

von Gaudecker, Hans-Martin; van Soest, Arthur; Wengström, Erik

Working Paper

Heterogeneity in risky choice behaviour in a broad population

IZA Discussion Papers, No. 4022

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: von Gaudecker, Hans-Martin; van Soest, Arthur; Wengström, Erik (2009) : Heterogeneity in risky choice behaviour in a broad population, IZA Discussion Papers, No. 4022, Institute for the Study of Labor (IZA), Bonn, <https://nbn-resolving.de/urn:nbn:de:101:1-20090304791>

This Version is available at:

<https://hdl.handle.net/10419/35555>

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.

IZA DP No. 4022

Heterogeneity in Risky Choice Behaviour in a Broad Population

Hans-Martin von Gaudecker
Arthur van Soest
Erik Wengström

February 2009

Heterogeneity in Risky Choice Behaviour in a Broad Population

Hans-Martin von Gaudecker

VU University Amsterdam

Arthur van Soest

Tilburg University and IZA

Erik Wengström

University of Copenhagen

Discussion Paper No. 4022

February 2009

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Heterogeneity in Risky Choice Behaviour in a Broad Population^{*}

We analyse risk preferences using an experiment with real incentives in a representative sample of 1,422 Dutch respondents. Our econometric model incorporates four structural parameters that vary with observed and unobserved characteristics: Utility curvature, loss aversion, preferences towards the timing of uncertainty resolution, and the propensity to choose randomly rather than on the basis of preferences. We find that all four parameters contribute to explaining choice behaviour. The structural parameters are significantly associated with socio-economic variables, but it is essential to incorporate unobserved heterogeneity in each of them to match the rich variety of choice patterns in the data.

JEL Classification: C90, D81

Keywords: risk aversion, loss aversion, uncertainty resolution, field experiments

Corresponding author:

Arthur van Soest
Tilburg University
P.O. Box 90153
5000 LE Tilburg
The Netherlands
E-mail: avas@uvt.nl

^{*} Financial support from the Dutch Science Foundation (NWO), the Swedish Institute for Banking Research (Bankforskningsinstitutet), and from the European Union under grant HPRNCT-2002-00235 (RTN-AGE) is gratefully acknowledged. We thank the team of CentERdata, especially Marika Puumala, for their support with the experiments, as well as Morten Lau and Joachim Winter for very helpful remarks on the experimental design. The analysis benefitted from comments by James Banks, Syngjoo Choi, Erik Sørensen, Peter Wakker, and seminar participants in Mannheim, Amsterdam, Rotterdam, Tilburg, Paris, at the 8th RTN-AGE meeting in Paris, the XIIth FUR conference at LUISS in Rome, the 2006 ESA meetings in Nottingham and Tucson, the 10th DNB Research Conference in Amsterdam, and the 3rd Nordic Conference on Experimental and Behavioural Economics in Copenhagen. The Online Appendix is available at <http://staff.econ.vu.nl/hgaudecker>

1 Introduction

We describe and analyse an experiment to elicit risk preferences in a representative sample from the Dutch population with more than 1,400 individuals. The main innovation of our work is that we estimate a structural model of the distribution of preferences in the population, distinguishing parameters for risk aversion, loss aversion, and a preference for the timing of uncertainty resolution. The rich nature of the data allows us to account for heterogeneity in all these parameters using a random coefficients model. In addition, we model heterogeneity in the tendency to make random optimization errors that explain why reported choices often are in conflict with fully rational decision making.

Analysing decisions under risk is an important theme in economic research. Economic theory has generated numerous models for it, starting with the early work of von Neumann & Morgenstern (1947). The class of models that combine linearity in probabilities with a non-linear utility function remains the workhorse in much of modern economics. Starting with Allais (1953), however, many violations of this basic model have been documented. Starmer (2000) reviews this literature. One particularly persistent finding in experiments is a greater sensitivity to losses than to gains of similar size. This was widely popularised through prospect theory (Kahneman & Tversky 1979) and defined formally as a kink of the utility function at the reference point by Köbberling & Wakker (2005). Heterogeneity in utility curvature has been documented in many studies. Important examples include Hey & Orme (1994), Dohmen, Falk, Huffman, Sunde, Schupp & Wagner (2005), or Andersen, Harrison, Lau & Rutström (2008). Nearly all studies that incorporate loss aversion concentrate on estimating mean or median parameters. Recent exceptions include Fehr & Götte (2007) and Gächter, Johnson & Herrmann (2007) who assume linear utility on each side of the reference point. A novel aspect of our analysis is the joint estimation of individual-specific parameters measuring the utility curvature and the strength of the kink in a structural model.

A difference between many real life economic decisions and most economic experiments is that the timings of the decision and of the resolution of uncertainty are different. For example, this is a characteristic of many asset allocation or insurance decisions. The timing of uncertainty resolution may matter for two reasons: A planning advantage of early resolution (Kreps & Porteus 1978) and anticipatory feelings such as hope and anxiety in case of late resolution (Wu (1999), Caplin & Leahy (2001)). Preferences for the timing of uncertainty resolution have received considerable empirical attention in the last decade (Chew & Ho (1994), Ahlbrecht & Weber (1996), Lovallo & Kahneman (2000), Noussair & Wu (2006), Eliaz & Schotter (2007), van Winden, Krawczyk & Hopfensitz (2008)). This literature has concentrated on estimating average parameters and produced mixed results.

In our experiment, all payoffs take place three months after the experiment. The timing of uncertainty resolution varies within subjects between the beginning and the end of this period.

This allows us to identify, at the individual level, preferences for the timing of uncertainty resolution. We estimate a specification of the Kreps & Porteus (1978) model that explicitly distinguishes parameters for utility curvature, loss aversion, and preferences for the timing of uncertainty resolution.

Most theories of utility are deterministic in nature. Provided knowledge of the relevant parameters, they imply a unique choice from any set of options except for very special cases of indifference. Wherever repeated or sufficiently similar choices allow a violation of revealed preference conditions, a significant fraction of subjects will violate them. Prominent studies demonstrating this fact include Hey & Orme (1994) and Choi, Fisman, Gale & Kariv (2007). Explanations for this include changing tastes, lack of motivation, and difficulties of understanding the choice problem. Empirical models have therefore added noise to the process that generates observed choices. There has been considerable progress on modeling this noise. Key contributions include Harless & Camerer (1994), Hey & Orme (1994), Loomes & Sugden (1995), Ballinger & Wilcox (1997), Loomes, Moffatt & Sugden (2002), and the papers in a special issue of *Experimental Economics* (Starmer & Bardsley 2005).

We build on this literature and estimate individual-specific error propensities along with the preference parameters. This allows us to assess how much confidence we can have in the estimates of individuals' preference parameters. Put differently, our estimates yield signal-to-noise ratios for any observed set of choices.

Recently, there has been an increased interest in using economic experiments to draw inference on the distribution of economically important preference parameters in a broad heterogeneous population (Bleichrodt, Pinto & Wakker (2001), Harrison, Lau & Williams (2002), Andersen et al. (2008), among many others). There has been a growing concern that the standard recruitment procedure – an experimenter inviting college students via emails or posters – restricts socio-demographic variation too severely to allow for meaningful inference on the broad population of interest.

We conduct our experiment on a heterogeneous subject pool that is representative of the Dutch population. Moreover, we model heterogeneity in all preference and error parameters, as a function of observed as well as unobserved characteristics. Our experiment is based on the CentERpanel, an Internet-based household survey in the Netherlands. It has been used before for related questions: Donkers, Melenberg & van Soest (2001) and Booij & van de Kuilen (2008) use hypothetical questions to estimate risk preference functionals and analyze their associations with observed covariates. Huck & Müller (2007) provide evidence that violations of Expected Utility Theory in the Allais paradox occur more frequently in some population subgroups than in others. We also briefly compare our experimental results with a parallel laboratory experiment with students from Tilburg University.

Our main results can be summarised as follows. First, utility curvature and loss aversion

turn out to be the key determinants of individuals' choices under risk. Second, the influence of preferences towards the resolution of uncertainty is less important and the median subject is almost neutral between early and late resolution. However, this is not true for everyone and the large preference heterogeneity that we document is likely to stand behind the mixed evidence accumulated so far. Third, while many people exhibit consistent choice patterns, some have very high error propensities. Errors are much more prevalent for the general population than for the students population represented in the laboratory. Fourth, parameters for utility curvature and loss aversion vary with socio-demographics, in a way that implies that students are not representative for the broader population. Fifth, the variation in preferences induced by our rich set of socio-demographic variables is small compared to the variance ascribed to unobserved heterogeneity. This implies that controls for individual preferences contain useful information in addition to socio-demographics and that it would be desirable to make them more widely available for empirical work based upon micro-data from socio-economic surveys.

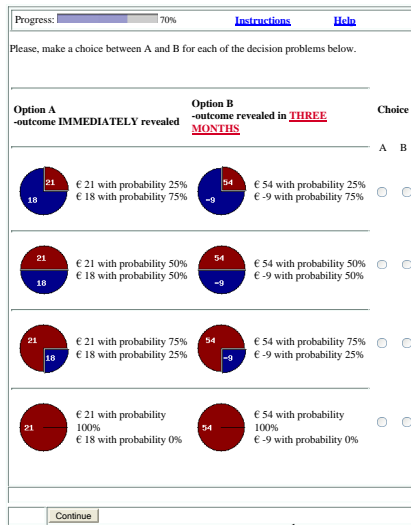
The remainder of this paper is organised as follows. In the next section, we describe the experimental design and motivate the subsequent analysis by means of aggregate features of the data and by describing the choice behaviour of five “benchmark” subjects, highlighting some common and some extreme examples of choice behaviour. In Section 3 we describe the theoretical and empirical models. Section 4 contains our results, describing average preference parameters and their socio-demographic correlates, the nature of unobserved heterogeneity, and the implications of the choices of the five selected subjects for their preference parameters. Section 5 concludes.

2 Data and Experimental Setup

We implemented the experiment in the CentERpanel, a Dutch household survey that is administered via the Internet. In order to avoid selection problems due to lack of Internet access, respondents without a computer are equipped with a set-top box for their television set (and with a TV if they do not have one). The panel consists of roughly 2,000 households who are representative of the Dutch population in terms of observable characteristics. It has rich background information on important demographic and socio-economic variables. Respondents are reimbursed regularly for their costs of using the Internet and we used the existing reimbursement system to make payments to the participants in the experiment.

We use data on 1,422 individuals in the CentERpanel and provide a description of the covariates and the construction of the sample used for estimation in the Online Appendix. Parallel experiments were conducted at Tilburg University's economics laboratory with a total of 178 student participants. While these are not the focus of this paper, we provide

Figure 1: Screenshot of Sheet 5, First Screen



some comparisons to results from these experiments.¹

2.1 Experimental Design

Our experiment uses an adapted version of the well-established multiple price list format, applied earlier by, for example, Binswanger (1980), Tversky & Kahneman (1992), Holt & Laury (2002), and Harrison, Lau & Rutström (2007). Andersen, Harrison, Lau & Rutström (2006) provide a detailed description and we limit ourselves to a brief introduction. Each subject is shown four pairs of lotteries such as the ones presented in Figure 1. We call the components of these pairs option ‘A’ and option ‘B’, where option ‘B’ always involves at least as much risk as option ‘A’. Subjects may opt for either option in each of the four choice tasks. The payoffs of both options do not change, but the probabilities of the high payoff in each option vary from 25 % to 100 % as one moves down the screen. The table is designed such that the expected value of option ‘A’ starts out higher but moves up slower than the expected value of option ‘B’.

Utility-maximising participants switch at some point from option ‘A’ to option ‘B’ or choose option ‘B’ throughout the screen. If such consistent behaviour is observed, the subject is routed to a screen containing lotteries with the same payoffs, but a finer probability grid. Andersen et al. (2006) recommend using this method and call it “iterated multiple price list”, others have used the term “chained method” for similar strategies (Wakker & Deneffe 1996). The grid now consists of steps of 10 percentage points located roughly between the subject’s highest choice of ‘A’ and his lowest choice of ‘B’ on the first screen. For a set of lotteries that differ only in the probabilities we use the term “payoff configuration”. All subjects were given

¹von Gaudecker, van Soest & Wengström (2008) perform an extensive analysis comparing the laboratory and the Internet experiments.

Table 1: Characteristics of the Seven Payoff Configurations

Payoff Configuration	Uncertainty Resolution, A	Payoff Low, A	Payoff High, A	Uncertainty Resolution, B	Payoff Low, B	Payoff High, B
1	early	27	33	early	0	69
2	early	39	48	early	9	87
3	early	12	15	early	-15	48
4	early	33	36	late	6	69
5	early	18	21	late	-9	54
6	early	24	27	early	-3	60
7	late	15	18	late	-12	51

Note: These values were shown in the high incentive and hypothetical treatments. For the low incentive treatment they were divided by three. The order was randomised.

the seven payoff configurations listed in Table 1. For each configuration, subjects made either eight or four decisions, depending on whether their answers on the first screen were consistent or not. Our data therefore constitute an unbalanced panel of 28 to 56 binary choices for each respondent.

A modification compared to previous studies is the inclusion of pie-charts as a graphical tool to help describing the probabilities of the outcomes. Pilot experiments showed that this supplement to the verbal descriptions of the decision tasks was appreciated by subjects who were not familiar with probability judgements. Moreover, we restricted the number of decision tasks per screen from the usual 10 to 4 to avoid the need for scrolling. Subjects could not go back to revise decisions on previous screens. We tested the instructions thoroughly. Unlike in typical laboratory settings, there was no experimenter assisting the respondents and answering questions they might have. In order to compensate for this, subjects had access to the instructions and specially designed help screens throughout the experiment.

Participants were allocated to one of three different incentive treatments. One of these was entirely hypothetical, with lotteries based upon the payoff configurations in Table 1. The two treatments with real incentives involved a participation fee paid to everyone who completed the experiment. Additionally, for one in every ten participants in these two treatments, one lottery was randomly selected and played out, and the payoff of that lottery was paid out.^{2,3}

²Andersen, Harrison, Lau & Rutström (forthcoming) find no significant difference between choices made in the case if one in ten participants is paid or if all participants are paid, although this may be different for other experimental designs (Baltussen, Post, van den Assem & Wakker 2008).

³The iteration introduces a distortion of incentive compatibility for some subjects. Sufficiently risk-tolerant participants (i.e. those whose preferences imply a switch point at or below 75%) had an incentive to switch

In the high incentive treatment, the completion fee was 15 Euros, with lottery payoffs listed in Table 1. In the low incentive treatment, payoffs and participation fee were one third of those of the high incentive treatment. We allocated subjects to one of two randomly determined orderings of the seven sheets.

Some of the lotteries included negative outcomes and we use a zero lottery payoff as the natural reference point. We avoided negative overall payoffs by setting the participation fees equal to the maximum losses that could be incurred. It should be noted that this has the potential drawback that loss aversion might be underestimated, since some respondents may combine payoff and participation fee and use a different reference point so that they never can experience a loss. Two facts suggest that this not so important. First, the high degree of narrow framing documented by Rabin & Weizsäcker (forthcoming) in a related experiment implies that most subjects will evaluate each screen separately. Second, we still find strong evidence of loss aversion and the less risky option ‘A’ always has weakly positive outcomes, so that subjects were able to avoid losses altogether.

All payoffs were made three months after the experiment. The timing of uncertainty resolution was set to either directly after the experiment or just before the payment was made. We emphasized to subjects that this concerned only the timing of uncertainty resolution and not the timing of payment: an entire screen of the introduction was devoted to this aspect and the timing of uncertainty resolution featured prominently at the top of each payoff configuration (see Figure 1).

2.2 Heterogeneity in the Choice Data


In this section we highlight the features of our experimental data that motivate the specification of our empirical model. We first present some descriptive statistics of the complete data set. Then we describe the choices of five exemplary individuals to show the need of incorporating heterogeneity and the possibility to make inconsistent decisions in the econometric model.

A convenient way to summarise the experimental choices are “switch points”, i.e. the probability at which a subject switches from choosing the safer option ‘A’ to the more risky option ‘B’ (see, for example, Holt & Laury (2002)). Figure 2 contains the average of these “switch points” for each of the seven payoff configurations, defined as the highest probability corresponding to an ‘A’ choice that is lower than the minimum probability at which ‘B’ is chosen.⁴ Following Holt & Laury (2002), all payoff configurations are designed in such a

at a higher probability for the high payoff because this would raise expected payoffs in the second step. This does not appear to be a major problem as the iteration procedure was not announced and we do not observe learning in the sense of increasing switch points over the course of the experiment.

⁴This is one of several ways to handle monotonicity violations. Some alternatives are explored in von Gaudecker et al. (2008). They lead to the same ranking of payoff configurations.

Figure 2: Mean “Switch Point”, by Payoff Configuration



Note: The numbering of the payoff configurations (PC) conforms to those in Table 1. “Switch points” are defined as the highest probability corresponding to an ‘A’ choice that is still lower than the minimum probability with a choice of ‘B’. Alternative ways to handle monotonicity violations lead to the same ranking of payoff configurations.

way that “switch points” would all be 40% for individuals who are risk neutral, have no loss aversion, are indifferent to the timing of uncertainty resolution, and make no errors. The much higher values are a clear indication that the average participant is quite different from this risk neutral benchmark respondent. Average “switch points” are typically higher if option ‘B’ involves the possibility of a loss (payoff configurations 3 and 5-7). The timing of uncertainty resolution for option ‘B’ (after three months in payoff configurations 4, 5, and 7, and immediate otherwise) does not influence average choices in a clear way. In particular, configurations 4 and 5 where the uncertainty is revealed early in option ‘A’ and late in option ‘B’ have the lowest and highest average switch points.


Optimisation errors in the sense of dominance violations (choosing ‘A’ at 100% chance of the high outcome or choosing ‘B’ at 0%) or monotonicity violations (switching back and forth) within a payoff configuration are found in 34.7% of the payoff configurations (where a maximum of one violation is counted per payoff configuration).⁵

The empirical model will take these findings into account, but in addition, the averages do not show the enormous amount of heterogeneity in individual choices. This is why we also describe the choices of five selected individuals, labelled respondents R1–R5. We will also come back to these five examples later, to analyze what our model estimates and their choices imply concerning their preference parameters and their tendencies to make mistakes.

The pattern of R1, described in Panel A of Figure 3 bears some similarity to the average figures. We see risk aversion for all configurations. The most risky choices are made in payoff configurations 1 and 2, characterised by early resolution and nonnegative payoffs. More safe choices were made in configurations 3 (with a probability of a negative payoff in option ‘B’ and

⁵The dominance violations are described in detail by von Gaudecker et al. (2008), who also show that the high numbers are due to the composition of the sample (with more low educated and older people than the usual lab experiment) rather than the Internet environment.

Figure 3: Choices Made by Selected Individuals



early resolution of uncertainty in both options) and 4 (late resolution of option ‘B’, positive outcomes only). R1 makes inconsistent decisions in configurations 5 and 7 (nonmonotonicities) and 6 (different choices when faced with the same decision twice). The latter could in principle also reflect indifference between the two options and is therefore not necessarily an inconsistency. However, preference parameters will have continuous distributions in our econometric model and the probability of drawing a respondent who is indifferent between two options is always equal to zero.

The other four benchmark respondents have choice patterns that are less common but not exceptional. The choices of R2 (Figure 3B) are very consistent and mostly safe – in all cases where the outcome set included a nonzero probability of a negative or zero payoff, R2 avoided this possibility by choosing the safer lottery. R3 (Figure 3C) makes the same choices as R2 for payoff configurations 5 and 7, but the choices on the other screens suggest that these are generated from different underlying preferences: In the other configurations the choices of R3 are much riskier than those of R2. This suggests that R3’s safe choices in these configurations

are driven by a combination of moderate risk aversion, loss aversion, and a preference for early resolution of uncertainty.

R4 in Figure 3D shows the opposite to R2: this subject exhibits risk-loving behaviour on all screens. The substantial number of such patterns in the data justifies a specification of heterogeneous preferences that can accommodate this kind of behaviour. Finally, R5 in Figure 3E makes seemingly idiosyncratic choices without any clear pattern. Although few individuals behave like this, there are many intermediate cases with a pattern suggesting an error tendency between that of subjects R2 and R5. The wide range of choice patterns with inconsistencies makes careful modelling of heterogeneity as well as errors important (cf., e.g., Moffatt (2005) for a similar motivation). At the end of the day, the estimates will tell us about the informational content of observed choices for individual preferences for every respondent in our sample. In contrast to other studies, our data suggests no clear cut-off point to exclude some individuals a priori from the sample (see, e.g., Abdellaoui, Barrios & Wakker (2007) or Choi et al. (2007)).

3 Theoretical Framework and Empirical Model

In this section, we first lay out the utility specifications that form the basis of our econometric analysis. We start with an expected utility of income specification, which incorporates loss aversion. In Section 3.2, we introduce a temporal component in a two-period model. While all payments are made in the second period, uncertainty may be resolved either in the first or the second period. We use a parsimonious version of the Kreps & Porteus (1978) model to allow for preferences towards the timing of uncertainty resolution. The last step consists in developing an empirical model that allows for sufficient heterogeneity (Section 3.3).

3.1 A Simple Model of Choice Under Risk

We start from a standard expected utility formulation with an exponential utility function:

$$u(z, \gamma) = -\frac{1}{\gamma}e^{-\gamma z} \quad (1)$$

where $z \in \mathbb{R}$ denotes a lottery outcome and $\gamma \in \mathbb{R}$ is the Arrow-Pratt coefficient of absolute risk aversion. We prefer exponential utility over power utility because it lends itself better to incorporating the timing of uncertainty resolution, as discussed below. In the Online Appendix, we compare the results to alternative specifications using power utility and find a better fit for exponential utility.

The first extension of (1) is to incorporate loss aversion, following prospect theory (Kahneman & Tversky 1979) and in line with the widely recognised stylised fact that “losses loom larger than gains” (see, e.g., Starmer (2000) for a review). In line with the literature, we augment

(1) with a loss aversion parameter $\lambda \in \mathbb{R}_+$:

$$u(z, \gamma, \lambda) = \begin{cases} -\frac{1}{\gamma}e^{-\gamma z} & \text{for } z \geq 0 \\ \frac{\lambda-1}{\gamma} - \frac{\lambda}{\gamma}e^{-\gamma z} & \text{for } z < 0 \end{cases} \quad (2)$$

The degree of loss aversion is measured by the ratio of the left and the right derivatives of the utility function at zero, as suggested by Köbberling & Wakker (2005). Prospect theory's original utility function is concave for gains and convex for losses. In contrast to this, (2) assumes the same type of curvature on the whole real line. This is primarily motivated by some recent empirical results that call prospect theory's original utility curvature findings for mixed gambles into question (Baltussen, Post & van Vliet 2006). Ideally, we would estimate separate parameters for the gains and loss domains, but our experiment does not have enough variation in negative outcomes to do this. In the Online Appendix, we present estimates based on alternative functional form assumptions that include prospect theory type preferences. The results suggest that assuming the same type of curvature on the gains and loss domains is preferred to the original prospect theory specification. The second reason behind using (2) is that this specification can much easier be built into a Kreps-Porteus environment capturing preferences for early or late uncertainty resolution.

3.2 Preferences towards the Timing of Uncertainty Resolution

In order to model preferences towards the timing of uncertainty resolution, we adopt the general framework of Kreps & Porteus (1978). In line with our experimental setup we consider a two-period setting. All decisions are made in the first period and payments made in the second period. The outcome of a gamble is either revealed in period 1, directly after all choices have been made (early resolution), or at the time of the payments in period 2 (late resolution).⁶

Assume that agents first calculate period 2 utility for all outcomes based on a function $v(z, \cdot)$, where z is the payoff and the dot replaces the preference parameters. Thereafter, agents are assumed to use a continuous and strictly increasing weighting function $h(\cdot)$ to calculate their first period utility, with period 2 utility v as its argument. The period 1 utility of a degenerate lottery that gives a certain outcome in period 2 is then given by $h(v(z, \cdot))$. The evaluation of nondegenerate lotteries hinges on the timing of uncertainty resolution: let

⁶Note that our motivation for modelling uncertainty resolution timing preferences is based on anticipatory utility. Caplin & Leahy (2001) argue that this concept is not captured well by the Kreps-Porteus model in settings where decisions take place in both periods because its temporal consistency axiom is frequently violated. In our setup without decisions in period 2, however, it is general and provides an attractive way of incorporating static and dynamic lottery characteristics.

$V(\pi)$ denote the period 1 utility of a lottery π with payoffs in period 2. Then V is given by:

$$V(\pi) = \begin{cases} \mathbb{E}[h(v(z, \cdot))] & \text{for early uncertainty resolution} \\ h(\mathbb{E}[v(z, \cdot)]) & \text{for late uncertainty resolution} \end{cases} \quad (3)$$

Note that the expectations operator is always applied to the quantity that is known at the end of period one. If uncertainty resolves early, the decision-maker applies the weighting function to the utility of the specific outcomes of π . If the outcome of π remains uncertain until the second period, he applies the weighting function to its expected value. Kreps & Porteus (1978) show that h is strictly convex (concave) if and only if the decision maker always prefers early to late (late to early) resolution, and is linear if the decision maker is indifferent. We choose the following parsimonious “power function” specification of the weighting function h :

$$h(v(z, \cdot)) = -S(-Sv(z, \cdot))^{\rho^{-S}} \quad (4)$$

with $\rho \in \mathbb{R}_+$ and where S is the sign operator given by:

$$S = \begin{cases} 1 & \text{for } \gamma \geq 0 \\ -1 & \text{for } \gamma < 0. \end{cases} \quad (5)$$

For $\rho > 1$, $h(\cdot)$ is convex and early resolution is preferred to late resolution. Indifference is obtained for $\rho = 1$, and late resolution is preferred for $\rho < 1$.

We model the second period utility function as a slightly modified version of (2):

$$v(z, \gamma, \lambda, \rho) = \begin{cases} \max\{-\frac{\lambda}{\gamma}, 0\} - \frac{1}{\gamma}e^{-\gamma\rho^S z} & \text{for } z \geq 0 \\ \max\{-\frac{\lambda}{\gamma}, 0\} + \frac{\lambda-1}{\gamma} - \frac{\lambda}{\gamma}e^{-\gamma\rho^S z} & \text{for } z < 0 \end{cases} \quad (6)$$

The building blocks of (4) and (6) seem complicated because of the necessity to accommodate both types of utility curvature. The term $-\frac{\lambda}{\gamma}$ is added for risk lovers to assure that the weighting function $h(\cdot)$ can be applied, i.e. it guarantees that $v(z, \gamma, \lambda, \rho)$ is always greater than zero for $\gamma < 0$. Including ρ^S in the exponent serves to retain the interpretation of $\gamma > 0$ as the coefficient of absolute risk aversion for early resolving lotteries on the positive domain. For such lotteries $V(\pi)$ collapses to $\mathbb{E}[u(\pi)]$ given in (2) if the subject is not risk loving. This implies that the distinction between risk aversion and uncertainty resolution timing preferences is identified for risk averse subjects if there are gambles on the positive domain and the timing of uncertainty resolution is varied.⁷

⁷This distinction is only approximately true for gambles with negative outcomes because of the additive term $\frac{\lambda-1}{\gamma}$ in (6). For risk lovers, the inclusion of $-\frac{\lambda}{\gamma}$ distorts the interpretation of γ by the same token. For the parameter values that we estimate, the magnitudes of the distortions are small. The Köbberling & Wakker (2005) definition of loss aversion remains valid for period 2 utility.

3.3 Econometric Implementation

Based on this specification of utility, we formulate structural econometric models that can be estimated by maximum likelihood. The models allow for individual heterogeneity in preference parameters and in the tendency to make errors. The heterogeneity can be captured by observed characteristics (“observed heterogeneity”) or not (“unobserved heterogeneity”). Assume that individual $i \in \{1, \dots, N\}$ faces $j \in \{1, \dots, J\}$ dichotomous choices between two binary lotteries $\pi_j^A = (A_j^{low}, A_j^{high}, A_j^{late}, p_j^{high})$ and $\pi_j^B = (B_j^{low}, B_j^{high}, B_j^{late}, p_j^{high})$. Each lottery is characterised by a low and a high outcome and whether uncertainty resolves late or not. Each pair of lotteries shares a common probability of the high outcome. Let $Y_{ij} = 1$ if the individual opts for π_j^B and $Y_{ij} = 0$ otherwise. Define the difference in certainty equivalents of the two lotteries in decision task j as:

$$\Delta CE_{ij} = CE(\pi_j^B, \gamma_i, \lambda_i, \rho_i) - CE(\pi_j^A, \gamma_i, \lambda_i, \rho_i),$$

where $CE(\pi_j^k, \gamma_i, \lambda_i, \rho_i)$, $k = A, B$ is the period one certainty equivalent of lottery π_j^k given the utility function defined by (3), (4), and (6) with the individual-specific parameters γ_i , λ_i , and ρ_i . It is straightforward to derive an exact analytical expression for $CE(\cdot)$ under our functional form assumptions, see the Online Appendix.

A perfectly rational decision maker would choose π_j^B if and only if $\Delta CE_{ij} > 0$. As a first step to allow for stochastic decision making, we add so-called Fechner errors (see, e.g., Loomes (2005)) to the CE comparison and model the individual’s choice as:

$$Y_{ij} = \mathbb{I}\{\Delta CE_{ij} + \tau \varepsilon_{ij} > 0\}, \quad (7)$$

where $\mathbb{I}\{\cdot\}$ denotes the indicator function. We assume that the ε_{ij} are independent of each other and of the random coefficients driving the utility function, and follow a standard logistic distribution. The parameter $\tau \in \mathbb{R}_+$ governs the individual’s probability to make this type of “mistakes”; the probability of such a mistake falls with the absolute value of ΔCE .

The use of certainty equivalents in (7) leads to a meaningful interpretation of ΔCE in monetary terms. This means that τ has an intuitive interpretation. For example, if the difference in valuations of two lotteries (i.e. the “cost” of making an error relative to the individual utility function parameters) is $\Delta CE = 10$ Euros and $\tau = 4$, the probability to choose the higher-valued lottery is 0.92. For $\Delta CE = 1$ Euro, this probability is only 0.56. Using certainty equivalents facilitates comparisons across subjects – using utility differences directly as, for example, Hey & Orme (1994), implies that the scale of the errors depends on the preference parameters and makes comparisons between subjects more difficult.

In addition to adding the errors $\tau \varepsilon_{ij}$, we allow for the possibility that subjects choose at random in any given task, following Harless & Camerer (1994). The propensity to do so is governed by the individual specific “trembling hand” parameter $\omega_i \in [0, 1]$ and the probability

of the observed choice Y_{ij} of individual i in choice situation j , given all the individual specific parameters, is given by:

$$l_{ij}(\pi_j^A, \pi_j^B, Y_{ij}, \tau, \gamma_i, \lambda_i, \rho_i, \omega_i) = (1 - \omega_i) \Lambda \left((2Y_{ij} - 1) \frac{1}{\tau} \Delta \text{CE}_{ij}(\pi_j^A, \pi_j^B, \gamma_i, \lambda_i, \rho_i) \right) + \frac{\omega_i}{2}, \quad (8)$$

where $\Lambda(t) = (1 + e^{-t})^{-1}$ stands for the cumulative standard logistic distribution function.

For the sake of a parsimonious and easily interpretable model, we restrict τ to be the same for all individuals, while allowing subjects to vary in their probability to make random choices. Alternative error specifications are possible in principle (see the references in Section 1) but beyond the scope of the current paper. For example, one might argue that τ should also be individual specific, but in practice it appears to be difficult to estimate heterogeneity in τ and ω separately (although both are identified, in theory).

We use a random coefficients model in order to estimate the distribution of the individual-specific parameters γ_i , λ_i , ρ_i , and ω_i in the population. This is a natural way of incorporating observed and unobserved heterogeneity directly.⁸ It has a better econometric justification and is much easier to do in the case of multiple parameters and not very many choices for each respondent than the alternative of first estimating the parameters for each individual separately (Hey & Orme 1994) and then regressing the results on socio-demographic characteristics (Dohmen et al. 2005). In contrast to the finite mixture model of Harrison & Rutström (forthcoming), we use a continuous distribution of the parameters of interest. The reason is that finite mixture models have difficulties handling a large number of potential values for the parameters and a small set of values seems insufficient to explain the very heterogeneous choice behaviour illustrated in Section 2.2. Our modeling of unobserved heterogeneity is similar to that of Conte, Hey & Moffatt (forthcoming); a difference is that we include observed heterogeneity in addition to unobserved heterogeneity.

In order to work with a concise notation, define

$$\eta_i = g_\eta(X_i^\eta \beta^\eta + \xi_i^\eta), \quad \eta_i \in \{\gamma_i, \lambda_i, \rho_i, \omega_i\} \quad (9)$$

where η_i denotes one of the four individual specific parameters, X_i^η are $1 \times K^\eta$ vectors of regressors, β^η are $K^\eta \times 1$ parameter vectors, and ξ_i^η are the unobserved heterogeneity components of the parameters. The first element of each X_i^η contains 1. The functions $g_\eta(\cdot)$ are used to impose the theoretical restrictions on the individual specific parameters. For γ , this is just the identity function; for λ and ρ , it is the exponential function, guaranteeing that these parameters are positive. For ω , it is the logistic distribution function, guaranteeing that ω is always between 0 and 1. We write $g(X_i \beta + \xi_i)$ for the vector of these four functions.

⁸See Bellemare, Kröger & van Soest (2008) for a similar modelling approach

We assume that $\xi_i = (\xi_i^\gamma, \xi_i^\lambda, \xi_i^\rho, \xi_i^\omega)'$ follows a jointly normal distribution independent of the regressors. The regressor matrix contains a dummy for the hypothetical high incentives treatment to capture the potential effect of giving hypothetical versus real payoffs. Preliminary estimations showed that the difference between low and high incentive treatments is better captured by a multiplicative specification than by adding a low incentive dummy to X . For the low incentive treatment, we therefore multiply all slope coefficients as well as the standard deviations of the unobserved heterogeneity terms by the same parameter $\beta_{\text{low incentive}}^\eta$.⁹

Defining $\xi^* = (\Sigma')^{-1}\xi$, where $\Sigma'\Sigma$ is the covariance matrix of ξ , we can express the likelihood contribution of subject i as:

$$l_i = \int_{\mathbb{R}^4} \left[\prod_{j \in J_i} l_{ij}(\pi_j^A, \pi_j^B, Y_{ij}, \tau, g(X_i\beta + \xi^*)) \right] \phi(\xi^*) d\xi^* \quad (10)$$

where l_{ij} is the probability given in (9) and $\phi(\cdot)$ denotes the probability density of $N(0, I)$. The log likelihood is given by the sum of the logs of l_i over all respondents in the sample and can be maximised by standard methods to obtain the maximum likelihood estimates. The integral in equation (10) does not have an analytical solution and we approximate it using standard simulation techniques. In particular, we employ Halton sequences of length $R = 1000$ per individual (Train 2003). We employ the BFGS algorithm with numerical derivatives to maximise the likelihood function. The variance-covariance matrix of the parameter estimates is based on the outer product of gradients. Standard errors for transformed parameters are calculated using the delta method.

4 Results

We present our results in four stages. First, we show that our model's average parameter estimates can explain the stylised facts in Figure 2. Second, we describe the estimated population distributions of the structural parameters. Third, we investigate how much of the total heterogeneity is accounted for by observed factors. Fourth, we move to the individual level and illustrate how the choices of respondents R1-R5 translate into preference parameter estimates. Finally, we compare the results of the Internet experiment to those from the parallel laboratory experiments.

⁹The multiplicative specification was selected by first considering separate models for each incentive treatment. This issue would not arise if we were only interested in average parameters without modelling heterogeneity.

4.1 Aggregate Features of the Data

The first columns of Tables 2-5 in the Appendix contain the estimates for the four parameter vectors β^η , $\eta \in \{\gamma, \lambda, \rho, \omega\}$ for models that contain only an intercept and treatment variables. The parameter vector for γ shows a very precisely estimated concave utility function. It is essentially the same in the high incentive and the hypothetical treatments but substantially larger in the low incentive treatment. An increase in the estimate of an exponential utility coefficient when the payoff scale is reduced was also found by Holt & Laury (2002). Whereas they focused on finding a functional form that provides a good approximation to aggregate behaviour over varying stakes, we are mainly interested in heterogeneity across subjects at a given level of payoffs. As the comparisons to power utility formulations provided in the Online Appendix show, exponential utility is much better suited for this purpose than the commonly considered alternative of power utility.

For the remaining parameters, it is important to note that the values in Tables 3-5 are all on the original parameter scale – the constant terms are simply $g_\eta(\beta_1^\eta)$. The first element in each table therefore represents the median (preference or error) parameter in the high incentive treatment. For the hypothetical treatment, the tables show $g(\beta_1^\eta + \beta_{\text{hypothetical}}^\eta) - g(\beta_1^\eta)$, the (partial) effect of moving from the high incentive to the hypothetical treatment on the median parameter value. For the low incentive treatment that enters multiplicatively, we simply report $\beta_{\text{low incentive}}^\eta$.

The estimated median parameter for loss aversion is 2.38 in the high incentive treatment, well in line with previous estimates. It is larger in the hypothetical treatment, probably because of the show-up fee paid in the real treatments – some subjects may have (partly) taken this into account so that they do not see the negative payoff as a loss. The loss aversion coefficient is lower in the low incentive treatment. The median coefficients for ρ are precisely estimated and not significantly different from 1 in any treatment. Hence the median subject appears to be close to indifferent for the timing of uncertainty resolution.

The magnitudes of the preference parameters are difficult to interpret directly. We follow Choi et al. (2007) and report the risk premia $\text{RP}(\pi) = \mathbb{E}[\pi] - \text{CE}(\pi)$ of standardised gambles implied by certain parameter constellations. To remain in line with the range of our payoffs, we use the gambles $\pi^1 = (25, 65, .5)$ and $\pi^2 = (-15, 25, .5)$, with both early and late resolution of uncertainty. Since $\text{RP}(\pi) \in [-20, 20]$ for both lotteries, the risk premia are directly comparable and their difference illustrates the impact of the loss aversion coefficient. The magnitudes of γ and λ are substantial in terms of the implied risk premia: They amount to 5.97 Euros for π^1 and to 10.65 Euros for π^2 ; the impact of ρ is negligible.

Comparing risk premia for π^1 is also a way to compare our estimates to those found in the literature. Holt & Laury's (2002) power-expo function estimates based on a wide range of payoffs imply a risk premium of about 2.79 Euros. The CRRA interval of their median

subject’s choices in their 20x treatment (with payoffs in the same range as ours) leads to an interval between 1.90 Euros and 3.20 Euros. Choi et al.’s (2007) estimates for the median subject imply risk premia of 3.34 Euros or 5.50 Euros for the Gul-CRRA and Gul-CARA cases, respectively. In a representative sample of the Danish population, the estimates of Harrison et al. (2007) using a CRRA functional lead to risk premia of 3.15 Euros. Our own estimates of about six Euros are slightly higher than the Choi et al. (2007) estimates using an exponential type utility function and substantially higher than the other estimates that use power utility type functionals (or, in the case of Holt & Laury (2002), more general functionals that are close to power utility for the estimated parameter values). The higher risk premia for the exponential utility seem to be a general characteristic in these settings. We also see the same drop in risk premia as in Choi et al. (2007) when we estimate a prospect theory type model based upon a power utility function (see the Online Appendix, in particular Table 21). The risk premium for the median subject is then 3.28 Euros, closely in line with the findings in the other studies.

As shown in Table 5, the median random choice propensity (ω_i) is about 8.3%. This can be related to the estimate for τ shown in Table 6, which is about 4.1. Taken together, the parameters imply that if $\Delta CE = 10$ Euros, the probability to choose the higher-valued lottery is 0.88; if $\Delta CE = 1$ Euro, this probability is only 0.55. These results correspond to the substantial error rates found in other studies using non-student samples. See, for example, de Roos & Sarafidis (2006) or Huck & Müller (2007). Comparing lab and Internet data, von Gaudecker et al. (2008) conclude that the high error rate is indeed due to the composition of the sample rather than the Internet environment. The random choice probabilities are the same in the hypothetical treatment and in the low incentive treatment. They are higher in the low incentive treatment, but the parameter τ is much smaller there. The implication for the overall tendency to make a suboptimal choice therefore varies with the characteristics of the lotteries involved.

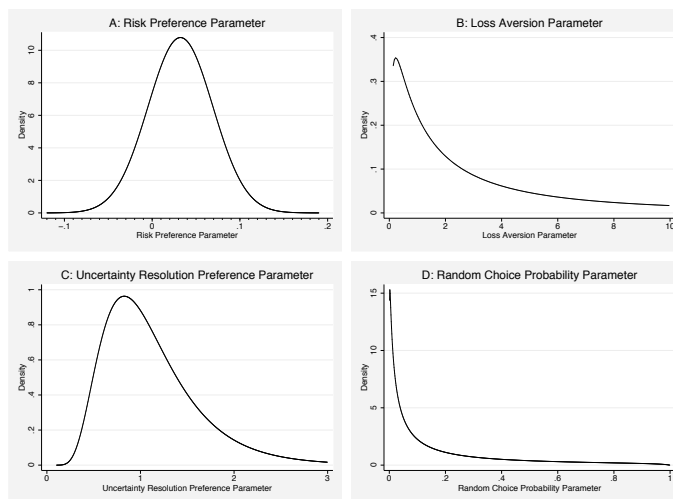
4.2 Distribution of Preferences and Errors

Figure 4 depicts the estimated population distributions of the four random parameters in our model for the high incentive treatment. The graphs are based upon the first columns of Tables 2-5, i.e. they account for observed as well as unobserved heterogeneity.¹⁰

Given the heterogeneity in individual choices that we showed in Section 2.2, the large dispersion in the four parameter distributions is hardly surprising. The large mass in the right tail of the distribution of λ implies that relying on a benchmark value between two

¹⁰Table 6 presents the underlying estimates for Σ . We only report the results imposing a diagonal structure on Σ ; Substantive results are very similar for a general variance-covariance matrix, but the estimates of the off-diagonal elements of Σ were very inaccurate; see the Online Appendix, Tables 8–11.

Figure 4: The Distributions of Preference and Error Parameters in the Population



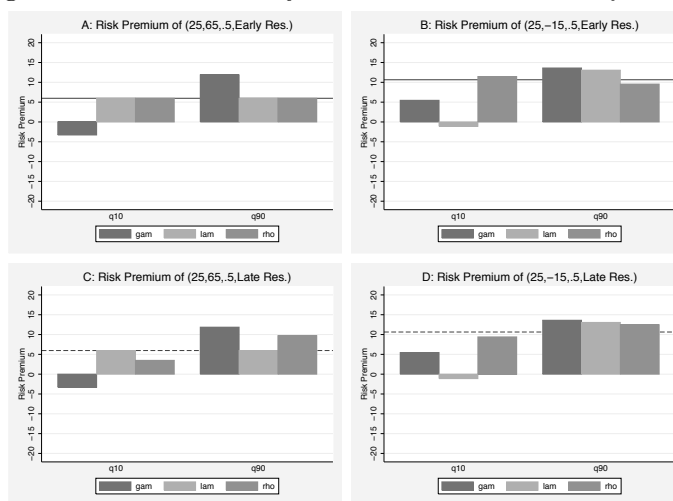
Note: Depicted are estimated parameter distributions taking observed and unobserved heterogeneity into account, based on the model with a minimal set of covariates (first columns of Tables 2–6). Treatment effects are netted out.

and three for the loss aversion parameter in model calibrations, something which seems to be common in the literature (see e.g. Benartzi & Thaler (1995)), is hard to justify. The estimated distribution of ω shows that most choices are driven by utility comparisons and therefore provide information on preferences. Still, a nonnegligible fraction of the populations often chooses at random.

Since the parameters themselves are difficult to interpret, we again consider the risk premia. Figure 5 plots the risk premia for $\pi^1 = (25, 65, .5)$ and $\pi^2 = (-15, 25, .5)$ at various quantiles of the parameter distributions, showing that the variation in preference parameters also induces substantial heterogeneity in terms of risk premia. The horizontal lines in Panels A-B plot the risk premia for the median preference parameters and for early resolution of uncertainty. The two sets of bars show what happens if each parameter separately is set to its 10th or 90th percentile.¹¹ For π^1 , the median subject demands a risk premium of about six Euros (see Panel A). Changing γ to its 10th percentile gives a negative risk premium of -3.30 Euros, while at the 90th percentile the risk premium is 11.90 Euros. The risk premium of π_1 does not depend on λ or ρ . The picture is quite different for $\text{RP}(\pi^2)$, as shown in Panel B. The baseline risk premium is now 10.65 Euros and the largest heterogeneity stems from the loss aversion parameter. The risk premium becomes negative when λ is set to its 10th percentile. Because of the already high median value of λ , the effect of increasing it to a higher value than the medium is much less. Finally, we see that due to the additive term in the definition of $v(z, \gamma, \lambda, \rho)$ in (6) for $z < 0$, there are slight changes in the risk premium

¹¹More detailed figures are given in Table 21 in the Online Appendix.

Figure 5: Risk Premia by Preference Parameter Quantiles



Note: Solid (dashed) lines depict the risk premia for the early (late) resolving lottery, evaluated at the median parameter estimates. The bars depict the risk premia when setting the parameters one at a time to their 10% and 90% quantiles.

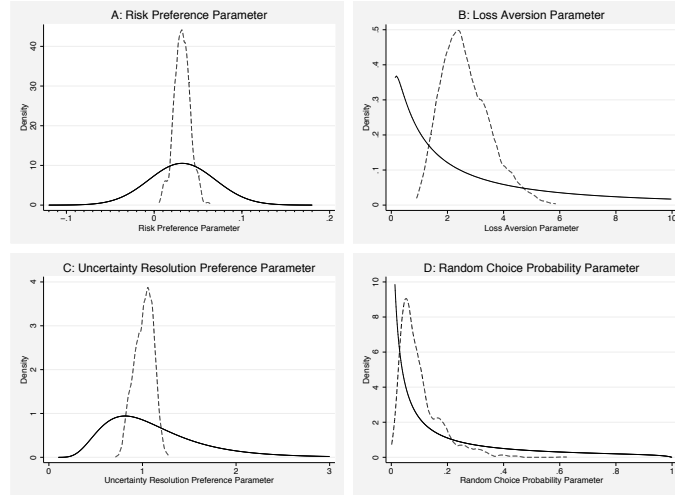
when ρ is changed, in spite of the fact that only early resolution of uncertainty is considered.

Panels C and D depict the case of late resolution. The median subject is almost indifferent with respect to the uncertainty resolution timing: risk premia are only 0.13 Euros lower than for early resolution. The only substantial difference with respect to the early resolution case concerns the impact of changing ρ : Moving it to the 10th or 90th percentile now has an impact of up to 3.70 Euros on the risk premia. This effect is substantial, but smaller than the effects of heterogeneity in other parameters found above.

4.3 Can Observed Factors Account for Parameter Heterogeneity?

An important question for empirical applications is the extent to which observed covariates can account for heterogeneity in structural parameters. In our case, we can phrase this question in terms of the reduction in variation of ξ_i^η if we control for typically available covariates in Equation (9). We can check this by comparing the distribution implied by the observed covariates only (setting ξ_i^η to zero) to the overall distribution. From Figure 6, it is evident that our large set of controls accounts only for a small part of the total variation in the preferences. For example, 90% of the conditional medians for the risk preference parameter γ lie in a range that accounts for less than 30% of the distribution of γ including the unobserved heterogeneity component. Only 1% of the conditional medians of ω imply a random choice propensity larger than .34, compared to 21% of the values of ω themselves. Furthermore, we note that for the overall distribution of parameters, it hardly makes a difference whether we account for observables or not. As Figure 13 in the Online Appendix demonstrates, the

Figure 6: The Distributions of Preference and Error Parameters in the Population: Comparing Total Heterogeneity and Observed Heterogeneity



Note: Solid lines are estimated parameter distributions taking observed and unobserved heterogeneity into account. Dashed lines neglect the unobserved part, they are kernel density estimates over the socio-demographic group means. Both are based on the model accounting for all covariates (second columns of Tables 2–6). Treatment effects are netted out.

two sets of graphs are virtually identical. The individual choices thus contain much more information than what is captured by socio-demographic groups.

Notwithstanding the small degree of overall heterogeneity that can be accounted for by observables, the dashed lines in Figure 6 also make clear that there are important differences between socio-economic groups. This is reflected in a large number of significant coefficients in the second columns of Tables 2-5. The associations broadly reflect those found in the literature and we just briefly highlight some of them.

Our results support the common finding that women are more risk averse than men (Croson & Gneezy forthcoming). We further find a positive age and a negative education gradient for risk aversion. The associations with income and wealth do not reveal a clear pattern. Being the household’s financial administrator is associated with lower risk aversion. These findings are in line with those of existing studies, cf. Donkers et al. (2001) or Dohmen et al. (2005); Harrison et al. (2007) find hardly any significant effects but this may be due to their moderate sample size; Benjamin, Brown & Shapiro (2006) find a negative association between risk aversion and cognitive skills, which is consistent with our findings for education. Loss aversion is higher among women and its age pattern peaks between 35 and 44 years of age and then declines. It also declines with household income, other variables are not significant. Some of the results for loss aversion contradict those of Gächter et al. (2007). Their experimental design, utility specifications, and sample selection procedure are very

different, however, so it is difficult to pin down a precise reason for the divergent results. The elderly are on average less averse to late uncertainty resolution than younger age groups.

The largest differences between socio-demographic groups are found for the error parameter. A plausible explanation is differences in numeracy, see Banks & Oldfield (2007). For low-educated persons older than 65 with low income and wealth, we estimate a random choice propensity of 40% at the median. Significantly fewer errors are made by the young and highly educated subjects. Point estimates for income and financial literacy are insignificant, but error rates decrease with wealth. We also included dummies for the time used to complete the experiment. They have the expected effect: those who completed the experiment rapidly have higher error rates and those who take a lot of time make fewer errors.

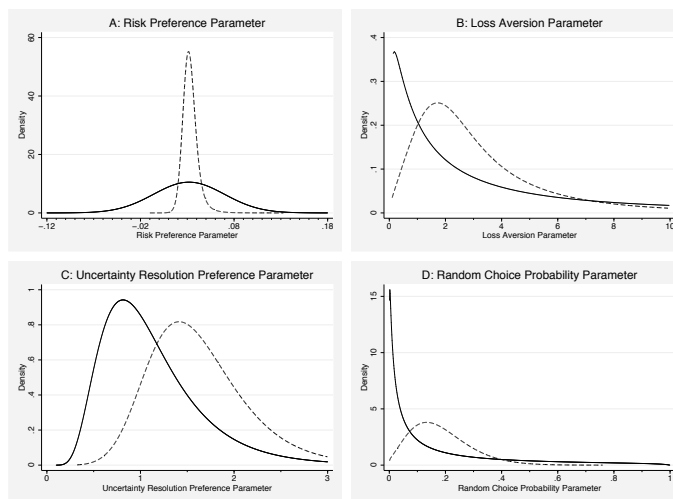
4.4 Choices and Preferences at the Individual Level

In this section we show how much information the choices in the experiment provide on the subjects' preference parameters. We do this for the five benchmark individuals described in Section 2.2. Our approach is similar to that of Revelt & Train (2000). The (“posterior”) distribution of preference parameters of each respondent is derived, conditional on observed choices of that respondent and given the estimated (“prior”) distribution of the preference parameters given individual characteristics X_i . Conditioning on the subjects' observed characteristics increases the prior distribution's chances to provide a good fit to the choices. Panels A-D of Figures 7-11 contain plots of both the prior distribution $F\left(\eta_i \mid X_i, \hat{\beta}, \hat{\Sigma}, g_\eta\right)$ and the posterior distribution $F\left(\eta_i \mid Y_i, X_i, \hat{\beta}, \hat{\Sigma}, g_\eta\right)$. The general picture that emerges is that the posterior estimates of γ_i are much more accurate than for the other parameters.

Consider Figure 7 with the posteriors for respondent R1. Her risk aversion parameter is very likely near the mean of the unconditional distribution – the 10% and 90% quantiles of the marginal posterior distribution imply risk premia for π^1 of 4.61 Euros and 7.67 Euros, respectively.¹² The same quantiles for the prior distribution (given the individual characteristics only, and not the observed choices in the experiment) imply risk premia of -3.14 Euros and 11.59 Euros. Hence knowing this subject's choices helps enormously to determine her risk aversion parameter. For the loss aversion parameter, the distance between the first and ninth deciles shrinks from more than 14 Euros in the prior to less than six Euros when conditioning on the subject's choices. Panel C shows that there is less than a 10% chance that individual R1 prefers early to late resolution while the corresponding prior probability for someone with her socio-economic characteristics is more than 40%. Finally, Panel D reveals that her random choice propensity is in the medium range. For all four parameters, conditioning on the choices makes the implied parameter ranges much tighter than conditioning on covariates

¹²The other parameter values are set to their posterior medians. Tables 22–25 in the Online Appendix contain the full details, including the risk premia for respondents R1-R4 at various quantiles of their distributions.

Figure 7: Choices and Preference Parameter Distributions of Respondent R1



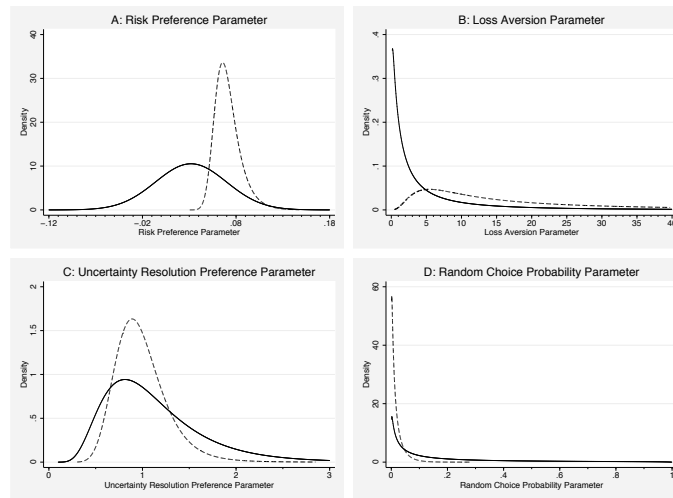
Note: Solid lines are the estimated parameter distributions for respondent R1’s socio-demographic group (female, age 45-54, higher secondary education or intermediate vocational training, household income between 22,000 Euros and 40,000 Euros, wealth above than 200,000 Euros, not financially knowledgeable, financial administrator, medium duration). Dashed lines are the marginal distributions of parameters conditional on the choices shown in the first panel. Graphs are based on estimates in Tables 2–6, second columns / panel.

only.

As expected from her choices, subject R2’s risk aversion parameter is in the upper range of the distribution. The risk premia for π^1 implied by the parameters at the posterior first and ninth deciles are 9.37 Euros and 12.54 Euros, respectively. The chance that her loss aversion parameter is less than the group median of 4.4 is slightly above 10%. The 90% quantile is 75 which makes sense since in all gambles with nonzero probability of a negative outcome this individual chose the safe option. The effect of this on the risk premia seems small – those for π^2 vary between 13.73 Euros and 14.47 Euros as λ moves from 4.3 to 75 – but it should be noted that even with infinite loss aversion, the risk premium will not become larger than 15 (eliminating the potential loss completely). The conditional distribution of ρ almost tracks the population distribution, implying that individual 2’s choices provide very little information on her value of ρ . Panel E shows that since her choices are highly consistent, her “trembling hand” error propensity ω is likely to be low: With 95% probability, it is below 7%.

Individual R3 (Figure 9) made moderately risk averse choices if payoffs were nonnegative and resolution of uncertainty was early. For late resolution and potentially negative payoffs, she never made a risky choice. The posterior distributions reflect this choice pattern: Her utility curvature parameter is below average with the first and ninth decile at .019 and .031 and corresponding risk premia for π^1 of 3.74 Euros and 5.87 Euros. The loss aversion parameter is large and more precisely estimated than for R2 since some of R3’s choices imply

Figure 8: Choices and Preference Parameter Distributions of Respondent R2




Note: Solid lines are the estimated parameter distributions for respondent R2's socio-demographic group (female, age 35-44, higher vocational training, household income between 22,000 Euros and 40,000 Euros, wealth above 200,000 Euros, not financially knowledgeable, not the financial administrator, medium duration). Dashed lines are the marginal distributions of parameters conditional on the choices shown in the first panel. Graphs are based on estimates in Tables 2–6, second columns / panel.

an upper bound on λ . The first and ninth deciles are at 6.59 and 28.22. The posterior median of ρ is 1.8 with the first and ninth decile at 1.34 and 2.44. This implies a preference for early resolution of uncertainty: the risk premium of 4.77 Euros for π^1 in the early resolution case rises to between 7.94 Euros if uncertainty resolves late. Even at the first decile of ρ , it is still 6.40 Euros.


Figure 10B shows that R4 is risk-loving ($\gamma < 0$). He is very likely to be neither loss averse nor does he prefer early resolution of uncertainty. All risk premia are substantially negative for this subject. He has a low tendency to make suboptimal choices. The posterior distributions of R4's parameters are quite narrow and directly reflect his choices. On the other extreme, individual R5's choices are hardly informative about her preferences (Figure 11). The marginal distributions conditioning on her choices are just as dispersed as the unconditional distributions. All we can say is that her behaviour is probably driven by a tendency to choose randomly instead of on the basis of the economic model – Panel D shows that her ω is probably close to 1.

Figure 9: Choices and Preference Parameter Distributions of Respondent R3



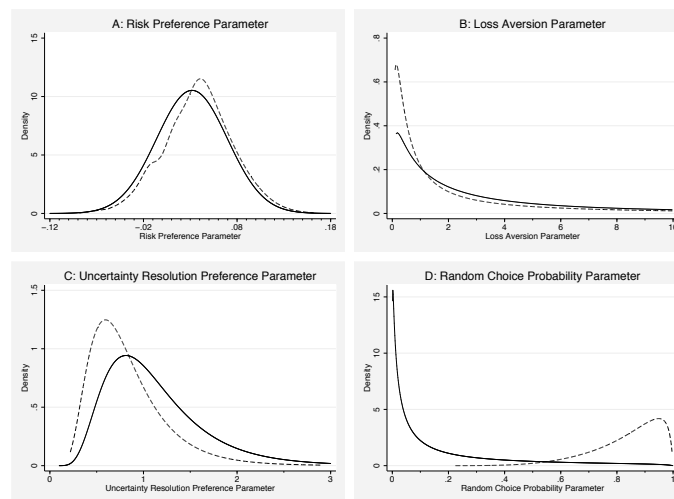
Note: Solid lines are the estimated parameter distributions for respondent R3's socio-demographic group (female, age 18-34, higher secondary education or intermediate vocational training, household income below 22,000 Euros, wealth between 10,000 Euros and 50,000 Euros, financially knowledgeable, financial administrator, medium duration). Dashed lines are the marginal distributions of parameters conditional on the choices shown in the first panel. Graphs are based on estimates in Tables 2–6, second columns / panel.

Figure 10: Choices and Preference Parameter Distributions of Respondent R4




Note: Solid lines are the estimated parameter distributions for respondent R4's socio-demographic group (male, age 45-54, primary or lower secondary education, household income above 40,000 Euros, wealth between 51,000 Euros and 200,000 Euros, not financially knowledgeable, financial administrator, short duration). Dashed lines are the marginal distributions of parameters conditional on the choices shown in the first panel. Graphs are based on estimates in Tables 2–6, second columns / panel.

Figure 11: Choices and Preference Parameter Distributions of Respondent R5



Note: Solid lines are the estimated parameter distributions for respondent R5's socio-demographic group (female, aged at least 65, primary or lower secondary education, household income below 22,000 Euros, wealth above 200,000 Euros, not financially knowledgeable, financial administrator, medium duration). Dashed lines are the marginal distributions of parameters conditional on the choices shown in the first panel. Graphs are based on estimates in Tables 2–6, second columns / panel.

Figure 12: Comparing the Distributions of the Preference and Error Parameters from the CentERpanel and Laboratory Experiments



Note: Both lines are estimated parameter distributions taking observed and unobserved heterogeneity into account. They are based on the models accounting for all available covariates (second and fourth columns of Tables 2–6). Treatment effects are netted out.

4.5 Comparing the Internet and Laboratory Experiments

Figure 12 illustrates the differences in the parameter distributions between the Internet and the laboratory samples. There are two salient differences: First, heterogeneity in the general population is much greater for all parameters except ρ . Second, the estimated median parameters imply substantially lower risk premia in the laboratory than in the Internet sample. In particular, they amount to about 3.50 Euros for π^1 (compared to 6 Euros in the Internet sample) and 7.50 Euros for π^2 (Internet: 10.65 Euros). These findings confirm that preference parameters found with student samples do not describe the preferences of the general population, as has been noted before by e.g. Andersen et al. (forthcoming) for utility curvature.

Considering the quantitative importance of the differences in heterogeneity, our estimates imply that 95% of the laboratory subjects have $\gamma \in [-.018, .049]$, whereas only slightly more than two thirds of Internet subjects have a γ in this range. For the loss aversion parameter, only 1% of subjects in the student sample have $\lambda > 10$, compared to more than 20% in the Internet sample. Finally, while 88% of the participants in the laboratory have a random choice probability below 10%, this fraction is only slightly more than one half for the Internet sample. This can be fully explained by the differences in the sample compositions: The error rates among the young and highly educated subjects in the CentERpanel are close to those estimated in the laboratory.

The laboratory based estimates of the structural parameters are presented in columns 3

and 4 of Tables 2–5. The lab based estimates for the medians of γ and λ are substantially lower than those based upon the Internet sample. We find no significant differences between Internet and lab estimates of uncertainty resolution timing preference. The median propensity to choose at random in the lab is close to zero, in line with the relatively small errors that are typically observed for student samples (Hey & Orme 1994). We find similar differences between hypothetical and real incentives treatments in the lab as over the Internet, except for the loss aversion coefficient, where the difference vanishes in the lab experiment. The fact that we paid subjects a show up fee in the hypothetical laboratory treatment but not in the hypothetical treatment Internet group might explain this – the participation fee is partly taken into account, even if the lotteries are purely hypothetical. The coefficients indicating differences between low and high incentive treatments for the lab experiment are largely in line with those for the Internet experiment. An exception is the large difference for ω : its estimate for the high incentive treatment in the laboratory is quite small.

5 Conclusions

We have described a large scale experiment on decision making under risk using a representative sample of a broad population. We have analyzed the experimental data using a structural empirical model, disentangling preference parameters for utility curvature, loss aversion, and preference for the timing of uncertainty resolution, and allowing for several types of errors.

Our model requires a number of specification choices. To save space, we have only presented the results of our preferred specification. In the Online Appendix, we further motivate our specification choices, showing that our model performs better than several alternatives. We also show that our substantive results are generally robust to the specification choices we have made. For example, not accounting for early or late resolution preferences (setting ρ equal to 1) or changing the assumptions on the utility curvature to prospect theory type preferences does not affect the estimates of the average risk premia for our benchmark lotteries very much. Using power utility rather than an exponential utility function leads to a drop similar to the one observed by Choi et al. (2007).

Our main finding is that risk preferences in the population are very heterogeneous, and only a small part of this heterogeneity can be captured with standard covariates such as age, gender, education, income, and wealth. The structural modelling approach combined with the rich data appears to be a useful tool to handle the heterogeneity. Our four main parameters of interest (the three preference parameters and the tendency to choose purely at random instead of on the basis of utility maximisation) are modelled as random coefficients, and we find substantial dispersion in all of them. For example, even though we find that the timing of uncertainty resolution does not matter much for the median respondent, our estimates

imply that there are groups in the population that clearly prefer early resolution and other groups that prefer late resolution. This heterogeneity may well explain the mixed evidence for preferences over temporal lotteries that has been accumulated so far in the literature.

Our structural model is particularly appropriate to analyse the informational content of each subject's choices for the parameters of that subject. This is shown by comparing the posterior distribution of some subjects given their choices to the prior distribution (given covariates only). We find that the choices are generally very informative about individual preference parameters, except, as expected, in cases where the choice data suggest that the subjects' choices are probably completely random. Importantly, the model is able to handle intermediate cases and does not require to exclude individuals from the sample a priori. Furthermore, the individual propensity to choose at random in experiments is likely to be informative about the quality of choices in other domains as well.

References

- Abdellaoui, Mohammed, Barrios, Carolina & Wakker, Peter P.** (2007). 'Reconciling Introspective Utility with Revealed Preference: Experimental Arguments Based on Prospect Theory', *Journal of Econometrics* 138(1), 356–378.
- Ahlbrecht, Martin & Weber, Martin.** (1996). 'The Resolution of Uncertainty: An Experimental Study', *Journal of Institutional and Theoretical Economics* 152, 593–607.
- Allais, Maurice.** (1953). 'Le Comportement De l'Homme Rationelle Devant Le Risque: Critique Des Postulats Et Axiomes De L'École Américaine', *Econometrica* 21, 593–607.
- Andersen, Steffen, Harrison, Glenn W., Lau, Morten I. & Rutström, E. Elisabet.** (2006). 'Elicitation Using Multiple Price List Formats', *Experimental Economics* 9(4), 383–405.
- Andersen, Steffen, Harrison, Glenn W., Lau, Morten I. & Rutström, E. Elisabet.** (2008). 'Eliciting Risk and Time Preferences', *Econometrica* 76(3), 583–618.
- Andersen, Steffen, Harrison, Glenn W., Lau, Morten I. & Rutström, E. Elisabet.** (forthcoming). 'Preference Heterogeneity in Experiments: Comparing the Field and Laboratory', *Journal of Economic Behavior & Organization* .
- Ballinger, T. Parker & Wilcox, Nathaniel T.** (1997). 'Decisions, Error and Heterogeneity', *Economic Journal* 107(443), 1090–1105.
- Baltussen, Guido, Post, Thierry, van den Assem, Martijn J. & Wakker, Peter P.** (2008). Random Task Incentive Systems in Risky Choice Experiments. Mimeo, Erasmus University Rotterdam.

- Baltussen, Guido, Post, Thierry & van Vliet, Pim.** (2006). ‘Violations of Cumulative Prospect Theory in Mixed Gambles with Moderate Probabilities’, *Management Science* 52(8), 1288–1290.
URL: <http://mansci.journal.informs.org/cgi/content/abstract/52/8/1288>
- Banks, James & Oldfield, Zoë.** (2007). ‘Understanding Pensions: Cognitive Function, Numerical Ability and Retirement Saving’, *Fiscal Studies* 28(2), 143–170.
- Bellemare, Charles, Kröger, Sabine & van Soest, Arthur.** (2008). ‘Measuring Inequity Aversion in a Heterogeneous Population Using Experimental Decisions and Subjective Probabilities’, *Econometrica* 76(4), 815–839.
- Benartzi, Shlomo & Thaler, Richard H.** (1995). ‘Myopic Loss Aversion and the Equity Premium Puzzle’, *Quarterly Journal of Economics* 110(1), 73–92.
- Benjamin, Daniel J., Brown, Sebastian A. & Shapiro, Jesse M.** (2006), Who is “Behavioral”? Cognitive Ability and Anomalous Preferences. Mimeo, Harvard University.
- Binswanger, Hans P.** (1980). ‘Attitudes Towards Risk: An Experimental Measurement in Rural India’, *American Journal of Agricultural Economics* 62, 395–407.
- Bleichrodt, Han, Pinto, José Luis & Wakker, Peter P.** (2001). ‘Making Descriptive Use of Prospect Theory to Improve the Prescriptive Use of Expected Utility’, *Management Science* 47(11), 1498–1514.
- Booij, Adam S. & van de Kuilen, Gijs.** (2008), A Parameter-Free Analysis of the Utility of Money for the General Population under Prospect Theory. University of Amsterdam, CREED Working Paper.
- Caplin, Andrew & Leahy, John.** (2001). ‘Psychological Expected Utility Theory and Anticipatory Feelings’, *Quarterly Journal of Economics* 116, 55–79.
- Chew, Soo Hong & Ho, Joanna L.** (1994). ‘Hope: An Empirical Study of Attitude toward the Timing of Uncertainty Resolution’, *Journal of Risk and Uncertainty* 8, 267–288.
- Choi, Syngjoo, Fisman, Raymond, Gale, Douglas & Kariv, Shachar.** (2007). ‘Consistency and Heterogeneity of Individual Behavior under Uncertainty’, *American Economic Review* 97(5), 1921–1938.
- Conte, Anna, Hey, John D. & Moffatt, Peter G.** (forthcoming). ‘Mixture Models of Choice under Risk’, *Journal of Econometrics* . forthcoming.
- Croson, Rachel & Gneezy, Uri.** (forthcoming). ‘Gender Differences in Preferences’, *Journal of Economic Literature* .

- de Roos, Nicolas & Sarafidis, Yianis.** (2006), Decision Making under Risk in Deal or No Deal. Mimeo, University of Sydney. Available at SSRN: <http://ssrn.com/abstract=881129>.
- Dohmen, Thomas, Falk, Armin, Huffman, David, Sunde, Uwe, Schupp, Jürgen & Wagner, Gert G.** (2005), Individual Risk Attitudes: New Evidence from a Large, Representative, Experimentally-Validated Survey. IZA Discussion Paper 1730.
- Donkers, Bas, Melenberg, Bertrand & van Soest, Arthur.** (2001). ‘Estimating Risk Attitudes Using Lotteries; a Large Sample Approach’, *Journal of Risk and Uncertainty* 22(2), 165–195.
- Eliaz, Kfir & Schotter, Andrew.** (2007). ‘Experimental Testing of Intrinsic Preferences for NonInstrumental Information’, *American Economic Review* 97(2), 166–169.
- Fehr, Ernst & Götte, Lorenz.** (2007). ‘Do Workers Work More If Wages Are High? Evidence from a Randomized Field Experiment’, *American Economic Review* 97(1), 298–317.
- Gächter, Simon, Johnson, Eric J. & Herrmann, Andreas.** (2007), Individual-Level Loss Aversion in Riskless and Risky Choices. IZA Discussion Paper 2961.
- Harless, David W. & Camerer, Colin F.** (1994). ‘The Predictive Utility of Generalized Expected Utility Theories’, *Econometrica* 62(6), 1251–1289.
- Harrison, Glenn W., Lau, Morten Igel & Rutström, E. Elisabet.** (2007). ‘Estimating Risk Attitudes in Denmark: A Field Experiment’, *Scandinavian Journal of Economics* 109(2), 341–368.
- Harrison, Glenn W., Lau, Morten Igel & Williams, Melonie B.** (2002). ‘Estimating Discount Rates in Denmark: A Field Experiment’, *American Economic Review* 92(5), 1606–1617.
- Harrison, Glenn W. & Rutström, E. Elisabet.** (forthcoming). ‘Expected Utility Theory and Prospect Theory: One Wedding and a Decent Funeral’, *Experimental Economics* .
- Hey, John D. & Orme, Chris.** (1994). ‘Investigating Generalizations of Expected Utility Theory Using Experimental Data’, *Econometrica* 62(6), 1291–1326.
- Holt, Charles A. & Laury, Susan K.** (2002). ‘Risk Aversion and Incentive Effects’, *American Economic Review* 92, 1644–1655.
- Huck, Steffen & Müller, Wieland.** (2007), Allais for All: Revisiting the Paradox. Mimeo, University of Tilburg. Available at: <http://center.uvt.nl/staff/muller/Allais071009.pdf>.

- Kahneman, Daniel & Tversky, Amos V.** (1979). ‘Prospect Theory: An Analysis of Decision under Risk’, *Econometrica* 47, 263–291.
- Köbberling, Veronika & Wakker, Peter P.** (2005). ‘An Index of Loss Aversion’, *Journal of Economic Theory* 122, 119–131.
- Kreps, David M. & Porteus, Evan L.** (1978). ‘Temporal Resolution of Uncertainty and Dynamic Choice Theory’, *Econometrica* 46, 185–200.
- Loomes, Graham.** (2005). ‘Modelling the Stochastic Component of Behaviour in Experiments: Some Issues for the Interpretation of Data’, *Experimental Economics* 8(4), 301–323.
- Loomes, Graham, Moffatt, Peter G. & Sugden, Robert.** (2002). ‘A Microeconomic Test of Alternative Stochastic Theories of Risky Choice’, *Journal of Risk and Uncertainty* 24(2), 103–130.
- Loomes, Graham & Sugden, Robert.** (1995). ‘Incorporating a Stochastic Element into Decision Theories’, *European Economic Review* 39, 641–648.
- Lovall, Dan & Kahneman, Daniel.** (2000). ‘Living with Uncertainty: Attractiveness and Resolution Timing’, *Journal of Behavioral Decision Making* 13(2), 179–190.
- Noussair, Charles & Wu, Ping.** (2006). ‘Risk Tolerance in the Present and the Future: An Experimental Study’, *Managerial and Decision Economics* 27(6), 401 – 412.
- Rabin, Matthew & Weizsäcker, Georg.** (forthcoming). ‘Narrow Bracketing and Dominated Choices’, *American Economic Review* .
- Revelt, David & Train, Kenneth.** (2000), Customer-Specific Taste Parameters and Mixed Logit: Households’ Choice of Electricity Supplier. University of California at Berkeley, Economics Working Paper E00-274.
- Starmer, Chris.** (2000). ‘Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk’, *Journal of Economic Literature* 38(2), 332–382.
- Starmer, Chris & Bardsley, Nicholas.** (2005). ‘Introduction to the Special Issue: Exploring the Error in Experimental Economics’, *Experimental Economics* 8(4), 295–299.
- Train, K.E.** (2003), *Discrete Choice Methods with Simulation*, Cambridge University Press.
- Tversky, Amos V. & Kahneman, Daniel.** (1992). ‘Advances in Prospect Theory: Cumulative Representation of Uncertainty’, *Journal of Risk and Uncertainty* 5(4), 297–323.

- van Winden, Frans, Krawczyk, Michal & Hopfensitz, Astrid.** (2008), Investment, Resolution of Risk, and the Role of Affect. University of Amsterdam, CREED Working Paper 2008.
- von Gaudecker, Hans-Martin, van Soest, Arthur & Wengström, Erik.** (2008), Selection and Mode Effects in Risk Preference Elicitation Experiments. IZA Discussion Paper 3321.
- von Neumann, John & Morgenstern, Oskar.** (1947), *The Theory of Games and Economic Behavior*, 2nd edn, Princeton University Press, Princeton.
- Wakker, Peter P. & Deneffe, Daniel.** (1996). ‘Eliciting von Neumann-Morgenstern Utilities when Probabilities Are Distorted or Unknown’, *Management Science* 42(8), 1131–1150.
- Wu, Geore.** (1999). ‘Anxiety and Decision Making with Delayed Resolution of Uncertainty’, *Theory and Decision* 46, 159–199.

A Tables with Estimation Results

Table 2: Estimated Risk Attitudes (γ)

Covariate	CentERpanel		Laboratory	
	(4),(6)	(4),(6)	(4),(6)	(4),(6)
Constant	0.0322*** (0.0010)	0.0317*** (0.0028)	0.0181*** (0.0022)	0.0126*** (0.0027)
Female		0.0080*** (0.0015)		0.0064* (0.0034)
Age 35-44		-0.0004 (0.0022)		
Age 45-54		0.0046** (0.0023)		
Age 55-64		0.0035 (0.0026)		
Age 65+		0.0150*** (0.0026)		
Hi Sec Educ / Int Voc Train		-0.0064*** (0.0018)		
Higher Voc Train		-0.0020 (0.0021)		
University		-0.0164*** (0.0026)		
Income EUR 22k-40k		-0.0013 (0.0016)		
Income EUR 40k+		0.0006 (0.0023)		
Wealth EUR 10k-50k		0.0088*** (0.0022)		
Wealth EUR 51k-200k		-0.0028 (0.0020)		
Wealth EUR 201k+		-0.0033 (0.0024)		
HH Financial Admin		-0.0034** (0.0017)		
Financially Knowledgeable		-0.0003 (0.0018)		
Hypothetical Treatment	-0.0004 (0.0015)	0.0008 (0.0016)	0.0013 (0.0033)	0.0031 (0.0036)
Low Incentive Treatment †	2.78*** (0.0869)	2.77*** (0.0873)	2.84*** (0.275)	2.84*** (0.253)

Note: Number of Observations is 1,422 (CentERpanel) and 178 (Laboratory), respectively. Estimation follows (10) based on the utility functions in the column headers. The left-out categories are: High incentive treatment, all payoffs non-negative, male, age 18-34, primary / lower secondary education, net annual household income below 22,000 Euros, total wealth below 10,000 Euros, not being the household's financial administrator, not being financially knowledgeable (self-rated).

† The low incentive treatment enters multiplicatively.

Table 3: Estimated Loss Aversion Parameters (λ)

Covariate	CentERpanel		Laboratory	
	(4),(6)	(4),(6)	(4),(6)	(4),(6)
Constant	2.38*** (0.159)	2.94*** (0.543)	1.67*** (0.243)	1.27 (0.182)
Female		0.617* (0.329)		0.766** (0.302)
Age 35-44		1.14** (0.541)		
Age 45-54		-0.752* (0.405)		
Age 55-64		-0.0463 (0.482)		
Age 65+		-0.673 (0.447)		
Hi Sec Educ / Int Voc Train		0.205 (0.391)		
Higher Voc Train		0.146 (0.435)		
University		0.663 (0.586)		
Income EUR 22k-40k		-0.378 (0.312)		
Income EUR 40k+		-1.15*** (0.388)		
Wealth EUR 10k-50k		0.193 (0.489)		
Wealth EUR 51k-200k		-0.0747 (0.367)		
Wealth EUR 201k+		-0.0697 (0.429)		
HH Financial Admin		-0.183 (0.299)		
Financially Knowledgeable		-0.418 (0.297)		
Hypothetical Treatment	1.31*** (0.317)	1.37*** (0.441)	0.0302 (0.382)	0.0418 (0.243)
Low Incentive Treatment †	0.861*** (0.0290)	0.839*** (0.0433)	0.862** (0.0702)	0.934 (0.0529)

Note: Number of Observations is 1,422 (CentERpanel) and 178 (Laboratory), respectively. Estimation follows (10) based on the utility functions in the column headers. Regression coefficients are transformed back to the original scale. In other words, the constant is defined by $g_\lambda(\beta_1^\lambda)$. The other values are partial effects of setting the dummy variables to one, given the reference value defined by the left-out categories. These categories are: High incentive treatment, all payoffs non-negative, male, age 18-34, primary / lower secondary education, net annual household income below 22,000 Euros, total wealth below 10,000 Euros, not being the household's financial administrator, not being financially knowledgeable (self-rated).

† The low incentive treatment enters multiplicatively.

Table 4: Estimated Uncertainty Resolution Preferences (ρ)

Covariate	CentERpanel		Laboratory	
	(4),(6)	(4),(6)	(4),(6)	(4),(6)
Constant	1.01 (0.0250)	1.04 (0.0849)	0.931 (0.0753)	0.891 (0.0812)
Female		0.0179 (0.0457)		0.321*** (0.0620)
Age 35-44		-0.0054 (0.0686)		
Age 45-54		0.0069 (0.0694)		
Age 55-64		-0.134* (0.0730)		
Age 65+		-0.207*** (0.0763)		
Hi Sec Educ / Int Voc Train		-0.0446 (0.0563)		
Higher Voc Train		-0.0903 (0.0620)		
University		-0.0100 (0.0811)		
Income EUR 22k-40k		0.0274 (0.0521)		
Income EUR 40k+		-0.0609 (0.0685)		
Wealth EUR 10k-50k		0.0895 (0.0744)		
Wealth EUR 51k-200k		0.0402 (0.0609)		
Wealth EUR 201k+		0.0444 (0.0742)		
HH Financial Admin		0.0588 (0.0517)		
Financially Knowledgeable		-0.0377 (0.0529)		
Hypothetical Treatment	-0.0964** (0.0405)	-0.109** (0.0460)	-0.177* (0.104)	-0.189* (0.0988)
Low Incentive Treatment [†]	0.999 (0.0015)	1.00 (0.0056)	1.00 (0.0167)	0.992 (0.0285)

Note: Number of Observations is 1,422 (CentERpanel) and 178 (Laboratory), respectively. Estimation follows (10) based on the utility functions in the column headers. Regression coefficients are transformed back to the original scale. In other words, the constant is defined by $g_\rho(\beta_1^\rho)$. The other values are partial effects of setting the dummy variables to one, given the reference value defined by the left-out categories. These categories are: High incentive treatment, all payoffs non-negative, male, age 18-34, primary / lower secondary education, net annual household income below 22,000 Euros, total wealth below 10,000 Euros, not being the household's financial administrator, not being financially knowledgeable (self-rated).

[†] The low incentive treatment enters multiplicatively.

Table 5: Estimated Random Choice Probabilities (ω)

Covariate	CentERpanel		Laboratory	
	(4),(6)	(4),(6)	(4),(6)	(4),(6)
Constant	0.0832*** (0.0083)	0.102*** (0.0216)	0.0049 (0.0053)	0.0027 (0.0038)
Female		0.0197* (0.0118)		0.0034 (0.0045)
Age 35-44		0.0201 (0.0182)		
Age 45-54		0.0479** (0.0203)		
Age 55-64		0.141*** (0.0326)		
Age 65+		0.293*** (0.0502)		
Hi Sec Educ / Int Voc Train		-0.0445*** (0.0133)		
Higher Voc Train		-0.0535*** (0.0140)		
University		-0.0622*** (0.0155)		
Income EUR 22k-40k		-0.0107 (0.0121)		
Income EUR 40k+		-0.0217 (0.0149)		
Wealth EUR 10k-50k		-0.0362** (0.0144)		
Wealth EUR 51k-200k		-0.0249** (0.0124)		
Wealth EUR 201k+		-0.0316** (0.0134)		
HH Financial Admin		0.0005 (0.0116)		
Financially Knowledgeable		-0.0136 (0.0115)		
Short Duration		0.0822*** (0.0207)		-0.0015 (0.0022)
Long Duration		-0.0485*** (0.0130)		0.0032 (0.0064)
Hypothetical Treatment	0.0070 (0.0106)	0.0049 (0.0127)	0.0043 (0.0054)	0.0017 (0.0036)
Low Incentive Treatment †	1.11*** (0.0078)	1.06*** (0.0082)	3.77*** (0.0150)	2.74*** (0.0147)

Note: Number of Observations is 1,422 (CentERpanel) and 178 (Laboratory), respectively. Estimation follows (10) based on the utility functions in the column headers. Regression coefficients are transformed back to the original scale. In other words, the constant is defined by $g_\omega(\beta_1^\omega)$. The other values are partial effects of setting the dummy variables to one, given the reference value defined by the left-out categories. These categories are: High incentive treatment, all payoffs non-negative, male, age 18-34, primary / lower secondary education, net annual household income below 22,000 Euros, total wealth below 10,000 Euros, not being the household's financial administrator, not being financially knowledgeable (self-rated).

† The low incentive treatment enters multiplicatively.

Table 6: Standard Deviations of the Random Coefficients, τ , and Log-Likelihoods

	CentERpanel		Lab	
	Few Covariates	Full Covariates	Few Covariates	Full Covariates
σ_γ	0.037 (0.001)	0.037 (0.001)	0.020 (0.001)	0.020 (0.001)
σ_λ	1.530 (0.042)	1.596 (0.049)	0.959 (0.080)	0.755 (0.081)
σ_ρ	0.452 (0.024)	0.456 (0.025)	0.415 (0.088)	0.401 (0.075)
σ_ω	1.957 (0.090)	1.811 (0.084)	2.729 (0.717)	2.858 (0.864)
τ	4.072 (0.068)	4.011 (0.069)	4.200 (0.131)	4.278 (0.134)
$\tau_{\text{Low Inc}}^\dagger$	0.281 (0.008)	0.286 (0.008)	0.255 (0.015)	0.264 (0.015)
Log-Likel	30234.4	30079.6	3507.1	3496.9

Note: Number of Observations is 1,422 (CentERpanel) and 178 (Laboratory), respectively. Estimation follows (10) based on the utility function defined by (4) and (6) in the top corners of the individual panels. The entries for σ are the standard deviations of the untransformed normal distributions of the random coefficients.

[†] The low incentive treatment enters multiplicatively.