Essays on Behaviour Under Risk

SERGIO ALMEIDA DE SOUSA, BSC, MSC

THESIS SUBMITTED TO THE UNIVERSITY OF NOTTINGHAM

FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

JULY 2010

ABSTRACT

This thesis consists of three essays on behaviour under risk. First, I investigate experimentally three related questions: (1) the effects of small-scale changes in wealth on risk attitudes; (2) whether potential changes in risk attitudes induced by such wealth increment are affected by (a) by the span of time this small-scale change in wealth has been anticipated for, and (b) the form taken by the wealth increment. There are three major results. One, whether risk attitudes are affected by a small-scale change in wealth depends on the form taken by the wealth increment. Two, that failure in replicating "house" money effect suggests that people may treat windfall money differently from earned money. Three, that the attitudes to risk are stable over the span of time we investigate. Second, I investigate how cognitive ability relates to consistency of behaviour under risk. Individual behaviour can be consistent in several forms. I find that individuals with higher cognitive ability display more consistent behaviour – in terms of choice and displayed type of risk preferences. Yet, in contrast to some recent studies, I find that individual measures of attitudes toward risk are not associated with cognitive ability. Third, I investigate the efficacy of a punishment mechanism in promoting cooperative behaviour in a public goods game when enforcement of punishment is uncertain. Numerous experimental studies have found that a sanctioning system can promote cooperative behaviour. But they rely on perfect enforcement of punishment. I find that a sanctioning system can no longer promote cooperative behaviour in a public goods game when punishment enforcement is a low-probability event.

Acknowledgements

I wish to express my deeply-felt thanks to my thesis supervisors, Professor Robin Cubitt and Professor Martin Sefton, for their thoughtful guidance. They have tirelessly mentored and warmly encouraged throughout my thesis. Their advice and comments helped me enormously in improving this thesis – any remaining error is mine though. But it would be a understatement to say their contribution is confined to this thesis, for I myself have learnt from and because of them in several ways; and I feel this is a much greater contribution. I also thank to Professor Horst Zank and Professor Simon Gächter for helpful discussion and several suggestions. I would like to show my gratitude to Ana Neves for her support and numerous discussions that helped me to improve my arguments on several issues. I am also grateful to the Brazilian Ministry of Education for funding.

To my parents

CONTENTS

List of Figures

List of Tables

Chapter 1

Chapter 1

Introduction

This thesis consists of three essays presenting results of experiments studying behaviour under risk. These essays are almost entirely self-contained and each addresses different issues. Yet, behaviour under risk is a unifying theme of this thesis so that this introduction provides an overview of the economic literature on this topic. It is by no means exhaustive, but tries to organise this enormous literature with the aim of pinpointing where and how each essay contributes to the body of research on decision-making under risk.

The significance of this topic is hardly disputable. Risk is so prevalent in a vast number of aspects of economic and social activity that it seems only natural that a great deal of theoretical and applied work carried out by economists has been devoted to it. While decision-making under risk is a research area of interest in its own right, the widespread use of game theoretic analysis and the search for microfoundations for macroeconomic models has made models of decision-making under risk a key building block of models across several fields of economics.

In the light of the huge literature on the topic and its importance, it is perhaps surprising that economists are still "hunting" for a good theory of choice under risk (Starmer, 2000) – a combination, ideally, of predictive power, simplicity and tractability. Such search has been motivated in large part by results from waves of experimental studies testing the axioms which the standard analysis of decision behaviour under risk, based on von Neumann and Morgenstern Expected Utility Theory (EUT), relies on. The assessment of EUT based on the accumulated evidence has been mixed. While most of the first wave of experimental studies testing EUT suggested that people's choices contravene even key axioms of the theory (for a review, see Camerer, 1995), tests of other competing theories have not decisively established their empirical superiority over EUT[1]. This seems to account, to some extent, for why EUT still remains a central framework to much applied theory. Nevertheless, the violations that have come to light spawned refinements of the standard model and the development of alternative theories of decision under risk, such as prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992), rank-dependent utility (Quiggin, 1982; Chew, 1983) and regret theory (Bell, 1982; Loomes & Sugden, 1982), to name but a few. The bulk of these new theoretical developments could well be subsumed into categories according to how the basic elements of a risky choice, its outcomes and associated probabilities,

[1]For empirical tests and a somewhat mixed assessment of competing decision theories see for example Harless & Camerer (1994) and Hey & Orme (1994).

are each dealt with[2]. There is a large number of functional specifications for each of these elements and how a rational[3] individual combines them prior to reaching a choice. Some theories may even have a stochastic version, in which some form of randomness is introduced in the process of choice; the modelling of the source of randomness itself has given rise to a variety of models. In fact, Loomes & Sugden (1995) show how a given base model such as EUT can have very different implications depending on how the stochastic component is introduced. All these new developments make this literature vast[4].

Empirical work of experimental nature on individual choices under risk is equally vast. This is hardly surprising given, on the one hand, the rich interplay between theory and data and, on the other hand, an ongoing interest in unfolding the determinants and correlates of people's attitudes towards risk. It seems natural, therefore, to divide this empirical literature, or most of it, into two broad categories: *theory testing* and *fact finding* studies. The first would consist of laboratory and field experiments primarily aiming to test fundamental assumptions of EUT and alternative theories; this would include, for instance, tests for transitivity of preferences (e.g. Loomes *et al.* , 1991; Birnbaum & Schmidt, 2008), and the independence axiom (e.g. Starmer & Sugden, 1989b; Loomes, 1991). The second category would consist of studies seeking to examine how a wide variety of aspects

[2]One could also classify these alternative theories according to properties of EUT they relax.

[3]The issue of rationality is rather fuzzy in normative and descriptive models. See Gilboa (2009) for a discussion.

[4]Camerer (1995) and Starmer (2000) offer surveys.

affect individuals' risky choices; this would include, for instance, investigations upon effects of features of a decision setting on risk-taking behaviour (Goeree *et al.* , 2003; Chetan *et al.* , 2007; Bothner & Stuart., 2007) as well as studies investigating risk-taking involving single and compound gambles (Benartzi & Thaler, 1999; Klos *et al.* , 2005).

We are ready to acknowledge that this classification is not problem free. Because experimental work often features a theoretical framework to make sense of observed data, this dual distinction could have trouble to provide an unambiguous classification, having to rely on subjective interpretations of the work. Yet, we use the term *fact finding* in a very broad sense; it refers to empirical work that purely generates data, but also to empirical work that is informed by theory but not primarily designed to test fundamental axioms and general behavioural implications of theories of decision under risk. Thus the empirical work of Starmer (1999), for example, would be seen as theory testing, for it investigates particular forms of non-transitivity in risky choices implied by prospect theory; while the work of Harrison *et al.* (2007) would be seen as fact finding, for while it tests an EUT account of effects of background risk on attitudes towards risk, it is primarily an investigation of how risk preferences are affected by the nature (monetary and nonmonetary) of the prizes involved. But whatever terminology one uses, it is undeniable that the empirical strand of the economic literature on behaviour under risk has gone beyond theory testing *à la* Allais (1953) or Kahneman & Tversky (1979).

Indeed, this empirical literature has been exploring various aspects that influence behaviour under risk. They center, in our view, on three main genres. First, *characteristics of the risk* (lottery) such as size and nature of prizes, knowledge and distribution of probabilities, presence or absence of losses, and so forth. Experimental investigations on effects of stake size (e.g. Holt & Laury, 2002), real and hypothetical prizes (Harrison *et al.* , 2005a), and studies on ambiguity (e.g., Halevy, 2007) are examples of this. Second, *characteristics of the decision setup* such as frames of choice (most on common-consequence effects e.g. Humphrey, 2000), elicitation methodology (Andersen *et al.* , 2006b), learning opportunities (Loomes *et al.* , 2003) and so on. Third, *characteristics of the decision-maker*; papers investigating the effects on risk attitudes of gender (e.g., Agnew *et al.* , 2008; Borghans *et al.* , 2009), income (e.g., Gertner, 1993b; Bosch-Domènech & Silvestre, 2003) and cognitive skills (e.g., Cesarini *et al.* , 2009) are examples of this. We then argue that the contributions of each essay in this thesis fit into these two latter lines of research in the literature.

Chapter 2: On characteristics of the decision setup and decision-maker

The first essay, presented in the next chapter, is a "hybrid" case for it examines the role of frames and characteristics of the individuals on their risky choices. This essay has two primary purposes.

First, it attempts to test experimentally the effects of small-scale changes in wealth on risk attitudes. Assumptions about the effects of changes to wealth on individuals' attitudes to risk plays a key role in empirical and theoretical results in a broad range of topics in economics. While it is appealing to assume that wealthier individuals are willing to take more risks, the empirical evidence as to how changes in wealth affect risk-taking behaviour is mixed at best[5]. Furthermore, most of this evidence is based on non-experimental data involving a cross-section of individuals. While econometric techniques could deal with the potential problem of endogeneity involving observed risky behaviour and wealth levels, these cross-sectional studies can arguably not provide an accurate account of wealth effects if preferences are heterogeneous. We design an experiment to examine this issue using an environment where both risks taken and changes of wealth are controlled.

Second, it further investigates whether and how potential changes in risk attitudes induced by a small-scale wealth increment are affected by the form taken by such increment. We examine two forms or frames that we term "inside" and "outside". The "inside" framing effect refers to changes in risk preferences when a small-scale change of wealth is given "inside" a lottery. That is, a subject's risk preference is elicited through certainty-equivalent of two lotteries, say L and L'. But L is just a shifted version of L', whereby a common amount of Δw is added to its prizes. The "outside" framing effect, in turn, refers to changes in risk preferences when Δw is simply given to the subject outside the lottery. In this case, the "outside

[5] I will review this evidence more fully in Chapter 2.

money" framing effect is captured by differences in risk preferences elicited before and after Δw is administered. Several studies investigate the effects on one's degree of risk aversion when the size of stakes is increased (e.g. Binswanger, 1980; Holt & Laury, 2002; Harrison *et al.* , 2005a), but little has been said about whether and how the form taken by an increase in the payoff levels affects individuals' degree of risk aversion. Furthermore, according to rational choice models, these framing effects, whatever they are, should not be different. But while there is evidence that people's choices are affected by changes of frame (see, e.g., Tversky & Kahneman, 1986), little has been said about whether the effects of a monetary gain on attitudes to risk are equivalent across the frames considered here. Part of the design of our experiment examines these issues.

Chapter 3: On characteristics of the decision-maker

The third chapter examines how cognitive ability is associated with some aspects of decision-making under risk – it would thus fall into the part of the literature examining effects of characteristics of the decision-maker on risky choices.

There is a recent literature on the relationship between cognitive ability and economic preferences[6]. Most of these studies examine time and risk preferences through laboratory experiments, comparing observed preferences across groups with different levels of cognitive ability. For the most part[7], these laboratory experiments use a multiple-price-list method, sim-

[6] We review this literature in Chapter 3.

[7] An exception is Frederick (2005). We come back to this later.

ilar to Holt & Laury (2002), to elicit subjects' attitudes towards risk. By this elicitation method, a subject faces a task consisting of a set of pairwise choices displayed in a table, with each involving a choice between a lottery and a risk-free sum of money. Each of such task provides a measure of risk preferences. In general, existing experimental studies have presented subjects with few risk-elicitation tasks, collecting only a few decisions from each subject.

Although this practice is understandable given other features in the design of these experiments, it may not necessarily reveal an accurate representation of individuals' attitudes to risk: responses to a single task with this format may include a significant "noise" component that one cannot even out when there is lack of repetition. Therefore, while the existing experimental studies suggest that people with higher cognitive ability are less risk averse relative to those with lower cognitive ability, they tell us little about whether this association holds when individuals are given some learning opportunity (by experience) so that their risk preferences become better reflected in their choice decisions. In Chapter 3, we address this issue by examining the connections of cognitive ability with risky choice behaviour in a repeated setting. We take advantage of the repeated nature of risky choices to, in addition, examine how several forms of consistency in individual choices relate to cognitive ability.

Chapter 4: On characteristics of the decision setup

Finally, the essay on chapter 4 examines how contribution and punishment behaviour in a public goods game is affected by the introduction of a new element of uncertainty into the environment – measurable uncertainty over punishment enforcement. This places this essay in the part of the literature examining how characteristics of the decision setup affect risky choices. This relies on the view that a public good game with punishment opportunities can have several sources of risk, which is hardly controversial.

Several experiments have brought evidence that the level of contributions in public good games in one-shot and repeated public good games is far from theoretical predictions[8]. Yet, it has also been shown that these contributions display a steady decay pattern when the game is repeatedly played (Mark Isaac *et al.*, 1985) [9]. Since then numerous experimental studies have found that a sanctioning system can induce individuals to adopt and sustain cooperative behaviour[10]. But the commonly used experimental design relies on two assumptions: perfect monitoring and perfect enforcement – features that most sanctioning systems outside the laboratory do not have. This abstracts away from important sources of risk in this decision setting. In chapter 4, we address this issue by relaxing one of this assumptions: we introduce a sanctioning system that is no longer "perfect" regarding the enforcement of punishment. By doing so, we introduce an additional source of risk into the decision setting, with the aim of testing

[8]For a review of the "first wave" of such experiments see (Ledyard, 1995).

[9]One interpretation is that contributions are partly due to some sort of misunderstandings that eventually vanishes when the game is repeated (see Andreoni, 1995a), although more recent evidence suggests that this might be due to "imperfect" conditional cooperation in the sense people do not match others' contributions Fischbacher & Gächter (2009).

[10]We review this literature in chapter 4.

whether such imperfect (monetary) sanctioning system can still promote cooperative behaviour as documented by Yamagushi (1986), Ostrom *et al.* (1992) and Fehr & Gächter (2000).

The rest of this thesis is organised as follows. Chapter 2, 3 and 4 present the above mentioned essays. Chapter 5 concludes; in it, we summarise the the major findings, discuss some limitations and propose further extensions to our research.

Chapter 2

Chapter 2

SMALL-SCALE CHANGES IN WEALTH AND ATTITUDES TOWARD RISK

2.1 Introduction

The primary purpose of this chapter is to describe an experimental attempt to test the effects of small-scale changes in wealth on risk attitudes. We also explore how the framing of the wealth change and the span of time this wealth change has been anticipated for influences attitudes to risk.

Understanding the attitude to risk of economic agents is a goal that has long been pursued by many economists. Much of the theoretical and empirical effort to analyze risk-taking behaviour has been influenced by the Arrow-Pratt measures of risk aversion for von-Neumann-Morgenstern utility functions. For an expected utility maximizer with a utility function $u(.)$ defined over wealth w, Arrow and Pratt interpret the functions $R_A(w) = -u''(w)/u'(w)$ and $R_R(w) = wR_A(w)$ as local measures of absolute and relative risk aversion, respectively. An individual is then characterised as a decreasing, constant, or increasing absolute (relative) risk averter depending upon whether $R'_A(w)(R'_R(w))$ is less than, equal to, or greater than zero. Each possibility describes how changes in wealth affect one's willingness to take a given risk.

Assumptions made about the sign of such wealth effects on risk aversion underpin empirical and theoretical results in a broad range of topics in economics. Ogaki & Zhang (2001), for instance, point out how strikingly different empirical tests of the risk sharing hypothesis involving household consumption models can be when estimation methods are based on preferences that allow relative risk aversion to vary with the level of wealth. Models dealing with phenomena as diverse as life-cycle savings (Weil, 1993), portfolio choice (Hadar & Seo, 1990), and asset pricing (Gollier, 2001), make predictions that are very sensitive to the way risk attitudes are affected by changes in wealth. How risk aversion varies with wealth has also implications for Samuelson's fallacy of large numbers (Samuelson, 1967) and Rabin's calibration theorem (Rabin, 2000), paradoxes that have been the object of considerable attention[1]. Samuelson's paradox refers to a pattern of choice that rejects positive mean gambles, such as an even chance to win \$200 or lose \$100, but accepts one hundred of such gambles in a row. Samuelson regarded that choice behaviour as inconsistent with Expected

[1]Regarding Samuelson's paradox: see, on the empirical front, Redelmeier & Tversky (1992); Haubrich (1998); Benartzi & Thaler (1999); Gneezy *et al.* (2003); Klos *et al.* (2005); Chen & Corter (2006); and, on the theoretical front, see Nielsen (1985); Ross (1999); Peköz (2002); Hammarlid (2005). Regarding Rabin's theorem, see Rubinstein (2001); Watt (2002); Wakker (2005); Bombardini & Trebbi (2005); Cox & Sadiraj (2006); Palacios-Huerta & Serrano (2006).

Utility Theory (EUT). Assuming that the single bet is unacceptable at all wealth levels, he proved a theorem stating that the initial rejection should imply a rejection of any sequence of such bets. But rejection of a gamble at all wealth levels is an assumption that, as showed by Ross (1999), holds only for a limited class of utility functions, namely, those displaying constant absolute risk aversion. Such utility functions describe individuals whose attitudes towards risk are the same across wealth positions. A similar claim has been made by Cox & Sadiraj (2006) and Palacios-Huerta & Serrano (2006) regarding the validity of Rabin's demonstrations that risk aversion over modest stakes within EUT implies absurd risk aversion over large stakes gambles. They point out that Rabin's striking results rely on the assumption that a given risk is *consecutively* rejected across a wide range of wealth levels, which in a sense amounts to saying that risk aversion does not vary with wealth.

Despite the analytical importance of the characterization of absolute and relative risk aversion, there is mixed empirical evidence as to the effects of changes in wealth on attitudes toward risk. Ogaki & Zhang (2001), Guiso *et al.* (1996) and Rosenzweig & Binswanger (1993), for instance, find evidence in support of the decreasing relative risk aversion hypothesis, while Szpiro (1986), using data on insurance, finds empirical support for constant relative risk aversion. Barsky (1997) and Donkers *et al.* (2001), instead, find evidence that risk aversion increases with wealth, while Binswanger (1980) finds that changes in wealth have no significant effect on risk aversion. Even though various methodology-related arguments may be given to explain that discrepancy, it is debatable whether these econometric studies have fully provided evidence on the way attitudes to risk are affected by changes in wealth. Most of the existing results are based on data involving choice behaviour among individuals of different wealth levels[2]. But inferring how risk aversion varies with wealth from cross-sectional observations may not be accurate when preferences are heterogeneous.

At first sight, a data set containing measures of risk attitudes at various wealth positions of an individual (i.e. a long panel) could fully overcome that concern. However, wealth is likely not exogenous to attitudes to risk: unobservable risk-driven choices can underly the changing of wealth positions. Thus, econometric estimates would still have to address the problem of endogeneity that could confound estimation. An alternative approach would be a laboratory experiment, where wealth can be exogenously manipulated. Though this method cannot produce, under incentivised conditions, an extensive map of individuals' wealth states onto their risk attitudes, it can produce evidence that complements econometric studies by providing careful controls of risks taken and changes of wealth experienced. While several experimental investigations (e.g., Harrison, 1986; Holt & Laury, 2002; Bosch-Domenech & Silvestre, 1999; Bosch-Domènech & Silvestre, 2003) have brought evidence about attitudes toward scaled-up risks given

[2]An exception is Eisenhauer (1997), who uses a long sample of aggregate time series data from the U.S. and finds evidence that absolute risk aversion increases with wealth – which is in contrast with the above mentioned studies based on cross-sectional analysis.

subjects' initial wealth level, contributions that test for effects of changes in wealth on attitudes toward a given risk are scarce[3].

This chapter attempts to elicit experimentally the sensitivity of risk attitudes to small-scale changes in wealth. We elicit attitudes to risk through a multiple price list method at two different times, say t_0 and t_1. A subgroup of subjects (treatment group) is awarded money between t_0 and t_1. Another sub-group (control group) is not awarded any money, and their choices are used to detect changing patterns of risk attitudes elicited at t_1 relative to t_0 that cannot be attributed to changes in wealth, induced by the experimenter.

Payment for their decisions was then made at the end of the experiment, in cash, according to the random lottery incentive system. This random-lottery procedure, by which several decision problems are faced but the subject is paid the outcome of only one of them, has been extensively used. It allows an incentivised elicitation of individual choices in multiple-task settings avoiding income effects Lee (2008). The random-lottery system provides an incentive-compatible elicitation mechanism both under EUT and PT, in the sense that subjects are incentivized to report genuine valuations to the lotteries they face which reflect their true preference ordering over the pairwise set of options.

This chapter also studies how effects of small-scale gains on risk attitudes are influenced by how far in advance such gains are anticipated. This

[3]See Wolf & Pohlman (1983) and Levy (1994).

issue has practical importance. Consider, for example, the random-lottery system; this is a procedure by which several decision problems are faced but the subject is paid the outcome of only one of them. The random-lottery system provides an incentive-compatible elicitation mechanism both under Expected Utility Theory and Prospect Theory, in the sense that subjects are incentivized to report decisions that reflect their true preference ordering over the set of options in each decision problem. This is a a common practice in experimental investigation of economic behaviour in multiple-task settings. It aims to address a concern that subjects' behaviour across rounds may be contaminated by wealth/payoff effects if they are paid for their decisions in all tasks (see Holt & Laury, 2002; Cubitt *et al.* , 1998a; Lee, 2008). The validity of paying for all tasks and, consequently, the motivation of the random-lottery system, relies to some extent on assumptions about (1) the argument of the utility function (accumulated winnings?) and (2) the speed with which subjects update such argument in light of earnings in previous tasks. If, for instance, individual behaviour acts on preferences represented by a utility function defined on deviations from a given reference point, and this reference point depends on the individual's endowments, then paying individuals for decisions in all tasks (assuming they are payoff-wise similar) would contaminate behaviour with endowment effects. But this would be true only if individuals' reference point for one task varies according to the payment they receive from other tasks – an assumption that has not yet been tested and both treatments in our experiment shed some light on.

Examining whether and how quickly individuals change their behaviour as money is earned in the earlier periods of an experiment has also theoretical importance – in particular, for reference-dependent theories. Theories of reference-dependent preferences, most notably, Prospect Theory (PT), postulate that individuals evaluate the outcomes of an economic prospect by contrasting them to a reference-point. While an individual's reference point has been traditionally interpreted in this literature as corresponding to her current endowment (e.g., Kahneman & Tversky, 1979; Tversky & Kahneman, 1992), applications often adopt the simplified assumption that such reference point is a zero quantity of such endowment (e.g., Harless & Camerer, 1994, p.1255; and Andersen *et al.* , 2006a, p.18). Koszegi & Rabin (2006) challenge these assumptions, building up a reference-dependent model in which the expectations about outcomes rather than wealth levels are used as a reference-point. They claim that an expectation-based reference-point helps their model to accommodate a wide variety of observed behaviour that has been found irreconcilable under standard formulations of the major decision theories (EUT and PT). We address some of those issues by running the above mentioned baseline experiment under two treatment conditions, which differ in the time elapsed between t_0 and t_1 (a few minutes in one treatment and one week in the other). By doing so, our overall experiment design also allows us to examine (a) the short-term stability of preferences over lotteries, and (b) the different assumptions as to reference-point determination and adjustment.

Finally, our experimental design also allows us to examine the effect of what is termed here the "inside" and "outside" frames on attitudes to risk. The "inside" framing effect refers to changes in risk preferences when a small-scale change of wealth is given "inside" a lottery. That is, a subject's risk preference is elicited through certainty-equivalent of two lotteries, say L and L'. But L is just a shifted version of L', whereby a common amount of Δw is added to its prizes. The "outside" framing effect in turn – also known as "house money effect" – refers to changes in risk preferences when Δw is simply given to the subject[4].

In this case, the "outside money" framing effect is captured by differences in risk preferences elicited before and after Δw is administered. An underlying principle of rational choice models is that different frames of a given choice problem should not induce an individual to make different decisions if the variations in frame leave the consequences of the choice problem unchanged. But while it is no longer novel that people's choices are affected by some changes of frame (see, e.g., Tversky & Kahneman, 1986), little has been said about whether the effects of a monetary gain on attitudes to risk are equivalent across such frames. These "Inside-Outside" money effects, if found, could be of significant relevance for (a) theoretical modeling (e.g., of risk preferences over different stakes and over different

[4]There is a subtle distinction here. The "house money effect" (Thaler & Johnson, 1990) is a change of risk preferences in a *particular* direction that is induced by money given prior to risky choices; the "outside" money effect does not postulate a particular direction for how money given given to an individual prior to risky choices would affect his decisions.

wealth levels), and (b) practical experimental design questions (e.g., balance between task incentives and show-up fees). Part of our experiment addresses this issue.

The rest of this chapter is organised as follows. Section 2.2 reviews the experimental evidence on risk attitudes. Section 2.3 describes the experimental design. Section 2.4 derives the predictions of expected utility for treatment effects on subjects' choices. Section 2.5 presents the results. Section 2.6 concludes.

2.2 Experimental evidence on attitudes toward risk

The study of individual decision-making under risk has been the object of interest of much experimental work (Camerer, 1995; Starmer, 2000). While much of this work was initially related to an attempt to challenge the EUT paradigm as a description of how people actually make decisions under risk, it has been extended over the years to investigate the determinants of risk behaviour, its measurement and correlates. The resulting experimental literature is massive and every year general and specialised journals bring more to this body of work. In spite of the broad range of risk aversion-related issues investigated by this literature, it could be said that much of this empirical literature falls into four major categories:

(i) Methods of elicitation of risk aversion

(ii) Theory-testing using the probability triangle

(iii) Individual correlates of risk-taking behaviour

(iv) Wealth effects on risk aversion

This classification is essentially based on the major purpose underlying most of the empirical work over the last twenty years on risk preferences. It provides a fairly comprehensive taxonomy that helps to make some sense of the nature of the empirical work (mostly experimental) that has been carried on this topic. In what follows, we briefly review a representative sample of the branch of the literature in each of those categories.

The first category refers to experimental investigations that look at whether and how measured individual risk attitudes differ across different methodologies of risk elicitation. Using a within-subject design, Isaac & James (2000), for instance, have found that estimates of risk attitudes elicited through first-price auction and Becker-DeGroot-Marschak (BDM) mechanisms are considerably different. But these results should be interpreted with some caution as some potential confounds were not controlled for[5] and there are differences in the way bid data from each elicitation procedure is collected[6].

[5]For instance, order effects: 40 rounds of first-price auction-based elicitation are *always* conducted first, and only after that, 4 rounds of a BDM procedure are conducted. Also, and perhaps more importantly, a subject's position alternates between procedures, being framed as a seller in the BDM and as a buyer in the first auction. Loomes *et al.* (2002) provide some evidence that attitude to risk indeed change over a large sequence of tasks.

[6]Estimates of risk aversion for bid behaviour in the BDM procedure are based on average-bid only of the last two BDM bids (Isaac & James, 2000, p.181), whereas the procedure to estimate risk coefficients for bid behaviour in the first-price auction is based on a transformation of a parameter estimate of a linear Nash equilibrium bid function.

Another example of how adopted methodological frames can influence elicited risk attitudes is Harrison *et al.* (2007). They primarily investigate whether using non-monetary commodities as lottery prizes instead of monetary rewards affects risk attitudes; a group of subjects is also assigned to a treatment condition in which the non-monetary outcome (rare coins) has some background risk. The experiment is run in the field using subjects (numismatists) who have experience with the commodity that replaces the conventional lottery outcomes. While it departs from the conventional laboratory experiment regarding the subject pool used, they use a multiple price list design to elicit risk attitudes. This is a common approach to elicit measurements of risk aversion from choices subjects make when confronted with a list of paired lottery-choice decisions. They found that replacing money by a non-monetary outcome does not itself significantly change elicited measures of risk aversion. Behaviour within their sample is mostly risk averse, which is in line with other studies that also use this risk elicitation procedure. Yet they found that subjects, as predicted by standard EUT, tend to behave in a more risk averse way when the outcome value involves some uncertainty[7].

The second category refers to a stream of the literature that has used the Marshak probability triangle to examine competing theories of choices under risk. While designed by Marschak (1950) as a geometrical illustration of rational choice involving risk, it was popularised by Machina (1987) as

[7]As coins used as outcomes of the lotteries do not have a grade certification and, as a result, cannot have their retail value easily assessed by collectors.

a device that can be used to illustrate preferences over a certain family of lotteries[8]. The Marshak triangle is a unit triangle in which each point represents a probability distribution $p(p_1, p_2, p_3)$ over a given set x_1, x_2, x_3 of outcomes, where $x_1 > x_2 > x_3$ and p_i is the probability of winning x_i. The vertices of the triangle represent sure outcomes – degenerate lotteries with probability mass concentrated on a single outcome – while the points lying on the triangle edge lines represent binary lotteries – one of the outcomes has no probability mass.

The Marshack probability triangle has been used to great effect in the study of choice under risk. Since lotteries and preferences over them have a diagrammatic representation, the probability triangle framework can be used to compare individuals's risk aversion characteristics, as well as to construct tests of alternative theories of choice under risk. In fact, much of the experimental literature using the probability triangle has been devised to investigate axioms of the expected utility theory. Several experimental studies have found evidence of violations of such axioms (for a review, see Machina (1987); Starmer (2000); Camerer (1995)). It is an implication of EUT, for instance, that risk aversion, which is measured by the slope of the indifference curves, is constant across the triangle. Yet the diagrammatic representation in the triangle of individuals' decisions in simple choice problems systematically departs from such representation. Many of these violations have been documented in what are termed as common con-

[8]Those involving a probability distribution over up to three outcomes.

sequence effects – patterns of choice in pairs of choice problems involving lotteries that violate the EUT's independence axiom[9].

In these studies, however, risky choices involve lotteries located on the corners and edges of the triangle; and there is some evidence that violations are mitigated when the risky choices involve lotteries that if represented in the triangle no longer lies on its boundaries. Conlisk (1989) shows, for example, that violations as observed in the examples used by Maurice Allais are less frequent when the choice problem (1) is re-phrased as a three-step problem, and (2) involves lotteries lying on the interior of the probability triangle. Discussions and empirical examinations of these effects can be found in Kahneman & Tversky (1979); Camerer (1995); Wu & Gonzalez (1998); Starmer (1992); Humphrey (2000) and Starmer & Sugden (1989a); Carlin (1992); Hey & Orme (1994); Cubitt *et al.* (1998b), among others. In sum, these experimental studies using the triangle have not only shown under which conditions systematic violations of the independence axiom are to be observed, but also that subjects may have variable levels of risk aversion as the probability of winning the middle prize changes – phenomena that violates the standard expected utility theory.

The work that falls into the third category we alluded to earlier is devoted to identifying demographic variables that may importantly influence individual attitudes to risk. Several experiments on risk aversion have

[9]It is worth noting that in most of these studies risk attitudes are captured by comparison of individuals' choices across different probability distributions over a fixed set of prizes; in our experiment, however, risk attitudes are captured by the curvature of an individual's utility function implied by her choices.

pointed out, for example, that women are more risk averse than men (Eckel & Wilson, 2004; Eckel & Grossman, 2007), whereas Schubert *et al.* (1999) have found that this gender gap in risk propensities no longer exists when the context of decisions is manipulated, and instead of abstract gambles, decisions are framed as investment and insurance problems. Risk aversion is also found to be significantly related to age and height (Dohmen *et al.* (2005)). While most of this work has been confined to investigate the effects of demographic variables that are more commonly and easily observed/measured, experimental findings in Frederick (2005) suggest that cognitive ability may also play a role in individuals' risk behaviour. This has been confirmed by Dohmen *et al.* (2007), who find that subjects who perform better in a 10 question IQ-type test tend to be less risk averse than those who perform worse.

Another important theme in experimental studies on risk aversion is how people's willingness to take risks is affected by changes in wealth. This may cover both the case where subjects, from a given initial wealth position, assess lotteries with different magnitudes of payoffs ("inside case"), and the case where a exogenous variation in wealth is administered and a given risk is evaluated at each of these wealth positions ("outside case"). The "inside case" refers to assessment of risk-taking behaviour of an individual with initial wealth w in two risky scenarios. For instance: in one, her risky choices involve the lottery $L(1/4, 100; 3/4, 50)$; in the other, her risky choices involve a lottery $L'(1/4, 150; 3/4, 100)$, which is just L with a common amount of 50 added to its prizes 100 and 50. The "outside case", in turn, refers to assessment of risk-taking behaviour of an individual in two wealth scenarios. In one scenario, an individual with initial wealth w faces a risky problem involving, say, the lottery $L(1/4, 100; 3/4, 50)$; in the other, she will face a risky problem involving the same lottery L, but before doing so, she will be endowed with a certain amount of money, say 50. Depending on problem framing, these situations should be viewed as equivalent[10]. Whereas in the "inside case" it is as if subjects assess different positions along the wealth scale, in the "outside case" the risk is given and subjects are supposed to assess it from two different wealth positions.

But studies in this category have been confined to the "inside" case in which it is investigated how risk aversion displayed by subjects in real lottery decisions is affected by the size of the payoffs at stake. Kachelmeier & Shehata (1992), Holt & Laury (2002) and Harrison *et al.* (2005a), for example, have found that elicited risk aversion measures increase when the size of lottery stakes involved in such decision tasks is *scaled up*, which is clear evidence of increasing relative risk aversion[11]. Indeed, attitudes to risk do not seem neutral to changes of wealth. Bosch-Domenech & Silvestre (1999) hypothetically endow subjects with several amounts of money. These amounts range from \$3.50 to \$103 US dollars. Then they ask subjects whether or not they want to buy fair insurance against losing the sum

[10]This is the case for an expected utility maximiser whose utility is defined on final wealth.

[11]In Holt & Laury (2002), this effect is only found with real incentives.

received, which has probability of 20%. The sum each subject actually receives at the end is randomly drawn from the list of amounts. Analysing the hypothetical decisions of subjects, they observe many patterns regarding insurance purchase, but find a significant positive association between the decision to insure and the size of the endowment. Assuming that insurance buying is monotonically related to risk aversion, their results seem to confirm that risk preferences change according to the size of payoffs at stake.

While these positive "inside" wealth effects provide some empirical support for utility functions exhibiting increasing relative risk aversion (IRRA), they still leave unanswered the question of whether and how behaviour toward a given fixed risk is altered by changes in one's wealth level. Studies on this are scarce. One of the few attempts to investigate this in the lab is done by Levy (1994). He uses a portfolio allocation-type of decision problem repeated over 10 periods, allowing subjects to accumulate their earnings at each period. He finds evidence that the more "wealth" subjects have, the more they are willing to take risks. Interestingly, he also finds that the proportion of the subject's "wealth" allocated to risky assets does not decrease, as dictated by IRRA, as wealth becomes larger. It is not entirely clear how robust these results are; not just because of the contrast with the preceding findings regarding relative risk aversion (RRA), but also because the regression analysis was structured in such a way that within-subject noise is not controlled for. Moreover, and regardless of robustness, it is doubtful if subjects integrate money earned throughout an one-hour task into their conception of wealth. Therefore, the observed results may not be accounted for by a utility function exhibiting decreasing absolute risk aversion (DARA), but by a "house money" effect (Thaler & Johnson, 1990). Besides, by allowing subjects to accumulate earnings across trading periods, his design re-introduces the problem with field data: endogeneity of risk-taking behaviour. One of the goals of our experiment is thus to fill this gap, using a design that administers a carefully controlled "exogenous" small-scale change in wealth.

The analysis carried out and the treatment conditions employed help to distinguish our experiment from the others, above mentioned, on this issue. Firstly, because we perform a within-subject analysis, we can control for changes in risk attitudes induced by noisy behaviour rather than by the stimuli – an aspect not addressed by Levy (1994). Secondly, we pay *all* subjects for their decision, thus having a sample in which all subjects were provide incentives to think carefully about their decisions. In Thaler & Johnson (1990), for instance, subjects are uncertain about whether even a single of their choices will be for real: only a few subjects (roughly 5% of the sample) will be selected to play out one of their gamble decisions[12]

Furthermore, our design allows us to explore other issues. First, how the span of time this small-scale change in wealth has been anticipated for influences attitudes to risk. We do so by contrasting before- and after-

[12]Subjects are, in this case, still incentivised; but the the large likelihood of having no decision for real may dilute subjects' incentives.

increment measures of attitudes to risk of group of subjects who received the increment news either a few minutes or a week before the second risk-elicitation stage of the experiment. Second, whether the framing of the small-scale increment – the above mentioned "inside" and "outside" frames – can affect sign and magnitude of its effect on subjects' risk attitudes. Third, and finally, by having the same set of risk tasks being faced on two different occasions, we can also use a test-retest approach to examine the short-term stability of risk preferences – an issue of importance to the reliability of the elicitation method itself.

2.3 Experimental Design

We now turn to a detailed description of our design. Because the elicitation of individual's attitudes toward risk is the building block of our experiment, we shall first present the method used to this end. Then, we describe the sequence of task stages the experiment consists of, followed by a description of the treatments. After that, we describe the payment procedures and the mechanisms used for incentive-compatible elicitation of responses.

2.3.1 Risk-elicitation procedure

We propose to elicit risk attitudes employing a variant of the Multiple Price List (MPL) procedure used by Holt & Laury (2002)[13]. The format of the MPL method devised by Holt-Laury has been widely used in risk-elicitation laboratory experiments (e.g., Andersen *et al.* , 2006b; Harrison *et al.* , 2007; Offerman & Schotter, 2009) and involves an easily understandable task. Measurement of risk aversion is based on ten pairwise choice problems presented altogether in a table, one per row. Each problem is to choose between a lottery A, say a p chance to win x_a and $1-p$ to win y_a (where $y_a < x_a$), and a lottery B, say a p chance to win x_b and $1-p$ to win y_b (where $y_b < x_b$), in which p is systematically varied from 1/10 to 1 when proceeding down the table. Then, because the difference between payoffs for A is much larger than the difference between payoffs for B (i.e., $x_a - y_a >> x_b - y_b$), one should cross over to the lottery B at some point when going down the table.

Here, we depart from Holt-Laury in two respects: first, we use more than ten pairwise problems; second, the pairwise choice problem at each row is to choose between a fixed lottery and an amount of money with certainty, say M, that is systematically varied from row to row by a constant amount, say δ. For example, consider that the fixed lottery is L, a lottery with p

[13] Another way of eliciting risk attitudes is by asking subjects' selling and purchase prices for a lottery through auction procedures (e.g., Harrison, 1986). While this method yields a point estimate of certainty-equivalents, the way the pricing task is framed can considerably alter the implied risk attitudes (see Holt & Laury, 2002, p.1644). Auction mechanisms have also been reported to generate biased elicitation as subjects tend to overbid to ensure they will be the winner bidder (see, e.g., Krahnen *et al.* , 1997).

chance to win x and $1 - p$ chance to win y, where $x > y$. At the first row, the decision problem an individual faces is to choose between x for sure and L. At the second row it would be a choice between $x - \delta$ for sure and L. At the third row it would be a choice between $x - 2\delta$ for sure and L, and so on until the sure thing equals y. Note that the sure amount at the last row is equal to the worst possible payoff yielded by the lottery L. As δ is kept fixed, the number of decision rows depend upon the range of prizes of the lottery option. Our statistical analysis of treatment effects shall control for that; as we discuss in section 2.5.2., the variation of decision rows is a change of framing across some risk tasks that, however small, may influence elicited measures.

In order to clarify the elicitation procedure underlying each risk task, consider the following example in Figure 2.1 of what a risk-elicitation task will look like. The task consists of eliciting the cash equivalent of the lottery $L(8.00, 1/5; 4.00, 4/5)$, where the fractions indicate the probabilities of winning, and the integer numbers indicate the winning[14]. To this end, the subject would face the following set of pairwise problems in a table:

Each decision row on the screen constitutes a choice problem, which is to choose between option A, a sure amount of money, or option B, the lottery. They are asked to indicate their preferences for each choice problem. As one proceeds down the table the sure amount of money decreases and becomes

[14]Probabilities in the lotteries are replaced by numbers. For example: if the lottery has two prizes, A and B, with probability 0.3 and 0.7, respectively, then A is paid if the chip drawn is numbered 1 to 30, whereas B is paid if the chip is numbered 31 to 100. For more details on how lotteries are implemented see section 3.4.

Figure 2.1: Illustration of a risk elicitation task

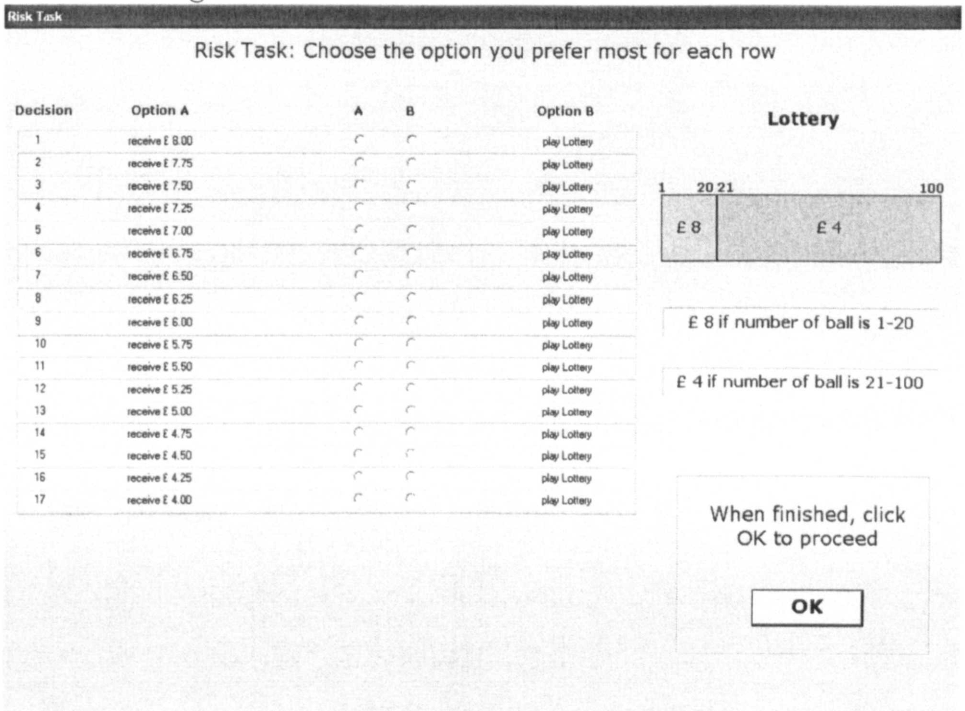

less and less attractive when compared to the expected value of the lottery (in this case £4.80). Note that in this Table, a risk-neutral individual, for instance, should take option A for the first thirteen rows cross over to the risky option when option A offers £5.00. A mildly risk-averse individual is expected to choose option A for the first thirteen rows, switching over to the lottery at some row thereafter. Even an extreme risk averter is expected to switch over to the lottery at some point – at the bottom row, to be precise, since the worst lottery prize is at least as good as the sure money at that row.

Note that, provided a subject starts by choosing A and switches once, task responses can be reduced to a closed switching interval within which the certainty-equivalent of the lottery option falls into[15]. For instance, if a

[15]Note that asking subjects for reporting an indifference point – in our risk task, an

subject crosses over to the risky option when the sure option offers £6.00, choosing the lottery thereafter, then we know that the sum of money that is regarded as good as the lottery lies between £6.00 and the sum offered in the next row, which is £5.75. We shall use the switching interval midpoint as our operational concept of the observed certainty-equivalent[16].

It is relatively common in this type of task to have some subjects switching back and forth between options as they proceed down the menu of choices. Our software, though, did not permit a subject to have multiple switch points. When one chooses option A, say £5.50, over option B, the lottery, the computer assumes that option A is also preferred over the lottery whenever it is offering a sum larger than £5.50, filling-in the buttons accordingly. Likewise, when the lottery option is chosen over a given amount of money, say £4.00 the computer also assumes that the lottery is preferred to the sure amount when it is less than £4.00. Before proceeding to a new risk task, subjects could change their choices and adjust their switching point as many times as they wished.

Some may argue that this device, by forcing a single switch point, is forcing a monotonicity that subjects's preferences may not have[17]. But

amount of money which makes the subject indifferent between receiving it with certainty or playing the lottery option – would turn out being a similar multiple-price list type of task if the truthfulness of reported indifference point were to be tested under incentivised conditions.

[16]This interval is quite narrow (0.25), which makes the midpoint of the switching interval a more refined estimate of subjects' money-equivalent point of lottery option sin each risk task. We keep this variation between sure amounts of money from decision row to decision row constant across all risk tasks.

[17]There is some evidence, yet, that an enforced single switching point does not yield significantly different elicited values relative to a multiple-price-list procedure without such feature (Andersen *et al.* , 2006b).

this feature has four potential advantages. First, it may help to alleviate boredom; subjects who understood it realise that they do not necessarily need to pick an option at every decision row. Second, it gives complete flexibility while embodying a feature that those who understand and take the task seriously would want to obey. Third, it may also simplify the decision problem, helping subjects to focus attention on the provision of a switch point that is as accurate as possible. Fourth, and last, it allows a more refined elicitation of certainty-equivalent from the *entire* sample[18] by eliminating the appearance of non-useable responses, since they violate monotonicity.

As we shall discuss later, task responses straightforwardly yield measures of risk aversion: once the switch point is elicited, simple indicators of attitudes to risk can be obtained either by calculating the bounds on the CRRA coefficient of an individual's utility function implied by the switch point (Holt & Laury, 2002), by computing risk premia (assuming the certainty-equivalent is the midpoint of the switching interval), or by simply comparing that "switch" point to the one predicted under risk neutrality. While these transformations of responses are risk aversion degree-preserving, the last two approaches have the advantage of yielding indicators of risk aversion that are not conditioned upon particular functional forms for utility functions or particular models of preferences.

[18]Provided, of course, a subject's choices do not violate first-order stochastic dominance, which can happen if she prefers a given option over the other in all decision rows in a given risk task.

2.3.2 Experiment's task stages

Having given an account of the risk elicitation procedure, we now describe how the experiment is organised around it. The experiment consists of three stages: (1) first risk-elicitation stage, (2) cognitive stage, and (3) second risk-elicitation stage, respectively. In each of these stages subjects perform the following types of task: risk tasks in the risk-elicitation stages and a IQ-type of test in the cognitive stage. To avoid confusion, we use "risk task" to mean a whole table with decision problems of the kind illustrated in Figure 2.1; and "choice problem" to mean a single decision problem in each risk task, so defined.

Stage I: Risk tasks

In the first stage, the subjects face 6 risk tasks in a sequence. In each of them, they are confronted with several pairwise choice problems, each posed in a given row of a table. They are asked to indicate a preference for one of the two options in every decision row. To ensure an incentive compatible elicitation mechanism and avoid income effects, they are told that one of the rows from each risk task they perform in the experiment (they are unaware of how many there are) will be randomly selected at the end of the entire experiment, being informed that only one of the selected decision rows will be used to determine their earnings (more about this in section 2.4).

As illustrated above, each row in a given risk task is a choice between

Option A, a sure amount of money, and Option B, a lottery. Table 2.1 below presents the set of lotteries used in each of these risk tasks in the order they are presented[19]:

Table 2.1: Lottery option per risk-elicitation task

Lottery	Payoff 1	Pr(Payoff 1)	Payoff 2	Pr(Payoff 2)	EV	Rows
L1	8	0.2	4	0.8	4.8	17
L2	9	0.2	3	0.8	4.2	25
L3	6	0.4	3	0.6	4.2	13
L4	9	0.3	4	0.7	5.5	21
L5	16	0.2	10	0.8	11.2	25
L6	6	0.4	3	0.6	4.2	13

Note that the number of decision rows each risk task contains, which is listed in the rightmost column in Table 2.1, varies across tasks. This is so because we keep the variation of sure money offered across decision rows in the risk tasks constant and the range of prizes vary across tasks[20]. In each risk task, the ordered list of certain amounts of money starts at the highest prize of the lottery option and decreases by £0.25 at each row down until the Option A offer equals the lowest payoff of the lottery. Hence, the different number of decision rows stems from the different range of lottery prizes across risk tasks.

The lotteries we use have four noticeable characteristics. First, they are all binary lotteries. Second, they only involve strictly positive outcomes.

[19]In order to test for order effects, we randomised the (pre-defined) position in which L2 and L5 would appear in the sequence – either in the second or the fifth risk task. Thus, to be precise, roughly half of the subjects face the sequence as shown in Table 2, while the remaining subjects face the same sequence but with L5 and L2 in, respectively, the second and fifth risk tasks instead.

[20]We control, however, for potential effects of the number of rows on elicited measures of risk aversion in our analysis.

Third, two of them are identical – risk tasks L3 and L6. The reason for this is that by making subjects face the same risk task twice, we can investigate short-term stability of risk preferences. Fourth, payoffs offered by lotteries L2 and L5 differ by £7.00 (i.e. L5=L2+7), which exactly matches the small scale change in wealth induced by the experiment under some treatments, as explained below.

The Cognitive stage

After completing a sequence of six risk tasks, subjects are then asked to complete a timed cognitive test. They have twelve minutes. They are told that their answers to these questions have no effect on their earnings in the experiment.

The cognitive test has three purposes. First, to allow risk attitudes to be related to cognitive ability (see Chapter 3). Second, to allow the small-scale wealth increment to be framed as a reward for completing the test. The idea is then to use this test as an "endogenous" treatment administration route: depending on the treatment condition the subjects were randomly assigned to[21], they learned that a money reward, for submitting a complete set of answers to the test, is guaranteed at the end of the experiment. This way, we want to induce them to think that the reward was "earned" rather than received as a "gift" from experimenters. Third, to crowd out subjects' working memory: as the same lotteries will be faced in a later stage task of the experiment, by going through a cognitive test-type of task, subjects'

[21] We come back to this issue in the "Treatment" subsection below.

working memory is likely to be loaded with new information; this makes less likely that they will spot the equivalence between first and second round of risk tasks, which might cause them to guess that the experiment tests for consistency, and respond accordingly (see Bertrand & Mullainathan, 2001).

Stage II: Risk tasks

In this stage subjects are asked to complete the same sequence of six risk tasks they faced before – though subject are not told this. They are told that they work just like the risk tasks they completed before: each risk task consists of a set of choice problems, in each of which they face two options: Option A, to receive an amount of money, and Option B, to play a lottery. All they need to do is to indicate the option they prefer most for each decision problem. We shall refer to this stage as the second-risk-elicitation stage.

Whether this stage is performed straight after the cognitive stage rather than at a second session taking place one week later depends upon the time delay treatment condition the group of subjects is assigned to. This brings us to the description of the treatment conditions under which risk attitudes are investigated.

2.3.3 Treatments

The major interest of this experiment is in ascertaining whether a small-scale change in wealth affects risk attitudes. But two other relevant related issues are also investigated as a by-product of our experimental design. One is whether such wealth effects are affected by, or depend on, how long the change in wealth has been expected for. The other is what we term here the equivalence of "inside" and "outside" money. This refers to the effects on risk preferences of different forms of introducing an increment – commonly added to the set of prizes of a lottery ("inside"), or simply given to the subject ("outside").

To this end, we have two treatment conditions:

(i) **Increment treatment**: here we manipulate the the money reward, say Δw, that subjects are given for completing the *Cognitive stage*. Δw takes one of two values: £0 or 7.00, which will be denoted by *zero* and *nonzero* increment conditions. The experimentally induced increment is modest, but it is larger than the expected value of almost all lotteries used in the risk tasks.

(ii) **Time treatment**: here we manipulate the length of the delay, say Δt, between the *cognitive stage* and the *Stage II* (Risk tasks)[22]. Δt is either around three minutes or an entire week. Henceforth, they are denoted by *instantaneous* and *delayed* conditions. Note that for subjects assigned to a *nonzero* increment condition, we are manipulating for how long the increment has been anticipated for before subjects face the second stage of

[22]To be accurate, the delay is between the end of the *Cognitive stage* and the beginning of *Stage II*.

risk tasks[23].

A visual illustration of the decision setting in the experiment is given by the timeline in Figure 2.2.

Now, given the two variables we manipulate, subjects are randomly assigned to one of the four treatment conditions: I7 (Instantaneous +7), I0 (Instantaneous +0), D7 (Delayed +7), or D0 (Delayed +0). Table 2.2 indicates how the experimental manipulation of time and wealth increment varies across the four treatment conditions investigated.

Table 2.2: Treatment Conditions

Independent variables	$\Delta w = 7$	$\Delta w = 0$
Δt = "Instantaneous"	**I7**	**I0**
Δt = "Delayed"	**D7**	**D0**

Note that subjects assigned to the treatment conditions in which $\Delta w = 0$, $I0$ and $D0$, are used as a control group, as we can use their responses across stages to control for differences in risk attitudes elicited at Stages I and II that are genuinely induced by $\Delta w = 7$ from those differences induced by inherently imprecise preferences (Butler & Loomes, 2007), stochastic choices (Loomes & Sugden, 1995; Loomes, 2005), or even changes in individual circumstances[24].

Information sets

Subjects in all treatments cannot infer either Δw or Δt from their

[23]They receive this money, as well as earnings from the risk tasks, at the end of the experiment only.

[24]This latter possibility cannot be ruled out for the "Delayed" cases.

Figure 2.2: Timeline for a given lottery L

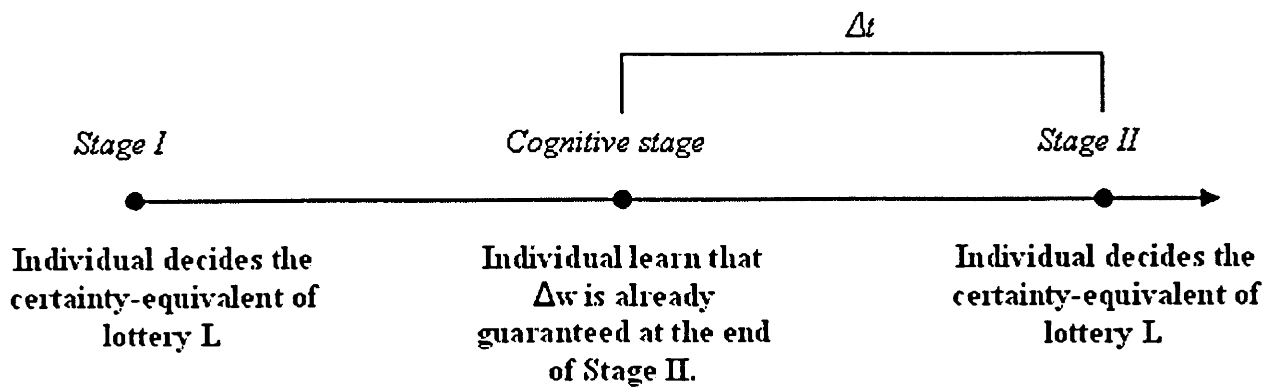

information until the end of the cognitive stage (see Section 2.5. for details). They also do not know about the existence of a second risk-elicitation stage until the very start of it – this can be a few minutes or a week after the end of the cognitive stage.

Inside-Outside money

To understand how the format taken by small-scale changes in wealth can affect risk attitudes, it will be instructive to examine the measure of risk aversion generated by choices in risk tasks L2 and L5. Lottery L5 is an increased version of L2: the prizes of L5 are precisely £7 larger than the prizes of L5. Contrasting the certainty equivalent of L2 and L5 in the first stage, before any manipulation of Δw takes place, provides insights into what we call "inside" money effect on risk attitudes, since when facing L5 it is *as if* subjects have been endowed with £7 relative to when they faced L2. On the other hand, in the second stage when subjects in some treatments are actually "endowed" with £7, to contrast responses to L2 and L5 before and after the increment allows us to test for the "outside" money effect; also known as "house money" effect, it refers to individuals' tendency to be relatively less reluctant to undertake risks after prior gains. Experimental evidence of "house money" effect has been found in individual (Thaler & Johnson, 1990) and market settings (Ackert *et al.* , 2006).

From a EU theoretical standpoint[25], however, the "inside" and "outside" effect should be equivalent; in terms of final consequences, $L2$ after the

[25]Specifically, for a EU maximiser who has a utility function defined on wealth or experimental income.

increment is not different from $L5$ before the increment. Thus, comparing the "inside" to the "outside" effect tests for a different thing: that the framing of the increment does not affect the individual's choice decisions.

2.3.4 Resolution of risk and payment

Payment is made according to the Random Lottery Incentive System. Subjects are informed prior to responding to the problems that, once they they have responded to all choice problems, one of the problems will be randomly selected and their winnings determined by the option they chose. This may involve a resolution of risk in the event the option chosen is a lottery rather than a sure thing.

This random-lottery procedure, by which several decision problems are faced but the subject is paid the outcome of only one of them, has been extensively used. It provides an incentive-compatible elicitation mechanism both under EUT and PT, in the sense that subjects are incentivised to attach genuine valuations to the lotteries they face which reflect their true preference ordering over the set of pairwise-options (Cubitt *et al.* , 1998a). Some may have concerns, though, that incentives may be importantly diluted when only one decision out of many is actually used to determine payment. However, several experiments have brought evidence that subjects' responses in random-lottery experiments are not significantly different from their responses in single-task experiments (Cubitt *et al.* , 1998a).

Even though subjects are unaware of the number of risk tasks they will perform, we have decided to use a two-stage randomization in order to maximize engagement in the task: subjects are told that the computer will pick one decision row from each risk task they perform, with the understanding that the computer will use a special randomization device that makes all decisions rows from the task equally likely to be chosen; so, for each risk task, a decision row is actually selected. Then, a physical device is used to determine the risk task, and so the decision row used to determine their earnings[26]. In what follows, we describe in more detail how the whole process is done.

After all tasks are completed, a table with 12 rows is displayed on the computer screen. Each row represents a risk task, and for each risk task there is a spin button: by hitting each of them, a decision row is selected at random from the relevant risk task. The screenshot in Figure 2.3 below shows the screen after subjects had hit spin for each case. Then to select one decision row, each subject draws a ball from a bag containing balls individually numbered from 1 up to the number of risk tasks they perform (12). Once they select the ball, and so the risk task to be used for real, the computer screen displays the two options from the relevant row and the choice the subject has made between them.

If they have chosen Option A, they receive the amount of money it

[26]We use a physical device (ball-drawing) to give transparency to the second-stage of the randomization process, as subjects could view computerized randomization as fixed. Such a suspicion would be unlikely at the first stage as subjects can usually see that, of the 12 rows chosen, some favour them and others do not. It is favourable when a relatively low numbered row is selected and unlucky when a high-numbered row is selected. Hence, Fig. 2.3 below illustrates the normal case where some spins are favourable and others not so favourable.

specifies, whereas if they have chosen Option B, they play the lottery. Risk is resolved by drawing a chip from a bag containing 10 numbered chips and receiving the payoff according to what the lottery specifies [27]. Note that subjects who happen to be assigned to one of the nonzero treatment conditions (I7 or D7) know beforehand (by the end of cognitive test) that £7.00 is already guaranteed. This money, nonetheless, is only to be paid at the end of the experiment along with earnings from the risk tasks.

2.3.5 Administration

A total of 138 subjects were recruited on a first-come first-served basis to take part in the experiment, divided in sessions involving 12-16 people at a time[28]. They signed-up for the experiment with the understanding that the experiment would have two sessions, one-week apart. They were also told that by signing-up for the first session of the experiment, they would be automatically signing-up for a second session, to take place one week later at the same time of the first session they choose to come.

We pre-randomised the combination of delay and increment treatment conditions to be assigned to each experimental session, so to all subjects in the session. Subjects in a given session were randomly seated at individual computer terminals in our laboratory. An individual ID number was entered for each subject, and this was used to record their decisions

[27] Probabilities in the lotteries are replaced by numbers. For example: if the lottery has two prizes, A and B, with probability 0.3 and 0.7, respectively, then A is paid if the chip drawn is numbered 1 to 3, whereas B is paid if the chip is numbered 4 to 10.

[28] Participants were recruited using ORSEE (Greiner, 2004).

Figure 2.3: Random risk task selection - Screen shot after subject has hit the “Spin” button for each risk task

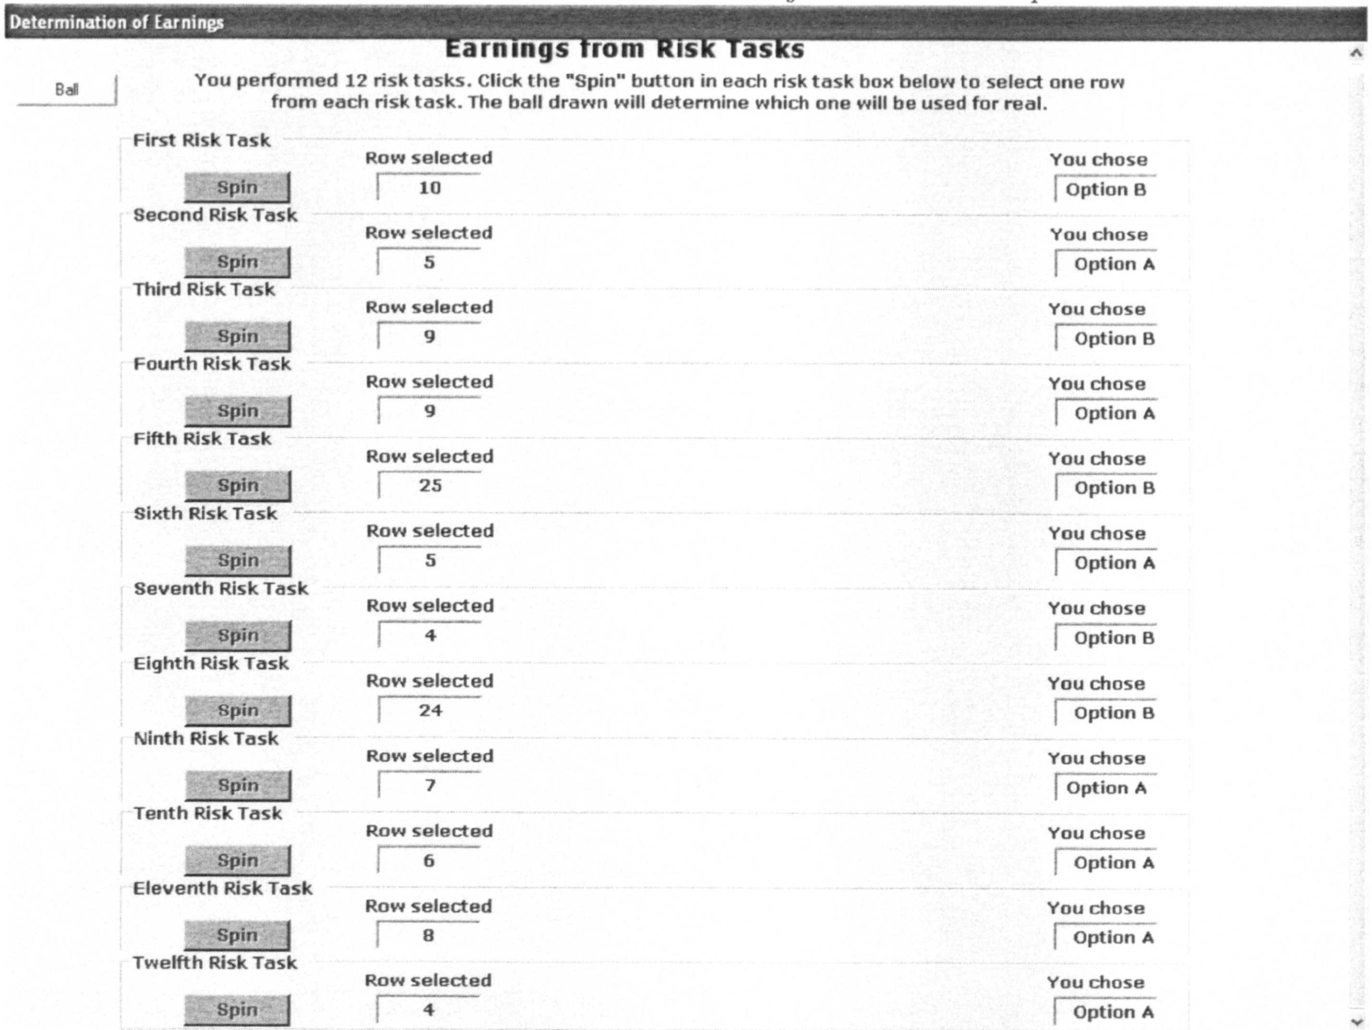

throughout their experiment. They were told at the beginning of the session that although there were many people in the room, their earnings would not depend on what others did. They were also reminded that all of them had signed up to a second session which would take place in one week's time. We told them that in the experiment they would be asked to complete risk tasks and multiple-choice tasks, without mentioning how many of them there were. Instructions for each task stage were handed out one at a time (see Appendix B). Subjects were asked to read them through with the experimenter, who read them aloud. They experienced a risk task trial round before the "real" ones; the main purpose of this was to demonstrate the feature of the software that "enforces" a single switch point. Throughout the session, there was an experimenter in the room to answer any questions and to ensure that subjects knew how to run the computer program used to present the risk tasks and the cognitive test.

While subjects took part in the experiment with the understanding that it would have two sessions, only subjects in the "Delayed" treatments, D7 and D0, had actually to return to the second session. [29]. For groups of subjects in the "Instantaneous" treatments, I7 and I0, the experiment was completed in one single session: they were told at the end of the first session that the experiment was completed and that there was no need to turn up for the second session – scheduled for one week later. Reminders of this were sent a few days later to all of them. Subjects in the "Delayed"

[29]They did not know this until the end of the first session, so there is no reason to believe that this could cause a selection bias.

treatments, in turn, were sent reminder e-mails about the second session they signed-up for[30]. In the second session, subjects were told on arrival that they would be asked to complete risk tasks just like the ones they completed in the first session. Instructions were handed out and subjects were asked to read them through with the experimenter, who read them aloud.

All subjects were paid at the end of the experiment. The average earnings for subjects in the "non-zero" increment conditions (I7 and D7) were £14.61, with payoffs ranging from £10 to £23. Among those in the "zero" increment conditions (I0 and D0) the average earnings were £6.70, with payoffs ranging from £3 to £16.

2.4 Theoretical predictions

This section presents theoretical predictions for the effect of our experimental treatments on subjects' task responses.

Our analysis will focus on predictions for three experimental manipulations. First, that of the money increment between risk-elicitation tasks. Second, that of the delay between the first and second stages of risk-elicitation tasks[31]. Our analysis will also focus on the "inside-outside" money issue, a "built-in" manipulation of the form taken by the money

[30]The large majority of subjects in the Delayed Treatments turned up for the second sessions; and the turn-up rates are relatively similar between the *zero* and *nonzero* increment conditions – 24 and 29 subjects out of 30, respectively.

[31]For subjects in the increment treatments, I7 and D7, this is also a manipulation of the delay between the announcement of the guaranteed money increment and its receipt.

increment.

To define a subject's task response, we must first describe the structure of their decision problem; despite some generality, it closely resembles the problem subjects face in our experiment. To begin, let L be a lottery which has two possible outcomes: x, with probability p, and y, with probability $1-p$. Both outcomes are positive. The decision problem has two stages. At stage I, the subject gives the certainty-equivalent of lottery L. Then an exogenous increment of either £0.00 or £7.00 is announced. This increment is to be received at the end of Stage II. Throughout our analysis we assume that there is no wealth-relevant news between stages. At stage II, the subject gives the certainty-equivalent of lottery L. Thus, a subject's task response at stage i ($i \in \{1, 2\}$), denoted by $T_i(L)$, is the sure amount of money that, if the task resulted in the receipt of it, would be regarded by the subject at stage i as exactly as good as the task resulting in play of L. Hence, by definition, $T_i(L)$ is the certainty equivalent of a lottery.

In describing the theoretical framework, we will confine attention to Expected-utility Theory. While Cumulative Prospect Theory (CPT) is one of the major alternative theoretical accounts of choice under risk[32], the features that make CPT more general than EUT and potentially better able to explain our data have a limited role in our experiment: *probability weighting* would not make much difference to our analysis of effects of the treatments; and because all lotteries are in the domain of gains, *loss*

[32]For a theoretical analysis of the relationship between risk aversion and the curvature of the utility function in CPT see Schmidt & Zank (2008).

aversion can only play a role if subjects' reference point is located near or above the high prizes of our lotteries.

2.4.1 Expected Utility Theory

We start by characterizing how risky outcomes are evaluated by a decision maker who obeys EUT. Assume now that she has a utility function $u(.)$ whose domain is $(\underline{w}, \overline{w})$, a nonempty interval of wealth levels. Assume that $u(.)$ is strictly increasing, time-invariant, and twice differentiable. This implies that $u(.)$ is a continuous function such that lottery L has a certainty-equivalent. The certainty-equivalent of L, $C(L, w)$, is defined as the amount of money m such that $m \sim L$ at wealth position w, where $\sim$ is a relation of indifference. The amount by which the expected value of L exceeds its certainty-equivalent, $E(L) - C(L, w)$, will be referred to as risk premium. The risk premium depends on w and on L, and henceforth shall be denoted by $\psi(L, w)$. So, if the lottery L has expected value $E(L)$, $\psi(w, L)$ is the maximum reduction in $E(L)$ that an individual with wealth w would accept to make herself indifferent between the lottery L and such amount with certainty, that is

$$u[w + E(L) - \psi(w, L)] = u[w + C(L, w)] = pu(w + x) + (1 - p)u(w + y).$$

By definition, the certainty-equivalent of L, $C(L, w)$, is the subject's task response at a given stage i to a risk task featuring L if her wealth is w. Trivially, then, $C(L,w) = T_i(L)$. Since $\psi(w,L)$ is linearly linked to $C(L,w)$ – for a given lottery, the higher $\psi(w,L)$, the lower $C(L,w)$ –, by providing a prediction as to how a change in wealth, say Δw, affects risk attitudes, as measured by the risk premium, we give a prediction as to how an increment of Δw will change a subject's task response at stage i, which we denote by $T_i(L)$[33].

2.4.1.1 Wealth effects

Thus, consider that the decision maker attaches the risk premium $\psi(w_0, L)$ to L when her wealth level is w_0. Let us assume for simplicity that even when her wealth level is w_0 but she then finds out that her wealth level is soon to be w_1, she attaches $\psi(w_1, L)$ to L as if her wealth level were w_1 – which is very much in the spirit of the asset integration axiom of EUT. Let us also assume that $w_1 = w_0 + \Delta w > w_0$.

Proposition 1 $\psi(w_1, L) \gtreqless \psi(w_0, L)$ *as the decision maker displays increasing, constant, or decreasing absolute risk aversion, respectively.*

Proof. See Appendix A ∎

We can conclude from Proposition 1 that, $T_2(L)$, a subject's task response to a given risk task at stage II after the increment has been administered, may equal, exceed, or fall short of, $T_1(L)$ depending on the form

[33]Hence, this prediction is only relevant to individuals assigned to *nonzero* increment conditions (I7 and D7). Section 4.1.2. examines the case of individuals in *zero* increment conditions (I0 and D0).

of (absolute) risk aversion embodied in $u(.)$. Unless we arbitrarily impose a uniform type of risk aversion over the interval $[w_0, w_1]$, there is, therefore, no unique prediction to how task responses will be affected by the the small-scale change in wealth administered in some treatments of the experiment.

Prediction 1 (Divergence in before- and after-increment task responses when $\Delta w = 7$): *For an expected utility maximizer with an utility function $u(.)$ defined over wealth, after-increment task response, $T_2(L)$, may equal, exceed, or fall short of, her before-increment task response, $T_1(L)$, depending on the form of absolute risk aversion embodied in $u(.)$.*

Thus, comparison of task responses (certainty-equivalents) between stages provides an experimental test of the form of absolute risk aversion for the scale of change in wealth considered here.

2.4.1.2 Time effects

The effect of our delay treatments can be divided into two distinct effects: the pure effects of time elapsed between risk elicitation stages, and the effect of time delay on the wealth effect – that is, the interaction of money increment treatment with the delay treatment.

With regard to time effects, EUT is mute implicitly assuming time-stability, that is, that an individual's risk preferences are not perturbed by a short passage of time. We know, however, that some experimental studies report that people, when asked to state their preferences over pairwise choices on two different occasions, reveal different choices on each of them (Hey & Orme, 1994; Camerer, 1989; Loomes, 2005). We shall use the choices from the control group (zero increment conditions) of our treatment to provide a further test of the time-stability of subjects' preferences over a few minutes ("Instantaneous" condition) and a week ("Delayed" condition).

With regard to delay treatment effects, it is easy to check that from the assumption that $u(.)$ is time-invariant, wealth effects on risk attitudes, whatever they are, are not influenced by the time dimension involved: for a given form of absolute risk aversion, the certainty-equivalent (task response) attached to a lottery L at two different wealth positions, say w_0 and w_1, have a relationship with each other that remains unchanged irrespective of the length of time elapsed between the moment at which w_0 and w_1 are each actually held. Therefore, under EUT a change in task responses across periods should not be influenced by the time elapsed between period 2 and period 3[34]. We conclude this subsection by stating the prediction for the increment/delay treatment interaction.

Prediction 2 (Delay effects on before- and after-increment task responses): *Divergence in before- and after-increment task responses, $T_1(L)$ and $T_2(L)$, respectively, should not be altered by the time elapsed between such periods.*

[34] Assuming there is no wealth-relevant news between periods.

2.4.1.3 "Inside-Outside" money effects

We now consider predictions for the effects of the form taken by the gain of £7.00 on task responses. Loosely put, the "outside" case refers to the effect of the £7.00 increment on task responses, while the "inside" case refers to the effect of adding £7.00 to the lottery prizes on task responses.

To make things more concrete, consider the lotteries L_2 and L_5. L_2 gives £9.00 with probability p and £3.00 with probability $1-p$. L_5 gives £16.00 with probability p and £10.00 with probability $1-p$. For this subsection, we confine our attention to these two lotteries because their payoff difference exactly matches the gain of £7.00. The certainty equivalent of these two lotteries, which will be denoted by $C(L_2)$ and $C(L_5)$, is the risk free amount that gives the same expected utility as taking the lottery. Thus, since $u(.)$ is continuous and the expected value of a given lottery $L(p,x;1-p,y)$ is given by $pu(x)+(1-p)u(y)$, the certainty-equivalent of L_2 and L_5 can be defined as follows:

$$C(L_2) \equiv u^{-1}\big(pu(9)+(1-p)u(3)\big) \tag{2.1}$$

$$C(L_5) \equiv u^{-1}\big(pU(16)+(1-p)u(10)\big) \tag{2.2}$$

Recall that in our decision setting a risk task is faced on two different occasions, Stage I and Stage II, and that the outside increment of £7.00 is administered right at the end of Stage I. Each risk task has a lottery option, and an individual's task response is the sure amount of money that the subject regards as exactly as good as playing that lottery – that is, the certainty-equivalent of the lottery. Let then $T_i(L_j)$ be the task response of an expected utility maximiser for lottery L_j at stage i, where $i \in \{1, 2\}$ and $j \in \{2, 5\}$. Then, it follows trivially that

$$T_1(L_2) = C(L_2) \tag{2.3}$$

$$T_1(L_5) = C(L_5) \tag{2.4}$$

The set of task response comparisons is illustrated in the table 2.3 below. Let us first consider the differences in task responses to L_5 and L_2 within a given stage. This is illustrated in Figure 2.4, for the case of $i - 1$.

Lotteries L_5 and L_2 are evaluated at different points of the wealth/income domain, and differences in their certainty equivalents (task responses, by definition) will depend on the properties of the individual's utility function $u(w)$ regarding her willingness to take risks as she becomes wealthier.

Proposition 2 *Suppose $u(.)$ is concave (convex). Suppose $C(L)$ is the certainty equivalent of the lottery $L(x, p; x', 1-p)$. Let $L'(x+k, p; y, x'+k, 1-p)$ be a lottery constructed from L by increasing each prize of L by the amount of $k > 0$. Let $C(L')$ denote the certainty equivalent of L'. Then $C(L') -$*

Table 2.3: Task response comparisons: within- and between-stage differences

Lotteries / Stage	**I**	**II**	**Between-stage differences**
L_2	$T_1(L_2)$	$T_2(L_2)$	$T_2(L_2)$ - $T_1(L_2)$
L_5	$T_1(L_5)$	$T_2(L_5)$	$T_2(L_5)$ - $T_1(L_5)$
Within-stage differences	$T_1(L_5)$-$T_1(L_2)$	$T_2(L_5)$ - $T_2(L_2)$	

Figure 2.4: Within-stage differences in the certainty-equivalent of L_5 and L_2

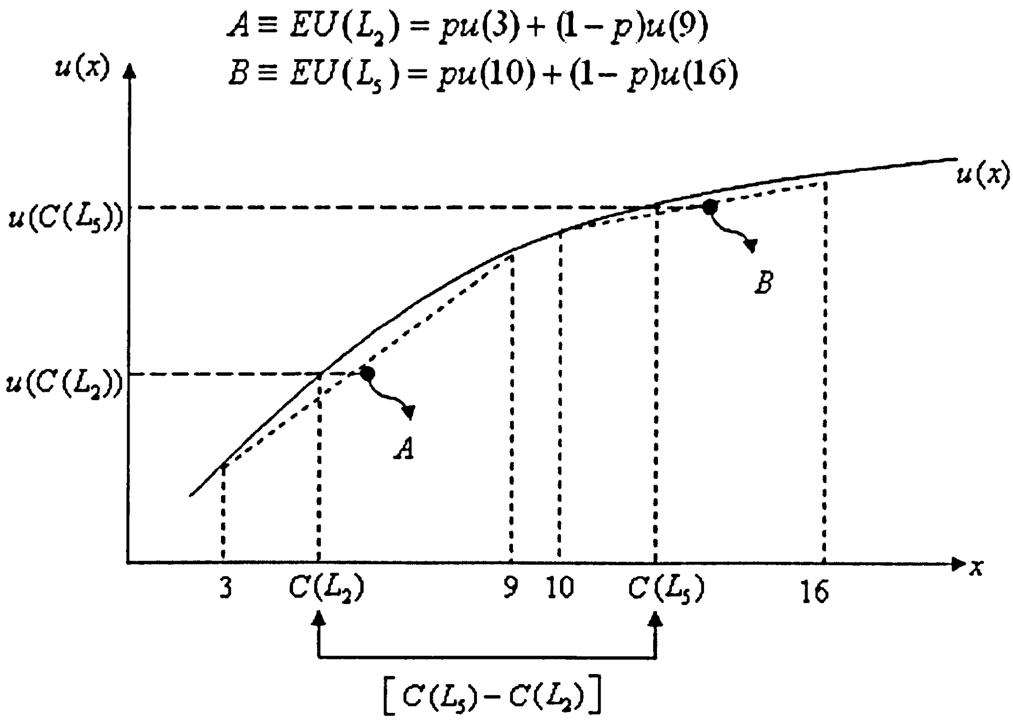

$C(L) \gtreqless k$ if, and only if, $u(x)$ exhibits decreasing/constant/increasing absolute risk aversion (increasing/constant/decreasing risk lovingness).

Proof. *See Appendix A* ■

Proposition 2 states an intuitive result: the difference of an individual's attitudes toward risk as a sure amount of money is added to all consequences of a given risk depends on properties of the utility function concerning risk-bearing behaviour as she becomes wealthier. The following prediction follows trivially from Proposition 2.

Prediction 3 (Inside money effect on task responses): *Differences in tasks responses to L_5 and L_2 in a given stage i, $T_i(L_5) - T_i(L_2)$ for*

$i \in \{1,2\}$, may equal, exceed, or fall short of the £7.00 common increase in prizes of L_2 depending on the type of absolute risk aversion/lovingness exhibited by the utility function.

Hence, there is no unique prediction for differences in task responses (certainty-equivalents) to L_2 and L_5 within a given stage. But the sign of such differences provide an alternative way of inferring properties of risk preferences as wealth changes.

An equivalent prediction holds for between-stage differences in task responses to a given lottery L_j ($j \in \{2,5\}$). To see why, consider how the subject responds at the second-stage (after-increment) to, for instance, L_2. The task is to set $T_2(L_2)$ such that receiving the outside increment of £7.00 and $T_2(L_2)$ for sure is exactly as good as receiving the outside increment and playing L_2. Thus, $T_2(L_2)$ solves

$$u\big(7 + T_2(L_2)\big) = pu(9+7) + (1-p)u(3+7). \tag{2.5}$$

By inverting both sides of (2.5), and using (2.2),(2.4) , it follows that

$$7 + T_2(L_2) = T_1(L_5) = C(L_5) \tag{2.6}$$

At stage I, in turn, task response to L_2 is the sure task reward that the subject regards as exactly as good as the task resulting in play of L_2. Trivially, then,

$$T_1(L_2) = C(L_2) \tag{2.7}$$

Note then that by combining (2.6) and (2.7), we can conclude that

$$T_2(L_2) - T_1(L_2) = [(C(L_5) - C(L_2)) - 7].$$

But the certainty-equivalent difference on right-hand side of this equality is just a different way of expressing the difference in tasks responses to L_5 and L_2 in a given stage i – a result that can exceed, equal, or fall short of 7 depending on whether risk preferences, as captured by the utility function $u(.)$, exhibit decreasing/constant/increasing absolute risk aversion (risk lovingness). Thus, we can make the following

Prediction 4 (Outside money effect on task responses for $\Delta w = 7$): *For an expected utility maximizer with a concave (convex) utility function $u(.)$ defined over wealth, between-stages differences in tasks responses to L_j ($j \in \{2, 5\}$), $T_2(L_j) - T_1(L_j)$, may equal, exceed, or fall short of 0 depending on the type (decreasing/constant/increasing) of absolute risk aversion (lovingness) exhibited by the utility function.*

Lastly, let us consider the prediction for a more interesting case: the across-stage differences, that is, the differences between task response to L_5 in the first stage (before increment) and task response to L_2 in the second stage (after increment).

When facing L_2 at stage II, the individual view the task as setting $T_2(L_2)$

such that receiving the outside increment of £7.00 and $T_2(L_2)$ (certainty-equivalent) for sure is exactly as good as receiving the outside increment and playing L_2. Thus, $T_2(L_2)$ solves (2.5), where the argument inside the utility function on the left-hand side of the equality is the certainty equivalent with certainty plus the £7.00 increment, and the argument inside the utility function on the right-hand side of the equality are the prizes of L_2 with the £7.00 increment added. It follows from (2.5), by inversion of both sides and using (2.2) and (2.4), that

$$7 + T_2(L_2) = C(L_5) \tag{2.8}$$

Note that neither (2.3), (2.4), nor (2.8) requires any assumptions about attitude to risk. Yet, combining all, they are sufficient to yield the following equality

$$T_1(L_5) - T_1(L_2) = T_2(L_2) - T_1(L_2) + 7$$

which, by eliminating $T_1(L_2)$ from both sides, can be reduced to the following:

$$T_2(L_2) + 7 = T_1(L_5) \tag{2.9}$$

which will be termed here as the *inside-outside equivalence condition*, whereby the sure sum which is just as good as getting the increment and playing L_2 should equal the sure sum which is just as good as playing L_5 without the increment. This is illustrated in Figure 2.5. Note that since $9+7=16$ and $3+7=10$, each figure, while representing the evaluation of different lotteries, is dealing with the same locus of points in the utility space.

From (2.9), we can now derive the following prediction:

Prediction 5 (Inside-Outside money equivalence condition): *For an expected utility maximizer, differences between task response to L_5 in the first stage (before increment) and task response to L_2 in the second stage (after increment) should be such that the following condition holds:* $T_1(L_5) - T_2(L_2) = 7$.

This prediction states that, within EUT, differences in task responses induced by an wealth increment of z should not be affected by the framing used to introduce such wealth increment.

2.5 Results

In this section, we present the results of our experiment, confronting the predictions for each experimental manipulation with the data. The repeated nature of our experiment will also allow us to analyse the short-term stability of risk preferences.

Our data sample set consists of 106 subjects, whom were about evenly divided among treatment conditions as follows: 26 assigned to I0, 27 assigned to I7, 24 to D0, and 29 to D7 [35].

[35]The sample used in our regression analysis, when the model used to estimate risk behaviour includes controls for treatment conditions and income class, is slightly different (102 subjects) since we excluded some subjects with missing income data.

Figure 2.5: Utility space: "Inside-Outside" equivalence condition

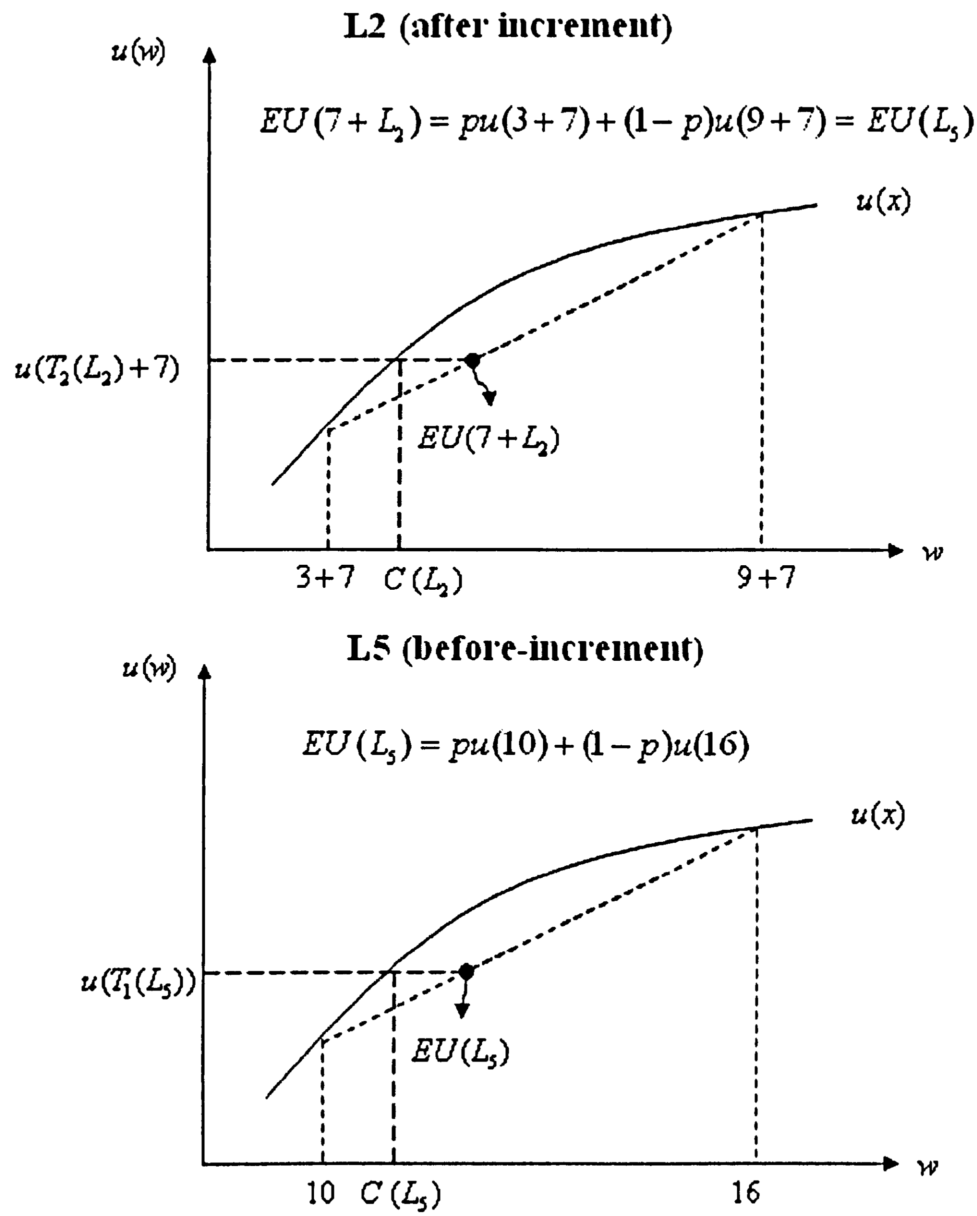

2.5.1 Elicited risk attitudes

We start by examining the overall distribution of risk attitudes elicited in the experiment.

A subject's attitude to risk in a given risk task featuring lottery L is measured here by the risk premium $R(L)$, which is the difference between the expected value of the lottery L $(E(L))$ and the certainty-equivalent the subject assigns to L $(C(L))$; i.e., $R(L) = E(L) - C(L)$[36]. By taking into account the expected value of each lottery, this measure is to some extent "normalised" across lotteries with different stakes, making individual's elicited risk preferences readily comparable across risk tasks.

Figure 2.6 provides charts with histograms of risk premia (in British Pounds) in each risk task. Data are pooled across treatments. Each risk task's histogram is overlaid with a normal density distribution with same mean and standard deviation of the data. The distribution of risk premia is clearly skewed to the left in all risk tasks. This indicates that a high proportion of subjects were displaying negative risk premia; hence, some degree of risk-loving behaviour.

Indeed, the majority of our subjects were systematically not risk averse throughout the risk tasks. Table 2.4 shows fractions of subjects in each distributional "class" of risk preference over the entire set of risk tasks. A subject is placed at class $[n, m]$ if she were risk averse in n risk tasks and risk neutral/loving in m, where $n + m = 12$. Very few displayed risk aversion in more than half of all twelve risk tasks. While the set of histograms in Figure 2.6 only tentatively suggests that our subject sample is characterised by some degree of risk-lovingness, Table 2.4 shows, for instance, that 77.36%

[36]From now on, $C(L)$ is taken to be the midpoint of the switching interval.

Figure 2.6: Histogram of risk premia, by risk task in all treatments

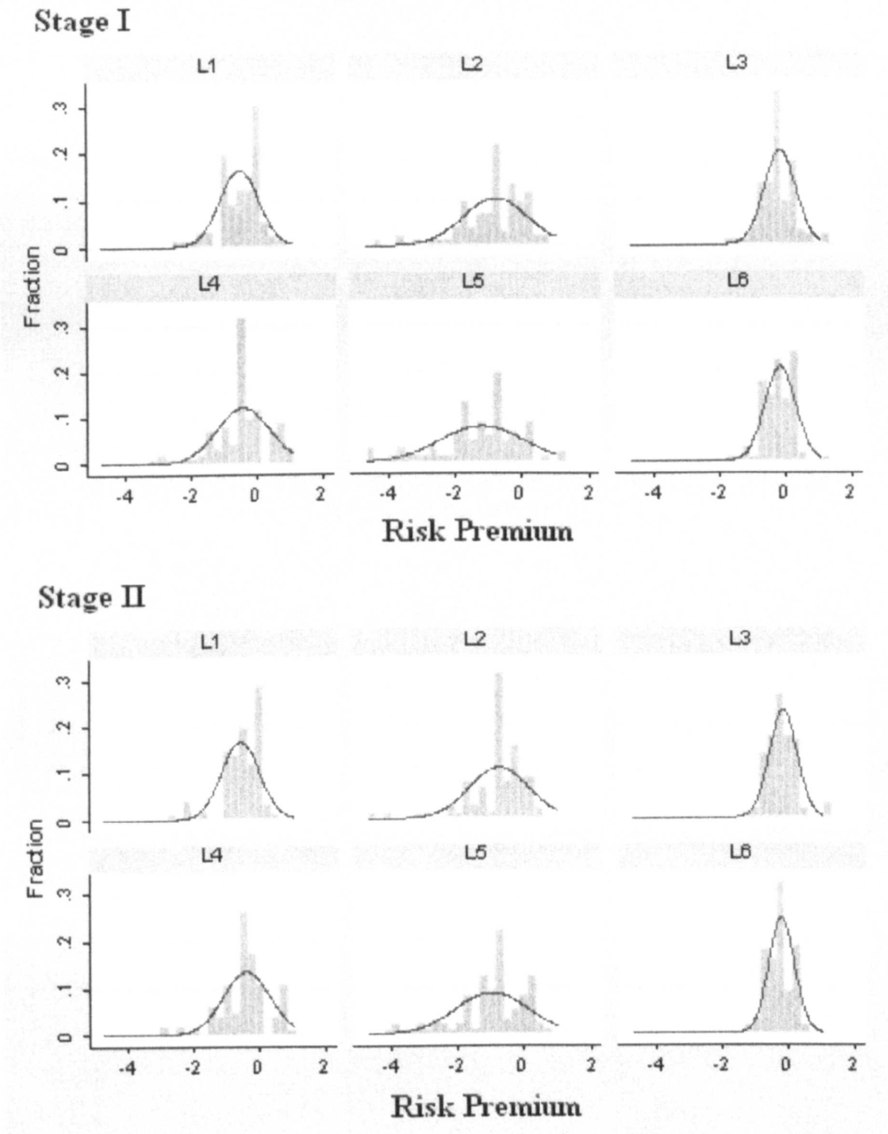

of all individuals in our experiment made either risk-neutral or risk-loving choices in at least 3/4 of all risk tasks. Less than 5% were systematically risk averse in more than half of the risk tasks.

We claim that raw data presented in Figure 2.6 and Table 2.4 support the following

Finding 1: *The great majority of subjects' choices exhibit non-risk-averse*

Table 2.4: Distributional classes of risk preferences in all risk tasks

Distributional class of risk preferneces	Frequency	%	Accumulated
[0,12]	47	44.34	44.34
[1,11]	14	13.21	57.55
[2,10]	13	12.26	69.81
[3,9]	8	7.55	77.36
[4,8]	11	10.38	87.74
[5,7]	6	5.66	93.40
[6,6]	2	1.89	95.28
[7,5]	1	0.94	96.23
[8,4]	-	-	-
[9,3]	2	1.89	98.11
[10,2]	1	0.94	99.06
[11,1]	-	-	-
[12,0]	1	0.94	100.00

Note: An individual is placed in class [n, m] if, considering all twelve risk tasks, she displayed risk aversion in n risk tasks out of twelve; and either risk neutrality or risk lovingness in m risk tasks.

behaviour in all risk tasks.

It is worth noting that this pattern of risk attitudes is in contrast to what is observed in most experimental studies that also elicit risk attitudes using the multiple-price list method and often similarly small stakes[37] (e.g., Gertner, 1993a; Holt & Laury, 2002; Andersen *et al.* , 2008, 2006b; Harrison *et al.* , 2007).

It is natural to ask if such distinctly different results are robust and not driven by (1) framing differences, or (2) by anchoring effects. We try here to address each of those worries.

There is some evidence to suggest that difference in risk attitudes does not stem from differences in the framing of risk tasks we adopted. Our

[37]An exception is Bombardini & Trebbi (2005), who find evidence of risk neutrality for small stakes.

risk tasks are indeed framed in slightly different ways, but they are likely simpler to understand. While each pairwise choice problem in our risk tasks involves a choice between a certain amount of money, which decreases as one moves down the table, and a fixed lottery, in other studies subjects are presented with a choice between two non-degenerate lotteries. Those studies closely follow Holt & Laury (2002), where the probabilities of the lottery prizes is the dimension being modified as one proceeds down the list of binary choice problems. Nevertheless, subjects in a laboratory elicitation of risk attitudes, which uses the same procedure used here, also showed a tendency toward risk-loving behaviour (Andersen *et al.* , 2006a)[38].

But one could still be concerned that the contrasting results in risk attitudes were produced by an anchoring effect towards the middle of the table. Since in most of our risk tasks the switch point for a risk-neutral individual lies at the bottom half of the table, responses of subjects who tend to *always* pick a middle of the table decision row to switch at would lead us to observe a great deal of risk-lovingness irrespective of her true risk propensities.

Indeed, if there is a strong bias to the middle row, this would matter for treatment comparisons because it would prevent any treatment effect. While our experiment was not designed to test for bias to the middle, there is no reason to think that a midpoint bias could generate spurious treatment

[38]When it is assumed, however, that cumulative income (earnings in the experiment) is the argument of the utility function, CRRA estimates are consistent with risk neutrality.

effects. If anything, bias towards the middle would make treatment effects hard to observe.

In sum, we found that the subject sample in our experiment is unusually risk loving. Since standard recruitment procedures were applied, this can hardly be attributed to a potential sample bias.

We now proceed to analyse the increment and time treatment manipulations.

2.5.2 Wealth effects

We start by testing whether risk attitudes across stages are affected by an increment of £7. There is no unique prediction if one looks at each task separately. Any sign of the wealth effect is consistent with EUT, as changes in risk attitudes induced by the increment ultimately depend on the form of absolute risk aversion embodied in the utility function. [39].

We investigate this question using a variety of statistical methods. We begin by performing unconditional tests on our measure of risk-taking behaviour in order to compare the group of subjects assigned to the nonzero increment condition with the group of subjects assigned to the zero increment condition.

We test the hypothesis that changes in risk premia of subjects in *nonzero increment* condition subjects are not significantly different from changes in

[39]It seems to be widely believed, though, that an utility function exhibiting non-increasing relative risk aversion would better describe observed risky behaviour. CRRA functions, for instance, are widely used on parametric estimations of risk attitudes (see, e.g., Harrison *et al.* , 2005a, 2007; Andersen *et al.* , 2008)

risk premia of subjects in a *zero increment* condition. Table 2.5 reports the results of Mann-Whitney tests. Tests are performed for each risk task[40], as it is of interest to see whether potential wealth effects on attitudes to risk are robust to risk tasks involving different lottery prizes and probabilities. The results do not show any systematic differences between those who knew £7 was guaranteed at the end of the experiment and those who were not expecting such extra gain.

Table 2.5: Effects of expected monetary gain on attitudes to risk by delay condition (within-subjects)

Risk Tasks	*Overall*	*Instantaneous*	*Delayed*
L1 *(£8,0.2;£4)*	z = 0.585 *p = 0.55*	z = 0.977 *p = 0.33*	z = 0.088 *p = 0.93*
L2 *(£9,0.2;£3)*	z = 1.542 *p = 0.123*	z = 0.012 *p = 0.99*	z = 0.213 *p = 0.83*
L3/L6 *(£6,0.4;£3)*	z = -1.93 *p = 0.23*	z = -3.020 *p = 0.00*	z = 0.926 *p = 0.35*
L4 *(£9,0.3;£4)*	z = -0.827 *p = 0.41*	z = -1.660 *p = 0.09*	z = 0.323 *p = 0.74*
L5 *(£16,0.2;£10)*	z = 0.045 *p = 0.96*	z = 1.278 *p = 0.20*	z = 0.766 *p = 0.44*

Note: Mann-Whitney two-sample test statistic and p-value reported in each entry. We test for the hypothesis that changes in attitudes to risk (variation in risk premia in a given risk task) across stages among *treated* ($\Delta w=7$) and *untreated* ($\Delta w=0$) subjects are not different. Tests are performed on aggregated sample and on sub-samples decomposed by delay condition. H_0: Δrisk premia for subjects in nonzero increment condition = Δrisk premia of subjects in zero increment condition.

Interestingly, by decomposing the sample by delay condition, we discover that the increment of £7.00 did have an effect on risky choices for subjects assigned to the *Instantaneous* time treatment – risk premium de-

[40] $L3$ and $L6$ are pooled as they are identical.

creased, on average, after the increment of £7.00; yet, this holds only for the identical risk tasks $L3$ and $L6$. A possible explanation for such differences between time treatments might be that the news of the £7.00 induces an immediate change in attitude but this has worn off by the following week. However, we have reasons to think this is not the case, as the increment effect is only detected in the identical risk tasks; this suggests that the differences are not capturing the effect of a small-scale change in wealth but, instead, a reduction in "noise" as result of practice. In sum, the unconditional analysis performed seem to support the following

Finding 2: *Elicited risk attitudes do not seem to be affected by small-scale changes in wealth (£7).*

Are these results robust to some individual controls? In order to examine that, we regress individuals' risk premia on individual and structural parameters of the experiment. With the panel data structure of our dataset, we can now look at the same issue not only exploiting the heterogeneity within a given subject's sequence of risk aversion measures, but also controlling for fundamental characteristics of the experiment and some observed demographics. To this end, we will implement the following panel data regression specification:

$$y_{it} = b_1 INCREMENT_i + b_2 DELAY_i + b_3 EXPECTVAL_{it}$$
$$+ b_4 ROWS_{it} + b_5 L1L5ORDER_i + +b_6 LOWINCOME_i$$
$$+ b_7 FEMALE_i + b_8 AGE_i + b_9 POSTGRAD_i + b_{10} OVSCORE_i + u_{it} \quad (2.10)$$

where y_{it}, the risk premium derived from subjects' choices in each risk task, is the dependent variable; the set of regressors mostly include dummies for characteristics of the experiment as well as for subject-specific characteristics:

1. $INCREMENT$ is a dummy variable for whether i received the money increment;
2. $DELAY$ is a dummy variable for whether i is assigned to one of the delayed conditions (D0 or D7);
3. $EXPECTVAL$ is the expected value of the lottery option in the risk task faced in period t;
4. $ROWS$ is the number of decision rows in the risk task i faces in period t;
5. $L2L5ORDER_i$ is a dummy for the order in which the risks involving lotteries $L2$ and $L5$ were faced[41];

[41] We randomised across subjects the order those two lotteries were faced. We did this to test for order effects in relation to the L2/L5 comparisons. Note that, while in

6. $LOWINCOME$ is a dummy equal to one if i said that her average monthly income is less than £1,000; We use this information to control for wealth effects due to income differences outside the lab.

7. $OVSCORE$ is the the overall score in the cognitive test;

8. $FEMALE$, and $POSTGRAD$ are two dummies: they are equal to one if i is female (postgraduate student), respectively.

9. AGE is the i's elf-reported age. u_{it} is a composite error term including a random intercept that captures subject-specific effect and a overall disturbance term assumed to be i.i.d over i and t.

We use a generalized least square random effects estimator to fit (2.10). In Table 2.6, we report the estimation results for this specification. The fact that the coefficient in front of $INCREMENT$ is not statistically significant suggests that risk attitudes, as measured by the lottery risk premium, are not influenced by the scale of the prior gain received. The estimates also reveal that the effect of scaling up the stakes of the lottery option in the risk tasks – reflected in its expected value – is to decrease elicited measure of risk aversion. We shall later follow up this question when analysing the "inside/outside" money framing effects.

Estimates showed that an increase in the number of rows in a risk task tended, on average, to reduce subjects' risk premia. The coefficient in front

a very moderate scale, this randomisation can also be seen as a partial test of order effects in our sequence of risk-elicitation tasks, as a full test for all possible sequence with which the tasks could be faced would be cost prohibitive.

Table 2.6: GLS estimates of a random-effects model, with the risk premium implied by the subject's choices as the dependent variable

Variable	Description	Estimate (Mean Effects)	Standard Error	p-value	Lower 95% Confidence Interval	Upper 95% Confidence Interval
Constant		0.45	0.81	0.58	-1.13	2.03
	Treatments					
delay	Delay condition	0.07	0.11	0.56	-0.16	0.29
increment	Increment condition	-0.12	-1.00	0.31	-0.35	0.11
evalue	Expected value	-0.05	0.01	0.00	-0.07	-0.03
rows	Number of rows (binary choices) in risk task	-0.04	0.01	0.00	-0.05	-0.03
L1-L5 order	Second and fifth tasks order	-0.05	0.12	0.65	-0.28	0.18
	Individual characteristics					
female	Female	-0.06	0.12	0.61	-0.29	0.17
age	Age	0.01	0.04	0.78	-0.06	0.08
postgrad	Taking some postgraduate education	-0.08	0.13	0.53	-0.35	0.18
IncLow	Lower level income	0.01	0.12	0.95	-0.23	0.24
ovscore	Overal score in cognitive test	-0.01	0.03	0.70	-0.08	0.06
σ_u	Standard deviation of individual effect	0.51				
σ_e	Standard deviation of residual	0.74				

Note: 1.188 observations based on 102 subjects. Delay condition takes value 1 if subject is assigned to the treatment conditions in which second stage takes place one week later; takes 0 otherwise. Increment condition takes value 1 if subject is assigned to the treatment conditions in which subjects learn, at the end of the first stage, that £7 is already guaranteed at the end of the experiment; takes 0 otherwise. L2-L5 order takes value 1 if L2 was faced before L5 and 0 otherwise.

of $ROWS$ is negative and statistically significant. Recall that risk tasks with more decisions rows have larger stakes, so the coefficient of the number of rows variable captures the effect of stake size on risk attitudes. This is consistent with the sign of the coefficient of the expected value variable, which also reflects the size of the lottery stakes. The remainder of the variables, including most demographic controls, are not statistically significant. In conclusion, our regression analysis confirms the result presented in *Finding 2*.

2.5.3 Time effects on choice decisions

We now consider the time treatment effects on the risk attitudes of subjects who received the increment of £7.00.

Recall that our time treatment conditions involve manipulating the length of the delay between the news of the reward of the £7.00 increment and second-stage of risk-elicitation tasks: a short delay of a few minutes in the "Instantaneous" treatment, and of one week in the "Delayed" treatment. For an expected utility maximiser with an utility function defined over wealth, £7.00 today and £7.00 in one week's time are, in lifetime terms, equivalent[42]. Thus, we would expect no effect of delay.

Table 2.7 reports statistics for t-test and non-parametric tests for whether second-stage responses of subjects in the "Delayed" and "Instantaneous" conditions are significantly different.

[42]Ignoring interest that might be earned in one week.

The hypothesis that there is not significant difference in risk premia between these treatment conditions cannot be rejected for any of the risk tasks. This supports the following:

Finding 4: *The length of delay between risk-elicitation task stages (one week) does not seem to induce different task responses.*

This finding hardly poses a challenge for a theory that has little to say about the influence of time delay on choice decisions involving identical problems. This result can, however, have relevant implications for an alternative theoretical account of our time treatment effects.

Consider, for instance, Prospect Theory, according to which a risky prospect is evaluated by contrasting its outcomes to a reference-point. On the assumption that a decision maker's reference-point is her expectations about outcomes[43], the length of the delay could induce different responses. One possible reason for that is the time-conditioned process of adjustment of expectations to new wealth-relevant information: even if the change in wealth induces an update in the reference-point (equated to her expectations), this may not take place immediately, and some time will be necessary to fully observe its effects. In this case, the length of the delay should induce a change in the coding rule that describes subjects' behaviour. Therefore, if one week seems a plausible amount of time to unfold the adjustment of expectations, the comparison of "Instantaneous" and "Delayed" treatments

[43]This is similar, but not precisely the same assumption advanced by Koszegi & Rabin (2006).

Table 2.7: Time effects on second-stage responses

	Time treatments		Tests Results	
Risk tasks	*Delayed*	*Instantaneous*	*Mann-Whitney*	*Kolmogorov-Smirnov*
L1 *(£8,0.3;£4)*	-0.56 (0.58)	-0.50 (0.54)	z = -0.298 *p = 0.76*	D = 0.107 *p = 0.99*
L2 *(£9,0.2;£3)*	-0.79 (0.95)	-0.48 (0.65)	z = -0.473 *p = 0.64*	D = 0.131 *p = 0.97*
L3 *(£6,0.4;£3)*	-0.23 (0.32)	-0.17 (0.34)	z = 0.977 *p = 0.33*	D = 0.250 *p = 0.35*
L4 *(£9,0.3;£4)*	-0.50 (0.64)	-0.30 (0.52)	z = -0.463 *p = 0.60*	D = 0.153 *p = 0.90*
L5 *(£16,0.2;£10)*	-0.91 (1.19)	-0.73 (0.98)	z = -1.29 *p = 0.20*	D = 0.177 *p = 0.80*
L6 *(£6,0.4;£3)*	-0.22 (0.34)	-0.13 (0.40)	z = -0.199 *p = 0.84*	D = 0.173 *p = 0.80*

Note: Risk premia means reported in columns of treatment conditions, with standard deviation reported in parantheses. The rightmost columns report parametric tests for time effects on second-stage responses from subjects assigned to the nonzero increment conditions. The null hypothesis is that second-stage risk measures of subjects (risk premium) who were assigned to the nonzero increment condition are not significantly different between "Instantaneous" and "Delayed" time treatment conditions. We test this null hypothesis for each risk task.

should provide a test of the speed of expectations adjustment – conditional, of course, on expectations being taken as the reference-point. But the above finding is open to several interpretations. It suggests, for instance, that subjects do not take the expected outcomes, in our case, the wealth increment announced at the end of stage II, as their reference-point. But it could also be the case that subjects do take their expectations regarding earnings in the experiments as their reference-point; but their expectations are either slowly (longer than one week) or instantaneously adjusted.

2.5.3.1 Time stability of risky choices

We now turn to look at whether choice decisions are stable between risk-elicitation stages, spaced either a few minutes or a week apart. This will be addressed via the test-retest component of our experiment design, as each risk task is faced on two different occasions by each subject. For this particular question, we use raw data: a subject's number of "safe choices", that is, the number of times the subject chose the sure money option before she switches from the sure thing to the lottery option.

Although very small, the length of the time interval over which stability of risk attitude is tested here seems to have relevant counterparts in real-world economic settings: many decisions involving risk in financial markets, which normally involve a great deal of risk, are made on a very short timescale. For example, it is rather common in these markets rapidly switch from buying to selling a given instrument on a time-interval of minutes.

While virtually every model, for instance, of portfolio allocation presumes that risk attitudes are stable over short time, experimental evidence on this question is mixed at best. Using the MPL method to elicit risk attitudes, Harrison *et al.* (2005b) has found that mean elicited CRRA coefficients differ between two experiments around six months apart, though differences in the mean coefficients are not statistically significant[44]. In

[44]But care must be taken in interpreting these results: differences in conditional

other laboratory experiments, though, it has been observed that as much as 25% of subjects reverse their decisions in pairwise choice tasks faced twice within a single experimental session (see Loomes & Sugden, 1998).

While we observed that there was some discrepancy between individual's responses to risk tasks at different stages, this turns out not to be statistically significant. In Table 2.8 we report the percentage of subjects whose risk attitudes, as measured by the number of safe choices, elicited in each risk task were equivalent between stages. We see, for instance, that when there was a short delay between stages ("Instantaneous" treatment), as much as half of subjects made precisely the same decision in some risk tasks.

Although the great majority of subjects did not make *identical* decisions in each stage a given risk task was faced, the distribution in both treatment is to some extent very much concentrated about a mean of zero – that is, no change in decisions across stages. Figure 2.7 shows the set of histograms of changes in decisions (in terms of number of "safe" decisions) across stages for each risk task L_i ($i = \{1, ..., 6\}$); for convenience, the histogram is overlaid with a normal distribution that has the same mean and standard deviation as the data.

However suggestive, inference based on simple comparison of raw individual responses can be misleading with respect to the distribution of

coefficient means, however small, can be misleading as within-subject reversals of risk attitudes may be difficult to perceive if changes across cross-sections happen to be mean-preserving.

Table 2.8: Stability of risk preferences between stages

Lotteries	*Instantaneous +0 (N=26)* %	*Delayed +0 (N=24)* %
L1 *(£8,0.3;£4)*	0.50	0.29
L2 *(£9,0.2;£3)*	0.11	0.21
L3 *(£6,0.4;£3)*	0.50	0.29
L4 *(£9,0.3;£4)*	0.38	0.17
L5 *(£16,0.2;£10)*	0.23	0.21
L6 *(£6,0.4;£3)*	0.42	0.50

Note: Numbers reported in the table are the percentage of subjects assigned to zero-increment treatment conditions whose responses (number of safe choices) were equal between stages, per risk task.

responses at each risk-elicitation stage. In Table 2.9 we present test statistics for the null hypothesis that risk attitudes across stages from subjects assigned to zero increment conditions are not significantly different. We report separate results for each delay condition in order to control for potentially differences in individual characteristics across treatment conditions samples. Two-sided Wilcoxon signed rank and t tests at 5 percent level suggest that there is no significant difference between the distribution of elicited risk attitudes: in most of the risk tasks, they were stable between stages in both delayed and instantaneous treatment conditions in which $\Delta w = 0$; We can only reject this assumption regarding the risk tasks involving lottery L3, and interestingly, when the length of delay between

Figure 2.7: Histogram of changes in number of safe choices across stages by risk task, by time treatment

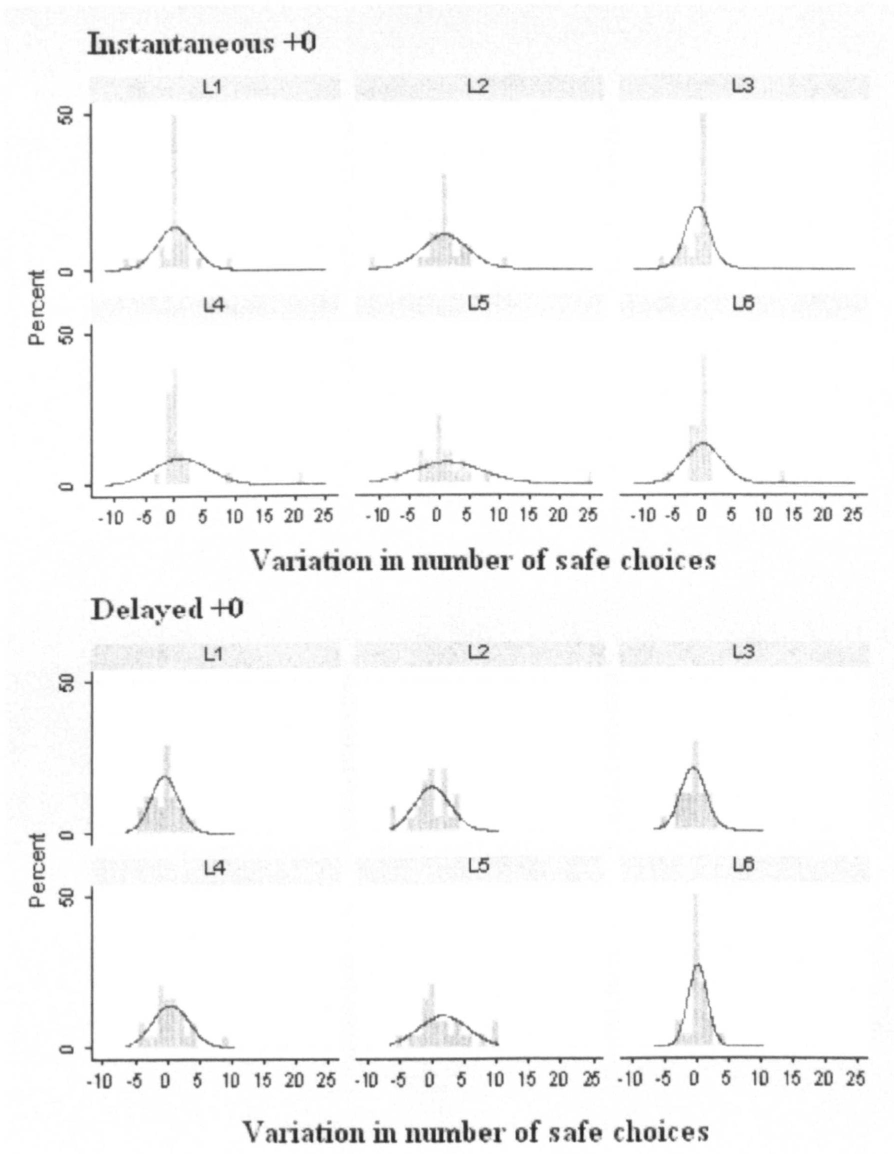

risk elicitation tasks was very short ($I0$ condition). As L3 have the lowest number of decision rows, this suggests that the format of the table may play a role in generating discrepancies in risk attitudes elicited at different periods of time. However, if this were the case in this particular instance, we should have also observed significant differences between responses to

L6, which just replicates L3. But this is not the case, as showed at the bottom row of Table 2.9.

We can thus state the following:

Finding 5: *Individual responses to risk elicitation tasks at Stage I do not differ from responses to the same risk tasks elicited at Stage II irrespective of the length of delay between them.*

2.5.4 "Inside" and "outside" money effects

We now investigate the effect of the form taken by the small-scale change in wealth on risk attitudes: the "inside" and the "outside" form. In the "inside" form, a money increment of £7.00 is added to the prizes of a baseline lottery; in the "outside" form, the money increment is given to subjects, who are then supposed to assess a fixed lottery from two different wealth positions – before and after the small-scale wealth increment.

Recall that such effects are examined through within-subject comparisons of responses to risk tasks with lotteries $L_5(0.2, 16; 0.8, 10)$ and $L_2(0.2, 9; 0.8, 3)$ (see Table 2.3). There are three relevant differences in task responses to these lotteries:

1. *Within-stage differences ("Inside money effect")*: $T_i(L_5) - T_i(L_2)$ $i \in \{1, 2\}$

2. *Between-stage differences ("Outside money effect")*: $T_2(L_j) - T_1(L_j)$ $j \in \{2, 5\}$

Table 2.9: Stability of risk preferences between stages by delay treatment

	Instantaneous		*Delayed*	
	Paired t test	Wilcoxon signed-rank test	*Paired t* test	Wilcoxon signed-rank test
L1 *(£8,0.2;£4)*	$t = -0.20$ $p = 0.84$	$z = -0.70$ $p = 0.49$	$t = 1.06$ $p = 0.30$	$z = 0.97$ $p = 0.33$
L2 *(£9,0.2;£3)*	$t = -1.43$ $p = 0.16$	$z = -2.16$ $p = 0.03$	$t = -0.22$ $p = 0.82$	$z = -0.49$ $p = 0.62$
L3 *(£6,0.4;£3)*	$t = 2.69$ $p = 0.01$	$z = 2.20$ $p = 0.03$	$t = 1.16$ $p = 0.26$	$z = 0.98$ $p = 0.33$
L4 *(£9,0.3;£4)*	$t = -1.11$ $p = 0.28$	$z = 0.11$ $p = 0.92$	$t = -0.87$ $p = 0.39$	$z = -0.66$ $p = 0.51$
L5 *(£16,0.2;£10)*	$t = -1.24$ $p = 0.23$	$z = -1.01$ $p = 0.31$	$t = -2.22$ $p = 0.04$	$z = -1.74$ $p = 0.08$
L6 *(£6,0.4;£3)*	$t = 0.32$ $p = 0.75$	$z = 1.86$ $p = 0.06$	$t = -0.55$ $p = 0.59$	$z = -0.94$ $p = 0.35$

Note: Wilcoxon signed rank sum test: The null hypothesis is that first and second-stage measures of risk aversion (number of safe choices) from subjects assigned to the zero increment condition are not significantly different.
H_0: 1st stage response ($\Delta w = 0$) = 2nd stage response ($\Delta w = 0$)

3. *"Inside-Outside" equivalence*: $T_1(L_5) - [T_2(L_2) + 7] = 0 \quad i \neq j$

The theoretical predictions regarding the sign and magnitude of the within- and between-stage differences are that they are all consistent with EUT: such differences can vary across individuals, with the sign of them depending on properties of the individual's $u(.)$ (i.e., if their utility function exhibit CARA, DARA, IARA). But theory gives a more precise prediction regarding the relationship between such differences. More specifically, it predicts an "Inside-Outside" equivalence, by which task response differences induced by an increment to wealth by Δw should not be affected by the framing used to introduce such increment. We showed that for this equivalence condition to hold, the after-increment "corrected" response to L_2 (when taking into account the outside increment) should equal the before-increment task response to L_5.

We first examine the within-stage differences. Table 2.10 reports statistics about the distribution of differences in responses, in terms of certainty equivalent, to L_5 and L_2 in each stage.

Table 2.10: "Inside" money effects: Within-stage differences in responses to L_5 and L_2

	Average	Standard deviation	Min	P25	P50	P75	Max
Stage I: T_1(L5) – T_1(L2)	7.43	1.59	3.5	6.75	7.25	8.00	12.50
Stage II: T_2(L5) – T_2(L2)	7.03	1.05	4.25	6.50	7.00	7.50	10.00

Note: Task response is the medium point of the elicited certainty-equivalent interval. Statistics of stage II differences are taking into account only subjects assigned to *zero* treatment conditions.

The distribution of within-stage differences in task responses is fairly symmetrical around the average, which is slightly larger than 7. When we consider the first-stage, for instance, differences in responses to L_5 and L_2 were larger than 7 for 78 subjects (74%). These patterns are kept similar when considering the second-stage differences in responses to L_5 and L_2. According to Proposition 2, this result is consistent with a convex utility function exhibiting increasing absolute risk lovingness.

This result has key implications for a recent literature over calibration critique of decision theories (Rabin, 2000; Cox & Sadiraj, 2006; Wakker, 2005; Safra & Segal, 2008). Motivated by the broad use of EUT to explain departures from risk neutrality when gambles are small, Rabin (2000) has proved an influential theorem showing that risk aversion over small gambles implies unrealistic levels of risk aversion over very attractive gambles with large stakes. This result has been extended to non-expected utility theories (Cox & Sadiraj, 2006; Safra & Segal, 2008). But it is known, however, that this calibration critique relies on the empirical assumption that a given gamble continues to be rejected over a wide range of wealth levels (see, e.g., Wakker, 2005). This pattern of risk aversion is only consistent with risk preferences that display constant absolute risk aversion. But our results suggest that this assumption may not have empirical validity; our subjects' pattern of risk preferences seem to change over what can be regarded as a narrow range of wealth levels, displaying, hence, everything but constant absolute risk aversion.

We now examine the between-stage differences, that is, differences in tasks responses before- and after increment to $L2$ and $L5$. Table 2.11 reports statistics about the distribution of such differences in the group of subjects assigned to nonzero increment conditions. We also report a break down of these statistics by time treatments. For ease of comparison of increment effects, we also present statistics on differences in responses to $L2$ and $L5$ across stages for subjects in zero increment treatments.

Average figures show that roughly half of the subjects were overall consistent with constant risk aversion: between-stage differences in responses to L_5 and L_2 were near zero. Yet, when considering average differences to the lottery with larger prizes, L_5, the outside increment seems to have induced risk aversion. A contrasting result, though, is suggested by results to the risk tasks with smaller stakes, L_2: among subjects in a non-zero increment condition ($I7$, $D7$), average difference in task responses across stages is positive, which is consistent with risk lovingness.

Yet, this contrasting pattern of differences in task responses was not induced by the outside increment. Using subjects assigned to the zero increment conditions as control groups, we test whether difference in task responses of subjects who received the outside increment ($I7$, $D7$) are significantly different from difference in task responses of subjects assigned to zero increment conditions ($I0$, $D0$). Table 2.12 reports statistics of Mann-Whitney two-sample tests (two-tailed) for lottery L_5 and L_2.

We now test whether the "Inside-Outside" equivalence condition holds, namely, that $T_2(L_2) + 7 = T_1(L_5)$. This equivalence condition states that,

Table 2.11: "Outside" money effects: Between-stage differences (before- and after-increment) in responses to L_5 and L_2

	Average	Standard deviation	Min	P25	P50	P75	Max
L_5 : $T_2(L5) - T_1(L5)$	-0.09	1.29	-3.75	-0.75	0.00	0.50	3.75
Instantaneous +7	-0.10	1.38	-3.00	-0.50	0.00	0.00	3.75
Delayed +7	-0.09	1.22	-3.75	-0.75	0.00	0.50	2.50
Instantaneous +0	-0.34	1.30	-3.75	-0.50	0.00	0.25	1.50
Delayed +0	-0.43	0.95	-2.50	-1.00	-0.125	0.25	1.25
L_2 : $T_2(L2) - T_1(L2)$	0.20	1.08	-1.50	-0.50	0.00	0.63	4.75
Instantaneous +7	0.28	1.35	-1.50	-0.50	0.00	0.75	4.75
Delayed +7	0.13	0.77	-1.50	-0.25	0.25	0.50	1.75
Instantaneous +0	-0.26	0.86	-2.50	-0.50	-0.25	0.00	2.50
Delayed +0	-0.03	0.68	-1.00	-0.50	0.00	0.25	1.50

Note: Task response is the medium point of the elicited certainty-equivalent interval. Statistics in the "lottery" row (L_5 or L_2) refer to between-stage differences from the sub-sample of subjects in *nonzero increment* treatment conditions. These statistics are broken down by time treatment conditions (I7 or D7) below the dashed line. Statistics about such between-stage differences from *zero increment* treatments (grey area) are included as a baseline.

Table 2.12: Are between-stage task response differences induced by outside money? Comparison of between-stage differences in responses to L_5 and L_2 between nonzero (treatment) and zero (control) treatment conditions

	Test statistics
Hypotheses	Mann-Whitney two-sample statistic
Instantaneous treatments	
$[T_2(L_5) - T_1(L_5) \mid I7] \vee [T_2(L_5) - T_1(L_5) \mid I0]$	$z = -0.209$ $p = 0.83$
$[T_2(L_2) - T_1(L_2) \mid I7] \vee [T_2(L_2) - T_1(L_2) \mid I0]$	$z = -1.596$ $p = 0.11$
Delayed treatments	
$[T_2(L_5) - T_1(L_5) \mid D7] \vee [T_2(L_5) - T_1(L_5) \mid D0]$	$z = -0.996$ $p = 0.32$
$[T_2(L_2) - T_1(L_2) \mid D7] \vee [T_2(L_2) - T_1(L_2) \mid D0]$	$z = -0.991$ $p = 0.32$

Note: Null hypothesis is that $[T_2(L5) - T_1(L5)]$ of subjects assigned to I7 (D7) is not significantly different from $[T_2(L5) - T_1(L5)]$ of subjects assigned to I0 (D0).

for an expected utility maximizer with an utility function defined on wealth, the sure sum which is just as good as getting the increment and playing L_2 at stage II should equal the sure sum which is just as good as playing L_5 before the increment. Figure 2.8 displays the histogram of differences between task response to L_5 at stage I and the outside increment-corrected task response to L_2 at stage II, which, if the equivalence condition holds, should be zero.

Considering that there is a wide range of values the differences between $T_1(L_5)$ and $T_2(L_2) + 7$ can take, it is worth noting that subjects for whom this difference takes the value zero, i.e. the equivalence condition holds, are relatively more frequent (21,43%) than subjects for whom the task response difference deviates from zero. Yet, we cannot accept the hypothesis that

Figure 2.8: Distribution of violations of the “Inside-Outside” equivalence condition

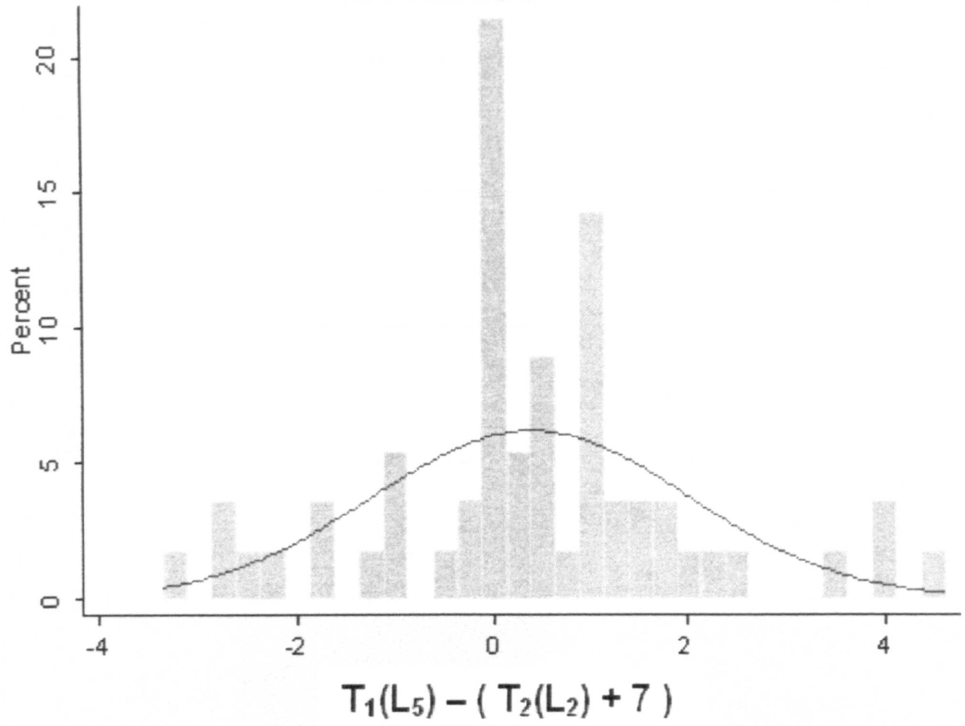

the equivalence condition is not violated. A Wilcoxon signed rank test of the condition, with 56 matched observations, yields $z = 2.041$, $p = 0.04$ (two-tailed). We conclude with the following

Finding 6: *The "Inside-Outside" equivalence condition does not hold. Differences between task response to L_5 at stage I (before increment) and increment-corrected task response to L_2 at stage II (after increment) are not, as predicted, equal to zero.*

Why is this equivalence condition violated? It is possible to come up with different theoretical accounts for that. One could argue, for example, that this violation stems from the presence of a "noise" component in subjects choices that behaves differently across different ranges of lottery prizes. But this would require the distribution of noise to be affected by the expected value of the lottery; otherwise, if noise has a distribution that holds across risk tasks, noise itself cannot predict a systematic tendency to violate the equivalence condition in a particular direction.

Alternatively, this violation could also be accounted for by a framing effect, inconsistent with EUT defined on overall wealth. To see how, let us start noticing that the "Inside-Outside" equivalence condition holds under the hypothesis that the decision-maker integrates the increment with the outcomes of the the lottery L_2 (*asset integration*); it is because of this hypothesis that task responses to L_2 after the increment is supposed to be framed in a way that makes it equivalent to the task response to L_5 before the increment, $T_1(L_5)$. The violation of the "Inside-Outside" condition can, then, be just a consequence of a failure of the asset integration hypothesis in describing how individuals assess a gamble: they simply do not merge prior gains with the direct consequences of taking the gamble.

2.6 Conclusions

We have investigated experimentally three related issues. First, the effects of small-scale changes in wealth on risk attitudes. Second, the effects of time on risk attitudes, which involves a twofold issue: an examination of different delay times on elicitation of risk attitudes in a repeated setting, and an examination of whether and how potential changes in risk attitudes induced by the small wealth increment are affected by how long such increment has been anticipated for – we compare the case in which the increment has just been earned to when it has been anticipated for a week. Third, how the frame adopted to introduce a small wealth increment affects attitudes to a given risk. We examine two "frames": the inside "inside" frame, when the increment of £7.00 is commonly added to the prizes of a lottery, and the "outside" frame, when the increment is simply given to the subject.

Regarding wealth effects, we have observed that risk attitudes do not seem to be systematically affected by the small-scale change in wealth (£7). The experimentally induced increment is modest, but it was larger than the expected value of almost all lotteries used in our experiment. Theoretically, and from a EUT standpoint, this result suggests that overall subjects display risk attitudes consistent with constant absolute risk aversion. This result is also entirely consistent with a narrow bracketing of problems, whereby individuals tend to evaluate new gambles they are offered in isolation from other wealth-relevant events. Barberis *et al.* (2006) show that this psychological feature can actually help a wide range of decision-making theories to account for departures from risk neutrality over small gambles. While risk aversion is far from widespread in our subject sample, we do find evidence that subjects' risky decisions ignore the small wealth increment they were given – and, more generally, their wealth level; instead, they see to base their decisions only upon the direct outcomes of the risk faced, therefore, adopting such "narrow bracketing".

Regarding the time treatments, and particularly the purely effect of delay times on elicited risky choices, subjects' decisions have tended to show, on average, a fair degree of short-term stability. We have had subjects facing risk tasks on two distinct points of time; some faced them a few minutes apart, while others faced them one week apart. We have found that risk task responses elicited at two different occasions do not systematically differ from each other regardless of the length of delay between them employed in our experiment. Indeed, as it has been pointed by some (e.g., Hey & Orme, 1994)), some of our subjects indeed do not give precisely the same answer to a risk task on the second occasion – which would be too "demanding" a test of stability. But differences have not been statistically significant and can be regarded as "noise". Thus, this result has suggested that subjects' risk preferences elicited through the multiple-price-list method tend to exhibit consistency at least over short period of time.

We have also examined the interaction of time and increment treatments. By doing so, we have checked whether how long the increment had been anticipated for would affect whether and how the increment would affect risk attitudes. The span of time was either a few minutes or a week. No effect has been found: comparing risky decisions before- and after-increment, we have found that there was no statistically significant difference between delay conditions. This has suggested that, regardless of how long the increment has been anticipated for, it is still ignored (narrow bracketing) in subsequent problems, hence, not inducing a different representation of the risky problem. Regarding the realm of reference-point theories (e.g., *Prospect Theory*), although our design have not provided a rigorous test of particular assumptions about reference-point determination, these results can be regarded as instructive; they suggest one of the following: either (1) that subjects do not take the expected outcomes, in our case the wealth increment announced before the second round of risk tasks, as their reference-point; or (2) that they do, but the expectations are slowly adjusted (a least longer than a week); or (3) that increment is embedded immediately in reference point.

Regarding the inside/outside money effect, we have found that the same amount of money, £7.00, given "inside" the lottery induces more risk lovingness, while money given "outside" does not induce changes in risk attitudes. This results is, in principle, entirely compatible with those from previous studies reporting *prior* monetary gains induce people to take more risks (e.g., Ackert *et al.* , 2006). Yet what we have found is not a "classic" house money effect – in fact, we have not replicated it in our experiment – but a "built-in" house money effect, whereby the "house money" is added to all potential outcomes of the lottery. This is essentially distinct from scaling-up (multiplicatively) the lottery' outcomes (Holt & Laury, 2002), which even has been shown to have an opposite effect on risk attitudes (Binswanger, 1980; Holt & Laury, 2002; Harrison *et al.* , 2005a).

On the other hand, different ways of framing the increment should not produce, in theory, different responses. According to standard EUT, for instance, an "equivalence condition" should hold, namely: the certainty-equivalent of a lottery L after the £7.00 increment is given should not differ from the certainty-equivalent of a lottery L', where L' is just a transformed version of L whereby £7.00 is added to lottery L's prizes. But when confronting this prediction with the data, we have observed that such equivalence condition does not hold with the difference of certainty-equivalents for $L5$ and $L2$ being larger than 7 (loosely put, $L5 = L2 + 7$).

This result unfolds two distinct theoretical implications. One is that the effects a wealth increment may have on individuals' risk preferences can be sensitive to the vehicle used to introduce it – frames matter. Our preferred explanation for this "framing" inconsistency is a simple one: when assessing a risk, individuals simply do not merge prior gains with the potential consequences of taking the risk. This is also consistent with the idea that decision makers are passive and accept the frames presented to them (Kahneman & Tversky, 1979). We readily acknowledge that this is no longer novel – it is just a violation of the "asset integration" hypothesis demonstrated in some experiments by (Kahneman & Tversky, 1979) that has been re-labeled as "narrow bracketing". But we have been able to go further on this "inside-outside" money issue. This brings us to the second theoretical implication: we have showed that the difference of certainty-equivalents for $L5$ and $L2$ (larger, smaller than, or equal to, 7) reveals information about properties of one's utility function regarding risk-bearing behaviour when wealth increases. The results suggest that anything but constant absolute risk aversion describes out subjects' choices. This may have implications for the for the practical consequence of the calibration critiques of decision theories. They are relying on an empirical assumption – rejection of a given lottery over a wide range of wealth positions – that holds for a class of utility functions that, our experiment suggests, do not capture an increment to wealth affects individuals' attitudes to a given risk.

Appendices

2.7 Appendices

2.7.1 Appendix A - Proofs

Proof. of **Proposition 1**

We provide a demonstration for the case where $\psi(w_1, L) > \psi(w_0, L)$. In terms of task response, in this case, it follows trivially that $T_2(L) \equiv C(L, w_1) < C(L, W_0) \equiv T_1(L)$. Proofs for the other cases use similar arguments and are therefore omitted.

We first prove that $\frac{\partial R_A(w)}{\partial w} > 0 \Rightarrow \psi(w_1, L) > \psi(w_0, L)$.

Assume that $\frac{\partial R_A(w)}{\partial w} > 0$ for all $w \in [\underline{w}, \overline{w}]$, where $0 \leq \underline{w} < \overline{w}$. Assume that $u(.)$ is monotone and strictly concave over $[\underline{w}, \overline{w}]$. Consider that u_0 and u_1 are the utility function evaluated at w_0 and w_1, respectively, where w_0, w_1 $\in [\underline{w}, \overline{w}]$. Since $\frac{\partial R_A(w)}{\partial w} > 0$ and $w_1 > w_0$, we can infer that the decision maker is more risk averse at w_1 than at w_0, that is, $-u_0''/u_0' < -u_1''/u_1'$. In this case, and at least over a closed ball with center w_1 and radius $r \geq w_1 - w_0$, we can see u_1 as a concave transformation of u_0, that is, $u_1 = \phi(u_0)$ where ϕ is a monotone and strictly concave function. Observe now that

$$u_1(w + E(L) - \psi(w_1, L)) = E[u_1(w + L)] \text{ (by risk premium definition)}$$

$$E[u_1(w + L)] = E[\phi(u_0(w + L))]$$

$$E[\phi(u_0(w+L))] < \phi(E[u_0(w+L)]) \text{ (by Jensen's inequality)}$$

$$\phi(E[u_0(w+L)]) = \phi(u_0(w+E(L)-\psi(w_0,L))$$

$$\phi(u_0(w+E(L)-\psi(w_0,L)) = u_1(w+E(L)-\psi(w_0,L),$$

This implies, by monotonicity of u_1, that $\psi(w_1, L) > \psi(w_0, L)$. This completes the first part of the proof.

We now have to prove that $\psi(w_1, L) > \psi(w_0, L) \Rightarrow \frac{\partial R_A(w)}{\partial w} > 0$. We do so using a simple argument. Let A be the statement that $\frac{\partial R_A(w)}{\partial w} > 0$ and B that $\psi(w_1, L) > \psi(w_0, L)$. Assume that $(\sim A)$ holds. If that is the case, then we know that it cannot be true that $-u_0''/u_0' < -u_1''/u_1'$. From the first part of the proof, we know then that $(\sim A)$ implies that $\psi(w_1, L)$ cannot be greater than $\psi(w_0, L)$. Thus, as $u(.)$ is strictly concave, it must be that $(\sim A) \Rightarrow (\sim B)$. Hence, $B \Rightarrow A$. This completes the proof. ■

Proof. of **Proposition 2**

We will prove the 'if" part the proposition ($\Rightarrow$) for the cases of concave and convex utility functions. But before we proceed with this part of the proof, we lay out the conditions of the problem for clarity.

Suppose $u : X \rightarrow \mathbf{R}$ is a utility function, where X is a convex subset of a real linear space, L. Suppose that $u(.)$ is continuous, increasing and twice-differentiable on X.

Let x and x' be elements of X $(x' > x)$, and let $p \in [0, 1]$. Then suppose that $L(x, p; x', 1-p)$ is a baseline lottery that gives x with probability p and x' with probability $1-p$. Since $u(.)$ is continuous and monotonically increasing, L has a certainty-equivalent, that is, there exists an element $c_1 \in X$ satisfying

$$u(c_1) = pu(x) + (1-p)u(x')$$

Now, let $L'(x+k, p; x'+k, 1-p)$ be a lottery constructed from L by increasing each prize of L by the amount of k, so that L' gives $x+k$ with probability p and $x'+k$ with probability $1-p$. For simplicity, assume that $x'+k \leq \sup X$. Again, L' has a certainty-equivalent, that is, there exists an element $c_2 \in X$ satisfying

$$u(c_2) = pu(x+k) + (1-p)u(x'+k)$$

We now prove the "if-part" of Proposition 2. We start with

Case 1: $u(.)$ is concave. *If $u(.)$ is concave then $c_2 - c_1 \gtrless k$ if, and only if, $u(.)$ exhibits decreasing/constant/increasing absolute risk aversion.*

Suppose $u(.)$ is concave. By Jensen's inequality, it follows that

$$u(c_2) < u(p(x+k) + (1-p)(x'+k)) \qquad (2.11)$$

$$u(c_1) < u(px + (1-p)x') \qquad (2.12)$$

Since $u(.)$ is an increasing function, $u(a) > u(b) \Rightarrow a > b$. Therefore, it follows from 2.11 and 2.12, respectively, that

$$c_2 < p(x+k) + (1-p)(x'+k) = E(L') \tag{2.13}$$

$$c_1 < px + (1-p)x' = E(L) \tag{2.14}$$

We know then that $\exists\, a, b \in \mathbb{R}^+$ such that $c_2 + a = E(L')$ and $c_1 + b = E(L)$. Thus

$$c_2 - c_1 = [E(L') - E(L)] + (b - a) = k + (b - a)$$

Let $R(M)$ denote the risk premium of a lottery M, defined as $EM - C(M)$, where $E(M)$ is the expected value and $C(M)$ is the certainty-equivalent of M, respectively. Note then that a and b are just the risk premium of L' and L, respectively. Let them $R(L) = b$ and $R(L') = a$. Thus, $c_2 - c_1 > k$ if $R(L) > R(L')$, that is, if the risk premium of L is larger than the risk premium of L'.

Note, however, that the evaluation of L' and L can be seen as equivalent to the evaluation of L at two initial levels of wealth, $w = 0$ and $w' = w + k$; and that the difference between the expected value of L and the amount of money, denoted by c, for which the individual is indifferent between the lottery L and the certain amount c is decreasing in w.

But it is a well known result that the risk premium of a given lottery is decreasing in wealth if and only if $u(.)$ exhibits decreasing absolute risk aversion (DARA). We can therefore state that

$c_2 - c_1 > k$ if $b > a$, that is, if $u(.)$ exhibits DARA.

Similarly,

$c_2 - c_1 = k$ if $b = a$, that is, if $u(.)$ exhibits constant absolute risk aversion.

$c_2 - c_1 < k$ if $b < a$, that is, if $u(.)$ exhibits increasing absolute risk aversion.

■

Case 2: $u(.)$ is convex. *If $u(.)$ is a convex function then $c_2 - c_1 \gtreqless k$ if, and only if, $u(.)$ exhibits incresing/constant/decreasing absolute risk lovingness..*

Suppose $u(.)$ is convex. By Jensen's inequality, it follows that

$$u(c_2) > u(p(x+k) + (1-p)(x'+k)) \tag{2.15}$$

$$u(c_1) > u(px + (1-p)x') \tag{2.16}$$

Since $u(.)$ is an increasing function, it follows from 2.15 and 2.16, respectively, that

$$c_2 > p(x+k) + (1-p)(x'+k) = E(L') \tag{2.17}$$

$$c_1 > px + (1-p)x' = E(L) \tag{2.18}$$

We know then that $\exists\, a, b \in \mathbb{R}^+$ such that $c_2 = E(L') + a$ and $c_1 = E(L) + b$. Thus

$$c_2 - c_1 = (E(L') - E(L)) + (a - b)) = k + (a - b)$$

In this case, note that a and b are just the negative symmetric of the risk premium of L' and L, respectively. Let them $-R(L) = b$ and $-R(L') = a$. Thus, $c_2 - c_1 > k$ if $-R(L') > -R(L) \rightarrow R(L') < R(L)$, that is, if the risk premium of L is larger than the risk premium of L'.

Note, however, that the evaluation of L' and L can be seen as equivalent to the evaluation of L at two initial levels of wealth, $w = 0$ and $w' = w + k$; and that the difference between the expected value of L and the amount of money, denoted by c, for which the individual is indifferent between the lottery L and the certain amount c is decreasing in w.

By the same argument used before, the evaluation of L' and L can be seen as equivalent to the evaluation of L at two initial levels of wealth, $w = 0$ and $w' = w + k$; and that the difference between the expected value of L and the amount of money, denoted by c, for which the individual is indifferent between the lottery L and the certain amount c is increasing in w. That is, since the individual is risk-lover, the higher her wealth, the higher the amount of money in excess of the expected value required, with certainty, to make her as good as playing L.

By symmetry, the risk premium of a given lottery of a risk-lover is increasing as the individual become wealthier if and only if $u(.)$ exhibits increasing absolute risk lovingness (IARL). Therefore, we can state that

$c_2 - c_1 > k$ if $b > a$, that is, if $u(.)$ exhibits IARL.

Similarly,

$c_2 - c_1 = k$ if $b = a$, hence, if $u(.)$ exhibits constant absolute risk lovingness.

$c_2 - c_1 < k$ if $b < a$, hence, if $u(.)$ exhibits decreasing absolute risk lovingness. ∎

2.7.2 Appendix B - Instructions

GENERAL INSTRUCTIONS

PIECE 1: COMMON TO ALL TREATMENT CONDITIONS

Welcome

Thanks for participating. Let me remind you that this experiment has two sessions. Today's session of this experiment should be complete within an hour. You signed up to a second session which will take place here in one week's time. I will say more about it at the end of today's session.

Important Note

Please do not communicate in any way with other participants during this experiment. Please remember to switch off your mobile. Also, please do not write on these instructions. If you have a question or problem at any point in today's session, please raise your hand and I will come to you.

The Experiment

This is a study of individual decision-making. Although there are many people participating in today's session of this experiment, your earnings will not depend on what others do. The money you receive for participating will depend partly on choices you make yourself and partly on chance.

In this experiment, you will be asked to complete risk tasks and multiple-choice tasks. We will give you instructions for the risk tasks now. Please, take a few minutes to read them through with me.

Risk Tasks

Piece 2: Common to all treatment conditions

Instructions

You will be required to complete several risk tasks. In each risk task you will face a set of choice problems, in which you have to choose between option A, which involves receiving a sure amount of money, or option B, which involves receiving an amount of money where the amount is determined by chance.

Please click on the "**Demonstration" button** on your screen to see what each risk task will look like. Please follow my directions now. Do not click on anything until told to do so.

INTERACTIVE DEMONSTRATION

Parts in Italic below were not on subjects' instructions; They were read out to them during on-screen demonstration

[The table you see on the screen is an illustration of what each risk task will look like].

[Observe that this table presents a set of 21 choice problems, each posed in a given row; each row is a choice between Option A, a sure amount of money, and Option B, a lottery. You must make a choice between the

options in every decision row. Observe that Option B, the lottery, is the same in each row, whereas option A varies, offering progressively smaller sums as you move down the table:]

[Let me explain now what a lottery is. As an example, look at two-colour bar on the right-side of your screen. This is a graphical representation of the lottery. The lottery has two possible prizes, in this example £7 and £2. We select the prize by drawing a numbered chip from a bag that contains 100 chips consecutively numbered from 1 to 100. £7 is paid if the chip drawn is numbered 1 to 35, whereas £2 is paid if the chip is numbered 36 to 100.]

[Let me explain now the following interactive feature.]

[Please click on the button corresponding to the Option A, offering £7.00 for sure, at the first decision row. This would mean that you prefer to receive £7.00 for sure to playing the lottery.]

[Now click on the the button corresponding to the Option A, offering £5.50 for sure, at the seventh decision row. This would mean that you prefer to receive £5.50 for sure to playing the lottery. Note that the computer automatically filled-in the buttons corresponding to the Option A up to £7.00. Why has it done so? When you chose option A at a certain row, the computer will then assume that you prefer option A whenever it is offering a sum larger than that at the row you clicked first. So, when you chose to receive £5.50 for sure to playing the lottery, the computer assumed that you also prefer £5.75, £6.00, and so on, up to £7.00, to playing the

lottery.]

[Click now please on the the button corresponding to the Option B, at the bottom decision row of the table. This would mean that you prefer playing the lottery to receiving £2.00 for sure.]

[Now click on the the button corresponding to the Option B at the decision row 15th. This would mean that you prefer playing the lottery to receiving £3.50 for sure. Note that the computer automatically filled-in the buttons corresponding to the Option B down to £2.00. The logic here is similar to before. When you choose option B at a certain row, the computer will then assume that you prefer option B whenever the amount of money offered by Option A is smaller than that at the row where you selected option B. So, when you chose to play the lottery over receiving £3.50 for sure, the computer assumed that you also prefer to play the lottery to receiving £3.25, £3.00 for sure, and so on, down to £2.00.]

[Note that you must give a choice for every decision row of the table. As it stands, in our example, if you have exactly followed my instructions, a choice is still to be made at decision rows 8 to 14. The risk task is not complete until an option is chosen at every decision row. So, choose any option you like at decision row number 11. If you chose option A, the computer will fill-in the option A at all decision rows up, whereas if you chose option B, the computer will fill-in the option B at all decision rows down. Choose any option you like for the decision rows at which an option was not chosen yet. Doing so, you have completed a risk task.]

[Before confirming your choices, you can, of course, change your mind and switch the option chosen at any row. The computer will adjust the choices accordingly.]

[It is important to observe that the computer fills in the button according to your choices. You can make any choices you like, subject to the constraint that you cannot switch between options A and B more than once in the same risk task.]

Now, I will give you a few minutes to practice yourself. To do so, close the window by clicking on the "Finish Demonstration" button. Then, click on the "Practice" button again to see a new instance of this risk task. If you have a question, please raise your hand and I will come to you.

PRACTICE

[Please, click on the "Finish Demonstration" button now and turn to the instructions now please]

How your earnings are determined

We will determine your earnings at the end of the experiment as follows. After you have finished all risk tasks in this experiment, the computer will then randomly pick one decision row on each risk task you performed. The following figure illustrates how the screen would look like at this point of the experiment if there had only been two risk tasks.

Figure 2.9: Determining your earnings: selection of decision row per risk task

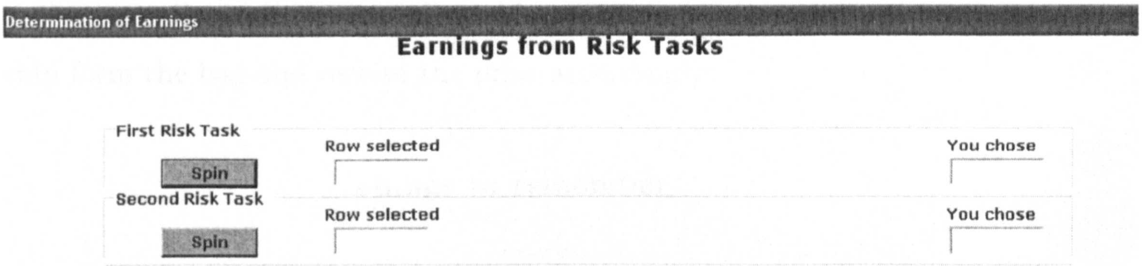

For each risk task you perform in this experiment, there will be a row like one of those on the figure above. Notice that there is a spin button for each task. You will be asked to hit all the spin buttons. As you do this, for a given task, the computer will randomly select one of the decision rows from that task. The computer will use a special randomization device that makes all decisions rows from the task **equally likely** to be chosen. So, note that for every risk task you perform, one decision row will be selected; and all are equally likely.

The set of selected decision rows will be displayed on the computer screen. The computer will then retrieve which option you chose at that row. We then select the option you chose from one of the selected rows. To do so, I will ask you to draw a ball from a bag, which contains balls individually numbered from 1 up to the number of risk tasks you performed. The number you draw will determine which risk task and its corresponding selected decision row will be used to determine your earnings. Once you select the ball, and so the risk task to be used for real, the computer screen will display the two options from the relevant row and the choice you made between them. If you chose Option A, you will receive the amount of money it specifies. If you chose Option B, you will play the lottery by drawing a chip form the bag and receive the prize accordingly.

Things to remember

(1) Though you make several decisions in this experiment, only one of them will be used to determine your earnings. (2) You will not know which decision will determine your earnings from the risk task until the end of the experiment. (3) Think carefully about your decisions in each risk task: every decision row of every risk task could turn out to be the choice that is for real for you.

Thanks for participating.

Please, click on the "Start" button and enter the password I am going to tell you now.

PIECE 3: COMMON TO ALL SUBJECTS

Instructions

We now ask you to complete a test consisting of 12 questions. You will have 1 minute to work on each question of this test. There will be a countdown timer on the top right-side of the screen.

Answers to the questions in this test have no effect on your earnings from the experiment. Nevertheless, please think carefully about the questions and answer them as well as you can in the time available.

Important Note: If you finish before time is called please wait for instructions for the rest of the experiment. Please DO NOT do anything with your computer until you are specifically instructed to do so.

Please, click on the "Start" button and enter the password I am going to tell you now.

MONEY REWARD!

PIECE 4: HANDED OUT TO NON-ZERO INCREMENT GROUPS ONLY.

{ $\rightarrow$ INSTANTANEOUS +7

[$\rightarrow$ DELAYED +7

{[

£7.00 is already guaranteed for you!

For completing the cognitive test, you will be given £7.00. This money will be paid in cash] today at the end of this session [next week at the end of this experiment's next session.].

{[Be assured that £7.00 is already guaranteed for you. You cannot lose this money or any part of it, provided you complete the experiment. This does not affect your earnings from the risk tasks. Therefore, £7.00 will be added to whatever you earn from the risk tasks.]}

Remaining part of the experiment

{We are now going to give you instructions for the remaining part of today's session of this experiment.}

[As you have noticed on the invitation to this experiment, it consists of two sessions: one today and the second one in one week's time.

Your earnings for the experiment, including the £7.00 reward, will be determined and paid at the end of the second session next week. It is, therefore, crucial that you turn up for the next session. Please, do so as you will not get paid if you do not turn up next week for completing the experiment.

Today's session is over.]

MORE RISK TASKS

PIECE 5.1: HANDED OUT TO INSTANTANEOUS GROUPS ONLY.

{ → INSTANTANEOUS +7

(→ INSTANTANEOUS +0

{(

Instructions

This is the last step of today's session of this experiment. Now, we ask you to complete some more risk tasks. They work just like the risk tasks you completed before.

In each risk task you will face a set of choice problems, in which you have to choose between Option A, to receive an amount of money, and Option B, to play a lottery. All you need to do is to indicate the option you prefer most, option A or option B, for each decision problem.

On Your Earnings

After you complete these risk tasks, we will determine your earnings for all risk tasks you have performed and you will be paid at the end of this session. Therefore, the money you earn in the risk tasks) as well as the £7.00 reward for completing the cognitive test (will be paid to you directly in cash at the end of this session.)}

{(Please, click on the "Start" button and enter the password I am going to tell you now.)}

INSTRUCTIONS

PIECE SET 5.2: HANDED OUT TO DELAYED GROUPS ONLY

[→ DELAYED +7

≪ → DELAYED +0

[≪

Welcome

You are in the second session of an experimental study of individual decision-making. You have already attended the first session one week ago. The entire session today should be complete within 30 minutes.

Today's Session

In today's session, you will be asked to complete some more risk tasks. They work just like the risk tasks you have completed during the first session. I will now remind you how that was.

Please click on the "**Demonstration" button** on your screen to see what each risk task will look like. Please follow my directions now. Do not click on anything until told to do so.≫]

INTERACTIVE DEMONSTRATION (Same parts in italic on Piece 2)

[≪ We will determine your earnings at the end of the experiment as follows. After you have finished today's risk tasks, the computer will then randomly pick one decision row on each risk task you performed in this experiment. The following figure illustrates how the screen would look like at this point of the experiment if there had only been two risk tasks.

Figure 2.10: Determining your earnings: selection of decision row per risk task

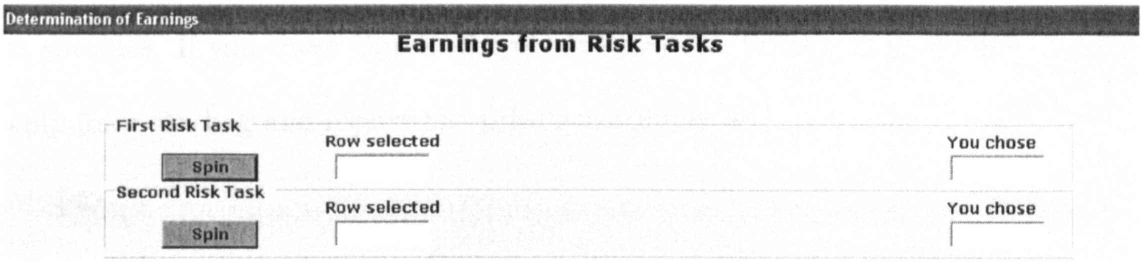

For each risk task you perform in this experiment, there will be a row like one of those on the figure above. Notice that there is a spin button for each task. You will be asked to hit all the spin buttons. As you do this, for a given task, the computer will randomly select one of the decision rows from that task. The computer will use a special randomization device that makes all decisions rows from the task **equally likely** to be chosen. So, note that for every risk task you perform, one decision row will be selected; and all are equally likely.

The set of selected decision rows will be displayed on the computer screen. The computer will then retrieve which option you chose at that row. We then select the option you chose from one of the selected rows.

To do so, I will ask you to draw a ball from a bag, which contains balls individually numbered from 1 up to the number of risk tasks you performed. The number you draw will determine which risk task and its corresponding selected decision row will be used to determine your earnings. Once you select the ball, and so the risk task to be used for real, the computer screen will display the two options from the relevant row and the choice you made between them. If you chose Option A, you will receive the amount of money it specifies. If you chose Option B, you will play the lottery by drawing a chip form the bag and receive the prize accordingly.≫]

[≪

Things to remember

(1) Though you make several decisions in this experiment, only one of them will be used to determine your earnings. (2) You will know which decision will determine your earnings from the risk task at the end of today's session. (3) Think carefully about your decisions in each risk task: every decision row of every risk task could turn out to be the choice that is for real for you. (4)≫ You have already earned £7.00 for completing the test. Thus, I will add £7.00 to the money you earn in the risk tasks.≪You will be paid in cash at the end of this session.≫]

[≪

Thanks for participating.

Please, click on the "Start" button and enter the password I

am going to tell you now. ≫]

2.7.3 Appendix C - Demographic questionnaire

We now ask you to complete a questionnaire consisting of 12 questions about yourself. Be assured that your responses are completely confidential. If you do not wish to answer a question in this questionnaire, you are free to omit it.

DIRECTIONS:

Please fill the ID Number box below. As for the questions, tick the box in the questions below whose option best describes you. ID NUMBER (one digit per cell)

... ...

QUESTIONS

1. What is your sex?

☐ Female ☐ Male

2. What is your age?

☐ years

3. What is your area of study?

☐ Social sciences
☐ Humanities
☐ Health-related sciences
☐ Natural sciences
☐ Other

4. What is the highest level of education you have completed?

☐ Secondary
☐ Undergraduate
☐ Postgraduate: Mphil/MSc

5. Which of the following ethnic groups is appropriate to indicate your cultural background?

☐ White
☐ Mixed
☐ Asian or Asian British
☐ Black or Black British
☐ Chinese
☐ Other ethnic group

6. Which of these categories best describes the average MONTHLY amount of money you have received for JUST YOURSELF over the last 12 months? This can include educational grant, educatinal loan, payments from a family member, income from part-time jobs, state benefits, etc. I'd like to remind you that anything you tell us is completely confidential.

☐ Less than £500
☐ £500 through £799
☐ £800 through £999
☐ £1,000 through £1,499
☐ £1,500 through £1,999
☐ £2,000 through £2,999
☐ £3,000 and greater
☐ Don't know

7. How well would you say you yourself are managing financially these days? Would you say you are (choose the option best describe you):

- [ ] Living comfortably
- [ ] Doing alright
- [ ] Just about getting by
- [ ] Finding it quite difficult
- [ ] Finding it very difficult
- [ ] Don't know

8. What is the highest level of education attained by the head of your family?

- [ ] Primary
- [ ] Secondary
- [ ] Undergraduate
- [ ] Postgraduate: Mphil/MSc
- [ ] Postgraduate: PhD

9. How much money would you ask for solving logical puzzles during 1 HOUR?

£ ______ Pounds

10. Different things can be important when deciding what type of occupation you want to follow. Please look at this table below and choose by clicking how important each of the following aspects are for you. How important is:

	Very Important	Important	Not Important	Not at all important	Don't know
Future job security	( )	( )	( )	( )	( )
High income	( )	( )	( )	( )	( )
Finding an occupation which is well respected	( )	( )	( )	( )	( )
Finding an occupation that leaves you with a lot of time for leisure	( )	( )	( )	( )	( )
A high degree of interaction with people	( )	( )	( )	( )	( )
Finding an occupation which makes a contribution to society	( )	( )	( )	( )	( )
Finding an occupation in which you can help others	( )	( )	( )	( )	( )

11. On a scale from 0 to 10, what is your willingness to take risks, in general, where 0 indicates "unwilling to take risks" and 10 indicates "very willing to take risks".

1	2	3	4	5	6	7	8	9	10

12. How much money would you ask for filling in forms during 1 HOUR?

£ ______ Pounds

***

Thank you for completing this questionnaire.

Chapter 3

Chapter 3

Cognitive Ability and Consistency of Behaviour Under Risk

3.1 Introduction

This chapter investigates experimentally whether violations from principles underlying standard models of rational choice (e.g. EUT) are attenuated as the cognitive ability of the subjects increases. To do this, it combines data on risk choices from the experiment presented in Chapter 2 with data from the cognitive test.

Individuals differ from one another in many aspects: gender, schooling, intelligence, cultural background, to name but a few. It is hardly controversial that differences in individual characteristics can be of substantial importance to understanding some phenomena. Sunden & Surette (1998), for instance, shows the importance of gender and marital status in understanding financial wealth accumulation for retirement. But in modeling economic behaviour, the diversity of individual characteristics has had a mixed role. A hallmark of twentieth century orthodox economics, for example, was to take preferences as a primitive, i.e. unexplained, even though preference differences were recognised as possible and, indeed, essential to motivate many of the main concerns of economic theory, e.g. general equilibrium theory, labor-leisure allocation, etc. Sometimes the issue of individual differences is just overlooked – like, for example, in representative agent models and many classes of game-theoretic models. The underlying principle in these cases is that people deploy similar toolboxes for making economic decisions and, regardless of personal attributes (e.g., age, intelligence, or gender), will respond to economic problems in similar ways. It seems a reasonable simplifying assumption especially when there is limited understanding of which individual characteristics are significant for economic behaviour at the macro level. But some models have tackled the question, introducing some form of heterogeneity (e.g. expectation formation, beliefs); and the introduction of heterogeneity has provided many insights. Grandmont (1992), for example, shows that an increasing degree of heterogeneity can have a regularising effects on aggregate demand. Yet individual characteristics have a neglected role as these models use a "black-box" approach to heterogeneity, whereby the source of it (Age? Cognitive skills?) is usually unspecified.(see, e.g., Morris, 1996; Xiong & Yan, 2006)

In applied work though, the possibility of differences in economic behaviour being linked to differences in individual characteristics has been examined more directly. It is natural in much of the current econometric practice, for instance, to control for observable individual characteristics. Failure to do so would produce biases in the estimated model if such characteristics influence the dependent variable. In a less instrumental perspective, demographics have also emerged as primary objects of interest in experimental economics. A number of studies have explored the links between certain individual characteristics and individual decision-making. Eckel & Grossman (1998) and Bolton & Katok (1995), for instance, investigate differences in pro-social behaviour between men and women in a dictator game; Andreoni & Vesterlund (2001) examine gender differences in altruistic behaviour. Harbaugh *et al.* (2001), in turn, investigate the role of age in choice rationality as measured by tests of transitivity. Gächter *et al.* (2009a) examine the incidence of framing effects between groups with different levels of seniority within the economics profession. Other studies have looked at the influence of individual ethnicity and cultural background on behaviour in ultimatum bargaining games (Roth *et al.* , 1991; Henrich, 2000; Henrich *et al.* , 2004) and public goods games (Gächter & Herrmann, 2009). Irrespective of the approach, these studies provide a general sense of relevance of some demographics to understand how people respond when faced with similar economic circumstances.

In this vein, economists have recently begun to look at cognitive ability. Since individuals engage in some form of reasoning when facing decision problems, it is only natural to inquire about how differences in cognitive ability are related to differences in choices in some economic domains. Frederick (2005), for instance, proposes a cognitive reflection test (CRT) – a three-item test to measure the disposition to resist reporting the response that first comes to mind. He finds that those who scored higher in the CRT are generally more patient and less risk-averse in the domain of gains. Despite using hypothetical choices to elicit time and risk preferences, Frederick's results have been confirmed in properly incentivized designs. Benjamin *et al.* (2006) and Sunde *et al.* (2009) find that individuals with higher intelligence test scores tend to be closer to risk neutrality and less impatient. Brañas-Garza *et al.* (2008), though, find no statistically significant difference in risk preferences between groups of subject with different levels of mathematical skills. Despite the contributions of these studies to our understanding of how differences in levels of intelligence affect human decision making, the literature on this issue is still relatively scarce and mainly focused on risk and time preferences.[1] This paper aims to contribute to this emerging literature by further examining whether and how cognitive ability relates to aspects of individual decision-making.

Using the experiment of the previous chapter, in which subjects make risky choices when facing a series of tasks, we are particularly interested in three aspects of such decision-making.

First, short-term *choice consistency.* We will use *choice consistency* to refer to a subject's choice switches in identical risk tasks. In our ex-

[1] Exceptions are Burks *et al.* (2009), who also investigate how cognitive ability relates to social awareness in Prisoner's Dilemma game and job attachment, and Charness & Levin (2005) study the link between bayesian updating skills and bids in a "winner's curse" experiment.

periment, subjects face a sequence of twelve risk tasks: four of them are identical and the other eight consist of four tasks, each played twice. An underlying principle in models of rational choice is that individuals should make the same choice when facing same circumstances. This has great normative appeal, but there is some evidence documenting a tendency to violate this principle[2]; these studies show that people tend to make different decisions when facing the same problem in different occasions (for a survey, see Bertrand & Mullainathan, 2001). Based on subjects' choice switches in these identical tasks, we introduce a simple index of individual consistency that captures dynamic aspects of their choice decisions. We shall thus extend the literature on this issue by examining whether *choice consistency* is influenced by cognitive ability.

Second, *type consistency*. We will use this term to refer to a subject's switches in her type (risk-averse, risk-lover). While models of stochastic choice often allow switches of type, standard theory restricts pattern of risk preferences across different tasks. As attitudes towards risk are captured by the curvature of the utility function, a risk-averse individual, for example, cannot be risk-loving over the *same* wealth/income range. Our experiment provides a rich decision environment to test for this principle as subjects proceed through sequence of several risk tasks. Using responses to this sequence of risk tasks, we introduce a simple index to characterise

[2]There are many competing explanations for that: an intrinsic stochastic component in people's decisions (Loomes, 2005), different time preferences, etc. It is beyond the scope of this chapter, though, to discuss the reasons behind time instability of preferences.

a subject's *type inconsistency* over all risk elicitation-tasks. We examine how it relates to individual's cognitive ability.

Third, *framing consistency*. This refers to the different frames adopted to introduce a small-scale change of wealth: either given "inside" the lottery, adding a common amount, say m, to all lottery outcomes, or given "outside" the lottery, by simply giving m to the subject. Several studies investigate the effects on one's degree of risk aversion when the size of stakes is increased (e.g. Holt & Laury, 2002; Harrison *et al.* , 2005a), but little has been said about whether and how the form taken by an increase in the payoff levels affects individuals' degree of risk aversion. We showed (Chapter 2) that within EUT and under the hypothesis of *asset integration*, the form taken by a monetary gain of m should not influence its effects, if any, on an individual's risk attitudes. This means that variations of risk attitudes should be consistent between the "inside" and "outside" frame. However, we have leaned in the previous chapter that whether the gain is "inside" or "outside" affects subjects' choices. We then examine here how this *framing inconsistency* relates to differences in cognitive ability as measured by our test.

We also examine how cognitive ability relates to attitudes towards risk. Obviously, this is not a new issue – several other studies have examined the influence of cognitive skills on risk preferences. Yet, as risk is elicited in a repeated setting and our measure of cognitive ability captures different forms of intelligence, this can be seen as robustness test of the general finding, namely, that those with higher cognitive ability are relatively less risk-averse (e.g. Burks *et al.* , 2009). Furthermore, the fact that we have unusually risk-loving subjects provides a chance to shed some light on a question that previous studies do not distinguish: will those with high cognitive ability be the most risk-loving or the closest to risk neutrality?

The remainder of the chapter is organized as follows. Section 3.2 presents the experimental design. Section 3.3 presents the indicators that seek to quantify attitudes to risk and of consistency of behaviour in our experiment and presents the results. Section 3.4 concludes.

3.2 Experiment Design

To study how cognitive ability affects attitudes towards risk and choice consistency, we designed a set of risk-elicitation tasks and a brief cognitive test. In what follows, we describe each component of the experiment. Since the risk-tasks and subject sample are the same used in the experiment described in Chapter 2, the following draws heavily from the experimental design section from the previous chapter (The reader is referred to that chapter for further details.).

3.2.1 Risk tasks

The question of how to elicit people's attitudes to risk is addressed via a Multiple Price List (MPL) procedure. We implemented a sequence of risk-elicitation tasks in each of which a subject faces a number of pairwise choice problems. Each problem is to choose between a given lottery L (a p chance to win x and $1-p$ to win y, where $x > y > 0$) and an amount of money with certainty; the certain money option is systematically decreased from x to y by a constant amount, say δ, when proceeding down the table. So, in the first row, the decision problem an individual faces is to choose between x for sure and L. In the second row it would be a choice between $x - \delta$ for sure and L, and so on until the sure thing equals y, the lower possible payoff yielded by the lottery L. Because the difference between the sure sum and the expected value of the risky option decreases and turns negative from some point on, even a very risk-averse individual is expected to switch over to the lottery at some row when going down the table.

Figure 3.1 below presents a screenshot of the set of pairwise problems presented to subjects in a given risk task. In this example, the task consists of eliciting the cash equivalent of the lottery $L(8.00, 1/5; 4.00, 4/5)$, where the fractions indicate the probabilities of winning, and the integer numbers indicate the winnings:

Each decision row on the screen constitutes a choice problem, which is to choose between option A, a sure sum, or option B, the lottery. Subjects are asked to indicate their preference for each choice problem. Under EUT, an individual will choose A in decision 1 and B in decision 17, switching from A to B at some point in between. Thus, provided a subject starts by choosing A and switches once, task responses can be reduced to a single number:

Figure 3.1: Illustration of a risk elicitation task

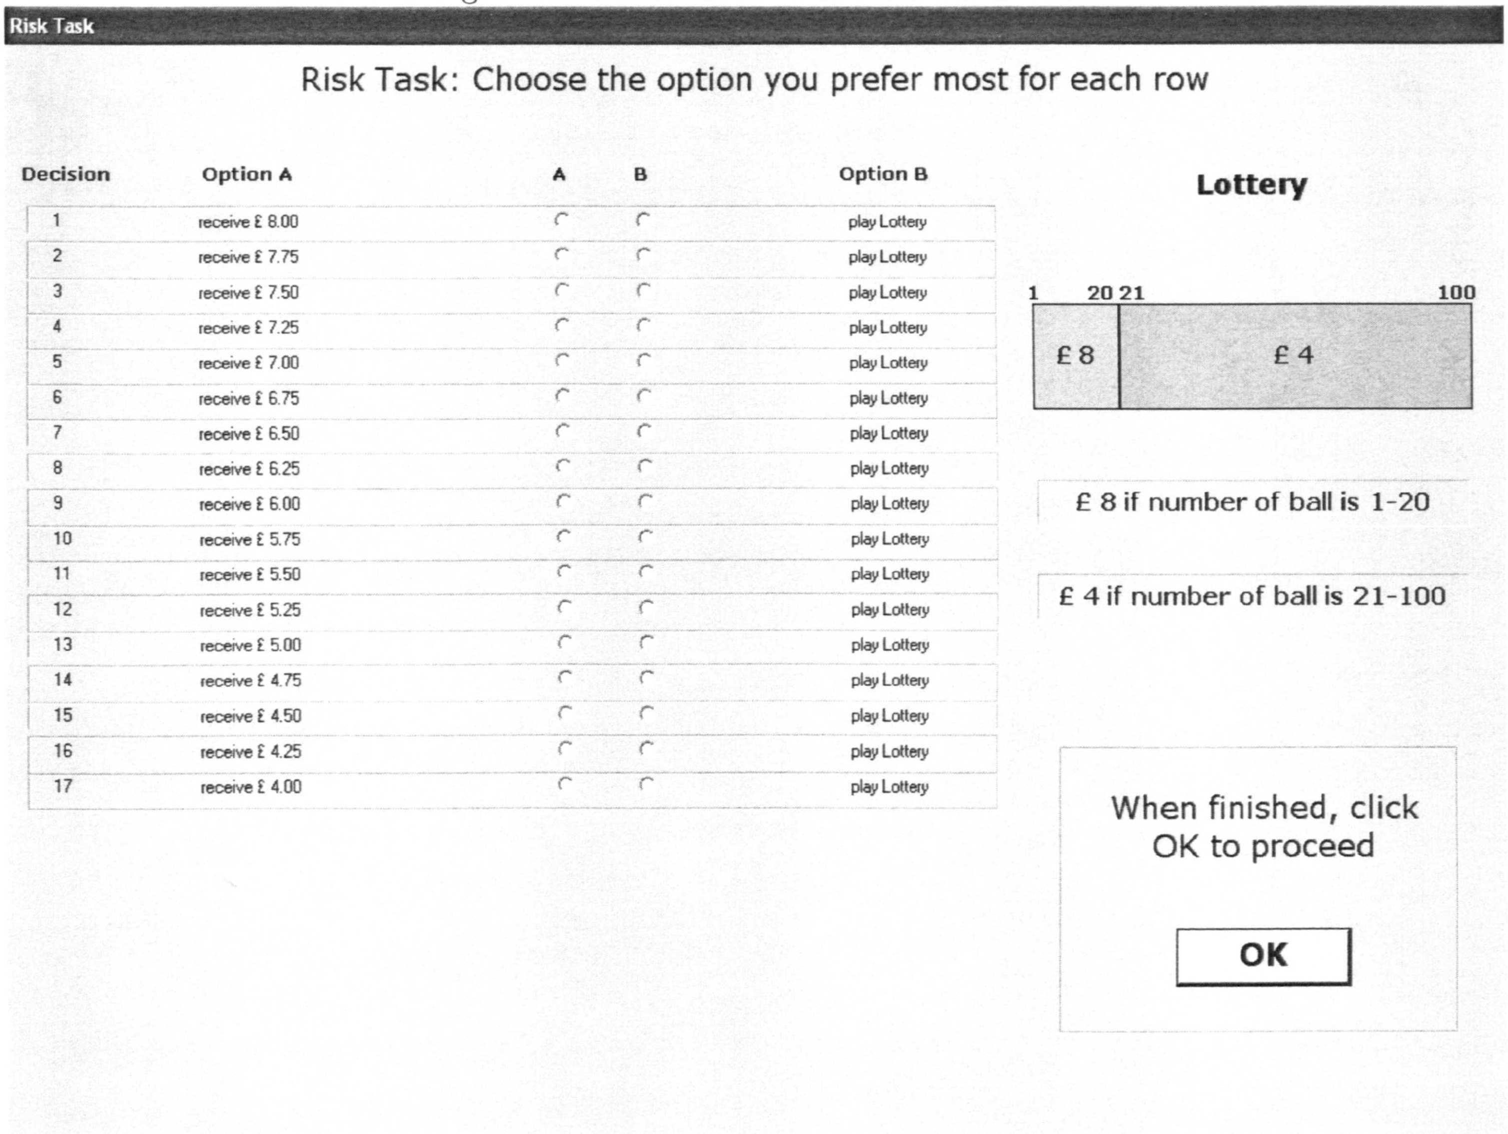

the decision row number at which the subject switched from option A to option B [3].

In our experiment, a subject faces a sequence of six risk tasks on two different occasions. They are not told though that the sequence of six risk tasks they face in those occasions are identical. For convenience, Table 3.1 below presents the set of lotteries used in each of these risk tasks in the order they are presented[4]:

Table 3.1: Lottery option per risk-elicitation task

Lottery	Payoff 1	Pr(Payoff 1)	Payoff 2	Pr(Payoff 2)	EV	Rows
L1	8	0.2	4	0.8	4.8	17
L2	9	0.2	3	0.8	4.2	25
L3	6	0.4	3	0.6	4.2	13
L4	9	0.3	4	0.7	5.5	21
L5	16	0.2	10	0.8	11.2	25
L6	6	0.4	3	0.6	4.2	13

The lotteries we use have four notable characteristics. First, they are all binary lotteries. Second, they only involve strictly positive outcomes. Third, two of them are identical – risk tasks L3 and L6. The reason for this is that by making subjects face the same risk task twice, we can use a test-retest approach to investigate short-term stability of risk preferences. Second, payoffs offered by lotteries L2 and L5 differ by £7.00 (i.e. L5=L2+7), which exactly matches the small scale change in wealth induced by the

[3]As explained in Chapter 2, the software allowed the subjects to economise on "clicking effort", and in doing so guaranteed a maximum of one switch. We view this feature as an advantageous one, as it prevents boredom and guarantees usable data while still leaving plenty of scope for cognitive ability to affect behaviour.

[4]For roughly half of the subjects. For the other half the order of L2 ad L5 was reversed. The purpose is to perform a small-scale test of order effects. We do not reject the hypothesis that there is no order effect on elicited risk attitudes (see Chapter 2).

experiment under some treatments, as described in Chapter 2. This manipulation will allow us to investigate the effects of the form taken by an increment on risk attitudes (here, the increment is either given "inside" or "outside" the lottery).

After completing the first sequence of six risk tasks and before facing the second one, subjects are asked to take a timed cognitive test – for completing it some subjects get a small-scale wealth increment. In what follows, we describe the components of the cognitive test.

3.2.2 Cognitive test

There has been a prominent effort to develop psychometric tests to measure individual reasoning abilities. While there is no clear conceptualization of intelligence (Sternberg & Detterman, 1986), several tests have been devised – if not to measure intelligence itself, to measure some closely related construct like scholastic aptitude, school achievement, etc (Neisser, 1996).

Like other psychometric tests, our cognitive test is a set of questions that seek to assess a range of reasoning skills. The test contains twelve questions divided into four sections: one on each of mathematical, verbal and sequential reasoning; plus the cognitive reflection test proposed by Frederick (2005) (see Appendix to this chapter for the complete test). Subjects are given one minute to complete each question included in the test as they are presented on the computer screen. In our investigation, we measure cognitive ability with scores obtained in this test. An individual's cognitive score is simply the total number of correct answers.

Besides allowing us to do an exploratory analysis of variation of choice behaviour in our risk tasks across groups with different levels of cognitive ability, the cognitive test has two additional purposes as explained in Chapter 2. First, to allow the small-scale wealth increment to be framed as a "earned" reward for completing the test and not as a "gift" from experimenters. Second, to load subject's working memory with new information and make less likely that they will spot the equivalence between first and second round of risk tasks. This might have caused them to guess that the experiment tests for consistency, and respond accordingly – not because their preferences are truly consistent but rather because they want to appear to be consistent (see Bertrand & Mullainathan, 2001).

Like many other tests, our test measures relatively abstract reasoning skills. The mathematical and verbal sections are very much like the GRE Quantitative and Verbal sections, requiring understanding of elementary mathematical concepts and working knowledge of vocabulary and grammar. The GRE is a test widely used by universities for admission and financial-aid in master's and doctoral programmes. A great deal of research has been done to investigate the reasoning abilities measured by the GRE, their predictive validity, and its correlation with scores in other tests of intelligence (see Williams, 1996, p.511-16). The mathematical section picks up problem solving skills involving basic understanding of arithmetic and algebra. The verbal section tries to measure ability to analyse parts of sentences, and recognise the relationship between words and literary concepts. The sequential reasoning section of our test, in turn, covers analysis of patterns and deductive reasoning in arithmetic and geometric context. The cognitive reflection test (Frederick, 2005) tries to measure one's ability to resist reporting intuitive answers that first spring to mind. Table 3.2 shows descriptive statistics for the cognitive test and its components.

Since it is apparent that one's performance in this test involves a command of skills developed over many years of education, performance in cognitive test is likely to be shaped by individual's successful school learning. While this depends on many personal characteristics other than intelligence, school performance is shown to be correlated with scores on psychometric tests to measure cognitive skills (Mayes *et al.* , 2009). Thus, our test, like other measures of scholastic achievement, is likely to be picking up individual's general cognitive ability as well. Furthermore, relative to the general population, our subjects are likely to be less heterogeneous in their schooling success; so differences in cognitive ability within our sample may be more influenced by "raw ability" than they would be in the general population.

On incentivisation of cognitive test

Performance in our cognitive tests has no effect whatsoever on subjects' earnings – a feature that this part of our experiment shares with

Table 3.2: Cognitive Ability Test: Descriptive statistics

Cognitive test component	***Mean***	***Min***	***Max***
Overall score	5.41	1.00	9.00
	(1.68)		
Quantitative score	0.92	0.00	3.00
	(0.77)		
Verbal score	1.07	0.00	3.00
	(0.73)		
Sequential score	2.50	1.00	3.00
	(0.59)		
CRT score	0.92	0.00	3.00
	(0.99)		

Note: The sample consists of 106 subjects. Numbers in brackets are standard deviations. Overall score coud range from 0 to 12. The test has four component tests – each one a three-item test. Scores in each one can range from 0 to 3.

other studies using psychometric tests. Let us, however, address here a potential criticism of this practice, namely, that the lack of financial incentive affects the validity of test score data – since, payoff-wise, test performance does not matter. The belief is that without financial incentives, variation in test performance, if any, would reflect individual intrinsic motivation differentials and not cognitive abilities.

This potential problem, however, is unlikely to be alleviated by monetarily rewarding subjects for correct answers; in fact, incentives could make things worse and arguably affect accuracy of test scores: by paying for correct answers, individual test performance differentials could then reflect an unknown mix of individual differentials in money-driven efforts exerted by subjects and cognitive abilities required to solve the questions. Furthermore, there is some evidence (Gneezy & Rustichini, 2000) that subjects, when regarding incentives as small, can perform in some tasks worse than subjects who were driven solely by their intrinsic motivation. Thus, it seems plausible that, even when not incentivised, test score variation to some extent reflects cognitive ability differentials.

3.2.3 Subject pool, parameters, and procedures

This experiment has a subject pool of 106 participants recruited from a database of volunteers.[5] Most are undergraduate students from different disciplines with a median age of 20 years and age range from 18 to 25 years.

[5]Participants were recruited using ORSEE (Greiner, 2004).

Slightly more than half of the subjects are women (52.83%).

The experiment was conducted in a computerized laboratory in several sessions. In each session, following their arrival, each subject received instructions explaining the risk tasks[6]. The instructions were read aloud while the students read them silently. Before the beginning of the actual experiment, subjects had the opportunity to face a risk-elicitation task for practice, knowing that it would not affect their earnings. The set of lotteries used in each risk-elicitation task and its parameters are presented above in Table 3.1. There were no time constraints to complete the risk tasks. But only when all participants finished making their choices, the set of instructions for the next part of the experiment was handed out.

After completing the first risk-elicitation stage, subjects were given instructions for the cognitive test (framed as questionnaire designed to assess their ability to perform certain type of reasonings[7]. They had twelve minutes to complete the test. The cognitive test phase was followed by the second risk-elicitation stage[8]. After this stage was completed, subjects were asked to complete a short questionnaire with survey questions and fill out a form with information for the administration of payments.

Payment for their decisions was then made at the end of the experiment, in cash, according to the random lottery incentive system. This random-

[6]Instructions for the cognitive test and second risk-elicitation stage were handed out one at a time when all subjects completed the first risk-elicitation stage.

[7]See Instructions – Appendix B of Chapter 2

[8]Which was immediately after the cognitive test for subjects in time treatments I0 and I7; but a week later for those in D0 and D7. We pool observations across treatments as time treatments had no effect on subjects' attitudes to risk.

lottery procedure, by which several decision problems are faced but the subject is paid the outcome of only one of them, has been extensively used. It allows an incentivised elicitation of individual choices in multiple-task settings avoiding income effects Lee (2008). The random-lottery system provides an incentive-compatible elicitation mechanism both under EUT and PT, in the sense that subjects are incentivized to report genuine valuations to the lotteries they face which reflect their true preference ordering over the pairwise set of options.

Thus, subjects were informed prior to responding to the elicitation tasks that, when they had completed all risk tasks, one would be randomly selected; one decision in the chosen task would be randomly selected. Their winnings is determined by the option they chose, after the resolution of the risk in the event that option was the lottery. As explained before, £7.00 was awarded for roughly half of the subjects for completing the cognitive test[9]. This was added at the end to their earnings in the risk tasks. On average, experimental sessions lasted around one hour and average earnings were £6.70.

3.3 Data Analysis

We use data from the two sequences of risk-elicitation tasks all subjects faced. Recall that subjects proceed through a sequence of six risk tasks

[9]This reward was announced only after the test was completed. Subjects assigned to the treatment I7 received it at the end of the session. Subjects assigned to D0 received it at the end of the second session one week later.

on two different occasions. There is an experimental manipulation of the time elapsed between those occasions that, in principle, could "contaminate" responses in the second sequence of risk-elicitation tasks, affecting any inference based on aggregation of the data between the two occasions. Kolmogorov-Smirnov tests of the equality of distributions reported in the previous chapter showed us that the time manipulation of our experiment did not have an effect, so that one cannot reject that distribution of risk attitudes in each risk task is statistically equivalent between groups assigned to different time treatments ("Delayed" and "Instantaneous"). There are differences, however, in the way choices from risk-elicitation stages are used to evaluate subjects' choice and type consistency. We will explain this in detail in Section 3.3.2. In what follows, we describe the data provided by the risk tasks and the procedures used to compute measures of risk aversion.

3.3.1 Measurement of risk aversion

There are different ways in which subjects' choices in each risk task can be used to identify individual estimates of risk aversion.

In a given risk-elicitation task, a subject faces a set of discrete binary choices. In each choice problem, a subject chooses between a sure amount of money and a given fixed lottery. These choices yield a *switching point* for each elicitation task, which is the number of times the subject prefers the sure thing before switching over to the lottery. In the case of a risk-elicitation task with, say, n decision rows, this indicator varies from 0, indicating that the subject prefers the lottery over all amounts given with certainty, to n, indicating that the subject prefers the sure thing over the lottery in all decision rows.

While there are several ways in which one can use subject's risky choices to infer a measure of her risk attitudes[10], we summarise the risk aversion of a subject by computing her risk premium for the lottery in each risk task. The risk premium for a lottery L is the certain amount of money an individual would forego in order to avoid the risk inherent in L. This is a simple and non-parametric method based on the subject's discrete choices. Recall that a subject's *switching point* in a given risk task provides a relatively narrow interval within each the amount of money that for her is as good as taking the gamble must fall into – that is, her certainty-equivalent of the lottery. For simplicity, assume that that such certainty equivalent is just the midpoint of such monetary interval. Now, the risk premium is calculated by subtracting the certainty equivalent from the expected value of the lottery. It can take a negative, positive or zero value depending on one's risk preferences.

We shall refer to a subject's *risk-score* as the average risk premium of the lottery options in the risk tasks faced by each subject[11]. The risk score

[10]For a discussion, see Chapter 2.

[11]We acknowledge that using a measure of central tendency to summarise subjects' risk attitudes has the drawback of throwing information away. But note that we average risk premia over the *same set* for all subjects, and use this measure to compare groups assigned to different treatment conditions – hence whatever problem this measure has, it is a "fixed effect" across the sample. When necessary, our analysis control for lottery characteristics (e.g. prize sizes, number of decision rows, etc).

can range from a positive number, indicating a degree of risk lovingness, to a negative number, indicating a degree of risk aversion. The risk score takes value zero in case of a consistently risk neutral individual. It can also take value 0 for someone who is risk-loving on some tasks and risk-averse on others. This brings us to the issue of consistency of risk preferences.

3.3.2 Measurement of inconsistency of risk preferences

A second important issue that this experiment allows us to address is the consistency of subjects' decisions. In our experiment, we consider two forms of consistency, which we term *choice* and *type* consistency.

The first one, *choice consistency*, relates to the (short-term) stability of subjects choice decisions in repeated risk tasks, an aspect addressed via the test-retest component of our experiment design. Recall that a subject faces the same set of six risk tasks on two different occasions, which are referred as first and second stages. But among this set of six risk tasks, two of them are equivalent– that is, within each stage, one of the risk-elicitation tasks, the third one, is faced twice. We shall combine subjects' decisions on those two occasions to create a measure of short-term consistency of their choices.

The second one, *type consistency*, relates to the variability of a subject's risk preferences over the course of the risk tasks. Since this involves subjects' decisions in all risk tasks, we will propose a measure that captures two aspects of such preference variability: first, the type consistency, that is, the frequency of each "type" of risk attitude over the set of risk tasks. For example, risk preferences of an individual who is risk-averse half of the time and risk-loving the other half should be deemed, *ceteris paribus*, as more inconsistent than an individual who displays risk-averse behaviour in 2/3 of the tasks and risk-loving in 1/3 of the them. The second dimension to be captured is the degree of preference change. For example, an individual who switches from extreme risk-lovingness to extreme risk-aversion should be considered, *ceteris paribus*, more inconsistent than an individual who switches from extreme risk-lovingness to moderate risk-aversion.

We propose an individual index for each form of consistency (*type*- and *choice-consistency*), later analysing the relationship of each one with cognitive ability. We start with the numerical characterisation of our indicator of *risky choice consistency.*

3.3.2.1 An index of *choice consistency*

Our choice consistency measure involves subjects' choice decisions in the risk-task $L3$. This task is faced four times, including the two times where it is faced as $L6$. The fact that they are all identical arises the question: which of these should be used as a reference to judge the consistency of a subject's choice decisions the other three times the risk task is faced? There are two problems with arbitrarily picking one of these tasks. First, that it is unclear which one captures the subject's "true" underlying preferences. Second, and more important, that using a task as a fixed "reference" for choice in this task overlooks the dynamic aspect of subject's decisions, that is, how a subject's choices may evolve over the set of identical tasks. We propose to quantify this type of inconsistency in our experiment with an index that captures this dynamic aspect. But because the overall consistency of all choices is taken into account, choice similarity and consistency are not necessarily the same; an individual who makes identical choices in $L3$ *within* each stage is not always more consistent that an individual who made different choices.

Denoting s_t as the number of "safe choices" made the t-th time ($t = 1, ..., 4$) the risk task was faced, the following set S describes all the six relevant pairwise comparisons: $S = \{(s_t, s_t + 1); t = 1, 2, 3\}$. A subject's *choice consistency index*, C, will then be denoted by

$$C = \sum_{p=1}^{3} \sum_{k=p+1}^{4} |s_p - s_k| \tag{3.1}$$

where the differences (in absolute value) in the inner summation term are the inter-tasks variations in choice behaviour over the identical tasks – that is, the difference between the number of safe choices made in each possible pairwise comparison of the identical risk task. Let us now give some examples.

Suppose two subjects, say A and B, who makes the following number of "safe" choices in the set of identical tasks, respectively: $\{6, 6, 9, 9\}$ and $\{4, 4, 9, 9\}$. The risky choice consistency index of A is 12, while B's index is 20 – A's risky choices are more consistent than B's. Suppose, instead, that

A has made identical choices, say $\{9, 9, 9, 9\}$, and B's choices were all but one identical, say $\{5, 5, 5, 6\}$. In such case, as one expects, A has a more consistent risky choice behaviour than B, as A's index is 0 while B's index of choice consistency is 3.

It is worth mentioning some properties of this index. First, that an individual's index of choice consistency depends on her decisions in each replica of the risk task $L3$; it depends neither on any other parameters, such as her coefficient of risk aversion, nor on any other risk task. Second, that the index is not more sensitive to upward shifts in choices than to downward shifts. Third, that the index is invariant to the order in which a given series of choices in the identical task are made. Fourth, that the index is not monotonic in the frequency of pairwise perfectly identical decisions. This property means that an individual who makes identical choices in three out of the four identical risk tasks need not be more consistent than an individual who made no identical choices in any pair of the risk task repeats. For example, let $S_i = \{1, 1, 1, 5\}$ and $S_j = \{1, 2, 3, 4\}$ be the sequence of "safe" choices of two individuals in each of the four occasions they face the risk task $L3$. The respective indices of risky choice consistency of i and j are $C_i = 12$ an $C_j = 10$; hence, j is more consistent than i, though i has made identical choices in all but one time risk task $L3$ was faced.

3.3.2.2 An index of *type consistency*

We now introduce a measure of a subject's *type consistency* over the course of the risk tasks. In a given risk task, a subject i is said to be a *risk-averse type* if her risk premium in that task is larger than zero. Likewise, i is said to be a *non-risk-averse type* if her risk premium in that task is smaller than, or equal to, zero. Obviously, this latter *type* encompasses risk-neutrality and risk-lovingness. The reason is simple; for one subject to be precisely risk neutral is a very strong requirement: in a risk task, there are several choice decisions that make one a risk-averse/risk-loving, but there is only one choice that makes her risk-neutral[12]. We believe that this strength makes the "merging" acceptable.

We purport to construct an index that captures two dimensions of consistency in one's risk *type*. First, the degree of uniformity in *type* throughout the set of risk tasks. Second, the degree of variation between *types*. To this end, consider first the number of times an individual displays risk aversion and non-risk aversion, R_A and $\tilde{R_A}$ respectively. Considering all risk tasks one faces in our experiment, we know that the *type* profile of an individual can be described by an element of $P = \{(R_A, \tilde{R_A}) | R_A + \tilde{R_A} = 12\}$, which is a set of ordered pairs representing an individual's distribution of types over her entire set of risky decisions. Then, for an individual with a *type* profile $(R_A, \tilde{R_A})$, let I_u be a measure of her *type* uniformity, where

[12]In fact, there is not really any choice that strictly implies risk neutrality, since we never really know where in the switching interval the subjects' true indifference lies.

$$I_u = \min(R_A, \tilde{R_A})$$

This measure takes value of zero, for example, for an individual who, type-wise, is perfectly uniform throughout all risk tasks. Yet this measure is not sensitive, as it stands, to the degree of variation between type switches. To see that, consider two individuals who are both risk-averse in half of the risk tasks – hence, non-risk-averse in half of the risk tasks. One of them, though, switches from extreme risk-lovingness to extreme risk-aversion, while the other switches from moderate risk-aversion to moderate risk-lovingness. While they would both have the same I_u, it is hardly disputable that such patterns of risk preferences should not yield equivalent indices of *type* consistency. This brings us to the second dimension of our indicator of *type* consistency – variability of choices when switching between types. We propose to formalise this dimension with a simple and widely used measure of variability: the variance-to-mean ratio. So let

$$I_v = \frac{\sigma^2}{|\mu|}$$

be a measure of an individual's (between-type) variability, where the numerator is the variance of risk-premia variation between choices associated to one *type* and the other; and the denominator is the absolute value of mean risk-premia variation between choices associated to one *type* and the other. This measure of variability is increasing on the dispersion of risk premia between risk tasks of one *type* relative to the other one.

We now combine these two measures, I_u and I_v to construct a simple index of *type consistency*. Let such index be denoted by T, such that

$$T = I_v \ln(1 + I_u) \tag{3.2}$$

This index is monotonically increasing in both measures of uniformity and dispersion of *type* in one's choices over the set of risk tasks. The virtue of combining I_u and I_v into a simple index is that this composite measure allows us a more refined ordering of individuals. For a pair of subjects who, for example, match each other in terms of uniformity of types over the risk tasks, the dispersion component of it will work as a "tie-breaker". For example: an individual who is *extremely* risk-averse half of the choices and *extremely* risk-loving the other half will be deemed as more *type* inconsistent than an individual who is *moderately* risk-averse half of the time and *moderately* risk-loving the other half. Note, thought, that because we intend to capture two dimensions, the ordering of individuals according to their *type* consistency does not necessarily preserve the ordering produced by each element T is composed of individually considered. That amounts to say that an individual who does not display much uniformity in the distribution of types over her choices is not necessarily more *type* inconsistent than all other individuals who were more uniform than her; final ordering depends on both uniformity and dispersion of choices made under each *type*.

3.3.3 Results

3.3.3.1 Cognitive ability and risk attitudes

This section begins by asking how individuals' measures of risk aversion are related to their performance in the cognitive test.

Figure 3.2 presents box plots with key distributional features of risk scores per cognitive score group – subjects are divided into two cognitive groups according to their score in the test: "High" score group, which includes subjects who scored above the median score, and "Low" score group, which includes those who scored below (or equal to) the median score. In each box, a line is drawn across the box at the median risk score value. The first and the third quartiles of the distribution of scores in each group are represented by, respectively, lower and upper "whiskers" in the plot lines used to outline the box.

By comparing the line drawn across the box at the median , one can see that subjects with higher level of cognitive ability do not have different attitudes towards risk relative to subjects in the low cognitive score group. Both cognitive groups have a similar pattern of distribution of risk scores, with the majority of subjects in both groups displaying a moderate degree of risk lovingness. Yet, low cognitive ability subjects are, on average, closer to risk neutrality than high cognitive ability subjects.

These differences are not statistically significant though. The result suggested by visual inspection of Figure (3.2) is indeed confirmed by a

Figure 3.2: Risk attitudes, by cognitive score

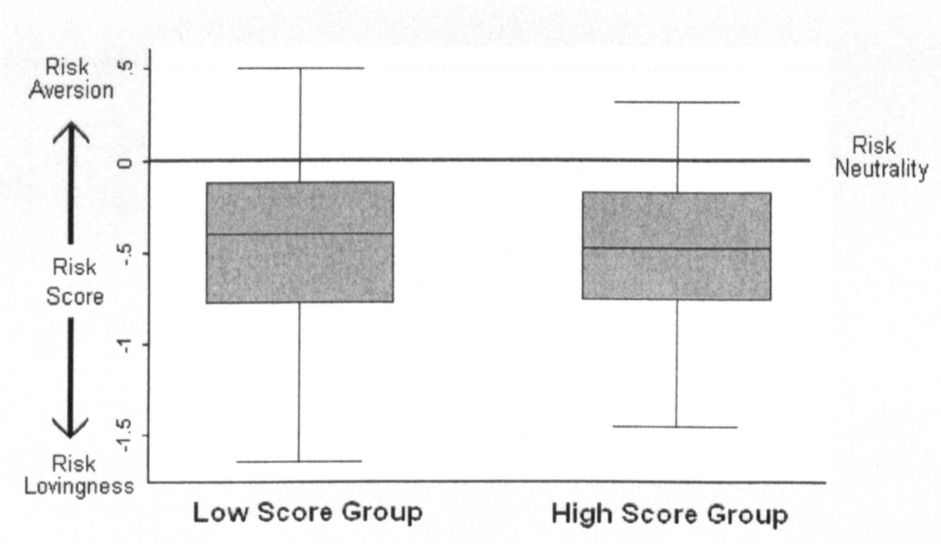

Kolmogorov-Smirnov test ($D = 0.105$, $p = 0.93$). Test results show that we cannot reject the hypothesis that the risk scores of low and high cognitive score groups are the same. One might wonder though if this result would hold in a less aggregated analysis. To check this, we test for differences in distribution of risk scores at a risk task level. Table 3.3 reports the average risk score per risk task in each cognitive group in both elicitation stages; the results of Mann-Whitney tests suggest that differences in risk score between cognitive groups are not statistically significant even at a risk task level. Hence, cognitive ability, as we measure it, is indeed not associated with risk attitudes in our experiment.

Surprisingly, these results hold true even when cognitive ability is measured by subject's performance in the Cognitive Reflection Test (CRT). Frederick (2005) shows that scores in this test correlate positively and significantly with scores in other tests of cognitive ability, claiming that the

Table 3.3: Average risk score per risk task, by cognitive group

	Average Risk score		*Statistical significance*
First stage[1]	*"Low Score"*	*"High Score"*	*Mann-Whitney Test*[a]
L1 (£8,0.2;£4)	-0.569 (0.654)	-0.462 (0.501)	Z = -0.670 (P = 0.503)
L2 (£9,0.2;£3)	-0.863 (1.141)	-0.642 (0.696)	Z = -0.531 (P = 0.595)
L3 (£6,0.4;£3)	-0.129 (0.532)	-0.210 (0.391)	Z = 0.853 (P = 0.393)
L4 (£9,0.3;£4)	-0.391 (0.917)	-0.457 (0.630)	Z = 0.822 (P = 0.411)
L5 (£16,0.2;£10)	-1.277 (1.395)	-1.107 (1.075)	Z = -0.652 (P = 0.514)
L6 (£6,0.4;£3)	-0.143 (0.515)	-0.240 (0.400)	Z = 1.000 (P = 0.317)
Second stage[2]			
L1 (£8,0.2;£4)	-0.710 (0.656)	-0.372 (0.391)	Z = -0.610 (P = 0.542)
L2 (£9,0.2;£3)	-0.804 (1.010)	-0.487 (0.575)	Z = -0.360 (P = 0.719)
L3 (£6,0.4;£3)	-0.198 (0.323)	-0.204 (0.344)	Z = 0.693 (P = 0.488)
L4 (£9,0.3;£4)	-0.434 (0.697)	-0.365 (0.467)	Z = 0.314 (P = 0.753)
L5 (£16,0.2;£10)	-0.844 (1.149)	-0.795 (1.029)	Z = 0.088 (P = 0.929)
L6 (£6,0.4;£3)	-0.181 (0.402)	-0.163 (0.339)	Z = 1.388 (P = 0.165)

Notes: "Low score" ("High score") are those subjects whose cognitive test sore is below (above) the median score. Standard deviations reported in parenthesis. P-values are two-tailed. * p<0.10, ** p<0.05, *** p<0.01. [a] H0: distribution of risk scores does not differ between the "high" and "low" score groups. [1] Based on observations from all subjects. [2] Based on observations from subjects in zero-increment treatments (I0 and D0) only.

CRT is a good predictor of time and risk preferences. We test the robustness of these results to an incentivised risk-elicitation procedure – a feature absent in Frederick's design. We do so by re-dividing the subject sample according to their score in Frederick's three-item Cognitive Reflection Test;

this is one of the sub-sections of the cognitive test subjects completed.

Following Frederick's analysis, we compare risk measures of two "extreme" groups: those who scored 0 out of 3 ("low" group) and those who scored 3 out of 3 ("high" group). No significant differences between these groups are found though[13]. A Kolmogorov-Smirnov test for risk scores in the low versus the high CRT score group yields $D = 0.433$, $p = 0.11$. Therefore, we cannot reject the hypothesis that those who scored higher on the CRT were not more or less risk-averse than those with the lowest scores.[14]

To examine the relationship between risk preferences and cognitive ability more rigorously, we conduct a regression analysis of the following basic econometric specification,

$$R_i = a_0 + a_1 SCORE_i + \beta' \mathbf{x}_i + \epsilon_i \tag{3.3}$$

where R_i denotes subject i's risk score, and $SCORE$ denotes the i's measure of cognitive ability. The vector $\mathbf{x}_i$ includes a set of dummy variables identifying subject i's parental education, income levels, gender and age. This will allow us to control for the role of some observable socio-demographic aspects in determining risk attitudes, providing a check of the robustness of the non-parametric test results.

[13]This result holds even when we enlarge the sample and redefine the "low group" as those who scored either 0 or 1 out of 3, and the "high group" as those who scored 2 or 3 out of 3. Kolmogorov-Smirnov test yields $D = 0.1133$, $p = 0.95$.

[14]Parametric tests (unreported) show that this result holds when we compare risk scores of subjects divided according to their performance on other components of our test – the quantitative, verbal and sequential, reasoning sections.

We estimate four econometric models. Model I is the baseline one, where the risk score, as measured, is used as dependent variable and the performance in the cognitive test (not including the CRT score) is used as a measure of cognitive ability. The other three models check whether the the baseline results are robust to using either different measures of cognitive ability or different measures of risk aversion as dependent variables. Model II uses the CRT score as the measure of cognitive ability. Models III and IV use a self-reported measure of willingness to take risks[15] (WTR) as dependent variable, while using either the cognitive test score or the CRT score as the measure of cognitive ability. This is to check whether the results of the baseline model are robust to using a simple survey question to measure subjects' willingness to take risks. Estimates of all models are based on standard linear regression.

Table 3.4 reports OLS estimates of the four model specifications. The second and third columns present the coefficient estimates of models I and II, respectively. The two remaining columns present the estimates of specification III and IV of the baseline model. The point estimate on the *SCORE* variable suggests that a better performance in the cognitive tests does not lead to a statistically significant change in attitudes towards risk. In models I-IV, the parameter estimates of the effect of cognitive ability on risk attitudes are slightly different from zero; in fact, the hypothesis that they

[15]Where 0 indicates "unwilling to take risks" and 10 indicates "very willing to take risk". This information was collected through a short socio-demographic questionnaire that subject were asked to answer at the end of the experiment.

are basically zero cannot be rejected. Similarly, in models II and IV in which cognitive ability is measured by the CRT, the estimated coefficients describing the effect of cognitive ability on risk scores are very similar to those found in the other models and equally not significantly different from zero. Thus, results based on WTR corroborate those based on risk score. Estimates across all four specifications confirm the non-parametric results shown before.

In summary, we find no relation between CRT scores and risk aversion. This is consistent with a recent study by Brañas-Garza *et al.* (2008), who find that individual computational capabilities are unrelated to risk attitudes. In general, demographic variable have little explanatory power. An exception is models III and IV where gender is significant. Based on those models, females are more risk averse than men.

One may wonder why these results differ from those found in (Frederick, 2005; Benjamin *et al.* , 2006; Sunde *et al.* , 2009), who report that risk aversion over small gambles is less common among individuals with higher scores in cognitive tests. It is not clear though where such differences are primarily stemming from. Our experiment differs from these in several aspects. The risk-elicitation procedure employed here, for instance, differs from Frederick (2005) in using an incentive-compatible mechanism to elicit risk attitudes, but uses the same measure of cognitive ability – along with others. Sunde *et al.* (2009) use the same multiple-price list method used here to elicit risk attitudes, but we use a different test to measure cognitive

Table 3.4: Determinants of individuals' Risk Score (OLS Regressions)

Dependent variable:	Risk Score	Risk Score	WTR	WTR
	(I)	(II)	(III)	(IV)
Constant	-0.469	-0.731	6.341*	6.277*
	(0.755)	(0.720)	(2.372)	(2.265)
Score	-0.017	0.065	-0.107	-0.205
	(0.047)	(0.057)	(0.149)	(0.180)
Female dummy	-0.035	-0.014	-0.618**	-0.621*
	(0.111)	(0.109)	(0.348)	(0.344)
Age	-0.004	0.010	0.017	0.004
	(0.035)	(0.036)	(0.111)	(0.112)
High income dummy	-0.084	-0.108	0.010	0.050
	(0.148)	(0.148)	(0.466)	(0.051)
Parental education dummy	-0.065	-0.086	-0.480	-0.404
	(0.119)	(0.120)	(0.375)	(0.378)
Observations	106	106	106	106
F(5,100)	0.18	0.41	0.98	1.14
R-squared	0.0089	0.0203	0.0466	0.0539

Note: Heteroskedasticity-robust standard error in parentheses. * Statistically significant at the 5% level. ** Statistically significant at the 10% level. SCORE is a subject's measure of cognitive ability. In models I and III, this is a subject's overall score in the cognitive test subtracted from her score in the cognitive reflection test (CRT). In models II and IV, this is a subject's score in the CRT only. "High income" dummy takes the value of 1 if the subject reported receiving an average monthly amount of money above one thousand pounds (answers 3-6 in Question 6 of Soci-demographic questionnaire (SDQ); see Appendix B of Chapter 2), and value of 0 otherwise. "Parental education" dummy takes the value of 1 if the subject reported that the hed of her family attained a postgraduate level of education (answers 4 and 5 in Question 8 of SDQ).

ability. This suggests that the relationship between cognitive ability and risk attitudes, and possibly other types of economic preferences, might be sensitive to intelligence and preference measurement methods used.

We summarize the results regarding the relationship between cognitive ability and risk preferences in

Result 1. *There is no difference in risk attitudes between "high" and "low" score groups. Risk aversion is not related to cognitive ability – even when we use Frederick (2005)'s cognitive reflection test to measure individuals' cognitive skills.*

3.3.3.2 Cognitive ability and *choice consistency*

We now ask whether individuals with higher levels of cognitive ability have greater choice consistency. Before proceeding to a parametric analysis of the data, we want to check the distribution of the index of choice consistency in each cognitive group.

The box-plot in Figure 3.3 shows an economical display of the distribution of the index of choice consistency in each cognitive group. The horizontal line in each box plot is the median value – 8 in the "low" cognitive group and 5 in the "high" one. The first and third quartile range defines the upper and lower boundaries of the box. The whiskers represent the range of more extreme values[16]. It is apparent that choice consistency

[16]By "extreme" we mean values lying on the interval between the upper quartile (inferior limit) and 1.5 times the interquartile range (superior limit).

has different patterns of distribution between cognitive groups.

Figure 3.3: Choice Consistency Index, by cognitive core

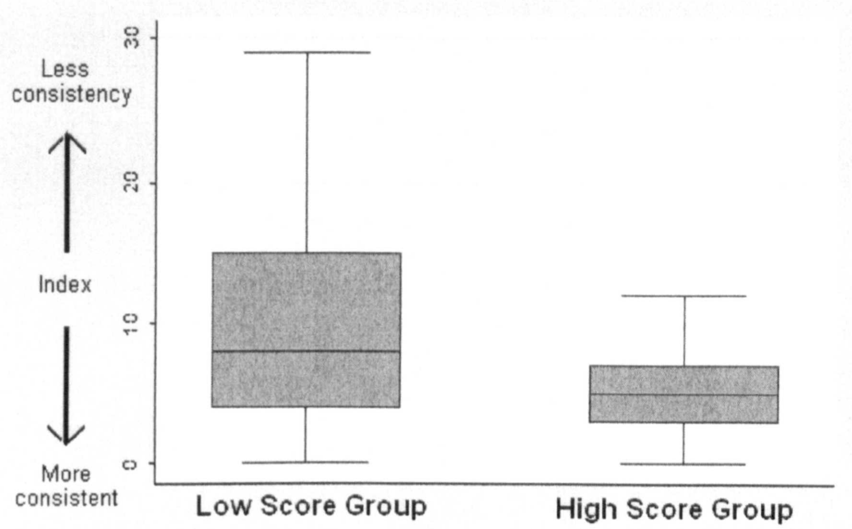

Indeed, choice consistency in the set of identical risk tasks is not the same across cognitive groups. Non-parametric tests show that we can reject the null hypothesis that both distributions of choice consistency are statistically equivalent[17]. A Mann-Whitney test yields $z = 2.804$, $P = 0.0050$, and a Kolmogorov-Smirnov test yields $D = 0.3157$, $P = 0.006$. We summarize this finding in the following

Result 2. *When a given risk task was faced on different occasions, individuals with higher cognitive ability displayed a more consistent pattern of choices than individuals with lower cognitive ability.*

Why do those with better cognitive skills display more choice consistency? Suppose a stochastic specification of choice is allowed, a subject's

[17]While the time period between the identical risk tasks varied across subjects, this result holds true even when we separate the subjects according to time treatment conditions they were assigned to ("Delayed" or "Instantaneous").

choice in a given task would be thought as a random variable equal to her true preference plus noise. Through this framework, Result 2 is to be interpreted as simply showing that individuals with higher cognitive ability make fewer mistakes in translating their preferences into choices – i.e. less noise. This is the account advanced by Burks *et al.* (2009) for a similar result[18]. While this is arguably a natural way to account for differences in choice precision in a stochastic framework, this account does not explain why choice decisions of individuals with lower cognitive ability are noisier than choice of individuals with higher cognitive ability. We offer two possible accounts for that.

One is psychology-based and takes into account the repetitive structure of this component of the experiment intending to test for consistency. If one thinks of "working memory" – the very short time over which we keep something in mind before dismissing it – as the cognitive device that enables subjects to retrieve information as they proceed through the experiment, it is natural to think that stronger memories will be associated to more consistent choices. Alexander & Smales (1997) and Engle *et al.* (1999) show some evidence that working memory capacity is positively correlated with cognitive skills. Individuals with better working memory would then demonstrate more accuracy in retrieving representations of recent events. Thus, in those risk tasks which are faced several times, individuals with

[18]They also find that individuals with higher cognitive ability make more consistent choice. But the type of consistency they examine, multiple switching in risk-elicitation task, is different from ours.

higher cognitive ability could more easily retrieve their previous choices and, therefore, demonstrate more consistency in their risky choices relative to individual with lower cognitive ability. Note that this does not necessarily rely on subjects spotting the similarity of these risk tasks; in fact, this argument is even consistent with a possible scenario in which all subjects spotted the similarity, since it still holds that some will retrieve their previous choices in similar scenarios more easily and precisely than others.

Another possible account is that those with lower cognitive ability choose without serious deliberation. Obviously, the effect of lack of engagement with risk tasks also extends to all parts of the experiment, in particular, the cognitive test. Psychometric tests, like the cognitive test used here, are designed in a way that there is a necessity of sustained cognitive effort in order to perform well; indeed, test-takers need to know certain logical rules; more importantly, to perform well they need to engage in a deliberate, slower, serial and effortful reasoning process rather than answer based on a plausible judgment that comes quickly to mind – that is, they need to use their "System 1 (reasoning) as opposed to their "System 2" (intuition) (Kahneman, 2003). This is cognitively costly and some subjects decided not to think that hard. Likewise, choice decisions in risk tasks did not involve much deliberation. Thus, the argument goes, the individuals who get lower scores in the cognitive test also display "noisier" behaviour in the risk tasks[19].

3.3.3.3 Cognitive ability and *type consistency*

We now focus our attention on the effect of cognitive ability on the consistency of subjects' *type* over the course of the risk tasks.

For each risk task, we classify a given subject into two types: risk-averse, if her risk premium for the lottery in that risk task is non-negative, or risk-loving, if her risk premium is negative. From a theoretical point of view (based on standard theory), displays of both types by a subject would be inconsistent, as there cannot be a unique utility representation of such preferences. Our index tries to measure the degree in which an individual violates this principle. Some could argue though that, if a stochastic specification of individuals' choice process is allowed, some choice pattern that yields a mixed sequence of risk preferences in terms of *type* may be compatible with *type* consistency; and depending on how the structure of such stochastic term is modelled, consistent but "noisy" choices may produce *type* reversals and some relatively extreme variations across the spectrum of *type* classes. But even in this stochastic specification, our index of *type* consistency capture both aspects[20].

So is there a difference in *type consistency* between high and low cogni-

[19]This account invites a discussion of why some subjects engage in more serious deliberation than others. There are many competing explanations (e.g. different understanding across cognitive levels), but it is beyond the scope of this chapter to examine them.

[20]In which case the type consistency index, so long there is some *type* mixing, would just be a more comprehensive measure of choice consistency in our experiment, since involving all risk tasks.

tive groups? Figure 3.4, plotting features of the index of type consistency's distribution, sheds light on this issue.

Figure 3.4: Index of type consistency by cognitive group

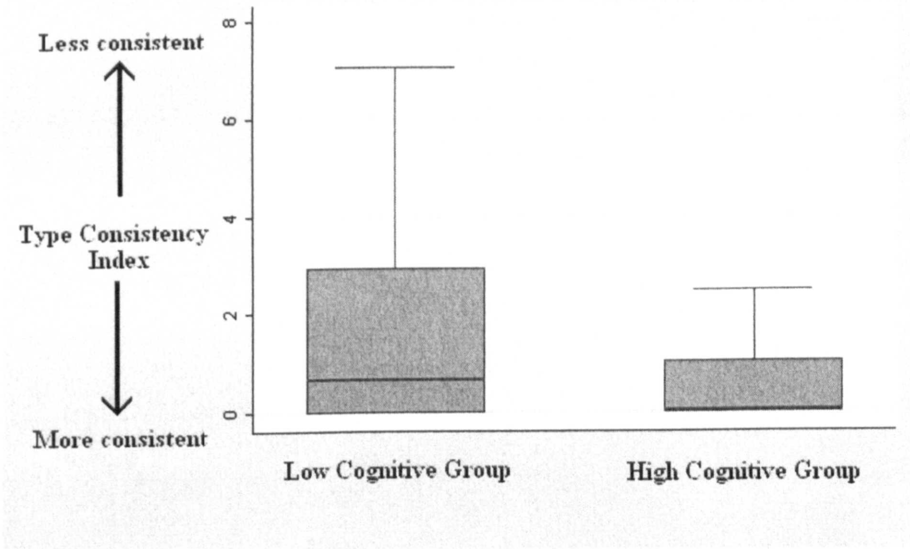

In each box, the red line drawn across each cognitive group box represents the median index value for that group. There are clear differences between cognitive groups. Not only the median index of type consistency is larger among subjects with relatively low scores, but also the interquartile range. These differences are statistically significant: a Mann-Whitney two-sample test yields $Z = 2.581$, $P = 0.0098$. We test whether these results are driven by the effect of differences in dispersion of choices, as measured by I_v. We do so by examining whether differences between cognitive groups in terms of *type* uniformity, as measured by I_u, are statistically significant. They are: a Mann-Whitney test yields $Z = 2.438$, $P = 0.0146$.

Result 3. *Individuals with higher cognitive ability displayed a more con-*

sistent pattern of risk attitude type than individuals with lower cognitive ability.

A first hypothesis to explain this result would not differ much from the explanation for differences in choice consistency in identical risk tasks. We have seen in the previous section that higher cognitive ability is associated with more consistent choices. It was advanced that this could possibly be related to the somewhat indirect effects of cognitive ability in reducing noise in subjects' choice decisions – either because of the way differences in cognitive ability affect "working memory", hence choice precision, or because of differences in deliberation when responding cognitive test and risk tasks alike.

But since our measurement of *type* consistency employs choices involving different lotteries, one could think that the above is not the most plausible theoretical interpretation for these results. It could be the case that the observed *type* reversals are actually capturing different curvature zones of individuals' utility function. Note, however, that all lotteries used have a narrow range of prizes, with almost all lotteries[21] featuring prizes between 3 and 9 British pounds; often, *type* reversals involve lotteries with similar range of prizes. For this reason, even if the utility function representing each subject's preferences has different concavity as we move along the wealth line, the index of *type* consistency is capturing risk preferences that cannot be simultaneously consistent. There is little reason, however, to

[21]Exception is $L5$ that gives a prize of either 10 or 16 British pounds.

think that such differences would be affected by reasoning skills.

3.3.3.4 Cognitive ability and framing consistency

Are individuals with higher cognitive ability less or more susceptible to framing effects than those with lower cognitive ability? According to rational choice models, they should not be. It is not that such a prediction can be derived from these models; they simply do not take into account the potential effect that demographic differences might have on what is regarded as rational decision-making. At least normatively, rational choice in these models mean choice that satisfy a certain set of principles.

One underlying principle of such models is that different frames of a giving choice problem should not induce an individual to different decisions if the variations in frame leave the consequences of choice problem unchanged – their decisions should be *consistent across frames.* Obviously, individuals' choice decisions should satisfy this principle regardless of levels of cognitive ability. But while there is plenty of evidence that, in a variety of problems, people's choices are affected by changes of frame (e.g. Druckman, 2001; Tversky & Kahneman, 1981), little has been said about whether individuals with different levels of cognitive skills are equally prone to such framing effects.

To examine this question, we make use of two risk tasks: $L2$ and $L5$. $L5$ is just the lottery constructed from $L2$ by increasing each prize of $L2$ by the amount of 7. By thinking of these lottery prizes as small-scale changes in wealth, as we increase the prizes in $L5$, we increase the individual's wealth level – but we do so through money given "inside" the lottery. So we denote risk premia differences between $L5$ and $L2$ in the first risk-elicitation stage, before the £7.00 increment is given, as the "inside money" framing effect. In turn, we denote risk premia differences between $L2$ across stages, before and after the £7.00 increment is given, as the "outside money" framing effect. For an expected utility maximiser whose utility is defined on final wealth, these framing effects should be equivalent as the consequences are wealth-wise identical between these frames. The only difference between them is that the small-scale wealth increment in one frame is given "inside" the lottery and in the other frame is given "outside" the lottery. Hence the variation in risk attitudes that may be induced by a small-scale change in wealth of, say, Δw, should be unaffected by the form the increment is framed; that is, individual choices should be consistent across frames.

In Table 3.5, we examine framing consistency for subjects in the high cognitive group versus subjects in the low one. Specifically, we report the means and standard deviation of risk premia variations between frames – a measure of framing inconsistency. Statistic tests and p-values from Kolmogorov-Smirnov tests of equality of distributions are reported in the third column. This test includes only individuals in the treatment conditions who were assigned to a non-zero increment condition ($I7$ and $W7$).

We do observe framing effects: 72% of subjects in the high cognitive group and 84% in the low cognitive group show some degree of inconsistency

Table 3.5: Framing consistency between cognitive groups

Inside/Outside framing variations	High score group (>average)	Low score group (<average)	Kolmogorov-Smirnov test
Stage I	-0.40 (1.70)	-0.34 (1.72)	D=0.308 (P=0.14)
Stage II	-0.10 (2.30)	-0.08 (1.75)	D=0.139 (P=0.95)

Note: Stage I(II) is the risk premia variation between frames, where the "inside" money framing effect is based on the risk premia differences between L2 and L5 in the first (second) stage. That is, Stage 1(2) = $[R_{1(2)}(L5) - R_{1(2)}(L2)] - [R_2(L2) - R_1(L2)]$, where $R_s(L)$ is a subject's risk premium for lottery L at risk-elicitation stage s. Standard deviations are reported in parentheses below each mean value.

between the framings used to introduce a small-scale change of wealth. Yet differences of framing consistency between cognitive groups are not statistically significant.

The framing inconsistency of most subjects' choices tells us that they tend to evaluate the risk tasks in terms of a minimal account, which would include only the sums that could be won in each risk task and exclude money won before. Tversky & Kahneman (1981) propose that this "narrow" framing makes the decision-making easier by simplifying the problem evaluation and reducing cognitive strain. We find that this mode of framing is adopted by individuals with different levels of cognitive ability. There can be at least two interpretations of this result. One, that this mode of framing reflects what seems to be a "natural" and intuitive way of assessing a choice problem i.e., choosing according to your preferences over the direct consequences. Another, that overriding the "narrow" focus underlying this framing inconsistency is too "complex" – depending on how "broad" it is assumed to be, it requires merging outcomes from other domains with the consequence of the current choice problem and calculating the joint distribution of them. Hence the deliberate deployment of cognitive operations that this requires is, in principle, only accessible to individuals with very high levels of cognitive ability. While subjects who were perfectly consistent between the "inside" and "outside" frames were in the high cognitive group, our sample size is too small to allow more definitive inference.

We summarize the results regarding the relationship between cognitive ability and risk preferences in

Result 4. *The way a small-scale wealth increment is framed -either given "inside" a given risk or "outside" it – affects subjects' risky choice regardless of their levels of cognitive ability. Thus, framing inconsistency is robust across cognitive score groups.*

3.4 Conclusions

We present a set of experimental results which, built on design presented in Chapter 2, explore the relationship between cognitive ability and consistency of behaviour under risk. An important innovation over previous work on this topic is that several dimensions of behavioural consistency are investigated.

First, we used a test-retest approach, where a given risk task is faced on different occasions, to examine *choice consistency*. We introduced an index that reflects the degree of choice variation over the repeated risk tasks; it also captures "dynamic" aspects of an individual's choice decisions in this set of tasks. We find that individuals with higher cognitive ability display greater consistency in their choice behaviour. Our preferred interpretation for this result is simply that higher cognitive ability help individuals to state their preferences with more precision, reducing noise in their observed choices.

Second, we examine the association of cognitive ability with what we termed *type consistency.* In order to do so, we introduce an index that intends to quantify this form of behavioral consistency based on *type* classification of subjects' choices in all risk tasks of the experiment. We think of type consistency as a type of bi-dimensional feature of the way risk preferences evolve in a repeated elicitation setting, reflecting uniformity but also variability of risk attitudes. Because both dimensions are somewhat related to how "noisy" is the process whereby individuals translate their preferences into choices, we expect individuals with higher cognitive ability to display relatively more *type* consistency. Cognitive ability through its effects on memory, deliberation, etc would reduce the propensity to error. It is not startling then that we find that individuals with higher cognitive ability also display more *type consistency.*

Third, we used different forms to represent a small-scale change in wealth to investigate *framing consistency.* This increment would take place under two frames: an "inside" and an "outside" frame. In the "inside", the small increment is incorporated into the outcomes of a risk task; in the "outside", it would be given to the subject. While there is no theoretical constraint to the sign of the effect of such increment on risk attitudes, theories of rational choice predict that such effect should be consistent between frames. While subjects in our experiment violated that principle, we do not find that differences in the patterns of violation are related to differences in cognitive ability. Under some hypotheses regarding the argument of the utility function, framing consistency in our context would require overriding a "narrow framing" approach, whereby lottery outcomes are not mentally merged with the wealth increment. We then conjectured that only individuals of a certain top percentile of the distribution of cognitive ability could, in principle, experience some success in mentally accessing the distribution of outcomes that a "broad framing" requires.

Perhaps the most surprising aspect of our set of results is the failure to replicate Frederick (2005)'s results regarding risk behaviour. The cognitive reflection test (CRT) proposed by Frederick is part of our cognitive test; and an important difference over his work is the incentivised elicitation of risk attitudes. Under incentivised conditions, we do not observe that individuals with higher cognitive ability are less risk-averse. In principle, this would cast doubts either on either the validity of the CRT in measuring cognitive ability or validity of elicitation of risk attitudes under non-incentivised conditions. But it is worth noting that we used a more comprehensive psychometric test and provided incentivised conditions to risk elicitation; yet we do not find any statistically significant association between risk attitudes and cognitive ability. While this seems to be in stark contrast to other recent studies (Burks *et al.* , 2009; Sunde *et al.* , 2009), this particular result should not be particularly startling. After all, individuals with higher cognitive ability benefit from a more efficient reasoning system; and there is nothing particularly complex in tasks used in laboratory-based elicitation of risk preferences that allow cognitive ability to arise as a source of differences in risk preferences. This is consistent with (Kahneman & Frederick, 2002; Stanovich & West, 2008)'s two-system based framework, according to which some experimental tasks will give cues to some subjects that a heuristically primed response needs to be overridden and an analytically derived choice substituted. Thus, the variability in the association of cognitive ability and performance in heuristic and biases tasks documented by Stanovich & West (2008) would derive from the failure of lower cognitive ability individuals in performing such overriding[22]. It is arguably the case that in expressing their risk preferences, individuals do not perceive that as a heuristically based answer that needs to be overridden by a analytic reasoning process. Therefore, cognitive capacity would have little bearing on the expression of preferences underlying a subject's display of risk aversion or risk lovingness in the type of task used in experimental investigations.

[22]Either because "mindware" is not available or, even when necessity for override is detected and mindware is available, the individual cannot carry out what they term sustained cognitive decoupling (Stanovich & West, 2008, p.687).

Appendix

3.5 Appendix: Cognitive Test

GENERAL COGNITIVE ABILITY ASSESSMENT*

Part I: Frederick's Cognitive Reflection test (JEP 2006) (3 questions)
Ability to spot erroneous features of intuitive answers that spring to mind when facing certain problems.

Part II: Quantitative Reasoning (3 questions)
Ability to use algebraic, arithmetical and geometrical methods to solve problems in a quantitative setting.

Part III: Sequential reasoning (3 questions)
Ability to use deductive and logical reasoning in a arithmetic and geometric context.

Part IV: Verbal Reasoning (3 questions)
Ability to analyze relationships between words and concepts and to interpret written information.

Total number of questions: 12
Time given: 12 minutes

*** This page was not shown to subects. The test section heading on the next pages also were not shows to subjects.**

Part I: Cognitive Reflection Test (Frederick, JEP 2006)

Solve each of the following problems and then write your answer.

1. A bat and a ball cost £1.10 in total. The bat costs £1.00 more than the ball. How much does the ball cost?

__5___pences.

2. If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets?

__5 ___minutes.

3. In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake?

___47______days.

Part II: Quantitative Reasoning

Section A: Problem solving
Solve each of the following problems and then choose the correct answer. Use the paper provided for any rough work.

4. A fish tank is half full of water. When 10 gallons are added, the tank is 6/8 full. What is the capacity of the tank in gallons?

(A) 30 gallons

(B) 40 gallons (X)

(C) 50 gallons

(D) 60 gallons

(E) 80 gallons

5. If a dealer had sold a stereo for £600, he would have made a 20% profit. Instead, the dealer sold it for a 40% loss. At what price was the stereo sold?

(A) £300 (X)

(B) £315

(C) £372

(D) £400

(E) £440

6. x and y are integers such that x + y < 11, and x > 6. What is the smallest possible value of x − y ?

(A) 1

(B) 2

(C) -2

(D) 4 (X)

(E) -4

Part III: Sequential Reasoning

7. Determine the number that should come next in the following series:

3 8 14 21 29 38 ?

(A) 46

(B) 42

(C) 51

(D) 54

(E) 48 (X)

8. Determine the missing square:

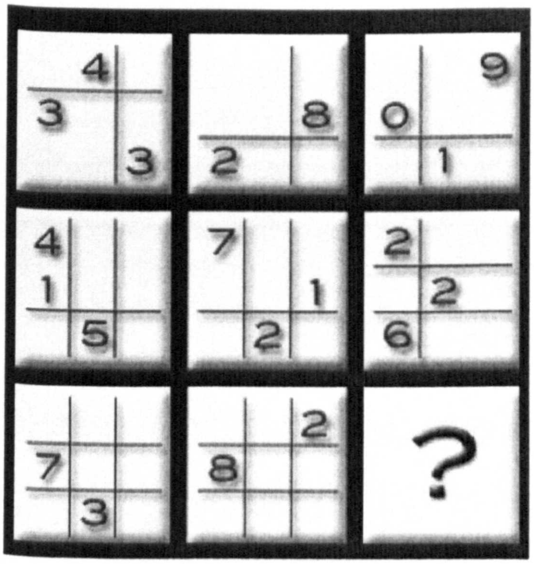

(A)	(B)	(C)	(D)
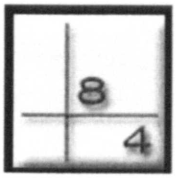		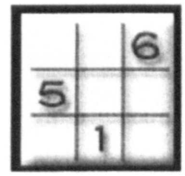	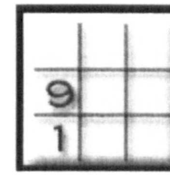
( )	( )	( )	(X)

9. Determine the missing square:

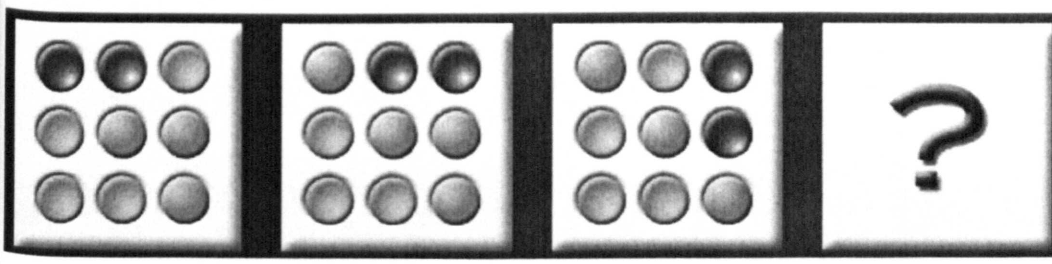

(A)	(B)	(C)	(D)
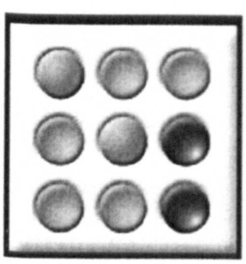	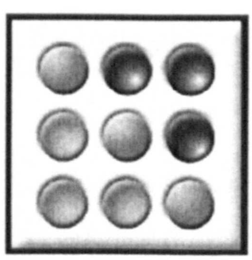	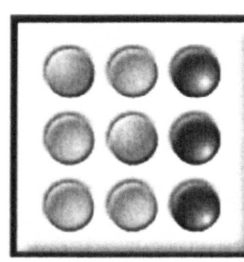	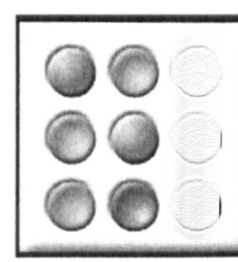
(X)	( )	( )	( )

Part IV: Verbal Reasoning

Section A: Analogies
Choose the answer key which contains a pair of words with a relationship most similar to the relationship between the pair of words in capital letters.

10. ARCHIVE : RECORDS

(A) arsenal : arms (X)

(B) locker : uniform

(C) box : shoes

(D) pantry : bread

(E) arsenide : death

Section B: Antonyms
Choose the answer key corresponding to the word with a meaning most nearly opposite to the meaning of the word in capital letters.

11. CENSURE

(A) proceed

(B) freedom

(C) praise (X)

(D) enclosure

(E) interest

Section C: Sentence completion
Choose the answer choice that contians the words that best complete the sentence.

12. To reach Simonville, the traveller needs to drive with extreme caution along the ____ curves of the mountain road that climbs ____ to the summit.

(A) jagged – steadily

(B) serpentine - steeply (X)

(C) gentle – precipitously

(D) shady – steadily

(E) hair-raising – languidly

Chapter 4

Chapter 4

Cooperation and Punishment under Uncertain Enforcement

4.1 Introduction

This chapter investigates experimentally the efficacy of a punishment mechanism in promoting cooperative behaviour when punishment enforcement is uncertain.

There has been a long-standing interest across many disciplines – situations where self-interested behaviour is at odds with collective interest[1] in behaviour in social dilemma situations. We commonly find ourselves facing such dilemmas. For instance, should we vote in a national election? It hardly seems individually rational to do so: it is costly and almost certainly has negligible impact on the final outcome. Yet, democratic political systems would breakdown if everyone refrained from voting. Quite often social

[1]See for example, in economics, Hardin (1968); Axelrod (1984); in psychology Dawes (1980); Messick & Brewer (1983); in biology Trivers (1971); Boyd & Lorberbaum (1987); in sociology Kollock (1998); Glance & Huberman (1993).

dilemmas appear in the use of public resources – water, forests, and health systems, to name a few – when self-interest of users conflicts with collective interests and the very sustainable use of such resources. A crucial question to social scientists and biologists alike is then how can self-interested individuals be induced to cooperate in social dilemmas?

Attempts to answer this question have led to a number of experimental studies on how to increase cooperation in social dilemmas. Many mechanisms have been investigated. Isaac & Walker (1988); Cinyabuguma *et al.* (2005); Guth *et al.* (2007) and Masclet *et al.* (2003), have showed that preplay communication, threat of expulsion, or even symbolic disapproval can all boost cooperative behaviour. Alternative mechanisms such as giving subjects an opportunity to penalise others financially can also effectively increase and maintain high levels of cooperation in repeated public goods game (Fehr & Gächter, 2000). This monetary sanction system has been receiving increasing attention (see e.g. Anderson & Putterman, 2006; Sefton *et al.* , 2007; Casari *et al.* , 2007; Ertan *et al.* , 2009).

Subsequent studies have confirmed that subjects are willing to pay from their own earnings to punish defectors (e.g., Fehr & Gächter, 2002; Masclet *et al.* , 2003; Noussair & Tucker, 2005; Bochet *et al.* , 2006; Sefton *et al.* , 2007; Gächter *et al.* , 2008); by doing so, they help to maintain contributions to the public good at high levels. Overall, they support the view that, at least under some circumstances, the existence of a sanctioning system can foster behaviour deemed as socially acceptable[2].

The laboratory setting used in these studies abstracts from many things. Take the uncertainty about the link between behaviour in the social dilemma and punishment. In principle, there could be uncertainty about, at least, three things: first, others' willingness to punish; second, whether our actions are being watched by others; third, whether willingness to punish can be translated into actual punishment decisions. There is no doubt that the first type of uncertainty is naturally present in the standard design: individuals face some uncertainty about whether others are willing to punish them, especially when it is costly to do so.

But the last two types of uncertainty have been largely neglected, as the commonly used experimental design relies on two assumptions: *perfect monitoring* and *perfect enforcement*. To see how this is the case, note first that in the standard public good design contributions are disclosed in every period, after which punishment opportunities are given. Hence, there is certainty of being monitored all the time throughout the game (*perfect monitoring*). Yet in many real world settings behaviour deemed as socially inappropriate escapes punishment simply because it is not observed. Note, also, that in these experiments there is no uncertainty about whether subject's demand for punishment will be satisfied: punishment decisions are always carried out (*perfect enforcement*). Most sanctioning

[2]There is some evidence that what is "socially acceptable" varies across societies, with high contributors in some cases getting more punished than low contributors; see Herrmann *et al.* (2008).

systems in modern societies, however, do not have this feature. Often, there are hindrances to punishment enforcement. For instance, individuals tasked with enforcing punishments can be corruptible, and anti-social behaviour, even if detected, could still not result in any penalty at all. Even when sanctions are decentralised and informal, individuals may have the willingness but not the ability to punish someone simply because an opportunity to do so will not arise. In either case, punishment is rarely perceived as certain. An experimental setting assuming perfect monitoring or perfect enforcement does not take into account those uncertainties, which could lead to a misleading assessment of the efficacy of punishment mechanisms in disciplining non-cooperators.

The aim of this study is then to relax one of those assumptions, isolating its effect on the deterrence force of punishment opportunities. We investigate, in particular, if a punishment mechanism can succeed in promoting cooperative behaviour in a public goods game when there is uncertain enforcement. While there exists a subjective element in individuals' perception of this uncertainty, assigning a probability distribution to enforcement of punishment could make the perceived uncertainty surrounding this event measurable (risk), controllable and, at least objectively, uniform across individuals.

Thus, to investigate the impact of uncertain enforcement on the common boosting effect of a punishment mechanism on cooperation, we designed an experiment that introduces measurable uncertainty into whether others' decision to punish a given player is actually carried out. To our knowledge only one previous study, by Walker & Halloran (2004), has investigated this. The authors compare cooperation in a one-shot two-stage punishment game (Fehr & Gächter, 2000) in which imposition of sanctions is certain to a two-stage game in which imposition is uncertain. They find that uncertainty does not change the level of cooperation or the willingness to punish in a significant way. We took a second look at this issue by examining it in a repeated setting and including different levels of uncertainty. Specifically, the experiment sought to investigate how "high" and "low" enforcement probabilities affect cooperation in a repeated-play public goods game, comparing behaviour in such uncertain environments to behaviour in an environment in which punishment enforcement is certain.

While we are primarily interested in the effects of uncertain enforcement on cooperation[3], such uncertainty may also affect individuals' willingness to punish free-riders. Its effects on punishment decisions are far from obvious. The reason is that backward- and forward-looking motives are likely to be driving punishment decisions. Experimental evidence suggests, for instance, that punishment is motivated by negative emotions triggered by past free-riding behaviour (Fehr & Gächter, 2002, p.139). There is also evidence that punishment tends to decrease over time (Nikiforakis, 2008, p.102), which suggests that the future matters: individuals presumably reason that the effectiveness of their punishment in enforcing cooperation

[3]Because we want to isolate the effect of uncertainty over enforcement, costs of punishment are not incurred unless it is enforced. We come to this point later.

is weakened as the game proceeds towards the end.

One can conjecture that uncertainty about punishment enforcement affect both backward- and forward-looking motives. Backward-looking motives because, as punishment may not be carried out, there may be many past episodes of free-riding behaviour that went unpunished; not because of unwillingness to punish, but because "luck" got free-riders "off the hook". Hence, through its effect on backward-looking motives, enforcement uncertainty could cause an increase in punishment – reflecting a delayed outlet of accumulated negative emotions caused by past free-riding behaviour. Uncertainty could also affect forward-looking motives because the anticipation that punishment may not be enforced could weaken its strategic use. Hence, through its effect on forward-looking motives, enforcement uncertainty could cause a decrease in punishment. Given these countervailing forces, it is not clear why punishment should be less effective. So whether and how uncertainty affects individuals' willingness to punish free-riding behaviour is also an empirical question that remains open to investigation, providing additional motivation for this study.

Further, this study can be seen as extending the current body of research on the "robustness" of the punishment mechanism used by Fehr & Gächter (2000, 2002). Recent papers have provided evidence that punishment may not help to maintain cooperation. Even when there is certainty over enforcement, the effectiveness of punishment in promoting cooperation is sensitive to (i) its price (Andersen *et al.* , 2006c), (ii) its payoff impact *per* *unit* of punishment (Egas & Riedl, 2008), (iii) whether individuals are given counter-punishment opportunities (Nikiforakis, 2008) and, to (iv) cultural differences regarding the strength of norms of civic cooperation (Herrmann *et al.* , 2008). The findings reported here add to this literature, furthering our understanding of under which circumstances a punishment mechanism can induce cooperation in social dilemmas.

The experiment has two major results. First, that the threat of punishment cannot raise and sustain high levels of contributions when punishment enforcement is perceived by the individuals as a low-probability event. The experimental results show that a relatively low probability of non-enforcement does not impair punishment to serve as an effective deterrent device, whereas a high probability of non-enforcement does. This indicates that there is more at work in sustaining cooperation than the simple existence of a sanction system. Second, that low contributors are more intensely punished when enforcement of punishment decisions is a low-probability event. Also, and curiously enough, punishment of free-riders and low contributors is generally more intense at the beginning and the end of the game. Thus, in contrast to Walker & Halloran (2004), we find that the existence of uncertainty over the imposition of sanctions has consistent implications on subjects' decision rules.

The remainder of this paper is organized as follows. Section 4.2 describes the experiment design. Section 4.3 presents the hypotheses to be tested. Section 4.4 reports the results. Section 4.5 concludes. An Appendix with the set of instructions is at the end of the paper.

4.2 Experimental Design

The design consists of a public good experiment with punishment with three treatment conditions. In one treatment (P100) there is certain enforcement. This corresponds to the standard case in the literature, in which punishment decisions are always enforced. The remaining two treatments differ according to the probability of enforcement of punishment decisions: one treatment with "high" probability of enforcement (P80), in which punishment decisions are carried out with probability 0.8; and the other with "low" probability of enforcement (P20), in which punishment decisions are carried out with probability 0.2. Thus, there is a chance in these two latter treatments that punishment decisions will not be actually carried out[4].

In each session, sixteen subjects are randomly partitioned into groups of four people. Composition of groups remains unchanged throughout the game – the so-called *partner matching* protocol. They play a public good game for ten periods. We use a *between-subject* design, so that in a session subjects are only exposed to one of the following three treatment conditions.

[4]We chose these probability values because we want to examine decisions in two enforcement settings that were rather contrasting; but not so much that the probabilities of enforcement were close to the endpoints of the unit interval.

4.2.1 Certain Enforcement Treatment (P100)

This treatment builds on the standard design for the public goods game with punishment, with three differences. First, while Fehr & Gächter (2000) frame contribution decisions as an investment into a group project, we frame them as investment into a Public Account. Second, they use a convex punishment cost function while we adopt a linear one. Third, in the current experiment, group members' contributions are identified by an ID number when disclosed on the computer screen; contribution are always listed in the same ID column position[5], rather than randomly reassigned every period. Of these, we believe this last feature is potentially a major distinction from the standard design; it allows participants to create a link between the actions of other group members across periods. There are two reasons for that. First, that by allowing individualization, we reduce the possibility for indiscriminate punishment and make interpretation of data more transparent. Second, that by allowing subjects to track group member's contributions, we can investigate the extent to which punishment decisions are influenced by contributions in previous rounds. This is particularly important when, as in the treatments P80 and P20 described below, the opportunity to punish is intermittent, but we allow it in P100 too, to avoid confounding possible effects of probability of punishment enforcement with information differences.

[5]Although players could track a particular co-player's contribution record, they have no way of identifying that person. This matters (a) for ethical treatment of subjects and (b) for elimination of confounds if, for example, subjects might respond to information contained in the name (e.g gender, nationality, etc).

At the beginning of each of the ten periods, each subject is endowed with a fixed amount of 20 Rubis (the experimental currency used). Each period unfolds in two stages. In the first stage, subjects are required to simultaneously decide how much of their endowment to invest in a Public Account, say c_i, and, consequently, how much of it to invest in a Private Account, $20 - c_i$. Each Rubi a player allocates to the Private Account has a return of 1 for that player. A Rubi allocated to the Public Account yields a return of 0.4 for *every* player in the group. At the end of the first-stage, each subject is informed of the group's total investment, her income from the Public Account and her first-stage earnings (π), which is given by:

$$\pi_i^1 = 20 - c_i + 0.4 \sum_{i=1}^{4} c_i \tag{4.1}$$

Note that the total return of investment in the Public Account depends on the total investment made by the entire group. While each Rubi allocated to the Public Account yields a marginal private return of less than 1, by investing in the Public Account players in a group may obtain earnings that exceed those associated with full investment in the Private Account. Investments in the Public Account, given its non-rivalness and non-excludability, can be seen as contributions to a public good.

In the second stage, participants are informed of the investment decisions of their group members and given the opportunity to punish each group member by assigning "deduction" points. Each deduction point costs the punisher one Rubi and reduces the punished players' first-stage income by 3 Rubis. Each subject can assign up to 10 "deduction points" to each other group member.

Additionally, it is imposed that a subject cannot have her first-stage income, π^1, reduced below zero as a result of the punishment given her by others. Nevertheless, as she always carries the cost of punishment she does, her period income may end up negative depending on the total number of "deduction" points received and assigned[6]. Subject i's end-of-period payoff is given by:

$$\pi_i^2 = \left\{ \begin{array}{l} \pi^1 - 3(P_{-i,i}) - P_{i,-i} \quad \text{if} \quad 3(P_{-i,i}) < \pi^1 \\ -P_{i,-i} \quad \text{if} \quad 3(P_{-i,i}) \geq \pi^1 \end{array} \right\} \tag{4.2}$$

where $P_{-i,i}$ stands for the number of deduction points imposed on subject i by other group members, and $P_{i,-i}$ stands for the total number of punishing points assigned by subject i to other group members.

4.2.2 Uncertain Enforcement Treatments (P80 and P20)

The other two treatment conditions involve a similar game to the one played in the above treatment condition. The difference now is that one stage is added after the second stage, which we refer here to as the "enforcement" stage. Recall that in the second stage, subjects are informed of the contribution decision of each other group member and are given the opportunity to punish them. In the "uncertain enforcement" treatment

[6]As in Fehr & Gächter (2000), Nikiforakis (2008) and others, each subjects is given a one-time lump-sum payment of 25 Rubis at the beginning of the experiment to pay for negative payoffs they might incurr during the experiment.

conditions (termed P80 and P20), they do so with the understanding that their "punishment"[7] decisions may not be carried out. They will be so with a probability p. They are told that this probability is the same for all 10 periods of the experiment. Note that it is as if their punishment decisions were delegated to a central authority that, depending on the state of the nature, may fail to implement their decisions. Thus, there were two states of the nature, say $S \in \{s_1, s_2\}$, and punishment decisions are enforced only when $S = s_1$, where $P(S = s_1) = p$. To investigate the effect of uncertain enforcement of punishment on cooperative (or punishing) behaviour, we ran "high probability" (P80) and "low probability" (P20) sessions, in which p is 0.8 and 0.2, respectively. The P100 treatment can be viewed as the particular case in which $p = 1$.

In each period, whether or not punishment decisions are enforced is decided at a group level as follows: for each group, a ball is drawn from a bingo cage with replacement. The bingo cage has balls numbered from 1 to 10. If the ball for a given group is numbered 9 or 10 in the P80 condition, or 3 to 10 in the P20 condition, then punishment decisions are not carried out[8]. In these cases, a subject's end-of-period earnings are equal to her earnings in the first stage. Otherwise, punishment decisions are implemented and the final earnings in the period are given by the equation in (4.2).

To avoid there being any communication of disapproval when punish-

[7]We did not use this terminology ("punishment") in the experiment.

[8]Doing these realisations separately for each group reduced the danger that P100 and P80 would actually be the same or that P20 would actually have no enforced punishment.

ment is not enforced (i.e., nonmonetary forms of punishment, see Masclet *et al.* (2003)), punishment assigned to each individual in a given group is not disclosed unless it is enforced. So only when punishment decisions are actually implemented are subjects informed of the total punishment points they received from the group. In a similar fashion, assigning punishment points will not have any cost to subjects if punishment is not to be enforced. This could correspond to a case where opportunity to punish, as opposed to willingness to, may simply not arise. More importantly, this feature of our design avoids that one's profile of punishment decisions be "contaminated" by her unwillingness to pay for something that may not happen.

Thus, the information disclosed at the end of each period depends on the enforcement state: in case punishment is not enforced, subjects are shown their final earnings, which in this case is equal to their earnings from the first stage. in case punishment is enforced, they are shown (a) the total cost of the punishment points they assigned, (b) the punishment points they received in total from the group, and (c) the associated reduction in their earnings along with their final earnings in the period. All subjects are also informed of their own accumulated earnings, which is equal to the sum of earnings over all previous periods.

In all three treatments conditions, the parameters of the experiment (endowment, the return rate from the Public and Private Accounts, group size, payoff functions, number of rounds) are publicly announced to the participants.

4.2.3 Administration

There are ninety six subjects in this experiment. None of them had previously participated in a public good experiment at the University of Nottingham[9]. The subjects signed up for one of six sessions. At that point, they only knew that the experiment would take up to 90 minutes. Treatment conditions were randomly allocated to sessions, with two sessions per treatment condition.

Sixteen subjects took part in each session. Following their arrival, each subject received instructions explaining the experiment[10]. The instructions were read aloud while the students read them silently. To ensure subjects' understanding of the game's structure and payoff determination, each of them was asked to complete a control questionnaire. The experiment only proceeded when all subjects had answered it correctly. The experiment was conducted using z-Tree (Fischbacher, 2007). Sessions took around fifty minutes to be completed. At the end of the experiment, subjects were asked to complete a short questionnaire about themselves. Their earnings were converted into Sterling Pounds and they were then paid in cash. The exchange rate was 1 Rubi = 2.5 pence. Participants earned on

[9]Participants were recruited using ORSEE (Greiner, 2004), a subject-recruitment software that allows us to exclude those in the subject database that have been recorded as participating in previous public good experiments. For elimination of confounds, we did not want subjects, in particular those assigned to P80 and P20 treatments, who have had previous experience of P100.

[10]Instructions are included as an Appendix.

average £8.51, which included a show-up fee of £2 and a one-time lump-sum payment of 25 Rubis.

4.3 Theory: Effects of uncertain enforcement on cooperation and punishment

We now present predictions for cooperative and punishment behaviour. We start considering a standard game-theoretic case in which players are of the same type: they are all strictly concerned with their material payoff. Then, we consider a mixed case in which some of the players have fairness concerns.

4.3.1 Cooperation: Homogeneous players

Assuming that individuals are monetary payoff maximizers and that this is common knowledge among them, they should contribute nothing to the Public Account. In the presence of punishment opportunities this still holds true. The threat of punishment is non-credible as this is a payoff-reducing action. Therefore, subgame perfection dictates that individuals would always be better off by not punishing at all. It is straightforward to see that the equilibrium outcome regarding punishment does not change when punishment enforcement is risky. In this case, the actual infliction of punishment is conditioned on a probability distribution over a set of states of nature. Even so, costly punishment would not be a credible action by self-interested payoff maximizers regardless of the status of enforcement. Since individuals do not punish, one should contribute nothing just as in the "certain enforcement" case. Thus, within the standard game-theoretic framework, zero cooperation and zero punishment would be the subgame-perfect equilibrium strategies in all enforcement conditions.

4.3.2 Cooperation: Heterogenous players

If individuals have other-regarding preferences and are motivated by more than their pecuniary payoffs then no punishment and full defection may not be an equilibrium outcome. Fehr & Schmidt (1999) show, for instance, that if some people care about payoff equity, full cooperation can be sustained as an equilibrium outcome in a public good game with punishment. The intuition behind such result is that individuals who care about disadvantageous inequality will be willing to punish defectors despite it being costly to them. Such a threat, given the information set of players, would be credible enough to sustain cooperation. While the Fehr and Schmidt model is consistent with a continuum of contribution profiles, it predicts full cooperation by using a refinement argument. In fact, a number of experimental studies have shown that individuals are indeed willing to pay to punish defectors, and that high levels of cooperation can be sustained in the presence of punishment. But uncertainty is likely to change the decision setting. To get an insight into this, we use a simple model for a two-player case.

Let us start with some preliminaries. Consider a game G played by two players. Each player has a type that determines the preferences she acts on. Player 1 is purely self-regarding ("selfish") whose utility function $u_1(.)$ is defined on her own payoff, say π_1, in the game. Player 2 is inequity-averse with a utility function $u_2(.)$ defined both on her own payoff and the other player's payoff in the game, say π_2 and π_1. $u_1(.)$ has a linear form defined by $u_1(\pi_1) = \pi_1$. $u_2(.)$ has a Fehr-Schmidt functional form (see Fehr & Schmidt, 1999, p.822) defined by

$$u_2(\pi_1, \pi_2) = \pi_2 - \alpha \max\{\pi_2 - \pi_1, 0\} - \beta \max\{\pi_1 - \pi_2, 0\}$$

With this in mind, let G be the following complete information public good game with three stages. In the first stage, players decide simultaneously whether or not to contribute to the public good. Each player has an endowment of e, so that $c_i \in \{0, e\}$ $(i \in \{1, 2\})$ is the discrete set of strategies each player can employ. Payoff at the end of this stage is given by

$$\pi_{i,C}(c_i, c_j) = e - c_i + r(c_i + c_j) \ , \ r \in (1/2, 1)$$

where r is the return to each player from contributions to the public good. In the second stage, each player is informed about the other player's contribution and decide simultaneously whether or not to impose a punishment on the other player. Let $p_i \in \{0, \rho\}$ $(i \in \{1, 2\})$ be the discrete set of strategies each player can employ in this stage. This stage's payoff is given by

$$\pi_{i,P}(p_i, p_j) = -[p_i + lp_j] \ , \ l > 1$$

where $l > 1$ is the punishment impact rate, which indicates the first-stage payoff deduction when the other player chooses to punish. Finally, in the third stage "nature" chooses whether to enforce players' punishment decisions; "nature" enforces punishment with probability $q \in [0, 1]$. Note that players move in the second stage without knowing what is nature's choice. Thus, the monetary payoff of a player i is simply $\pi_i = \pi_{i,C} + \pi_{i,P}$.

Now, what is the prediction for this game under the assumption that players are of different *types*? More specifically, how are decisions in the first stage affected by the probabilistic enforcement of punishment decisions? The prediction is summarized in the following:

Proposition 1

(1.1) *If $q = 0$, then it is a dominant strategy for both players to choose $c_i = 0$.*

(1.2.) *If $q > 0$, then punishment sustains an equilibrium in which both players contribute if (i) $\alpha > \frac{1}{l-1}$ and (ii) $q > q^* = \frac{e(1-r)}{lp}$.*

Proof. See Appendix B ∎

The crucial implication of the above results is that the threat of punishment can only induce self-regarding players to contribute (hence, sustain full cooperation in the game) if the probability q of punishment enforcement is sufficiently high. Otherwise, if q is too low, the self-regarding type will free-ride because expected punishment is too low to deter defection. In this case, free-riding is also the best response of the inequity-averse player. Based on Proposition 1, we conjecture the following

Hypotheses 1 (Contribution) *The presence of punishment opportunities will not raise contributions if enforcement is perceived as "weak" (low-probability event) to a sufficiently high proportion of subjects. The lower q, the more free-ride types will be, breaking down prospects of sustained cooperation.*

Now, what are the predicted effects of probability enforcement on punishment behaviour? We know of no formal hypothesis that has been put forward which would allow us to predict the direction of punishment enforcement probability effects in subjects' punishment behaviour; and our previous basic framework regards only cooperative behaviour. Yet, we conjectured in the introduction to this chapter that the effect of imperfect enforcement on punishment is ambiguous; it would ultimately depend upon how backward- and forward-looking elements influence punishment decisions.

4.3.3 Punishment with forward-looking dominance

Instead, if forward-looking motives dominate punishment decisions, then we conjecture that the more uncertainty over punishment enforcement, the less punishment will be observed. The intuition is that the anticipation that punishment may not be enforced would weaken its strategic use: the ability to shape future interactions. If uncertainty creates a hindrance to individuals' ability to influence the future, then the more the uncertainty, the less punishment would be exercised by individuals.

Hypotheses 2.1 (Punishment) *If punishment decisions are dominantly driven by forward-looking motives, then punishment points assigned (not necessarily implemented) to free-riders and low contributions will be higher in P100 and P80 than in P20.*

4.3.4 Punishment with backward-looking dominance

If backward-looking motives dominate punishment decisions, then we conjecture the following: the more likely it is that free-riders can escape punishment due to enforcement failure, the more intense will be the willingness to punish them. The intuition here is that "bygones are not bygones" and players might get more intensely punished in a given round t for failing to contribute in t *and* in rounds prior to t. By punishment intensity we mean the punishment points assigned to a player per deviation of her contribution from others' average. Note that with weak enforcement in a repeated setting, there will probably be players, especially in treatment P20, with a history of free-riding behaviour that went unpunished because of other players' punishment decisions were not enforced. Thus, if punishment is dominantly backward-looking and mainly directed towards free-riders, then more punishment will be directed towards free-riders in P20 than in P100 and P80 treatments.

Hypotheses 2.2 (Punishment) *If punishment decisions are dominantly driven by backward-looking motives, then punishment points assigned (not necessarily implemented) to free-riders and low contributions will be higher in P20 than in P80 and P100.*

If that is the case, one may wonder whether this extra punishment would not compensate the low probability of enforcement and lead to high contributions. Note, though, that a more intense *willingness* to punish a free-rider may have no bearing on cooperation, as willingness to punish may not translate into actual punishment. Thus, even if others are willing to punish free-riders in P20 more intensely than, say, in P80, this may not necessarily induce cooperation levels in P20 as high as in P80.

4.4 Results of the experiment

4.4.1 Cooperative Behaviour

We start by examining contribution patterns across treatments. Figure 1 presents box plots of contribution to the Public Account over the 10 periods for each treatment condition.

Figure 4.1: Average contribution, by enforcement condition

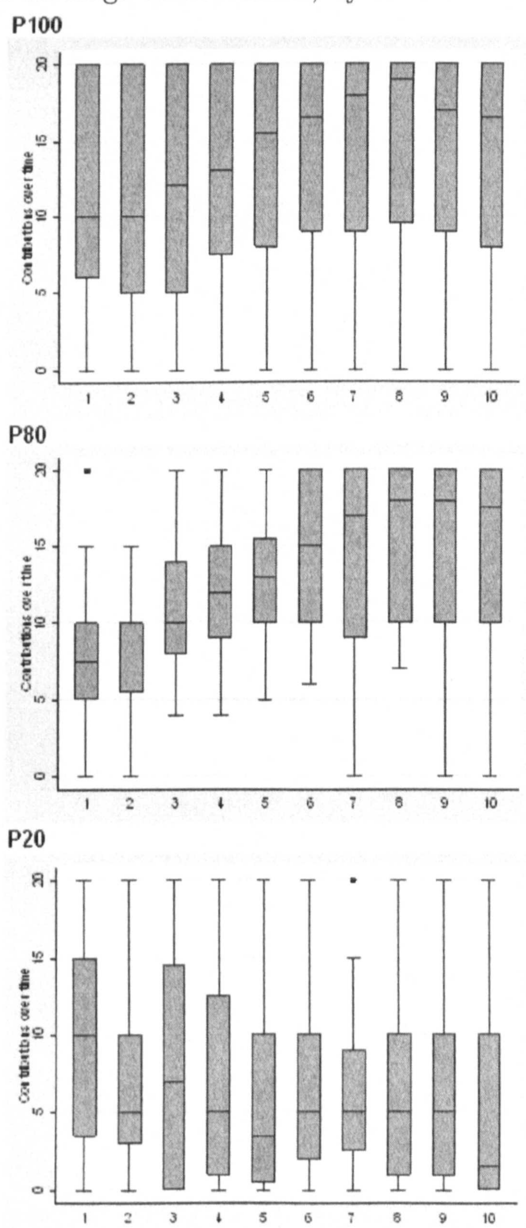

Each box plot describes key distributional features of the data. The median contribution value is shown as a line drawn across the box[11]. The variability in contribution is represented by first (lower hinge) and third (upper hinge) quartiles of the distribution in each period. Let this interquartile difference be H. The lower and upper adjacent values of the contribution in each period are shown as "whiskers" in the plot lines used to outline a box[12]. The upper adjacent contribution value represents the largest contribution between the upper hinge (uh) and the threshold value of $uh + 1.5 \times H$; the lower adjacent contribution value, in turn, represents the smallest contribution between the lower hinge (lh) and the threshold value of $lh - 1.5 \times H$. Dots outside the box plot identify contributions that lie unusually far from the main body of data[13].

In the box plots for each treatment, by following the line drawn across the box at the median, one can see the evolution of median contribution to the Public Account over the ten periods. Contributions under the P100 condition, for instance, are in line with previous experimental findings: they start at roughly half of subjects' endowments and keep increasing over time. This result confirms that the existence of punishment can improve cooperation over time. Additionally, and perhaps more importantly, it suggests that the ability of punishment to sustain cooperation is unaffected

[11]In the second period of P80, the line representing the median contribution seems absent of the box plot. This is because it coincides with the upper quartile (10).

[12]For contribution data in some periods, these adjacent values coincide with the first (P20, periods 3 and 10) or third quartile (e.g. all periods in P100) of the distribution of contributions.

[13]Any contribution which lies more than three times the inter-quartile range either lower than the first quartile, or higher than the third quartile falls into this category.

by knowledge of contribution histories.

There is, however, clear separation in contributions between the uncertain enforcement treatments. While median contributions in the P80 condition increase over time, closely following contributions in the certain enforcement condition, it is clear that contributions in the P20 condition are noticeably lower and on a divergent path compared to the P80 condition. While median contributions in the P20 condition start higher than contributions in the P80 condition, they keep decreasing from the second period on, while contributions in the P80 treatment increase over time. This suggests that the existence of punishment opportunities is not effective in raising contributions if enforcement is perceived as a "low" probability event.

This result, based on visual inspection, is indeed confirmed by non-parametric tests. We conduct pairwise Mann-Whitney tests between treatments for each period at a time in order to test for equality of distribution of mean contribution of groups between all enforcement treatments[14]. Test statistics are reported in Table 4.1.

Two features stand out in the test results. First, they show that, apart from the first period, there are no statistically significant differences in mean contribution of groups in the P100 and P80 treatments. Second, they also show that there are statistically significant differences between

[14]It is worth noting that the sample of observations from a given treatment is formed by the mean contribution of groups of players in a given treatment. This is so because individuals' contributions, while independent across samples, are not independent within treatments – which violates an assumption which the test relies on.

Table 4.1: Are groups' mean contribution different across enforcement treatments? Pairwise Mann-Whitney Tests

	Test Statistics		
Period \ Treatment Conparison	*P100 vs P80*	*P100 vs P20*	*P80 vs P20*
1	$z = 2.10$ $p = 0.03$	$z = 1.31$ $p = 0.19$	$z = -0.53$ $p = 0.60$
2	$z = 1.26$ $p = 0.21$	$z = 1.79$ $p = 0.07$	$z = 0.63$ $p = 0.53$
3	$z = 0.10$ $p = 0.92$	$z = 1.37$ $p = 0.17$	$z = 1.34$ $p = 0.17$
4	$z = 0.74$ $p = 0.46$	$z = 1.68$ $p = 0.09$	$z = 1.47$ $p = 0.14$
5	$z = 0.58$ $p = 0.56$	$z = 1.79$ $p = 0.07$	$z = 2.05$ $p = 0.04$
6	$z = 0.47$ $p = 0.63$	$z = 1.79$ $p = 0.07$	$z = 2.26$ $p = 0.02$
7	$z = 0.53$ $p = 0.60$	$z = 2.11$ $p = 0.03$	$z = 2.37$ $p = 0.02$
8	$z = 0.58$ $p = 0.56$	$z = 2.21$ $p = 0.02$	$z = 2.73$ $p = 0.00$
9	$z = -0.32$ $p = 0.75$	$z = 1.90$ $p = 0.05$	$z = 2.53$ $p = 0.01$
10	$z = -0.47$ $p = 0.63$	$z = 2.42$ $p = 0.01$	$z = 3.00$ $p = 0.00$
All periods	$z = 1.23$ $p = 0.21$	$z = 6.15$ $p = 0.00$	$z = 6.32$ $p = 0.00$

mean contribution of groups in P20 and either P100 or P80 treatments after the initial periods of the game (in most periods, at the 5% level of significance). Both features are more salient when considering test results involving data from all periods pooled together.

We need, however, to examine the robustness of these results. The no-parametric tests do not capture intertemporal dependencies in group contributions and may confound treatment effects. We then turn to a more formal analysis of the data; we do so by running a regression of individual contributions on treatment and individual variables. The panel structure of the data allows us to handle some degree of individual heterogeneity and obtain more consistent estimates of treatment effects.

We estimate an empirical model relating contribution to individual and structural parameters of the game that largely follows a common specification in these studies (e.g. Andersen *et al.* , 2006c; Nikiforakis, 2008). But our econometric specification also includes lagged variables that seek to capture recursive elements in contribution decisions. The underlying reason for this is hardly controversial: in repeatedly played games, individuals tend to reciprocate actions of other players; this produces behaviour that is largely reactive and influenced by past outcomes (see, e.g. Fischbacher *et al.* (2001), Frey & Meier (2004) and Gunnthorsdottir *et al.* (2007)). The model then has the following specification:

$$c_{i,t} = \beta_0 + \beta_1 \bar{c}_{-i,t-1} + \beta_2 \left(P^R_{i,t-1}\right) + \beta_3 \Sigma E_{i,t-1} + \beta_4 P80 + \beta_5 P20 + \mathbf{z}_i' \alpha + \mathbf{u}_{i,t} \tag{4.3}$$

where the $\bar{c}_{-i,t}$ is the average contribution of the other group members in period t, $P^R_{i,t}$ is the total punishment points actually *received* by individual i in period t – which is 0 if punishment decisions were not enforced. $\Sigma E_{i,t}$ is the number of previous periods in which punishment was enforced in the group i belongs to; this is meant to capture the effects of the particular sequence of enforcement experienced by i. $P80$ and $P20$ are dummy variables that equal one if individual i is taking part in the "high" or "low" probability of enforcement condition, respectively. Components of $\mathbf{z}$ will control for the variation strictly related to some subject-specific attributes (gender, ethnicity, etc). Dummy variables to control for group effects are included. $\mathbf{u}_{i,t}$ is a composite error term including a subject-specific random intercept and a purely random disturbance term which is assumed to be i.i.d. over i and t.

Table 4.2 reports the results of the generalized-least-squares regressions of the model in (4.3). Contributions are, on average, positively affected by retaliatory behaviour from others in the past: the actual number of punishment points received in the previous period as well as the number of periods in which punishment points were actually enforced have both significant and positive effect on contributions. Of interest in the results is the estimation of the "low" probability of enforcement treatment effects on contribution decisions, which in the case of this model consists of estimated value and significance of the parameter in front of the dummy variable $P20$. Even after controlling for the different enforcement conditions and group effects (interaction and sequence of enforcement experienced by groups), one can see that contributions from subjects in the low-probability of enforcement treatment are lower than contributions in both certain and "high" probability of enforcement conditions. $P20$ is, in fact, the only enforcement treatment whose effect on contributions is statistically significant. Thus, parameter estimates of the model support the raw results depicted in Table 4.2.

Therefore, as was apparent in Figure 1, there are significant differences in contribution estimates between "high" and "low" probability of enforcement conditions. The mere knowledge that sanctions may be imposed to punish those regarded as free-riders cannot induce cooperative behaviour if punishment enforcement is viewed as "weak". Based on the non-parametric and regression analysis one can conclude that the experimental data support the following:

Result 1. *The threat of punishment can only promote cooperative behaviour if enforcement is* perceived *as a high-probability event.*

4.4.2 Punishment Behaviour

The next issue to be examined is whether and how subjects' willingness to punish is affected by the possibility of not having their punishment decisions enforced. To get an intuition on this, we begin with some descriptive statistics.

Table 4.3 presents the frequency of individuals who assign no punishment. Two things are worth noting: first, that there is a considerable amount of "free-riding" behaviour on punishment efforts across treatments. In most periods, the option to punish is exercised by less than half of the subjects. Second, that there is more punishing of individuals in the first

Table 4.2: Do enforcement treatments affect contributions

Independent variable	Dependent variable: individual i's contribution in period t to the public good
Constant	3.514* (0.381)
Others' average contribution (t-1)	0.828* (0.019)
Punishment received (t-1) (= Enforcement state X Punishment assigned)	0.116* (0.049)
Number of previous periods with enforcement $\left(\sum_{k=1}^{t-1} E_{t-k}\right)$	0.021 (0.04)
Enforcement condition P80	-0.085 (0.235)
Enforcement condition P20	- 2.683* (0.919)
Female	- 0.854 (0.205)
Number of observations	864
Wald χ^2	4600.6*

Notes: The regression reports GLS estimates with individual random-effects. Values in parenthesis are standard errors. Estimates are heteroscedastic-consistent. * Significance at the 1% level. Dummy variables for groups are included.

Table 4.3: Fraction of subjects who assign no punishment points

Period	P100	P80	P20
1	0.31	0.53	0.57
2	0.65	0.66	0.63
3	0.50	0.53	0.63
4	0.53	0.59	0.63
5	0.63	0.66	0.56
6	0.69	0.63	0.72
7	0.66	0.53	0.72
8	0.69	0.56	0.69
9	0.56	0.63	0.72
10	0.66	0.66	0.63

period of the certain enforcement conditions than there is in the uncertain enforcement conditions.

Examining punishment points assigned to subjects, we find that individuals in the P100 condition were assigned more punishment points (2.19) on average than those who are in the P80 and P20 conditions (1.41 and 1.53). A Mann-Whitney test shows that this difference in punishment in the first period is significant at 5% level of significance (P100 versus P80: $p < 0.0224$; P100 versus P20: $p < 0.0507$). A natural question to ask is why subjects in P100 are imposing more sanctions relative to P80 and P20 conditions in the beginning of the game?

It is not that there is a great deal more of free-riding behaviour in the P100 relative to the other two conditions. In the first two periods, only one subject in the P100 contributes nothing to the Public Account, against five and four subjects in the P80 and P20 conditions, respectively. While the fraction of subjects who in the first period contribute less than the group average is slightly greater in P100 (59%) than it is in P80 and P20 conditions (56% and 71%), average contribution in P100 is actually higher (11.03) than it is in P80 and P20 treatments (7.78 and 9.03, respectively). Kolmogorov-Smirnov tests provide a second bit of evidence consistent with that: they show that one cannot reject the hypothesis that there is no significant differences in the distribution of deviations from others' average contribution between treatments in the first period (P100 versus P80: $p < 0.627$;P100 versus P20: $p < 0.964$; P80 versus P20: $p < 0.627$).

A possible interpretation of this first-period differences in punishment between treatments is that subjects in the certain enforcement condition are trying to discipline behaviour from the beginning by signalling "toughness" with free-riders and low-contributors. Yet, this strategic reputation building would be mitigated among subjects in P80 and P20 enforcement conditions. Because they know that their punishment decisions may fail to be enforced, they would be unwilling to accept the cost of enforced punishment as the potential "extra" cost of such strong signals early in the game may not be compensated by higher cooperation levels later in the game. This is likely to be the case of a forward-looking subject who believes that punishment will only work if it is enforced frequently, in which case it would be rational not to punish in P20 even though unenforced punishment is costless.

Graphics in Figure 4.2 show other interesting aspects of punishment behaviour in each enforcement condition. Each figure plots the average punishment points assigned to i by range of deviation from the others' average contribution over time. There are four classes of deviation: 1 (-20,-10], 2 (-10,0], 3(0,10], and 4 (10,20]. For example, the leftmost cluster of ten columns in each figure shows the average punishment assigned to individuals whose contribution is between 10 and 20 Rubis *less* than the average contribution from other group members (Deviation class 1). The other three ranges of deviation move towards a positive domain as one moves to the ten-sequence cluster of columns on the right-hand side of the x axis of the graph.

Visual inspection of these plots suggests three things. First, that there are some sort of first- and last-period effects. Note that the willingness to punish free-riders and low contributors (those in deviation class 1) is stronger at the beginning and at the end of the game. Second, that there is no "anti-social punishment" when enforcement is weak: individuals whose contribution is above the average contribution from the other group members – deviation classes 3 and 4 – are barely punished in P20. Third, that "negative deviators", especially in the middle rounds, are more intensively punished in P20 than they are in P100 and P80: individuals whose contribution falls short of the average seem, on average, to have more punishment points assigned to them in P20 treatment than in the certain and P80 enforcement conditions.

We now perform an econometric analysis of treatment effects on punishment behaviour. We regress the amount of punishment assigned to a player on lagged contribution treatment and structural parameters of the

Figure 4.2: Average punishment, by range of deviation from others's average contribution

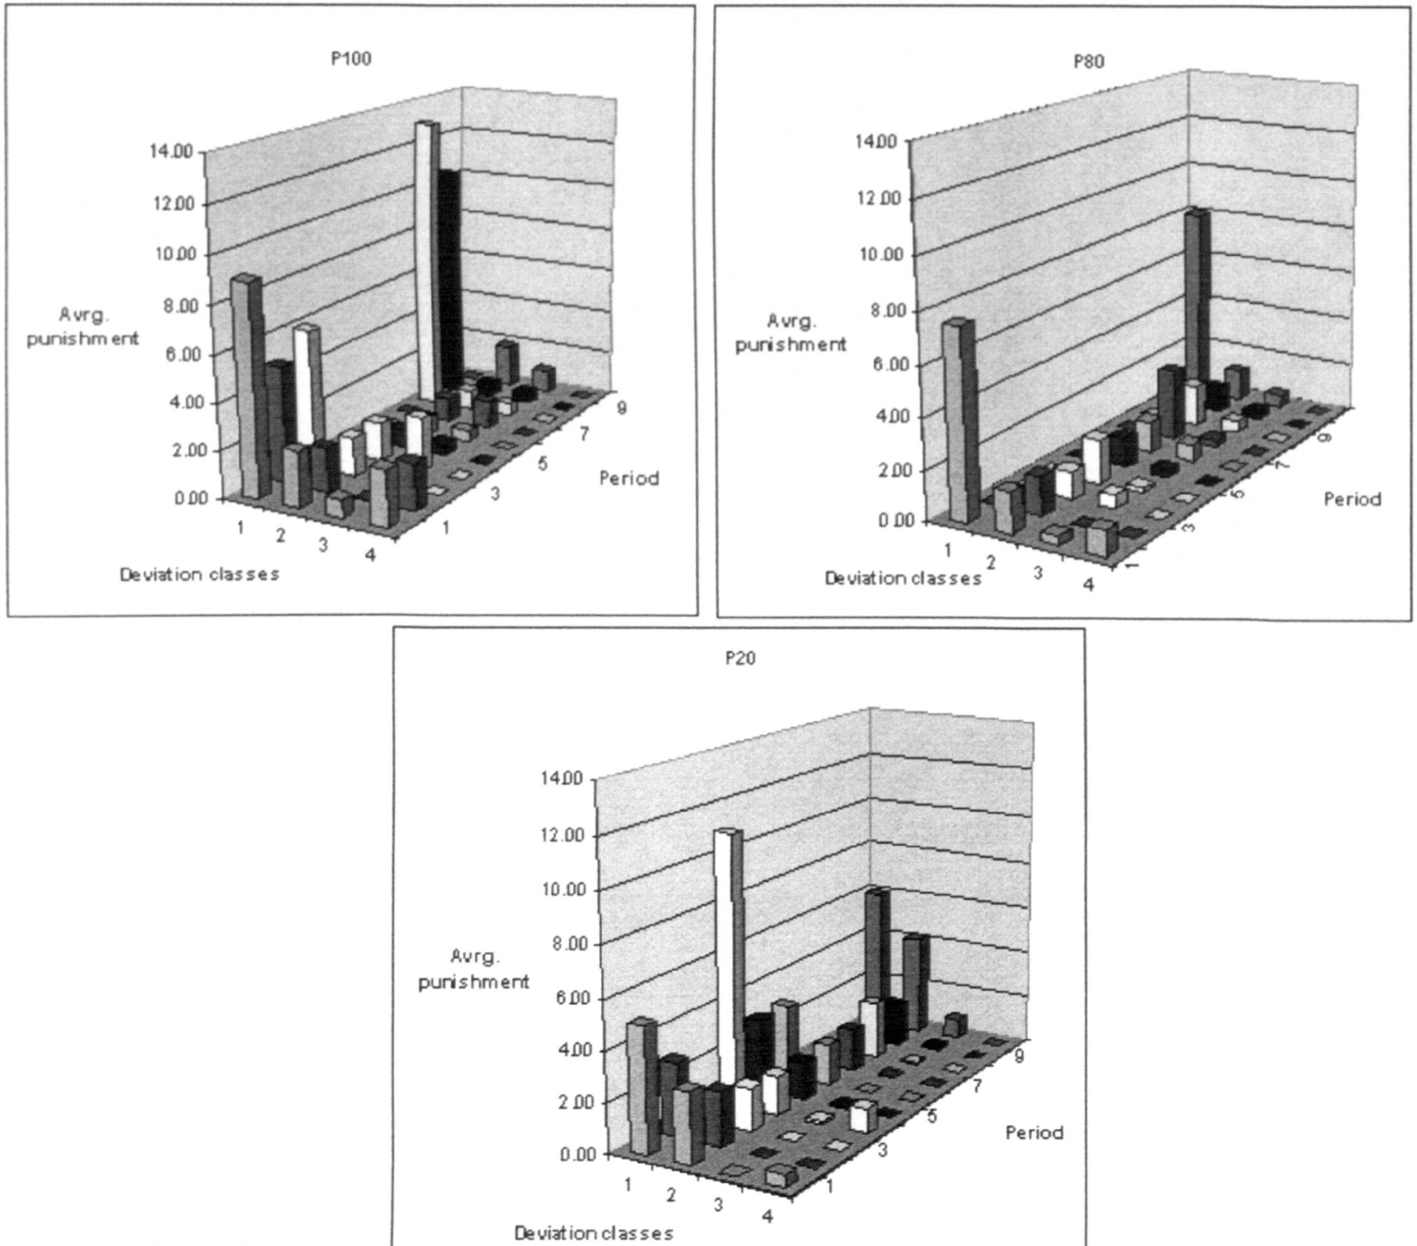

game. The general empirical model has the following form:

$$P_{i,t} = \beta_0 + \beta_1 \bar{c}_{-i,t} + \beta_2 POSDEV + \beta_3 NEGDEV + \dots \tag{4.4}$$
$$+ \beta_4 ANGER + \beta_5 P80 + \beta_6 P20 + \dots + \mathbf{z}'\alpha + \mathbf{u}_{i,t}$$

where $P_{i,t}$ represents the number of punishment points assigned to subject i, $\bar{c}_{-i,t}$ is the average contribution from other group members, $POSDEV$ and $NEGDEV$ are the absolute values of the deviation of i's contribution from other group members' average. We follow here (Fehr & Gächter, 2000), including them as separate regressors. One of those variables is zero depending whether i's contribution is either above (or equal) or below the others contribution. $ANGER$ denotes all the punishment points assigned i that have not been actually enforced over the previous periods. $P80$ and $P20$ are dummy variables that are equal to 1 if i is in the $P80$ or $P20$ enforcement treatments and 0 otherwise. Due to the random assignment of participants to treatment conditions, those dummies allow us to isolate the effect of enforcement conditions on subjects' willingness to punish. $\mathbf{z}$ is a vector of other dummies and interaction terms between treatment conditions and deviation from i's contribution from other group members' average that try to capture different levels of intensity of punishment assignment in each treatment condition. We include, for instance, a dummy regressor for the last period to capture last period effects on punishment decisions. $u_{i,t}$ is the compound error term. Parameter estimates of model 4.4 are presented in Table 4.6 column (1).

We also separate the data according to enforcement treatments and run separate regressions for each sub-sample of subjects. This allows us to examine our conjecture (see Section 3) that, because punishment is likely to be less frequent in P20 than in P8 and P100, subjects will assign punishment differently across enforcement treatments. These results are reported in Columns (2)-(4).

Beginning with the estimates of the general model in column (1), we notice that enforcement conditions do have an effect on punishment decisions: subjects in the uncertain enforcement conditions punish relatively less. We have conjectured that this effect has to do with the impact of uncertainty over punishment enforcement on the strategic value of punishment: players would be less inclined to punish if enforcement failure threats their ability to send a signal to free-riders on a consistent basis.

Looking across the treatments, there are other noticeable aspects influencing punishment decisions. First, we see that an increase in the group average contribution induces a reduction in punishment. This holds for all but the P80 treatment. Second, that punishment is mostly directed towards free-riders, those who contribute below the group average. These two results illustrate the elements of reciprocity in individuals' behaviour. Third, we find that "bygones are not bygones": the more punishment points towards a player ended up not being enforced in the history of the game –

Table 4.4: Do punishment decisions differ by treatment?

Independent variable	(1)	(2)	(3)	(4)
Other's average contribution	-0.012*	-0.039*	-0.006	0.015**
	(0.005)	(0.013)	(0.005)	(0.008)
Negative deviation	0.483*	0.457*	0.524*	0.410*
	(0.010)	(0.019)	(0.024)	(0.028)
Positive deviation	0.002	0.012	0.007	-0.052
	(0.005)	(0.014)	(0.009)	(0.020)
"Accumulated anger"	0.138*	---	0.183*	0.130*
	(0.012)		(0.027)	(0.014)
P80	- 0.270*	---	---	---
	(0.069)			
P20	- 0.505*	---	---	---
	(0.105)			
First period	0.002	0.081	-0.041	0.318**
	(0.058)	(0.139)	(0.060)	(0.173)
Last period	0.052	0.031	-0.063	0.973*
	(0.052)	(0.131)	(0.059)	(0.163)
Constant	0.541*	0.942*	0.129	-0.130
	(0.106)	(0.247)	(0.093)	(0.104)
Sample data	Pooled	P100	P80	P20
Wald $\chi 2$	2673.84*	731.79*	576.40*	478.70*
Number of observations	960	320	320	320

Notes: Dependent variable is the punishment points assigned in total to subject *i* at the end of a given period by other group members. Standard errors are in parentheses. * Significance at a p-level of 1%. ** Significance at a p-level of 10%. Estimates take into account error correlation within a subject's sequence of observations and correct for heteroskedastic error structure across panels. "Accumulated anger" consists of the number of punishment points in all previous periods that have been assigned to *i* but not enforced.

what we term "accumulated anger" –, the more punishment from others is directed to her. This can arguably indicate that punishment decisions are driven by emotions and not only by intertemporal concerns with material payoff.

All in all, these results seem to support the view that punishment is driven by a mix of backward- and forward-looking motives. The uncertainty over whether the willingness to punish one will be materialised over the course of the game seems, on the one hand, to weaken the strategic value of punishment in shaping future behaviour; on the other hand, because of the history of free-riding that goes unpunished, it also creates frustration and increasing "anger" towards those who have gotten "off the hook".

It should not come as surprise, therefore, that punishment in the first and the last periods is statistically significantly different from punishment over the other periods of the game in the P20 treatment. Since there is no strategic incentive to punish relatively more at the end of the game, this seems to suggest that individuals are pursuing some revenge for something they deemed as unfair during the game. Indeed, the last round of the game is the only round in which i can punish other group members without any danger of repercussions.

Thus, the results from the regression results suggest that the existence of uncertainty on whether punishment decisions will be carried out has statistically significant effect on punishment levels. The following result summarizes the findings of this section.

Result 2: *The willingness to punish free-riders is affected by the "uncertainty" over whether punishment will be actually enforced. In both uncertain treatments, individuals tend to punish less. There is a backward-looking element in punishment decisions as the more an individual has escaped being punished in the past, the more punishment is directed to her.*

4.4.3 Welfare Analysis

In addition to looking for differences in punishment behaviour across treatments, we now investigate how "uncertain" enforcement affects individuals' welfare. The key difficulty in addressing this issue is that aspects that are likely to affect individuals' utility in this experiment are not directly measured. For instance, there must be gains in utility from punishing a free-rider as much as there are losses in utility from not being able to punish a free-rider because of an enforcement failure. We sidestep this problem for a while, and following Nikiforakis (2008), we use individual earnings as a proxy measure for welfare. Using the certain enforcement treatment as a benchmark, we begin by examining whether earnings are increased in the "uncertain" enforcement conditions.

Table 4.5 provides an overview of how earnings look like in each enforcement condition. While average contributions are slightly lower in P80 than they are in P100, subjects in P80 have higher earnings on average than subjects in P100. We have seen that contribution levels are similar between P100 and P80 treatments, despite the fact that in P80 punishment decisions might not be enforced with a probability of 20%. While such a possibility weakens the threat of punishment, the degree of enforcement was sufficient to lead to an increase in contributions over time. Since punishment assignment is costless in P80 if punishment is not enforced, it should make intuitive sense then that subjects in P80 could benefit from higher contributions without necessarily incurring punishment costs in every period. As a result subjects in P100 have lower earnings than subjects in P80.

Table 4.5: Earnings by enforcement treatment

Treatment	Average contribution	Average earnings after contribution	Punishment-associated costs[1]	Average ings
P100	13.20	27.92	5.98	22.75
P80	12.67	27.60	3.11	24.54
P20	6.82	24.09	1.63	22.46

1 Average costs of punishment points given out (and enforced) to other group members and average deductions in first-stage payoff as a consequence of punishment received.

Let us now look at what happens in P20. Compared to conditions where the sanction system in place has stronger enforcement, earnings in P20 are lower. As discussed before, for most groups in P20 punishment enforcement occurred in few occasions irregularly spaced over the ten-period sequence. This created a "disbelief" in the enforcement system, encouraging more free-riding behaviour and, consequently, leading to a decline in contributions over time. Yet, note that average earning differences between P20 and P100 are not great; this highlights, as showed by others (e.g. Fehr & Gächter (2002); Sefton *et al.* (2007)), that in public goods experiments with certain enforcement, the benefit of higher contributions may be outweighed by punishment costs.

We now turn to a formal assessment of the relation between individual earnings and enforcement conditions. Table ?? reports regressions which examine treatment effects on individual earnings. Estimated results are reported in columns (1) and (2).

The first regression results show that none of the enforcement treatments have a significant effect on period earnings. This is in line with the general impression that average earnings do not show much variation across treatments (see Figure 4.3).

Figure 4.3: Average earnings

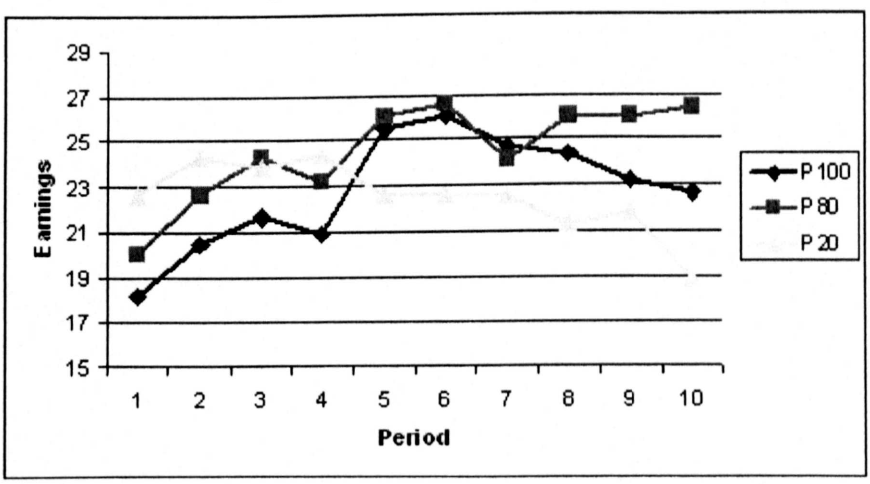

P100 and P80 follow similar trends in terms of contribution and punishment-associated costs yielding similar earning levels. In P20, average earnings in each period are higher at the first half of the experiment compared to P100 and P80, as in these latter treatments individuals are costly trying to induce higher contributions. Yet, earnings in P20 decrease in the second-

Table 4.6: Determinants of earnings (GLS Regressions)

Independent variable	(1) Model without period-treatment interactions	(2) Model with period-treatment interactions
P80	1.796	1.560
	(1.168)	(1.421)
P20	-0.290	4.881***
	(1.168)	(1.421)
Period	0.215***	0.514***
	(0.619)	(0.1039)
Period x P80	-----	0.0430
		(0.1470)
Period x P20	-----	-0.9403***
		(0.1470)
Constant	21.562***	19.917***
	(0.893)	(1.004)
Number of observations	960	960
Wald χ^2	15.84***	73.77***

Notes: Dependent variable in (1) and (2) is the earnings of subject i at the end of a given period. Regression results reports GLS estimates with individual random-effects. Standard errors are in parentheses. *p<0.10, ** p<0.05 and *** p<0.01.

half of the experiment as a result of the decay in contributions, while in P100 and P80 individuals are benefiting from higher contributions relative to the first half. The second regression adds interaction terms between period and enforcement treatments, $Period \times P80$ an $Period \times P20$. These variables try to capture time trends in P80 and P20 with respect to that in P100. The coefficient for the interaction term $Period \times P20$ is negative and significant. This just confirms what we have seen before: that there is a continuous decrease in average period earnings over time in P20 as a consequence of decline in contributions, whereas earnings are kept at higher levels in P100. Note, however, that earnings in P80 are higher with respect to that in P100. This is the result of a more significant increase in contributions and that in some period individuals in P80 need not incur any punishment cost.[15] Result 4 summarizes.

Result 4: *The highest welfare level, measured by accumulated earnings, is found in P80. While punishment enforcement is not certain in P80 like it is in P100, individuals in P80 condition benefit from higher contributions without incurring punishment costs in every period.*

[15]Not only less punishment points were assigned in P80 (326 in total) with respect to those in P100 (422), but as punishment failed to be enforced sometimes in P80, 77 out of 326 punishment points did not result in any cost.

4.5 Conclusions

Using a public goods experiment, we investigate whether cooperation prospects in a social dilemma situation can be affected when sanction opportunities are present but their enforcement is not certain. By not certain we mean that enforcement in each period of the game happens with a known probability p. The game is played under three treatment conditions, which differ only by the value of p (1, 0.8, or 0.2).

One of the findings is that punishment opportunities do not promote cooperative behaviour when enforcement is perceived as "weak" (treatment in which $p = 0.2$). In this case, average contributions start at around half of subjects' endowment and keep declining over time. This contrasts with the levels of cooperative behaviour observed in the treatment where punishment enforcement is perceived as "strong" (case in which $p = 0.8$): average contributions are raised and sustained at a high level. This result is somewhat comforting as it suggests that a sanctioning system with some degree of "imperfectness" can still induce cooperative behaviour in social dilemma situations. It also indicates that the deterrence effect of a sanctioning system operates through the perception it induces regarding either detection or enforcement likelihood. This result is in line, for example, with the evidence that income tax compliance increases when taxpayers are simply threatened to have their income reports "more closely examined" (see Slemrod *et al.* , 2001). Tax compliance, which is a form of cooperative behaviour, is promoted not by a threat of more severe punishment, but by a change in the likelihood of being detected.

Another finding is that punishment of free-riders and low contributors in general is more intense at the beginning and the end of the game. While this could be rationalized as a compromise between strategic (reputation building) and emotional (vindictiveness) components of individual's decision making, it is still unclear how to interpret these phenomena within a rational framework. Such end-of game effects, in particular, may have implications for the theoretical study of iterated prisoner's dilemma type of games as they hint at the existence of path-dependencies in the play of the game.

It is also observed that individuals in P80 condition benefit from higher contributions without incurring punishment costs in every period and, as a result, accumulated earnings in P80 are higher than in P100. Having said that, the P80 and P20 enforcement treatments have no significant effect on average earnings. This serves as an indication that, while failing to lead to high contributions, a punishment mechanism with "weak" enforcement might lead to similar earnings to the treatment where enforcement is certain.

The major finding of our experiment – that, put loosely, subject's perception of the likelihood of punishment enforcement matters – raises some interesting questions: in a repeated setting, can a strong threat of punishment deter individuals from deviating from a collective optimal course of action without it being ever 'demonstrated'? Can a history of "punishment" itself sustain cooperation in social dilemmas without a strongly credible threat of "punishment"? Can a threat of "punishment" efficiently induce and sustain high levels of cooperation in social dilemmas without a periodic demonstration of "punishment"? In sum, which "mixes" of threat and punishment history can induce cooperation?

It is worth noting that from a theoretic standpoint a threat should suffice. In the theory of infinitely repeated games, it is a classic result that it is possible to achieve a subgame-perfect equilibrium in which players achieve the highest payoff of all existing equilibria (Friedman, 1971). The core idea underlying this result is that a given player is persuaded to follow such perfect equilibrium strategy by *threatening* her with the strongest credible punishment. Punishment may not necessarily be history-dependent (Abreu, 1988). But while the perfect equilibrium strategy profile specifies punishment for deviations, the outcome path ends up not involving any imposition of punishment – the simple *threat* of punishment has a deterrence effect. Some may view results from Fehr & Gächter (2000) and many other studies as lending support to this: looking at first-period data, when there is no history of play, one can see that contributions are significantly higher in punishment treatments relative to no-punishment treatments. While in an experimental setting there is some degree of uncertainty over the size of punishment in terms of payoff, the simple threat of punishment often encourages pro-social behaviour.

But our experiment raises some questions. Its results suggest that the incentive constraints implicit in such punishment schemes may not rely only on credibility, but also on what we term here "punishment demonstration" – that punishment must be exercised upon subjects. We have found that an "imperfect" sanction system (in terms of enforcement) can achieve higher levels of pro-social behaviour by simply changing, through probability manipulation, subject's *perception* about the likelihood of sanction enforcement. It is unknown though to what extent the efficacy of punishment in inducing cooperative behaviour depends on perceived credibility of punishment threat (probability) and the factual history of the game. We view this as of theoretical and empirical relevance. Our experiment, like all other experimental studies on cooperation with sanction systems, threat and "punishment demonstration" are entangled. We intend to investigate the influence of these factors in further research.

APPENDICES

4.6 Appendices

4.6.1 Appendix A - Proofs

Proof. of **Proposition 1: Part 1.1.** If $q = 0$, then only the first-stage is payoff-wise relevant. In this stage, the strategy $c_i = e$ is strictly dominated by $c_i = 0$, since $\pi_{i,C}(c_i = 0, c\_i = 0) > \pi_{i,C}(c_i = e, c\_i = 0)$ for both players. The game has then a unique Nash equilibrium in which both players defect: $(c_1 = 0, c_2 = 0)$, given their payoff outcomes, is the pair of strategies that maximises the utility of the self-regarding player, $u_1(.)$, and the utility of the inequity-averse player, $u_2(.)$, for all $\alpha > 0$. ∎

Proof. of **Proposition 1: Part 1.2.**
If $q > 0$, then actions in the punishment stage may have payoff consequences for both players. For the self-regarding player, imposing no punishment, i.e. $p_1 = 0$, is a dominant strategy as $p_1 = 0 = \arg\max u_1(\pi_{i,C}(c_i, c\_i) + \pi_{i,P}(p_1, p_2)$ for all cooperative and punishment strategies of the inequity-averse player. Since this is common knowledge, it is easy to see that the inequity averse player will choose no punishment, i.e. $p_2 = 0$, at the second-stage of the game. Now, the inequity-averse player will also choose no punishment, $p_1 = 0$, when the profile of actions chosen in the first stage are $\{c_1 = e, c_2 = e\}$, $\{c_1 = 0, c_2 = 0\}$, $\{c_1 = e, c_2 = 0\}$; in all these cases, the inequity-averse player cannot be better off by choosing to punish, i.e. $p_2 = \rho$. For instance, choosing $p_2 = \rho$ following $\{c_1 = e, c_2 = e\}$ is dominated by $p_2 = 0$ since $r2e > r2e - (\rho + \beta(\rho + l\rho))$ for all $\beta > 0$. But the inequity averse player will punish following the first-stage pair of strategies

$\{c_1 = 0, c_2 = e\}$ if the final payoff when she assigns punishment to the self-regarding player is larger than the final payoff of not doing so, that is, if $\pi_1(c_2 = e, p_2 = \rho) > \pi_1(c_2 = e, p_2 = 0)$. This amounts to the following condition

$$(re-\rho)-\alpha max\{(e(1+r)-l\rho)-(re-\rho),0\} > re-\alpha max\{(e(1+r))-(re),0\}$$

which holds only if $\alpha > \frac{1}{l-1}$. Note that in this case the threat of punishment can only induce the self-regarding type to cooperate, and she will have no incentive to deviate from that, if the final payoff from cooperating is larger than the expected final payoff from free riding in the first stage and getting punished in the second stage

$$r(2e) > (1-q)[e(1+r)] + q[e(1+r) - l\rho]$$

which holds only if $q > q^* = \frac{e(1-r)}{l\rho}$. This completes the proof. ∎

4.6.2 Appendix B - Instructions

INSTRUCTIONS

CONDITION: CERTAIN PUNISHMENT ENFORCEMENT

WELCOME

Thank you for participating in this study. Your earnings in this study will be paid to you in cash at the end of the session. The instructions are simple. If you follow them carefully, you may, depending on your decisions, earn a considerable amount of money.

IMPORTANT NOTE

Please do not communicate in any way with other participants during this experiment. Please remember to switch off your mobile. If you have a question or problem at any point in today's session, please raise your hand and I will come to you.

THE EXPERIMENT

During this experiment Rubis will take the place of traditional monetary units. At the end of the experiment the total amount of Rubis you have earned will be converted into Pounds at the following rate:

$$1 \; Rubi \; = \; 2.5 \; Pence$$

Each participant receives a lump sum payment of 20 Rubis at the beginning of the experiment. This one-time payment can be used to pay for eventual losses during the experiment. However, you can always avoid losses with certainty through your own decisions. At the end of the experiment, your entire earnings from the experiment plus the lump sum payment will be converted to Pounds and immediately paid to you in cash.

This experiment is divided into **10 periods**. In each period, participants will be divided into groups of four. You will therefore be in a group with three other participants. During these 10 periods the group composition is constant. You are, therefore, grouped with the same people throughout the experiment. Your earnings will depend on the decisions that you and the other members of your group make. You will never learn whom you have been grouped with.

Each period consists of **two stages**. In the first stage, your task is to decide how many Rubis you want to invest in each of two investment accounts. One account is a Private Account, which only you benefit from. The second account is a Public Account, the benefits of which are shared equally by all members of your group. In the second stage you will be shown the amount invested in the Public Account by the three other members of your group. Your task is to decide whether you want to reduce their earnings from the first stage by distributing "deduction points" to them.

In what follows, these stages are described in more detail.

The First Stage

At the beginning of each period, each participant receives 20 Rubis. Your task is to decide how many of your 20 Rubis you want to invest in each of the two accounts mentioned above. To make your investment decision, you will type the amount of Rubis, a number between 0 and 20, you want to invest in the Public Account in the input field on the following input-screen:

Figure 4.4: 1^{st} stage - Investment decision

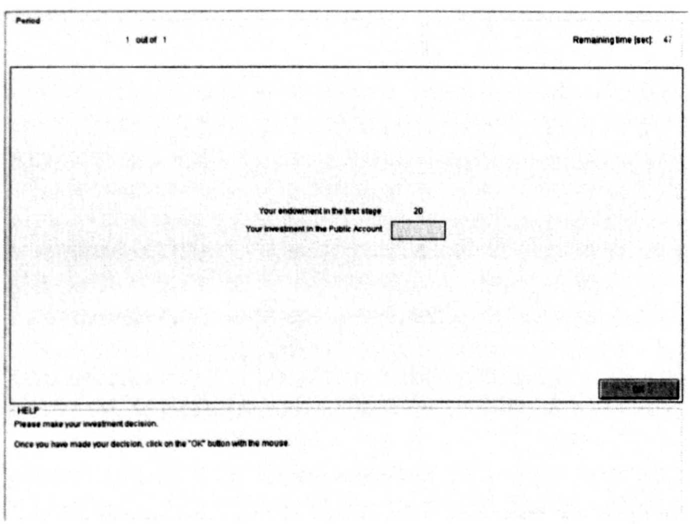

Those Rubis that you do not invest in the Public Account are automatically invested in the Private Account. Once you have made your decision, you must click on the red "OK" button to submit it, after which your decision cannot be changed. What you earn from your investment in the Public Account will depend on the total number of Rubis that you and the other three members of your group invest in the Public Account.

Your earnings in the first stage of a period consist of two parts, A and B:

A = Your return from your Private Account.

Your Private Account returns 1 Rubi for each Rubi invested. That is, for each Rubi invested in the Private Account you get 1 Rubi.

B = Your return from the Public Account.

To calculate this, we sum up all investments made in the Public Account in your group, multiply the sum by 1.6 and divide the result equally between the four members of your group.

Your investment in the Public Account also raises the earnings of the other group members. On the other hand, you earn Rubis for each Rubi invested by the other members in the Public Account. For each Rubi invested by any member you earn $\frac{1.6\times1}{4}$ = 0.4 Rubis. The income of each group member from the Public Account is calculated the same way. Everyone gets 0.4 Rubis from each Rubi invested in the Public Account by any group member. Therefore, you and the other group members receive the same amount from the total investment in the Public Account. The process is best explained by a number of examples.

Example: Suppose that you invest 0 Rubis in the Public Account but that the three other members of your group invest a total of 50 Rubis. Then your return from the Public Account would be $\frac{1.6 \times 50}{4} = 20$. Everyone else in your group would also earn 20 Rubis from the Public Account. Your total earnings in the first stage would be 40 Rubis (20 Rubis you kept in your Private Account + 20 Rubis from the Public Account).

Example: Suppose that you invest all your 20 Rubis in the Public Account but that the other three members invest nothing. Then your return from the Public Account would be $\frac{1.6 \times 20}{4} = 8$ Rubis. Everyone else in the group would also earn 8 Rubis from the Public Account. Your total earnings in the first stage would be 8 Rubis (0 Rubis you kept in your Private Account + 8 Rubis from the Public Account).

Example: Suppose that you invest 15 Rubis in the Public Account and that the three other members of your group invest a total of 50 Rubis. Then your return from the Public Account would be $\frac{1.6 \times 65}{4} = 26$ Rubis. Everyone else in your group would also earn 26 Rubis from the Public Account. Your total earnings in the first stage would be 31 Rubis (5 Rubis you kept in your Private Account + 26 Rubis from the Public Account).

Once you have confirmed your contribution, your decision can no longer be altered. The first stage is over only when all groups have made their decisions. After that, the second stage commences.

The Second Stage

In the second stage you will see how much each of the other group members invested in the Public Account and they will see your decision. At this stage you will have the opportunity to reduce the income of each group member by distributing "**deduction points**". The other group members can also reduce your income if they wish. How your earnings in the period are affected by decisions you and the other group members make in the second stage are described below.

If you assign deduction points to another group member, the earnings of this group member will be reduced by three times the amount of assigned deduction points. This means that if you assign one deduction point to another group member, her first-stage earnings will be reduced by 3 Rubis. If you assign 2 deduction points to a group member, her earnings will be reduced by 6 Rubis. If you assign 9 deduction points her earnings will be reduced by 27 Rubis, and so on. If you decide to assign 0 deduction points to a particular group member her earnings will not be changed by you. **You can assign a maximum amount of 10 deduction points to each other member.**

If you assign deduction points, you will also have costs. For each assigned deduction point, you will have costs of one Rubi. For example, if you assign 5 deduction points, you will have costs of 5 Rubis; if you assign 10 deduction points, you will have costs of 10 Rubis, and so on. If you assign 0 deduction points, you will have no costs from assigning deduction points. After all participants have made their decisions in the second stage, your final earnings for the period will be calculated as follows:

Total earnings at the end of the Period =

(Earnings from the 1st stage - 3 × (deduction points received from other group members) - (cost of deduction points you assigned to other group members)

Please note there is an exception to this: if the tripled amount of deduction points you have received exceeds the earnings from the first stage, the earnings after the second stage will be zero minus the costs of the deduction points you assigned to other group members. That is, no matter how many deduction points you have received, you cannot loose more than your first-stage earnings as a result of the deduction points assigned by other people. But independent of deduction points received, you always have to bear the full costs of deduction points you assign to other members. This means that your earnings at the end of the second stage can be negative. However, you can always avoid such losses with certainty by the decisions you make.

How do you make your decisions at the 2nd stage?

In the 2nd stage your task is to decide how many deduction points to assign to each of the other three group members. You enter your decision into a input screen like the one in Figure 2.

In the first row you see the amount endowed by each member of your group. In the second row you see how much of that endowment was invested in the Public Account. Your investment is displayed in the first column, under the heading "You", while the amounts invested in this period by the other group members are shown in the remaining three columns. Note that each column is always headed by the group member ID number (1,2,3 or 4), which means that the column position of each member on this screen is kept constant throughout the game.

Figure 4.5: 2^{nd} stage - Assignment of "Deduction points"

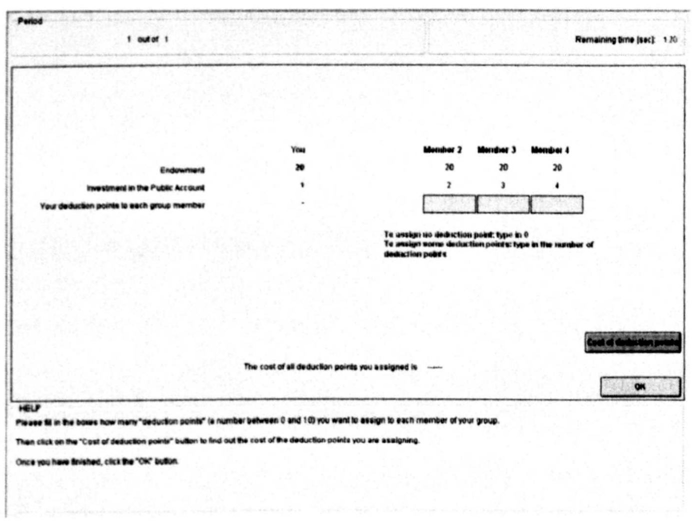

In the the last row, "Your deduction points", you have to make your decisions for the second stage. You must now decide how many deduction points you would like to assign to each of the other group members. You must type in the respective box a number between 0 and 10. You have to make an entry into each box. If you do not wish to change the income of a specific group member then you enter 0.

After all participants have made their decision, your income from the period will be displayed on an output screen like the following:

INSTRUCTIONS

CONDITION: UNCERTAIN PUNISHMENT ENFORCEMENT

WELCOME

Thank you for participating in this study. Your earnings in this study will be paid to you in cash at the end of the session. The instructions are simple. If you follow them carefully, you may, depending on your decisions, earn a considerable amount of money.

IMPORTANT NOTE

Please do not communicate in any way with other participants during this experiment. Please remember to switch off your mobile. If you have a question or problem at any point in today's session, please raise your hand and I will come to you.

THE EXPERIMENT

During this experiment Rubis will take the place of traditional monetary units. At the end of the experiment the total amount of Rubis you have earned will be converted into Pounds at the following rate:

$$1 \; Rubi \;\; = \;\; 2.5 \; Pence$$

Each participant receives a lump sum payment of 20 Rubis at the beginning of the experiment. This one-time payment can be used to pay for eventual losses during the experiment. However, you can always avoid losses with certainty through your own decisions. At the end of the experiment, your entire earnings from the experiment plus the lump sum payment will be converted to Pounds and immediately paid to you in cash.

This experiment is divided into **10 periods**. In each period, participants will be divided into groups of four. You will therefore be in a group with three other participants. During these 10 periods the group composition is constant. You are, therefore, grouped with the same people throughout the experiment. Your earnings will depend on the decisions that you and the other members of your group make. You will never learn whom you have been grouped with.

Each period consists of **two stages**. In the first stage, your task is to decide how many Rubis you want to invest in each of two investment accounts. One account is a Private Account, which only you benefit from. The second account is a Public Account, the benefits of which are shared equally by all members of your group. In the second stage you will be shown the amount invested in the Public Account by the three other members of your group. Your task is to decide whether you want to reduce their earnings from the first stage by distributing "deduction points" to them. **However, everybody's "deduction points" are only carried out with a probability of 80%(20%). Therefore, there is a chance**

that everybody's first-stage earnings will not be reduced even if "deduction points" were assigned to them.

In what follows, these stages are described in more detail.

The First Stage

At the beginning of each period, each participant receives 20 Rubis. Your task is to decide how many of your 20 Rubis you want to invest in each of the two accounts mentioned above. To make your investment decision, you will type the amount of Rubis, a number between 0 and 20, you want to invest in the Public Account in the input field on the following input-screen:

Figure 4.7: 1^{st} stage - Investment decision

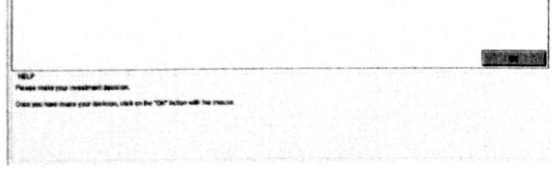

Those Rubis that you do not invest in the Public Account are automatically invested in the Private Account. Once you have made your decision, you must click on the red "OK" button to submit it, after which your decision cannot be changed. What you earn from your investment in the Public

Account will depend on the total number of Rubis that you and the other three members of your group invest in the Public Account.

Your earnings in the first stage of a period consist of two parts, A and B:

A = Your return from your Private Account.

Your Private Account returns 1 Rubi for each Rubi invested. That is, for each Rubi invested in the Private Account you get 1 Rubi.

B = Your return from the Public Account.

To calculate this, we sum up all investments made in the Public Account in your group, multiply the sum by 1.6 and divide the result equally between the four members of your group.

Your investment in the Public Account also raises the earnings of the other group members. On the other hand, you earn Rubis for each Rubi invested by the other members in the Public Account. For each Rubi invested by any member you earn $\frac{1.6 \times 1}{4} = 0.4$ Rubis. The income of each group member from the Public Account is calculated the same way. Everyone gets 0.4 Rubis from each Rubi invested in the Public Account by any group member. Therefore, you and the other group members receive the same amount from the total investment in the Public Account. The process is best explained by a number of examples.

Example: Suppose that you invest 0 Rubis in the Public Account but that the three other members of your group invest a total of 50 Rubis. Then your return from the Public Account would be $\frac{1.6 \times 50}{4} = 20$. Everyone else in your group would also earn 20 Rubis from the Public Account. Your total earnings in the first stage would be 40 Rubis (20 Rubis you kept in your Private Account + 20 Rubis from the Public Account).

Example: Suppose that you invest all your 20 Rubis in the Public Account but that the other three members invest nothing. Then your return from the Public Account would be $\frac{1.6 \times 20}{4} = 8$ Rubis. Everyone else in the group would also earn 8 Rubis from the Public Account. Your total earnings in the first stage would be 8 Rubis (0 Rubis you kept in your Private Account + 8 Rubis from the Public Account).

Example: Suppose that you invest 15 Rubis in the Public Account and that the three other members of your group invest a total of 50 Rubis. Then your return from the Public Account would be $\frac{1.6 \times 65}{4} = 26$ Rubis. Everyone else in your group would also earn 26 Rubis from the Public Account. Your total earnings in the first stage would be 31 Rubis (5 Rubis you kept in your Private Account + 26 Rubis from the Public Account).

Once you have confirmed your contribution, your decision can no longer be altered. The first stage is over only when all groups have made their decisions. After that, the second stage commences.

The Second Stage

In the second stage you will see how much each of the other group members invested in the Public Account and they will see your decision. At this stage you will have the opportunity to reduce or leave equal the first-stage income of each group member by distributing **"deduction points"**. The other group members can also reduce your income if they wish. However, it is not certain whether your decision and the decision from other group members of reducing someone's first-stage income will be actually carried out. This will be determined by chance. How your earnings in the period are affected by decisions you and the other group members make in the second stage are described below.

If you assign deduction points to another group member, the earnings of this group member may be reduced by three times the amount of assigned deduction points. This means that if you assign one deduction point to another group member, her first-stage earnings may be reduced by 3 Rubis. If you assign 2 deduction points to a group member, her earnings will be reduced by 6 Rubis. If you assign 9 deduction points her earnings may be reduced by 27 Rubis, and so on. If you decide to assign 0 deduction points to a particular group member her earnings will not be changed by you. **You can assign a maximum amount of 10 deduction points to each other member.**

If you assign deduction points, you may also have costs. For each assigned deduction point, you may have costs of one Rubi. For example, if you assign 5 deduction points, you may have costs of 5 Rubis; if you assign 10 deduction points, you may have costs of 10 Rubis, and so on. If you assign 0 deduction points, you will have no costs from assigning deduction points. **In each period, the deduction points decisions will be carried out with a probability of 80%(20%). This probability is the same for all 10 periods of this experiment.**

After all participants have made their deduction points decisions in the second stage, we decide whether the deduction points assigned are actually implemented. We decide this at a group level as follows: for each group, we will draw a ball from a bingo cage. The bingo cage has balls numbered from 1 to 10.

If the ball for your group is numbered 9 or 10 (1 to 8), then the deduction points you and the other members of your group have assigned ARE NOT carried out. In this case, your earnings in the period will be equal to your earnings in the first stage.

If the ball for your group is numbered from 1 to 8 (9 or 10), then the deduction points you and other members of your group have assigned ARE carried out.

In this case, your earnings in the period will be calculated as follows:

Total earnings at the end of the Period =

(Earnings from the 1st stage - 3 × (deduction points received from other group members) - (cost of deduction points you assigned to other group members)

Please note there is an exception to this: if the tripled amount of deduction points you have received exceeds the earnings from the first stage, the earnings after the second stage will be zero minus the costs of the deduction points you assigned to other group members. That is, no matter how many deduction points you have received, you cannot loose more than your first-stage earnings as a result of the deduction points assigned by other people. But independent of deduction points received, you always have to bear the full costs of deduction points you assign to other members. This means that your earnings at the end of the second stage can be negative. However, you can always avoid such losses with certainty by the decisions you make.

How do you make your decisions at the 2nd stage?

In the 2nd stage your task is to decide how many deduction points to assign to each of the other three group members. You enter your decision into a input screen like the one in Figure 2.

In the first row you see the amount endowed by each member of your group. In the second row you see how much of that endowment was invested in the Public Account. Your investment is displayed in the first column, under the heading "You", while the amounts invested in this period by the other group members are shown in the remaining three columns. Note that each column is always headed by the group member ID number (1,2,3 or 4), which means that the column position of each member on this screen is kept constant throughout the game.

In the the last row, "Your deduction points", you have to make your decisions for the second stage. You must now decide how many deduction points you would like to assign to each of the other group members. You must type in the respective box a number between 0 and 10. You have to make an entry into each box. If you do not wish to change the income of a specific group member then you enter 0.

Figure 4.8: 2^{nd} stage - Assignment of "Deduction points"

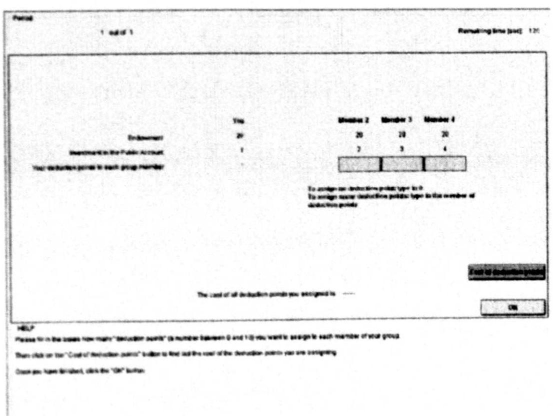

After all participants have made their decision, your income from the period will be displayed on an output screen like the following:

When you have finished reviewing your earnings for the current period you will click the grey "OK" button. When everyone is done, the experiment will proceed to the next period starting with stage one. Remember this experiment consists of 10 periods.

Figure 4.9: Period earnings display

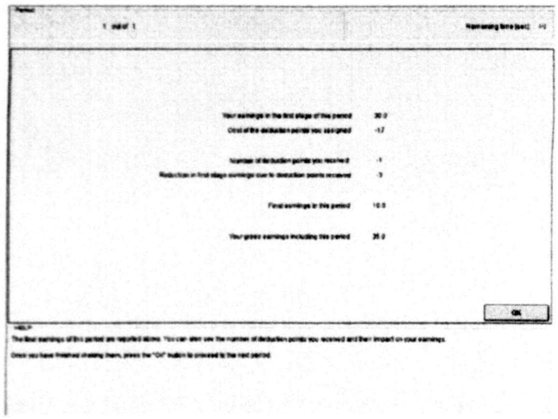

Thanks for participating.

If you have any questions, raise your hand.

Chapter 5

Chapter 5

CONCLUSIONS

In this chapter, we bring together the main results presented within previous chapters, discuss some limitations of our analysis and briefly discuss some extensions and follow-up questions for future research.

5.1 Main Findings

Chapter 2 presents evidence that endowing subjects with a relatively small sum of money prior to eliciting risk preferences does not affect their "baseline" attitudes to risk. This seems to be in stark contrast with findings from other studies reporting that subjects' willingness to accept a given actuarially fair gamble is increased when subjects had a prior gain[1] – the so-called "house money effect"[2] A natural question is why such an effect failed to be replicated in our experiment?

While a definite answer requires further experiments, there are reasons

[1]Thaler & Johnson (1990). For evidence of such effect in dynamic setting see e.g. Ackert *et al.* (2006).

[2]Harrison (2007) reports evidence of a house money effect in public goods games.

to believe that the source of the money – or the framing of it – matters[3]. The sum of money given to subjects in between the risk elicitation stages was administered in a way to induce them to think the money was a genuine earning rather than a windfall gain granted by the experimenter[4]. Indeed, there is some evidence that people may treat money they are promised differently from money they have earned. Cherry *et al.* (2002), for instance, show that when subjects in a dictator game made allocation decision over earned wealth, self-interested game theoretic behaviour was the norm. While it is an open question whether our treatment administration manages to legitimise the money given to subjects with effort, this result does raise the question of whether effects of changes in monetary income on risk preferences – or any other type of preference – are sensitive to the origin of the money. This may have serious implications for the reliability of results from laboratory experiments in which subjects are endowed with (any form of) money [5]. But not only this. It would also tell against the standard practice in economic theory of aggregating income from different sources.

Chapter 2 also showed that individuals may be willing to take more risks when they become "wealthier" – despite the increment to wealth be-

[3]Tangibility of the money may also be important. But this is an aspect "held constant" between our experiment and other experiments on house-money effects. In fact, it is common practice that earnings throughout a laboratory experiment be only handed over at the the end of it.

[4]The money was framed as a reward for completing the test. Subjects may have treated this as "earnings" even though the reward was not related to test performance.

[5]This would include, for example, studies on myopic loss aversion in investment decisions (e.g. Gneezy & Potters (1997), Haigh & List (2005) and Fellner & Sutter (2009). Yet if the origin of endowments is to play a role, this is more likely to happen, if at all, in experiments examining game situations.

ing small. But whether this is true, the evidence suggests, depends on the vehicle used to introduce the wealth increment. We measure subjects' attitudes towards a given risky prospect in two occasions: before and after a given small-scale increment to wealth, say Δw, was given; no statistically significant differences in these risk measures, hence wealth effects, are observed. However, subjects do change their risk attitudes when the wealth increment is merged with all possible consequences of the prospect: they exhibited more risk-lovingness in this situation. Because of the different forms used to introduce an exogenous change in wealth, we term this asymmetrical effect on risk attitudes the inside-outside framing effect.

Three observations about this effect are in order. First, that this asymmetrical framing effect suggests that individuals judge risk by its direct consequences not merging them with pre-existing outcomes – a violation of the "asset integration" hypothesis demonstrated in some experiments by Kahneman & Tversky (1979). Second, that while framing effects are no novelty in the literature[6], it is the first time to our knowledge that such behavioral asymmetry between the changes in risk attitudes induced by the "inside" money and changes induced by the "outside" money has been documented. It is worth noting that this result is not only in contrast to studies reporting "house-money" effects of money granted by the experiment on risk attitudes, but also with studies reporting an increase in risk aversion when lottery stakes are scaled-up (e.g. Holt & Laury, 2002). Third,

[6]See e.g. Kahneman & Tversky (1979), Andreoni (1995b), Cookson (2000) and Gächter *et al.* (2009b).

that this result has implications for calibration critiques of decision theories (see Rabin, 2000; Cox & Sadiraj, 2006). As shown by Wakker (2005), these critiques rely on an empirical assumption, namely, that attitudes to risk hold constant for a wide range of levels of wealth. Our result challenges this assumption. Although the existence of a wealth effect reported here relies on a "forced" integration of wealth increment to prizes of the lottery, this effect is not consistent with a utility function that exhibits constant absolute risk aversion.

We now turn to the main findings of Chapter 3. The primary result of this chapter is that individuals with higher cognitive ability show more choice consistency in a series of risk-elicitation tasks than individuals with lower cognitive ability. There is no relationship, however, between cognitive ability and "framing consistency", that is, violations of descriptive invariance. Such violations occur as a result of asymmetrical effects of different forms used to induce a small-scale wealth increment on attitudes to risk. Also, we find no association between performance in the cognitive test and individual differences in risk preferences.

The account we offer for each of these results hints at an important question, namely, when is an association between cognitive ability (CA) and task performance likely to arise? Regarding the association found between CA and choice consistency we advance that, given the repetitive structure of the design, this is either (a) due to differences in subjects' ability to retrieve information about their choices as they proceed through the experiment (correlated with cognitive ability), or simply (b) a consequence of an underlying lack of serious deliberation and engagement with tasks of the experiment. Framing inconsistency, in turn, would be a manifestation of a "narrow framing", whereby the increment received tends not to be mentally merged with lottery consequences prior to reaching a decision. The reasons why framing inconsistencies and risk preferences in our experiment are not associated with cognitive ability seem, in our view, closely related: individuals do not see framing consistency in our context and particular types of risk attitudes as normatively appropriate responses. In line with Kahneman (2003) and Stanovich & West (2008), we argue that it is irrelevant whether subjects differ in their cognitive resources if there are no situational cues suggesting that certain patterns of responses are more normatively-correct than others. Thus, if in performing a given experimental task individuals do not detect the existence of some sort of "correct answer", a relationship between cognitive ability and task responses is unlikely to be observed even if the ability to override a naturally primed response and replace it with an analytic response is related to cognitive ability.

Finally, Chapter 4 has two major results. First, that the threat of punishment cannot raise and sustain high levels of contributions when punishment enforcement is perceived by the individuals as a low-probability event. The experimental results show that a relatively low degree of uncertainty over enforcement does not impair punishment to serve as an effective deterrent device whereas a high degree of uncertainty does. This result furthers our understanding about under which circumstances a monetary sanction system can help to promote cooperation in social-dilemma situations. Second, that punishment behaviour is driven by forward- and backward-looking considerations. You find, for instance, that relatively more punishment was directed towards individuals who were "lucky" enough to not have punishment assigned to them actually enforced. This suggests that "bygones are not bygones" and that punishment decisions were driven by the "anger" being accumulated towards players with a history of free-riding behaviour that went unpunished in the past. This result points out the presence of inertial elements in the play of repeated games.

5.2 Limitations

There are some limitations to our analysis. First, in our measures of attitudes to risk. While the method we use (multiple-price-list) is transparent and incentive-compatible, it may be sensitive to the format of the multiple price list table[7]. Second, in the single set of wealth increment administered. We observe how risk behaviour responds to a given variation of wealth. There are, of course, other possible ranges of variation. Third, the cognitive test data has a relatively small number of observations. While one should not expect a uniform distribution of scores across the entire performance scale, there is a relatively small number of observations in the tails of our distribution. Another caveat concerning, in particular, the

[7] See Andersen *et al.* (2006c).

lack of association between cognitive ability and risk preferences relates to the restriction of range in our sample. Individuals with above average cognitive ability are likely to be overrepresented in our sample, which restricts the generality of our results. Fourth, there might be sequence effects in cooperation data that are not fully controlled for. While a "natural" randomisation (as opposed to a pre-selected one) *at a group level* of the enforcement status during the public good game in each session of P80 and P20 treatments creates variation, it produces a sample of enforcement sequences with a potentially unbalanced representativeness. Note, however, that this problem is worsened if we were to use pre-designed enforcement sequences in the sessions with uncertain enforcement condition.

5.3 Extensions

I now outline some extensions or closely-related questions for future research.

We have made the case that results in Chapter 2 suggest that subjects may treat money earned differently from money granted by the experimenter. A natural follow-up line of investigation is then to examine if the replication of the "house-money" effect in a simpler setting is susceptible to whether the money gained by subjects before facing up a given risk is legitimised with effort. Part of this line of enquiry could also investigate whether other departures from game-theoretic predictions, e.g. in public goods games, are diminished when subjects' endowment is earned. Thus, the idea is to legitimise with effort the endowment used in the allocation decisions. If it is a repeated setting, there are different ways to examine this issue. A possible mechanism to change the origins of the endowments and make them subjects' "own" earnings in this case is to make the earnings at the end of the first period the "endowment" of the subsequent ones. In fact, many social dilemma situations that we face repeatedly are likely to be intertemporally linked, for the outcomes of our decisions at one moment are likely to affect the subsequent decisions we make. Yet, some of the intertemporal structure of these decisions tend to be abstracted away in most public good experiments since endowments are "topped-up" at the outset of every period. A public good experiment in which subjects use their own earnings as an endowment surely involves a more complex design with linkages between periods. But it seems hard to avoid these complexities if one is to examine the role of intertemporal incentives in subjects' decisions.

In Chapter 3 we find that cognitive ability relates to some aspects of individual decision-making – for instance, individuals with relatively low cognitive ability have relatively noisier choices relative to individuals with high cognitive ability. This is somewhat surprising given that samples of university students are rather homogeneous in respect to this demographic characteristic. But our result reinforces the general message being sent off here and other similar studies: that is worth investigating whether differences in cognitive ability are associated with other types of economic behaviour. Another potentially interesting extension to Chapter 3 is to examine how cognitive ability is associated with ambiguity aversion. This relates to studies in which probabilities associated with outcomes are either not explicitly given or are only partially known – a situation that reflects many economic problems of interest.

Finally, in Chapter 4 we examine how patterns of cooperation and punishment behaviour respond to an introduction of a source of risk over punishment enforcement. A natural extension to it is to examine how *imperfect observability* affects cooperative behaviour. In this case, the uncertainty is placed at a different point of the timeline of the game: before the punishment stage takes place rather than after it. While with *imperfect enforcement* the (measurable) uncertainty is resolved after punishment decisions have been made, with *imperfect observability* the uncertainty is resolved at the end of the contribution stage, which determines the very existence of a punishment stage. This examines to some extent the sensitivity of cooperative behaviour to the framing of uncertainty, for both types of "imperfections" are logically equivalent – as whether one is to be punished for choosing a given contribution level is not affected by when this uncertainty is resolved in the remaining of the game period.

References

Abreu, Dilip. 1988. On the Theory of Infinitely Repeated Games with Discounting. *Econometrica*, **56**(2), 383–96.

Ackert, Lucy, Charupat, Narat, Church, Bryan, & Deaves, Richard. 2006. An experimental examination of the house money effect in a multi-period setting. *Experimental Economics*, **9**(1), 5–16.

Agnew, Julie R., Anderson, Lisa R., Gerlach, Jeffrey R., & Szykman, Lisa R. 2008. Who Chooses Annuities? An Experimental Investigation of the Role of Gender, Framing, and Defaults. *American Economic Review*, **98**(2), 418–22.

Alexander, J.R.M., & Smales, S. 1997. Intelligence, learning and long-term memory. *Personality and Individual Differences*, **23**(5), 815–825.

Allais, Maurice. 1953. Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica*, **21**(4), 503–546.

Andersen, Steffen, Harrison, Glenn, & Rutstrom, Elisabet. 2006a. *Dynamic Choice Behavior: Asset Integration and Natural Refer-*

ence Points. Working Paper 06-07. Economics Department, University of Central Florida.

ANDERSEN, STEFFEN, HARRISON, GLENN, LAU, MORTEN, & RUTSTRÖM, E. 2006b. Elicitation using multiple price list formats. *Experimental Economics*, 9(4), 383–405.

ANDERSEN, STEFFEN, HARRISON, GLENN, LAU, MORTEN, & RUTSTRÖM, E. 2006c. 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.

ANDERSON, CHRISTOPHER M., & PUTTERMAN, LOUIS. 2006. Do non-strategic sanctions obey the law of demand? The demand for punishment in the voluntary contribution mechanism. *Games and Economic Behavior*, 54(1), 1–24.

ANDREONI, JAMES. 1995a. Cooperation in Public-Goods Experiments: Kindness or Confusion? *American Economic Review*, 85(4), 891–904.

ANDREONI, JAMES. 1995b. Warm-Glow versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments. *The Quarterly Journal of Economics*, 110(1), 1–21.

ANDREONI, JAMES, & VESTERLUND, LISE. 2001. Which Is The Fair Sex? Gender Differences In Altruism. *The Quarterly Journal of Economics*, **116**(1), 293–312.

AXELROD, R. A. 1984. *The evolution of cooperation.* New York: Basic Books.

BARBERIS, NICHOLAS, HUANG, MING, & THALER, RICHARD H. 2006. Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing. *American Economic Review*, **96**(4), 1069–1090.

BARSKY, ROBERT B, ET AL. 1997. Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study. *The Quarterly Journal of Economics*, **112**(2), 537–79.

BELL, DAVID E. 1982. Regret in Decision Making under Uncertainty. *Operations Research*, **30**(5), 961–981.

BENARTZI, SHLOMO, & THALER, RICHARD. 1999. Risk Aversion or Myopia? Choices in Repeated Gambles and Retirement Investments. *Management Science*, **45**(3), 364–381.

BENJAMIN, DANIEL J., BROWN, SEBASTIAN A., & SHAPIRO, JESSE M. 2006 (May). *Who is "Behavioral"? Cognitive Ability and Anomalous Preferences.* Levine's Working Paper Archive. UCLA Department of Economics.

BERTRAND, MARIANNE, & MULLAINATHAN, SENDHIL. 2001. Do People Mean What They Say? Implications for Subjective Survey Data. *American Economic Review*, 91(2), 67–72.

BINSWANGER, HANS P. 1980. Attitudes toward Risk: Experimental Measurement in Rural India. *American Journal of Agricultural Economics*, 62(3), 395–407.

BIRNBAUM, MICHAEL, & SCHMIDT, ULRICH. 2008. An experimental investigation of violations of transitivity in choice under uncertainty. *Journal of Risk and Uncertainty*, 37(1), 77–91.

BOCHET, OLIVIER, PAGE, TALBOT, & PUTTERMAN, LOUIS. 2006. Communication and punishment in voluntary contribution experiments. *Journal of Economic Behavior & Organization*, 60(1), 11–26.

BOLTON, GARY E., & KATOK, ELENA. 1995. An experimental test for gender differences in beneficent behavior. *Economics Letters*, 48(3-4), 287–292.

BOMBARDINI, MATILDE, & TREBBI, FRANCESCO. 2005. *Risk Aversion and Expected Utility Theory: A Field Experiment with Large and Small Stakes*. unpublished manuscript. Harvard University.

BORGHANS, LEX, GOLSTEYN, BART H. H., HECKMAN, JAMES J., & MEIJERS, HUUB. 2009. Gender Differences in Risk Aversion and Ambiguity Aversion. *Journal of the European Economic Association*, **7**(2-3), 649–658.

BOSCH-DOMENECH, A., & SILVESTRE, J. 1999. Does risk aversion or attraction depend on income? An experiment. *Economics Letters*, **65**(3), 265–273.

BOSCH-DOMÈNECH, ANTONI, & SILVESTRE, JOAQUIM. 2003 (June). *Do the Wealthy Risk More Money? An Experimental Comparison*. Economics Working Papers 692. Department of Economics and Business, Universitat Pompeu Fabra.

BOTHNER, MATTHEW, JEONG-HAN KANG, & STUART., TOBY E. 2007. Competitive Crowding and Risk Taking in a Tournament: Evidence from NASCAR Racing. *Administrative Science Quarterly*, **52**(2), 208–247.

BOYD, R., & LORBERBAUM, J. P. 1987. No Pure Strategy is Stable in the Repeated Prisoner's Dilemma Game. *Nature*, 58–59.

BRAÑAS-GARZA, PABLO, GUILLEN, PABLO, & DEL PASO, RAFAEL LÓPEZ. 2008. Math skills and risk attitudes. *Economics Letters*, **99**(2), 332–336.

BURKS, STEPHEN V., *et al.* . 2009. Cognitive skills affect economic preferences, strategic behavior, and job attachment. *PNAS - Proceedings of the National Academy of Sciences of the USA*, **106**(19), 7745–7750.

BUTLER, DAVID J., & LOOMES, GRAHAM C. 2007. Imprecision as an Account of the Preference Reversal Phenomenon. *American Economic Review*, **97**(1), 277–297.

CAMERER, COLIN. 1995. *Handbook of Experimental Economics*. Princeton University Press. Chap. Individual Decision Making.

CAMERER, COLIN F. 1989. An Experimental Test of Several Generalized Utility Theories. *Journal of Risk and Uncertainty*, **2**(1), 61–104.

CARLIN, PAUL S. 1992. Violations of the reduction and independence axioms in Allais-type and common-ratio effect experiments. *Journal of Economic Behavior & Organization*, **19**(2), 213–235.

CASARI, MARCO, HAM, JOHN C., & KAGEL, JOHN H. 2007. Selection Bias, Demographic Effects, and Ability Effects in Common Value Auction Experiments. *American Economic Review*, **97**(4), 1278–1304.

CESARINI, D., LICHTENSTEIN, P., AN O. SANDEWALL, M. JOHANNESSON, & WALLACE, B. 2009. Genetic Variation in Financial Decision Making. *Journal of Finance*, **Forthcoming**.

CHARNESS, GARY, & LEVIN, DAN. 2005 (Mar.). *The Origin of the Winner's Curse: A Laboratory Study*. Levine's Bibliography 666156000000000602. UCLA Department of Economics.

CHEN, YUH-JIA, & CORTER, JAMES E. 2006. When Mixed Options are Preferred in Multiple-Trial Decisions? *Journal of Behavioral & Decision Making*, **19**, 17–42.

CHERRY, TODD L., FRYKBLOM, PETER, & SHOGREN, JASON F. 2002. Hardnose the Dictator. *American Economic Review*, **92**(4), 1218–1221.

CHETAN, DAVE, ECKEL, CATHERINE, JOHNSON, CATHLEEN, & ROJAS, CHRISTIAN. 2007. *Eliciting Risk Preferences: When is Simple Better?* unpublished. University of Massachusetts, Amherst.

CHEW, SOO HONG. 1983. A Generalization of the Quasilinear Mean with Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox. *Econometrica*, **51**(4), 1065–92.

CINYABUGUMA, MATTHIAS, PAGE, TALBOT, & PUTTERMAN, LOUIS. 2005. Cooperation under the threat of expulsion in a public goods experiment. *Journal of Public Economics*, **89**(8), 1421–1435.

CONLISK, JOHN. 1989. Three Variants on the Allais Example. *American Economic Review*, **79**(3), 392–407.

COOKSON, R. 2000. Framing Effects in Public Goods Experiments. *Experimental Economics*, **3**(1), 55–79.

COX, JAMES C., & SADIRAJ, VJOLLCA. 2006. Small- and large-stakes risk aversion: Implications of concavity calibration for decision theory. *Games and Economic Behavior*, **56**(1), 45–60.

Cubitt, Robin, Starmer, Chris, & Sugden, Robert. 1998a. On the Validity of the Random Lottery Incentive System. *Experimental Economics*, **1**(2), 115–131.

Cubitt, Robin P, Starmer, Chris, & Sugden, Robert. 1998b. Dynamic Choice and the Common Ratio Effect: An Experimental Investigation. *Economic Journal*, **108**(450), 1362–80.

Dawes, Robyn M. 1980. Social Dilemmas. *Annual Review of Psychology*, **31**, 169–193.

Dohmen, Thomas, Falk, Armin, Huffman, David, Sunde, Uwe, Schupp, Jürgen, & Wagner, Gert G. 2005 (Sept.). *Individual Risk Attitudes: New Evidence from a Large, Representative, Experimentally-Validated Survey*. IZA Discussion Papers 1730. Institute for the Study of Labor (IZA).

Dohmen, Thomas, Falk, Armin, Huffman, David, & Sunde, Uwe. 2007 (Apr.). *Are Risk Aversion and Impatience Related to Cognitive Ability?* IZA Discussion Papers 2735. Institute for the Study of Labor (IZA).

Donkers, Bas, Melenberg, Bertrand, & Van Soest, Arthur. 2001. Estimating Risk Attitudes Using Lotteries: A Large Sample Approach. *Journal of Risk and Uncertainty*, **22**(2), 165–95.

DRUCKMAN, J. 2001. Evaluating framing effects. *Journal of Economic Psychology*, 22, 96–101.

ECKEL, CATHERINE C, & GROSSMAN, PHILIP J. 1998. Are Women Less Selfish Than Men? Evidence from Dictator Experiments. *Economic Journal*, 108(448), 726–35.

ECKEL, CATHERINE C., & WILSON, RICK K. 2004. Is trust a risky decision? *Journal of Economic Behavior & Organization*, 55(4), 447–465.

ECKEL, C.C., & GROSSMAN, P. 2007. *Handbook of Experimental Results.* New York: Elsevier. Chap. Men, women and risk aversion: experimental evidence.

EGAS, MARTIJN, & RIEDL, ARNO. 2008. The economics of altruistic punishment and the maintenance of cooperation. *Proceedings of The Royal Society B-Biological Sciences*, 275(1637), 871–878.

EISENHAUER, JOSEPH G. 1997. Risk aversion, wealth, and the DARA hypothesis: A new test. *International Advances in Economic Research*, 3(1), 46–53.

ENGLE, R. W., TUHOLSKI, S. W., LAUGHLIN, J. E., & CONWAY, A. R. 1999. Working memory, short-term memory, and general fluid intelligence: a latent-variable approach. *Journal of experimental psychology*, 128(3), 309–31.

ERTAN, ARHAN, PAGE, TALBOT, & PUTTERMAN, LOUIS. 2009. Who to punish? Individual decisions and majority rule in mitigating the free rider problem. *European Economic Review*, **53**(5), 495–511.

FEHR, ERNST, & GÄCHTER, SIMON. 2000. Cooperation and Punishment in Public Goods Experiments. *American Economic Review*, **90**(4), 980–994.

FEHR, ERNST, & GÄCHTER, SIMON. 2002. Altrustic Punishment in Humans. *Nature*, **415**(10), 137–140.

FEHR, ERNST, & SCHMIDT, KLAUS M. 1999. A Theory Of Fairness, Competition, And Cooperation. *The Quarterly Journal of Economics*, **114**(3), 817–868.

FELLNER, GERLINDE, & SUTTER, MATTHIAS. 2009. Causes, Consequences, and Cures of Myopic Loss Aversion - An Experimental Investigation. *Economic Journal*, **119**(537), 900–916.

FISCHBACHER, URS. 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, **10**(2), 171–178.

FISCHBACHER, URS, & GÄCHTER, SIMON. 2009. Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Good Experiments. *American Economic Review*, Forthcoming.

Fischbacher, Urs, Gächter, Simon, & Fehr, Ernst. 2001. Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*, **71**(3), 397–404.

Frederick, Shane. 2005. Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, **19**(4), 25–42.

Frey, Bruno S., & Meier, Stephan. 2004. Social Comparisons and Pro-social Behavior: Testing "Conditional Cooperation" in a Field Experiment. *American Economic Review*, **94**(5), 1717–1722.

Friedman, James W. 1971. A Non-cooperative Equilibrium for Supergames. *Review of Economic Studies*, **38**(113), 1–12.

Gächter, Simon, & Herrmann, Benedikt. 2009. Reciprocity, culture, and human cooperation: Previous insights and a new cross-cultural experiment. *Philosophical Transactions of the Royal Society B - Biological Sciences*, **364**(March), 791–806.

Gächter, Simon, Renner, Elke, & Sefton, Martin. 2008. The Long-Run Benefits of Punishment. *Science*, **322**(5907), 1510.

Gächter, Simon, Orzen, Henrik, Renner, Elke, & Starmer, Chris. 2009a. Are Experimental Economists Prone to Framing Effects? A Natural Field Experiment. *Journal of Economic Behavior and Organization*, **70**(3), 443–446.

GÄCHTER, SIMON, ORZEN, HENRIK, RENNER, ELKE, & STARMER, CHRIS. 2009b. Are Experimental Economists Prone to Framing Effects? A Natural Field Experiment. *Journal of Economic Behaviour and Organization*, Forthcoming.

GERTNER, ROBERT. 1993a. Game Shows and Economic Behavior: Risk-Taking on "Card Sharks.". *The Quarterly Journal of Economics*, **108**(2), 507–21.

GERTNER, ROBERT. 1993b. Game Shows and Economic Behavior: Risk-Taking on "Card Sharks.". *The Quarterly Journal of Economics*, **108**(2), 507–21.

GILBOA, I. 2009. Questions in Decision Theory. *Annual Reviews in Economics*, Forthcoming.

GLANCE, N.S., & HUBERMAN, B. A. 1993. The outbreak of cooperation. *Journal of Mathematical sociology*, **17**(4), 281–302.

GNEEZY, URI, & POTTERS, JAN. 1997. An Experiment on Risk Taking and Evaluation Periods. *The Quarterly Journal of Economics*, **112**(2), 631–45.

GNEEZY, URI, & RUSTICHINI, ALDO. 2000. Pay Enough Or Don'T Pay At All. *The Quarterly Journal of Economics*, **115**(3), 791–810.

GNEEZY, URI, KAPTEYN, ARIE, & POTTERS, JAN. 2003. Evaluation Periods and Asset Prices in a Market Experiment. *Journal of Finance*, 58(2), 821–838.

GOEREE, JACOB K., HOLT, CHARLES A., & PALFREY, THOMAS R. 2003. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 45(1), 97–113.

GOLLIER, CHRISTIAN. 2001. Wealth Inequality and Asset Pricing. *Review of Economic Studies*, 68(1), 181–203.

GRANDMONT, J. M. 1992. Transformations of the Commodity Space, Behavioral Heterogeneity, and the Aggregation Problem. *Journal of Economic Theory*, 57(1), 1–35.

GREINER, BEN. 2004. *Forschung und wissenschaftliches Rechnen 2003.* Gottingen: Gesellschaft fur Wissenschaftliche Datenverarbeitung. Chap. An Online Recruitment System for Economic Experiments, pages 79–93.

GUISO, LUIGI, JAPPELLI, TULLIO, & TERLIZZESE, DANIELE. 1996. Income Risk, Borrowing Constraints, and Portfolio Choice. *American Economic Review*, 86(1), 158–72.

GUNNTHORSDOTTIR, ANNA, HOUSER, DANIEL, & MCCABE, KEVIN. 2007. Disposition, history and contributions in public goods experiments. *Journal of Economic Behavior & Organization*, 62(2), 304–315.

GUTH, WERNER, LEVATI, M. VITTORIA, SUTTER, MATTHIAS, & VAN DER HEIJDEN, ELINE. 2007. Leading by example with and without exclusion power in voluntary contribution experiments. *Journal of Public Economics*, **91**(5-6), 1023–1042.

HADAR, JOSEF, & SEO, TAE KUN. 1990. The Effects of Shifts in a Return Distribution on Optimal Portfolios. *International Economic Review*, **31**(3), 721–36.

HAIGH, MICHAEL S., & LIST, JOHN A. 2005. Do Professional Traders Exhibit Myopic Loss Aversion? An Experimental Analysis. *Journal of Finance*, **60**(1), 523–534.

HALEVY, YORAM. 2007. Ellsberg Revisited: An Experimental Study. *Econometrica*, **75**(2), 503–536.

HAMMARLID, OLA. 2005. When to accept a sequence of gambles. *Journal of Mathematical Economics*, **41**(8), 974–982.

HARBAUGH, WILLIAM T., KRAUSE, KATE, & BERRY, TIMOTHY R. 2001. GARP for Kids: On the Development of Rational Choice Behavior. *American Economic Review*, **91**(5), 1539–1545.

HARDIN, G. 1968. Tragedy of commons. *Science*, **162**(3859), 1243–1248.

HARLESS, DAVID W, & CAMERER, COLIN F. 1994. The Predictive Utility of Generalized Expected Utility Theories. *Econometrica*, **62**(6), 1251–89.

HARRISON, GLENN. 2007. House money effects in public good experiments: Comment. *Experimental Economics*, **10**(4), 429–437.

HARRISON, GLENN W., JOHNSON, ERIC, MCINNES, MELAYNE M., & RUTSTROM, E. ELISABET. 2005a. Risk Aversion and Incentive Effects: Comment. *American Economic Review*, **95**(3), 897–901.

HARRISON, GLENN W., JOHNSON, ERIC, MCINNES, MELAYNE M., & RUTSTRÖM, E. ELISABET. 2005b. Temporal stability of estimates of risk aversion. *Applied Financial Economics Letters*, **1**(1), 31–35.

HARRISON, GLENN W, LIST, JOHN A, & TOWE, CHARLES. 2007. Naturally Occurring Preferences and Exogenous Laboratory Experiments: A Case Study of Risk Aversion. *Econometrica*, **75**(2), 433–458.

HARRISON, G.W. 1986. An Experimental Test for Risk Aversion. *Economics Letters*, **21**(1), 7–11.

HAUBRICH, JOSEPH G. 1998. Bank diversification: laws and fallacies of large numbers. *Economic Review*, 2–9.

HENRICH, JOSEPH. 2000. Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining among the Machiguenga of the Peruvian Amazon. *American Economic Review*, **90**(4), 973–979.

HENRICH, JOSEPH, *et al.* . 2004. *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence in Fifteen Small-Scale Societies*. OUP.

HERRMANN, BENEDIKT, THÖNI, CHRISTIAN, & GÄCHTER, SIMON. 2008. Antisocial Punishment Across Societies. *Science*, **319**, 1362–1367.

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**(5), 1644–1655.

HUMPHREY, STEVEN J. 2000. The common consequence effect: testing a unified explanation of recent mixed evidence. *Journal of Economic Behavior & Organization*, **41**(3), 239–262.

ISAAC, R MARK, & JAMES, DUNCAN. 2000. Just Who Are You Calling Risk Averse? *Journal of Risk and Uncertainty*, **20**(2), 177–87.

ISAAC, R MARK, & WALKER, JAMES M. 1988. Communication and Free-Riding Behavior: The Voluntary Contribution Mechanism. *Economic Inquiry*, **26**(4), 585–608.

KACHELMEIER, STEVEN J, & SHEHATA, MOHAMED. 1992. Examining Risk Preferences under High Monetary Incentives: Experimental Evidence from the People's Republic of China. *American Economic Review*, **82**(5), 1120–41.

KAHNEMAN, DANIEL. 2003. Maps of Bounded Rationality: Psychology for Behavioral Economics. *American Economic Review*, **93**(5), 1449–1475.

KAHNEMAN, DANIEL, & FREDERICK, S. 2002. *Heuristics and Biases: The Psychology of Intuitive Judgment.* New York: Cambridge University Press. Chap. Representativeness revisited: Attribute substitution in intuitive judgment, pages 49–81.

KAHNEMAN, DANIEL, & TVERSKY, AMOS. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, **47**(2), 263–91.

KLOS, A., WEBER, E. U., & WEBER, M. 2005. Investment Decisions and Time Horizon: Risk Perception and Risk Behavior in Repeated Gambles. *Management Science*, **51**(12), 1777–1790.

KOLLOCK, PETER. 1998. Social Dilemmas: The Anatomy of Cooperation. *Annual Review of Sociology*, **24**(1), 183–214.

KOSZEGI, BOTOND, & RABIN, MATTHEW. 2006. A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics*, **121**(4), 1133–1165.

KRAHNEN, JAN PIETER, RIECK, CHRISTIAN, & THEISSEN, ERIK. 1997. Inferring risk attitudes from certainty equivalents: Some lessons from an experimental study. *Journal of Economic Psychology*, **18**(5), 469–486.

LEDYARD, JOHN O. 1995. *Handbook of Experimental Economics.* Princeton University Press. Chap. Public Goods: A Survey of Experimental Research, pages 111–194.

LEE, JINKWON. 2008. The effect of the background risk in a simple chance improving decision model. *Journal of Risk and Uncertainty*, **36**(1), 19–41.

LEVY, HAIM. 1994. Absolute and Relative Risk Aversion: An Experimental Study. *Journal of Risk and Uncertainty*, **8**(3), 289–307.

LOOMES, G., & SUGDEN, R. 1995. Incorporating a stochastic element into decision theories. *European Economic Review*, **39**(3), 641–648.

LOOMES, GRAHAM. 1991. Evidence of a New Violation of the Independence Axiom. *Journal of Risk and Uncertainty*, **4**(1), 91–108.

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, & SUGDEN, ROBERT. 1982. Regret Theory: An Alternative Theory of Rational Choice under Uncertainty. *Economic Journal*, **92**(368), 805–24.

LOOMES, GRAHAM, & SUGDEN, ROBERT. 1998. Testing Different Stochastic Specifications of Risky Choice. *Economica*, **65**(260), 581–98.

LOOMES, GRAHAM, STARMER, CHRIS, & SUGDEN, ROBERT. 1991. Observing Violations of Transitivity by Experimental Methods. *Econometrica*, **59**(2), 425–39.

LOOMES, GRAHAM, MOFFATT, PETER G, & SUGDEN, ROBERT. 2002. A Microeconometric Test of Alternative Stochastic Theories of Risky Choice. *Journal of Risk and Uncertainty*, **24**(2), 103–30.

LOOMES, GRAHAM, STARMER, CHRIS, & SUGDEN, ROBERT. 2003. Do Anomalies Disappear in Repeated Markets? *Economic Journal*, **113**(486), C153–C166.

MACHINA, MARK J. 1987. Choice under Uncertainty: Problems Solved and Unsolved. *Journal of Economic Perspectives*, **1**(2), 121–154.

MARK ISAAC, R., MCCUE, KENNETH F., & PLOTT, CHARLES R. 1985. Public goods provision in an experimental environment. *Journal of Public Economics*, **26**(1), 51–74.

MARSCHAK, JACOB. 1950. Rational Behavior, Uncertain Prospects, and Measurable Utility. *Econometrica*, **18**(2), 111–141.

MASCLET, D., NOUSSAIR, C., TUCKER, S., & VILLEVAL, M.C. 2003. Monetary and Nonmonetary Punishment in the Voluntary Contributions Mechanism. *American Economic Review*, **93**(1), 366–380.

MAYES, SUSAN DICKERSON, CALHOUNA, SUSAN L., BIXLERA, EDWARD O., & ZIMMERMAN, DENNIS N. 2009. IQ and neuropsychological predictors of academic achievement. *Learning and Individual Differences*, **19**(2), 238–241.

MESSICK, D. M., & BREWER, M. B. 1983. Solving social dilemmas: A review. *Review of personality and social psychology*, **4**, 11–44.

MORRIS, STEPHEN. 1996. Speculative Investor Behavior and Learning. *The Quarterly Journal of Economics*, **111**(4), 1111–33.

NEISSER, U. ET AL. 1996. Intelligence: knowns and unknowns. *American psychologist*, **51**(2), 77–101.

NIELSEN, LARS TYGE. 1985. Attractive Compounds of Unattractive Investments and Gambles. *Scandinavian Journal of Economics*, **87**(3), 463–73.

NIKIFORAKIS, NIKOS. 2008. Punishment and counter-punishment in public good games: Can we really govern ourselves. *Journal of Public Economics*, **92**(1-2), 91–112.

NOUSSAIR, CHARLES, & TUCKER, STEVEN. 2005. Combining Monetary and Social Sanctions to Promote Cooperation. *Economic Inquiry*, **43**(3), 649–660.

OFFERMAN, THEO, & SCHOTTER, ANDREW. 2009. Imitation and luck: An experimental study on social sampling. *Games and Economic Behavior*, **65**(2), 461–502.

OGAKI, MASAO, & ZHANG, QIANG. 2001. Decreasing Relative Risk Aversion and Tests of Risk Sharing. *Econometrica*, **69**(2), 515–26.

Ostrom, Elinor, Walker, James, & Gardner, Roy. 1992. Covenants With and Without a Sword: Self-Governance is Possible. *The American Political Science Review*, **86**(2), 404–417.

Palacios-Huerta, Ignacio, & Serrano, Roberto. 2006. Rejecting small gambles under expected utility. *Economics Letters*, **91**(2), 250–259.

Peköz, Erol A. 2002. Samuelson's fallacy of large numbers and optional stopping. *Journal of Risk and Insurance*, **69**(1), 1–7.

Quiggin, John. 1982. A theory of anticipated utility. *Journal of Economic Behavior & Organization*, **3**(4), 323–343.

Rabin, Matthew. 2000. Risk Aversion and Expected-Utility Theory: A Calibration Theorem. *Econometrica*, **68**(5), 1281–1292.

Redelmeier, D. A., & Tversky, A. 1992. On the framing of multiple prospects. *Psyhological Science*, **3**, 191–193.

Rosenzweig, Mark R, & Binswanger, Hans P. 1993. Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments. *Economic Journal*, **103**(416), 56–78.

Ross, Stephen A. 1999. Adding Risks: Samuelson's Fallacy of Large Numbers Revisited. *Journal of Financial and Quantitative Analysis*, **34**(3), 323–339.

Roth, Alvin E., Prasnikar, Vesna, Okuno-Fujiwara, Masahiro, & Zamir, Shmuel. 1991. Bargaining and Market Behavior in

Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study. *American Economic Review*, **81**(5), 1068–95.

RUBINSTEIN, ARIEL. 2001. *Comments on the Risk and Time Preferences in Economics.* Workin Paper. Tel Aviv University.

SAFRA, ZVI, & SEGAL, UZI. 2008. Calibration Results for Non-Expected Utility Theories. *Econometrica*, **76**(5), 1143–1166.

SAMUELSON, P. A. 1967. Risk and Uncertainty: A fallacy of large numbers. *Scientia*, **98**, 108–113.

SCHMIDT, ULRICH, & ZANK, HORST. 2008. Risk Aversion in Cumulative Prospect Theory. *Management Science*, **54**, 208–216.

SCHUBERT, RENATE, BROWN, M., GYSLER, M., & BRACHINGER, H.W. 1999. Financial Decision-Making: Are Women Really More Risk-Averse? *American Economic Review*, **89**(2), 381–385.

SEFTON, MARTIN, SHUPP, ROBERT, & WALKER, JAMES M. 2007. The Effect Of Rewards And Sanctions In Provision Of Public Goods. *Economic Inquiry*, **45**(4), 671–690.

SLEMROD, JOEL, BLUMENTHAL, MARSHA, & CHRISTIAN, CHARLES. 2001. Taxpayer response to an increased probability of audit: evidence from a controlled experiment in Minnesota. *Journal of Public Economics*, **79**(3), 455–483.

STANOVICH, KEITH, & WEST, RICHARD. 2008. On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology*, **94**(4), 672–695.

STARMER, CHRIS. 1992. Testing New Theories of Choice under Uncertainty Using the Common Consequence Effect. *Review of Economic Studies*, **59**(4), 813–30.

STARMER, CHRIS. 1999. Cycling with Rules of Thumb: An Experimental Test for a new form of Non-Transitive Behaviour. *Theory and Decision*, **46**(2), 139–157.

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, & SUGDEN, ROBERT. 1989a. Probability and Juxtaposition Effects: An Experimental Investigation of the Common Ratio Effect. *Journal of Risk and Uncertainty*, **2**(2), 159–78.

STARMER, CHRIS, & SUGDEN, ROBERT. 1989b. Violations of the independence axion in common ratio problems: An experimental test of some competing hypotheses. *Annals of Operations Research*, **19**(1), 79–102.

STERNBERG, R. J., & DETTERMAN, D. K. 1986. *What is Intelligence? Comtemporary viewpoints on its nature and definition.* Norwood NJ: Ablex.

SUNDE, UWE, DOHMEN, THOMAS, FALK, ARMIN, & HUFFMAN, DAVID. 2009. Are Risk Aversion and Impatience Related to Cognitive Ability? *American Economic Review*, Forthcoming.

SUNDEN, ANNIKA E, & SURETTE, BRIAN J. 1998. Gender Differences in the Allocation of Assets in Retirement Savings Plans. *American Economic Review*, **88**(2), 207–11.

SZPIRO, GEORGE G. 1986. Measuring Risk Aversion: An Alternative Approach. *The Review of Economics and Statistics*, **68**(1), 156–59.

THALER, RICHARD H., & JOHNSON, ERIC J. 1990. Gambling with the house money and trying to break even: the effects of prior outcomes on risky choice. *Management Science*, **36**(6), 643–60.

TRIVERS, R.L. 1971. The evolution of reciprocal altruism. *Quarterly Review of Biology*, **46**, 35–57.

TVERSKY, A., & KAHNEMAN, D. 1981. The framing of decisions and the psychology of choice. *Science*, **211**(4481), 453–458.

TVERSKY, AMOS, & KAHNEMAN, DANIEL. 1986. Rational Choice and the Framing of Decisions. *Journal of Business*, **59**(4), S251–78.

TVERSKY, AMOS, & KAHNEMAN, DANIEL. 1992. Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*, **5**(4), 297–323.

WAKKER, PETER P. 2005. *Formalizing Reference Dependence and Initial Wealth in Rabin's Calibration Theorem.* Workin Paper. Econometric Institute, Erasmus University Rotterdam.

WALKER, JAMES M., & HALLORAN, MATTHEW A. 2004. Rewards and Sanctions and the Provision of Public Goods in One-Shot Settings. *Experimental Economics*, **7**(3), 235–247.

WATT, RICHARD. 2002. Defending Expected Utility Theory. *Journal of Economic Perspectives*, **16**(2), 227–229.

WEIL, PHILIPPE. 1993. Precautionary Savings and the Permanent Income Hypothesis. *Review of Economic Studies*, **60**(2), 367–83.

WILLIAMS, WENDY M. 1996. Consequences of How We Define and Assess Intelligence. *Psychology, Public Policy, and Law*, **2**(3/4), 506–535.

WOLF, CHARLES, & POHLMAN, LARRY. 1983. The Recovery of Risk Preferences from Actual Choices. *Econometrica*, **51**(3), 843–50.

WU, GEORGE, & GONZALEZ, RICHARD. 1998. Common Consequence Conditions in Decision Making under Risk. *Journal of Risk and Uncertainty*, **16**(1), 115–39.

XIONG, WEI, & YAN, HONGJUN. 2006 (Dec.). *Heterogeneous Expectations and Bond Markets.* NBER Working Papers 12781. National Bureau of Economic Research, Inc.

YAMAGUSHI, TOSHIO. 1986. The Provision of a Sanctioning System as a Public Good. *Journal orpersonality and Social Psychology*, **51**(1), 110–116.