REPORT ON THE PENSIONS EXPERIMENT

A Report on an Experiment on Pensions
Financed by the
Department for Work and Pensions
and Implemented at the
Centre for Experimental Economics (EXEC)
at the
University of York

John D. Hey
LUISS, Italy and University of York, UK

21st October, 2007

CONTENTS

List of Tables

List of Figures

VOLUME 1: THE REPORT

1. Executive Summary

1.1. This is a report on an original research study, using the methodology of experimental economics, carried out for the DWP. Its purpose was twofold: to assess the potential of the experimental approach in assisting the DWP with questions relating to its research and policy agenda; and to explore the applicability of the experimental approach in understanding the possible impact of proposed policy measures within the field of pensions, and, in particular, the possible impact of proposed pension reforms like that of the introduction of Personal Accounts.

1.2. Whilst the methodology of experimental economics is well-developed in other contexts, its applicability in the field of pensions is relatively under-explored. Therefore, one of the prime interests of this study is how the methods might be adapted to the field of pensions. A major difficulty was the translation of the 'pension problem' into a form suitable for examination using experiments. Typically, an economics experiment involves an abstraction of the real-life problem into a form replicable in an experimental setting. It also involves an appropriate payment mechanism that replicates in some sense the real-life incentive of the participants. In the field of pensions, the real-life incentive is that of maximising life-time happiness.

1.3. This first experiment examined the possible implications of various possible policy measures designed to encourage increased pension contributions. The typical methodology of experimental economics is to run several 'treatments' of the experiment, one without the proposed policy measure and one with. The analysis then proceeds by comparing the behaviour of the participants in the two treatments. Clearly one of these two treatments should therefore be the *status quo* – namely, the situation prior to the introduction of the proposed policy measure.

1.4. This is where difficulties with the experimental approach first became manifest - since the *status quo,* particularly of the target population - is one of relatively little involvement with pension decisions. To reproduce this situation in an experimental setting is not an easy task. It was clear that the 'pre-policy measure' treatment of the experiment would not be an accurate reproduction of the *status quo*. A further difficulty is that in experiments, information relevant to the decision task is usually supplied in a clear and (relatively) easily digestible form to the participants. There is a clear conflict here – in real-life, and particularly for the target population, information is not easily obtainable and usually not in an easily digestible form.

1.5. The actual experiment was carried out in the laboratory and lasted a little under three hours. Participants were presented with a version of the 'pension problem' and were asked to take decisions concerning their contributions to pensions.

1.6. An important and relevant innovation of this project was the use of a quota sample of 174 participants from the target population (defined as those members of the UK

population believed to be under-saving and subsequently at risk of having insufficient private pension provision) as the participants in the experiment. Using participants from the target population implied some work on the presentation of the problem in a form that would be understandable. It is not clear that all these problems were overcome, though, in a sense this may not be a bad thing – as it reproduces the difficulties that this target population have in understanding the nature of their real-life pension problem.

1.7. Three different treatments of the experiment were run – with roughly one-third of the total of 174 participants in each treatment (Treatments 1, 2 and 3). Comparing behaviour in Treatments 1 and 2 allows an analysis of the possible effects of the introduction of a personal accounts scheme, while a comparison of Treatments 2 and 3 helps us to understand the possible impact of the introduction of a risk-free fund. The participants in the experiment were recruited by an agency.

1.8. In all treatments, allocations to pensions were higher than in the real-life *status quo* of the target population. This is almost certainly the consequence of two features of the experiment which are not present in real life: first, the regular prompting for pension decisions; and secondly, the provision of information about future pensions in a relatively easily digestible form (compared with information provision in real life). The experiment also showed that there were very strong gender affects and these varied from treatment to treatment. In Treatment 1 females allocated considerably more to pensions than males while the reverse was true in Treatments 2 and 3. These gender effects were observable with the help of demographic information that we collected from the participants. This information also helped us to identify a strong effect of risk attitude on behaviour – with the more risk-averse participants clearly choosing the least risky pension funds.

1.9. The *broad* findings that emerge from the experiment can be classified under two broad headings. First, one of the strongest messages emerging is that engagement with the 'pension problem' increases participation and involvement, and that the provision of information in a relatively digestible form aids decision making. Second, it was learned that the experimental *setting* is crucial, and that more thought needs to be given to this aspect if further experimental work is to be carried out: using a three-hour laboratory setting may not be the best way to understand real-life pension allocation behaviour.

1.10. The *detailed* findings of the experiment can be summarised as follows. Participants appeared to have understood the essential objective of the decision task in which they were engaged and tackled it with varying degrees of success. Compared to the *status quo* reference point (of almost zero pension allocations) participants performed well, even though they did not quite achieve the allocations that would be implied if they wished to maximise their expected earnings from the experiment. In comparing treatments, we saw that overall pension allocations were somewhat lower in Treatment 2 than in Treatment 1 and again somewhat lower in Treatment 3 than in Treatment 2. In all treatments, allocations to pensions had a generally downward sloping profile through time - as they should do to take advantage of compound interest. Nevertheless, in relative terms, in all treatments pension allocations were generally somewhat too low early in life and too high later in life. Females generally allocated less to pensions than males, and risk-averse participants allocated significantly less to pensions than those who were less risk-averse. Opting out of the saving scheme was almost minimal (as it should be) with very little variation across

treatments, gender or time. Comparing Treatments 1 and 2 there was very little difference in fund choice, though the introduction of a risk free fund in Treatment 3 caused a significant shift to safer funds. Rather obviously, risk-averse participants chose safer funds, as did females.

1.11. It is clear that the experiment was difficult. It was difficult for the experimenters to organise and to present the instructions and information needed by the participants in an easy-to-assimilate fashion. It was difficult for the participants to appreciate the nature of the decision task and to tackle it. Clearly participants were learning as they proceeded through the experiment. This learning process was helped by the instructions, the PowerPoint presentation, the Control Questions and the software. We cannot discount that learning also took place during the experiment itself, but it is difficult to know how one might test for this. There was no obvious convergence of the participants' strategies, though there was a modest reduction in the variability of their allocations through time. This could be taken as some modest sign of learning by doing.

2. Terms of Reference and Policy Background

2.1. **Terms of Reference** (taken from the Framework Agreement)

"This proposal is for the use of experimental methods in the investigation of a number of possible pension designs, concentrating on the likely responses of the British public in terms of their private savings decisions and their contributions to a National Pensions Savings Scheme. The methodology used will be that of implementing in the laboratory a simplified replication of the decision problem faced by individuals when deciding on their saving and pension provision. Of necessity, the experimental implementation will concentrate on the key issues and abstract from apparently irrelevant issues. Several different treatments will be implemented, with the objective of being able to infer from the experimental responses the differential responses of the British public to different possible features of a National Pension Saving Scheme. An essential feature of the proposed experimental implementation will be the use of a sufficiently heterogeneous set of subjects – so that inferences can be made about the aggregate and distributional responses of the entire British population."

2.2. Policy Background: Introduction to personal accounts (contributed by the DWP)

2.2.1. In December 2002 the Government set up the independent Pensions Commission to review the longer-term challenges faced by the pensions system and make recommendations for reform. The Commission published its conclusions in November 2005, setting out its proposals for meeting the challenge of providing a fair and adequate retirement income for all in retirement. Concurrently, the National Pensions Debate invited the public to comment on the emerging options for reform, culminating in the National Pensions Day in March 2006 in which members of the public across the UK were asked to consider and vote on the framework of the Pension Commission's proposals

2.2.2. Building on the Commission's report and the findings of the National Pensions Debate, the Government published proposals for pension reform in May 2006 in the White Paper 'Security in retirement: towards a new pensions system'. These proposals were designed to meet the 5 tests of personal responsibility, fairness, simplicity, affordability and sustainability and set out a new structure for the UK pension system for the long term. The proposed reforms include changes to the State pension system, primarily: increasing the State Pension age to 67 for both men and women, streamlining the contributions conditions, and changing uprating to be calculated in line with earnings rather than prices.

2.2.3. The second major aspect of the proposed reforms involves the introduction of a new national pensions saving scheme called personal accounts, to make it easier for more people to save more for their retirement.

2.2.4. Personal accounts are particularly targeted at the estimated 7 million people who are not saving enough to provide retirement incomes they are likely to consider

adequate. This group tends to be younger and on moderate to low incomes. They are also likely to be part-time workers and/or work for small employers.

2.2.5. There are persistent and powerful barriers to people taking long-term saving decisions, including inertia, financial myopia, the cost of pension saving and the complexity of the decisions involved. Personal accounts are designed to tackle these barriers, using automatic enrolment to combat inertia and with a combination of employer contribution, tax relief and low costs providing good incentives to save.

2.2.6. In future, individuals will be automatically enrolled into personal accounts, or comparable workplace pension schemes, if they are employees aged between 22 and State Pension age and earning above approximately £5,000 per year. Employees will contribute around 4 per cent of their earnings between approximately £5,000 and £33,500 a year; this will be matched by 3 per cent from their employer together with around 1 per cent from the State in the form of normal tax relief.

2.2.7. The key items of interest that motivated the implementation of the experiment in this particular form include:

2.2.7.1. What factors determine the decision to opt-out of the personal accounts scheme?
2.2.7.2. Is there any pattern to the choice of funds by the participants?
2.2.7.3. What is the effect of introducing a risk-free fund?
2.2.7.4. Is there an influence of risk-aversion on behaviour?
2.2.7.5. Are there any significant demographic variables influencing behaviour - such as gender, age and income?
2.2.7.6. Do participants come close to showing behaviour that might in some sense be considered good behaviour?
2.2.7.7. Did participants make as much money out of participating in the experiment as they could have done?

3. Developing an Experiment in the Field of Pensions

3.1. Preamble

3.1.1. This project was initiated following discussions within the DWP as to the possible use of experimental economics as an evaluation tool within its Social and Economic Research function. Until the initiation of this project, experimental methods had not been used within the DWP. It was thought that experiments might be of assistance in complementing other methods of enquiry. This section describes the methodology of experimental economics in general and the difficulties and challenges of applying them in the field of pensions in particular.

3.2. Background on the Methodology of Experimental Economics

3.2.1. Experimental economics has a well-defined methodology, which can be summarised as follows. Participants are recruited and presented with some kind of decision problem. They are usually paid for their participation, with their payment related to their performance in the experiment. This latter is intended to provide participants with some kind of 'appropriate' incentive for taking the decision problem seriously. Usually, different sessions of the experiment are run (normally with different participants in the different sessions). In the different sessions the problem presented to the participants differs. The experimenter can compare behaviour across the different sessions. This enables the experimenter to understand how changing the nature of the problem changes the way that participants respond to it. This, in turn, enables inferences to be made as to how participants are solving the decision problem posed to them. In addition, when the problem posed to the participants has a unique 'best' solution, the experimenter can see how close the participants are getting to this 'best' solution.

3.2.2. In experimental economics, the problem posed to participants is usually an economics problem – the intention being that of seeing how participants solve that problem. The information can be used to test economic theories or to generate new theories, that is, new explanations of behaviour. Applications of the experimental approach have been made across the whole field of economics, from individual behaviour, through games to markets; from microeconomics, through macroeconomics to international trade. **In the field of pensions, however, experiments are only just beginning to be run.** Some of the reasons for this will become apparent as this report proceeds. They include:

3.2.2.1. The inherent difficulty of the pension problem;

3.2.2.2. The difficulty of replicating in a laboratory experiment that lasts a few hours a decision problem which lasts over a lifetime;

3.2.2.3. The difficult in conveying information to the participants in a form that is reasonably easy to digest.

3.2.3. One main advantage of using experimental economics is that the problem presented to the participants can be posed in controlled conditions – the experimenter can control other factors that may influence the decisions of the participants (such as variations in the individual's employment status or income - these factors being irrelevant to the solution of the problem posed in the laboratory). Thus, the data obtained on behaviour is not influenced by these other factors. This is in stark contrast with data obtained using other techniques. Control is a key feature of the experimental approach. It enables the experimenter to investigate behaviour under controlled conditions – thus eliminating other factors which may influence behaviour. In a sense, it allows testing of theories under the conditions underlying the derivation of those theories.

3.2.4. Usually, in order to maintain this control (over the conditions under which experiments are run), experiments are typically implemented in an experimental laboratory. Such a laboratory is, in essence, a room with computer terminals at which the participants sit and perform the experiment. Usually, participants are screened one from the other, so that there are no uncontrolled interactions between the participants. However, and particularly so in recent years, experiments have been increasingly run outside the laboratory – the term used is 'in the field', though the

extent to which it is actually in the field varies from experiment to experiment. Such field experiments have certain advantages – they may be nearer to the real world – but they also have disadvantages, particularly that some of the control crucial to experiments may be lost.

3.3. Applying Experimental Economics in the Field of Pensions

3.3.1. This involves a number of stages. In order to simplify this section, we split up the process of applying the methodology of experimental economics in the field of pensions into three broad stages: (1) a first abstraction (eliminating some extraneous considerations); (2) a second abstraction (pinning down the details of the experimental implementation); (3) financial incentives (that is, how do we pay the participants to give them an appropriate incentive to take the experiment seriously?).

3.3.2. **Our first Abstraction - reducing the complexity of the problem.**

3.3.2.1. In essence, investigating some economic phenomenon using experimental economics involves replicating the decision problem in the laboratory or in the field. We study behaviour in the experiment and this enables us to understand behaviour in that context in real life. Of necessity, a replication must involve some kind of abstraction – we must take the key elements of the decision problem and replicate them in the experiment. This is, of course, exactly what theory does. So experiments follow theory in terms of its methodological abstraction. Let us examine this in the field of pensions.

3.3.2.2. **The real-life pension problem is very complicated.** People can put money into pensions at any stage of their working life; they can chose when to put money into pensions and how to put money into pensions; they usually also have to decide when to retire and then in what form to take their pension entitlements. Economists necessarily abstract from this problem, and frame it in a stylised way that seems, to the economist, to capture the key elements. Let us consider one such abstraction.

3.3.2.3. The usual starting point is to assume that life consists of a number of discrete periods - like years or months. In practice, of course, life is continuous and decisions about pensions can be made at any stage, but a discrete model is easier to handle and to describe. The number of periods in any person's life is necessarily random – as the date of death is not known. In each of the periods of a person's life, he or she may receive an income, and this income could be either determined by the individual or by outside forces. Before retirement (which itself could be decided by the individual) he or she must decide how to allocate his or her income – between saving in various forms (including pensions) and consumption in various forms. After retirement, the individual receives an income (which will partly depend on his or her past decisions about savings and pensions) and must decide once again how to allocate this between various forms of saving and various forms of consumption. If we are interested solely in pensions, we can further abstract from this problem, by proceeding as follows: we can assume that income and the retirement date are determined by outside forces, and that the only decision that the individual can make is how to allocate his or her income between consumption and pension allocations in the periods preceding retirement. This is the abstraction on which we based this experiment. We think that it captures the key elements of the pension decision. Clearly, if it does not, then our experiment will depart from the problem faced by real people in their real life.

3.3.3. **We now consider a second abstraction: how do we implement this in the laboratory?**

3.3.3.1. We now have an abstraction on which we can base our experiment. However, at this point we encounter a major problem – in real life the 'pension decision' occupies a whole life. How do we reproduce this in an experiment? Clearly in practical terms this is impossible, so we once again need to make an abstraction. Crucial to this is whether the passage of real time has an effect on the decisions. Economists would say that this is one important element – as it is clear that individuals evaluate differently incomes and expenditures at different points in time: individuals discount[1] incomes and expenditures in the future relative to incomes and expenditures in the present. If we want to incorporate such discounting in our experiment we need to implement an experiment which continues over a sufficiently long period of time – if not necessarily a lifetime. This necessarily makes a laboratory experiment lasting a few hours impossible. The alternatives are either to run an experiment in which participants return to the laboratory on several occasions over a period of time or one in which participants carry out the experiment at their homes over a period of time. Both these alternatives imply a loss of control: both of them because the circumstances of the participants change in an uncontrolled way over time, and the second because the experimenter does not know exactly who is responding to the experiment on the various occasions. However, both these alternatives have one clear advantage: one could set the experiment up in such a way that participants were not forced to take active decisions at any stage – thus more accurately reflecting the real-life pension allocation decision.

3.3.3.2. An alternative abstraction starts from the premise that the passage of real time is not the only key element in the pension decision – another is the allocation of income across the lifetime. Irrespective of whether individuals discount the future, they still have to decide how much of their incomes to allocate to pensions. This decision is usually necessary because income after retirement is typically lower than income before retirement – and therefore the provision of a pension is necessary to smooth consumption over the lifetime: if individuals do not make an adequate pension provision then their consumption after retirement is less than they would wish it to be. So, even if we ignore discounting, we still pick up **an important element of the pension decision – that of providing the appropriate smoothing of consumption over the lifetime[2].** This is the element that we focussed on in this experiment, and we decided to ignore the issue of the passing of real time and hence avoid the complications involved in running an experiment extending over a significant period of real time.

3.3.3.3. We therefore arrived at the following abstraction: individuals would live a random number of periods. Before retirement they would receive a given income and their decision problem would be that of deciding how much to allocate to pensions and how much to consumption; after retirement they would receive an income in the form of a pension (partly determined by their contributions to pensions before retirement and partly consisting of a basic pension), all of which would be devoted to consumption. There would be no significant passage of real time and hence no discounting[3].

[1] Appendix 10 contains a glossary of some of the more technical terms used in this report.

[2] Readers who are interested in pursuing the technical aspects of this consumption smoothing exercise might find it helpful to refer to Technical Appendix 7.

[3] An extended comment is relevant at this stage: even though an experiment may not last a significant amount of real time, it is still possible to include discounting in the experiment – by having an increased probability of the experiment

3.3.4. **We now need to discuss financial incentives - how do we get the participants to take the experiment seriously?**

3.3.4.1. At this stage, we should note that the abstraction at which we have arrived is very similar in some respects to what economists term the life-cycle model of consumption and savings. There have already been experiments on this model: see, in particular, Ballinger *et al* (2003), Carbone (2005), Carbone and Hey (2004) and Hey and Dardanoni (1998). These, however, were not set specifically in a pension context, and hence our experiment differs in certain respects from these earlier experiments; we will explain the differences shortly. However, we borrow one key aspect from these earlier experiments – the payment or incentives structure. It will be recalled from the discussion above that **economics experiments usually involve a payment to the participants – which provides an appropriate incentive to take the experiment seriously.**

3.3.4.2. What does that mean in this context? To answer this we must go back to the key issue involved in the life-cycle model – the smoothing of consumption over the lifetime. Why do individuals want to smooth consumption? The economist's answer is as follows: consumption provides happiness (or utility, to use the economist's term); moreover, the more an individual consumes the happier he or she is; however, while happiness (or utility) increases with consumption it does so at a decreasing rate[4]: an increase in consumption from £600 in a year to £700 in a year increases an individual's happiness, as does an increase in consumption from £700 in a year to £800 in a year – but the increase in happiness implied by the latter increase is less than increase in happiness implied by the former increase. The implication of this is that individuals want to smooth consumption: if presented with a choice between consuming £600 in one year and £800 in the following year, or consuming £700 in each of the two years, individuals prefer the latter. They prefer to smooth fluctuations in consumption.

3.3.4.3. How do we implement this in an experiment, particularly an experiment that lasts a short time in the laboratory? While we can divide the experiment up into periods and give the participants some income in each of the periods, we cannot actually get the participants to spend their income on consumption goods and get them to consume these consumption goods during the experiment. The problem is simply that the time spent by the participants in the laboratory is too short. Hence, we adopt an alternative procedure. We pay the participants an amount of money for participating in the experiment – that depends, in a way that we will explain shortly, on the decisions that they take. We tell them that their objective is to make as much money as they can out of the experiment, and we set up the payment mechanism in such a way that their payment is the experimental analogue of the lifetime happiness in the economist's story. Accordingly, we say that, in each period, the amount of their income that they 'consume' will be converted into real money, and we will pay them the total of the money equivalents of their consumption in all the periods for which the experiment lasts. So real money earned in each period of the experiment is the analogue of the 'happiness' that the individual experiences in that period (from consumption), and the total amount of money earned in the experiment is the analogue of the lifetime happiness of the individual. To make the analogy complete we need that this conversion (from consumption in the experiment to real money

---

finishing at the end of any one period. This was implemented in Hey and Dardanoni (1988), but introduces extra complexity into an already complex experiment.

[4] It is not being asserted that this is a fact - merely an empirical observation: if you repeatedly give an individual increases (of the same amount) in his or her consumption, most individuals report that each increase makes them happier, but for each one the increase in happiness is less than for the one before.

paid to the participant) is exactly like the conversion of consumption (in the economist's story) to happiness: the more the individual converts/consumes (in the experiment) the more real money he or she earns for that period in the experiment, but real money increases at a decreasing rate with consumption. Thus, for example: 600 units of income in the experiment converted into real money become 82 pence, 700 units become £1.17 and 800 units become £1.48. Thus, an increase in conversion/consumption from 600 to 700 increases earnings in the experiment by 35 pence, while an increase in conversion/consumption from 700 to 800 increases earnings by 31 pence. Thus, according to economic theory, participants have an incentive to smooth their conversion/consumption over their lifetime. One of the purposes of the experiment is to see whether they do.

3.3.4.4. The experiments referred to above have all implemented this particular version of the life-cycle model of consumption/saving – using this payment mechanism. Although the results are not directly relevant to the pension version, which we describe in more detail shortly, there are some features that are of interest. In particular, it was noted that consumption generally was too sensitive to fluctuations in income – so that participants did not achieve the appropriate amount of consumption smoothing[5]. This phenomenon closely mirrors evidence from elsewhere on savings behaviour (see, as an excellent reference on both theoretical and empirical evidence on consumption and saving, Deaton, 1992). It has been interpreted as a sort of myopia – in that participants are not able to understand the full opportunities for consumption smoothing. Rather than look at their lifetime as a whole, participants seem excessively concerned with the short-term future. This type of myopia it seems is also at work in pension decisions.

3.3.4.5. The experiments we have referred to above were primarily concerned with investigating the life-cycle model of consumption and savings. In these experiments, savings were a general and flexible way of transferring income through time, and hence of smoothing consumption. In a pension context we need to restrict savings to be of a particular form. Accordingly, in our experiments money saved, that is, money allocated to pensions, could not be withdrawn until after retirement – and then only in the form of an annuity[6]. We also imposed the constraint that, once allocated to a particular fund, the allocation would remain in the fund throughout the remaining working lifetime of the participant. These latter constraints constitute further simplifications of reality. However, in these pension experiments we supplied different funds in which pension allocations could be made – which differed in terms of their rates of interest. **The DWP were particularly interested in the issue of the fund choice of individuals.**

3.3.5. This completes the basic design of the experiment: a discrete random period world in which the individual receives a fixed and known income in every period before retirement; in which individuals have to decide how much of their income to 'consume' (convert into real money) and hence how much to allocate to pensions; in which there is a choice of pension funds into which to allocate any pension contributions; and in which the objective of the participants is to maximise the amount of money they earned from participating in the experiment. We now discuss the details of the experiment that we conducted.

[5] We do not really think it helpful at this stage to go into technical detail but an example may be useful. If individuals value consumption in all periods equally, if the interest rate is zero and if there is no uncertainty in income, then consumption should be the same in all periods. If there is *any* variation in consumption it is clearly excessive. In the experiments referred to in this section, there were reasons for some variation in consumption, but not as great as that observed.

[6] A definition of *annuity* can be found in Appendix 10.

3.4. The Specific Details of the Experiment

3.4.1. **There are usually different experimental sessions conducted in an experiment.** In the different sessions, there are different 'rules of the game' imposed. This enables us to see how the rules influence behaviour. In addition, of course, within any experimental session we can see how close behaviour is to what may be considered to be 'best' (as defined by economic theory) behaviour. We return to this latter point later, after we have described the various experimental sessions that we implemented. **In experimental economics jargon we refer to these different types of session as different *treatments*.**

3.4.2. The different treatments in an experiment are usually determined by the objectives of the experimenters: different treatments enable us to see how behaviour changes when we change the 'rules of the game' of the experimental problem. As section 2 of this report makes clear, the experiments were being designed in the context of a possible introduction of some kind of personal accounts scheme. As a consequence, certain features were of interest, including the effect of automatic enrolment on behaviour, the fund choice of individuals and the effect of the introduction of a risk-free fund. It is normally considered to be good experimental practice to have only one element different between two different treatments.

3.4.3. Our basic treatment, which we call *Treatment 1*, is one in which participants are regularly prompted to take conscious decisions (by the simple fact that they are sitting there in the laboratory), and in which the decision problem is that outlined above. More specifically their pre-retirement working life consisted of 15 periods in each of which they received an income of 1,000 experimental units. In each of these periods they had to decide how much to allocate to pensions and in which find to allocate it. We need now to add a little detail about the funds in which participants could invest their pension allocations. As we were particularly interested in the choice of funds, yet at the same time wanted to keep things as simple as possible, we allowed participants a choice of three funds which differed in their riskiness. All three funds had two possible rates of interest (each equally likely) and were as follows:

Fund 1: either 1.0% or 5.2% (average 3.1%)
Fund 2: either 1.5% or 4.3% (average 2.9%)
Fund 3: either 2.0% or 3.4% (average 2.7%)

3.4.4. It will be noted that as the average rate of interest rises going from Fund 3 to Fund 2 and then to Fund 1 and so does the riskiness of the fund (Fund 3 has rates on interest ±0.7% around its average, Fund 2 ±1.4% around its average and Fund 1 ±2.1% around its average). Considerable discussion took place at the design stage of this experiment as to whether we should include funds with possibly negative rates of return, but it was decided in the end that it was *relative* riskiness of the various funds that was of interest – rather than their *absolute* riskiness. In addition, given that we had promised participants that they would take away at least their participation fee of £30 from the experiment, we did not want to leave open the possibility that participants might lose money on the experiment itself. This ruled out negative interest rates.

3.4.5. After retirement, participants received a basic pension of 100 experimental units plus an extra pension arising from their pension allocations during their working life. The basic pension implies a replacement ratio of just 10% - we note that this is unrealistically low, but was designed so as to provide a strong incentive for allocating money to pensions.

3.4.6. *Treatment 2* involved the introduction of a stylised version of a possible variant of the personal accounts scheme – which we termed the 'savings scheme'. The key elements that we wanted to introduce were the following: (1) that participants would be by default enrolled into the scheme (and hence that they had to take a conscious decision to opt out); (2) that there would be a certain minimum pension allocation if they wished to remain in the scheme; and (3) that there would be a matching contribution made from elsewhere – if the participants remained in the scheme. The actual details of these elements were less crucial than their existence: we choose a minimum contribution rate of 10% (of income) and a matching contribution of 75% from elsewhere[7].

3.4.7. Finally *Treatment 3* introduced a further element of interest for the DWP – the possible effect of the inclusion of a risk-free fund in the set of available pension funds. This treatment therefore was exactly the same as Treatment 2 with the addition of an extra fund (Fund 4) with a certain rate of interest of 2.2%. Note that this rate of interest is lower than the average rate of interest on any of the other funds.

3.4.8. As a reference point for the subsequent discussion, we introduce a fictitious treatment which might be deemed as representing the *status quo*. This is a situation in which participants are not regularly prompted to take decisions. We call this fictitious treatment *Treatment 0* – the '0' signifying that we did not in fact implement it – for the reasons which we have already explained above.

3.4.9. Table 1 summarises the various treatments.

3.4.10. The Main experiment was preceded by three pilots: two with student subjects and the third with a sample of participants from the same target audience as the Main Experiment. The Main Experiment itself involved a total of 174 active participants, each of whom was paid a participation fee of £30 (and any expenses) plus their actual earnings from the experiment itself.

3.4.11. In each session of the experiment, participants were greeted by a member of EXEC, reminded that they would be paid the participation fee of £30 plus any earnings that they made from the experiment, and then taken to the laboratory. There they were given a set of written instructions. After they had read these, a PowerPoint presentation was played on their individual screens, after which they were given a set of Control Questions to ensure that they understood the instructions. Any wrong

[7] A word about this figure of 75% and the treatment of tax is necessary at this stage. Pension allocations are subjected to tax relief. Therefore, if the income of experimental units is considered as *post-tax* then the following implications would flow: participants opted out of the personal accounts scheme would have 25% added to their pension allocations (the tax rebate) while participants opted in would have 100% (25% from the tax rebate and 75% from the augmentation) added to their pension allocations. For simplicity, we decided to treat the income of experimental units as *pre-tax*, so that no tax rebate would be forthcoming in either case. Hence zero would be added if out of the scheme and 75% added if in the scheme.

answers were discussed with them and any questions they had were answered. They could then begin the experiment, carrying it out at their own speed (subject to a certain minimum time). At the end, they completed a brief questionnaire, asking for demographic information, were paid their earnings and were free to leave.

Table 1: The various treatments

	The experimenter provides a 75% matched contribution(1)	**Fund 1 High risk (3.1%)**	**Fund 2 Medium risk (2.9%)**	**Fund 3 Low risk (2.7%)**	**Fund 4 Guaranteed (2.2%)**
Treatment 0 *Scenario which 'artificially'(2) simulates a status quo of individuals not saving*	✘	✘	✘	✘	✘
Treatment 1 *Scenario where participants make saving allocation and fund choice decisions without personal accounts*	✘	✓	✓	✓	✘
Treatment 2 *Scenario where participants are auto-enrolled into a personal account and make saving allocation and fund choice decisions*	✓	✓	✓	✓	✘
Treatment 3 *Scenarios where participants are auto-enrolled into a personal account, make saving allocation and fund choice decisions with the choice of an additional fund offering a guaranteed rate of return*	✓	✓	✓	✓	✓

4. Broad Findings and Conclusions from the Experiment

4.1. Preamble

4.1.1. We divide this section into two sub-sections:

(1) What this particular experiment tells us about pensions;

(2) How it might assist the DWP in understanding pension behaviour and the possible impact of policy measures.

4.2. Conclusions Concerning the Experiment that we conducted: what this experiment tells us about pensions.

4.2.1. We have reproduced in the experimental economics laboratory an abstraction of the pension problem facing a typical individual. We have implemented this experiment on a quota sample of the part of the UK population which is currently at risk of entering retirement with inadequate pension provision. We implemented three different versions of the pension problem – with different specifications of the options open to the participants in the experiment. We have not only analysed the behaviour of the participants in absolute terms, but also have compared behaviour across the different versions of the experiment and with two possible theoretical accounts of behaviour.

4.2.2. **The abstraction that we implemented was a stripped-down version of the generic pension problem, in which participants received a constant income up to retirement and a basic pension after retirement, and in which participants were required to allocate their income before retirement between consumption and saving (towards an augmented pension).** Participants could choose between various funds in which to put their pension allocations; most of these funds were risky, as was the length of life of the participants in the experiment. Two of the versions of the experiment (Treatments 2 and 3) had, in addition, a savings scheme in which allocations made by the participants were matched with a 75% contribution from the experimenter. One of these (Treatment 3) version had a risk-free fund as one of the fund options. Treatment 1 had neither of these features.

4.2.3. Because of the nature of an experiment, participants were given information in a form that was relatively simple - compared to their real-world situation. In particular, participants could see (in both graphical and tabular) form, the precise implications in terms of likely future pensions of any decisions being considered. We think that the provision of information in this form had a material bearing on the behaviour of the participants. Although it was not particularly simple for participants to work out what they ‘should’ be doing, they had information and instructions which significantly informed them and pushed them in the right direction. This does not reflect accurately their actual circumstances in their real-world decision problem.

4.2.4. Participants in the experiment were given an incentive to consider their allocations seriously – with their payment for participation depending upon their decisions. One of the theoretical accounts of behaviour, with which we compared actual behaviour,

was an allocation strategy which maximised the expected payment to the participant for his or her participation in the experiment (we call this the risk-neutral model-optimal strategy - because a risk-neutral individual is interested in maximising his or her expected return and is not concerned about risk). The other theoretical account was one in which participants made no allocations to pensions throughout their lifetime (we call this the *status quo* strategy). This latter might be considered close to the real-life actual behaviour of the target population. **We found that actual behaviour in the experiment was much closer to the first of these two theoretical accounts than the second.**

4.2.5. If we compare behaviour across the various treatments we find that absolute earnings were highest in absolute terms in Treatment 2 and lowest in Treatment 1, both for males and for females. However, relative to the maximum possible earnings from the experiment, earnings for males were highest in Treatment 2, and lowest for Treatment 3. For females, earnings in these relative terms were roughly constant across treatments, though slightly higher in Treatment 1 and slightly lower in Treatment 3. These give some indication of how close participants came to maximising their 'lifetime happiness' in the underlying pension problem.

4.2.6. The actual allocations to pensions, in effectively all treatments and for both males and females, declined slightly through time – as they should do to take advantage of the effects of compound interest. However, they did not decline as fast as they do in the risk-neutral model-optimal strategy. We took this to be evidence of a type of myopia – but quite a subtle kind of myopia – not one that ignores the need to make pension provision, but one that does not consider the full life-time effects of pension accumulation. Interestingly, in all treatments and for both males and females, pension allocations were considerably higher than both those implied by the risk-neutral (see section 4.2.5 above and the glossary in Appendix 10) model-optimal strategy and the *status quo* strategy. We suspect that this may have partly a consequence of the risk attitude of the participants, and also a consequence of the experimental design – in that participants were forced to take conscious decisions about their pension provision – unlike in real life. **Indeed, this latter is one of the striking features of the experiment: prompting them to take active decisions, and providing them with full and relatively easy-to-digest information, implied that they took decisions that were remarkably good.** Given the design of this experiment, it is not possible to determine the relative importance of these two possible explanations – though a further experiment could be designed to shed light on this important point. A further possible reason for the rather modest decline in pension allocations through time could be the effect of learning: as time passed participants became more aware of the need and importance of making pension contributions. However, it is difficult to know how to test this learning hypothesis and how to disentangle it from other factors influencing behaviour.

4.2.7. In the treatments with the saving scheme (in which participants were automatically enrolled into the scheme and had to take a conscious decision to opt out) we noticed a certain amount of inertia in behaviour. Indeed the opt-out rate was remarkably low. However, this comment should be set against the risk-neutral model-optimal strategy - which says that it is never rational to opt out. Given that in the experiment participants opted out 4.4% of the time in Treatment 2 and 8.6% of the time in Treatment 3, it is not clear how to interpret this evidence. On the one hand, 4.4% and 8.6% are significantly greater than 0%; on the other hand, 4.4% and 8.6% are small and hence the stay-in rate is high. Moreover, given that the participant is obliged to

contribute at least 10% of his income when not opted-out, this means that a high proportion of the time in Treatments 2 and 3 participants were allocating more than 10% of their income to pensions. This should be contrasted with the fact that in Treatment 1 participants were allocating less than 10% of their income 28.8% of the time. Nevertheless mean contributions were higher in Treatment 1 than in Treatments 2 and 3. So the 10% floor in the saving scheme may not only have been a minimum but also a target.

4.2.8. Fund choice was quite clearly influenced by the risk-attitude of the participants with the more risk-averse (and hence more of the females than the males - see Cox and Harrison 2007) choosing the less risky funds. With the introduction of a risk-free fund in Treatment 3, there was a significant number choosing that fund and a general shift towards the less risky funds. A formal analysis shows strong treatment effects - with the safer funds chosen more often in Treatment 2 than in Treatment 1 and (rather obviously) much more often in Treatment 3 than in Treatment 2. Elderly people were more likely to choose safer funds, as were those who were risk-averse, independently of the treatment.

4.2.9. In all treatments there was a default fund, and it was noticed that participants were more likely to stay in the default fund in Treatment 2 than in Treatment 1, and that they are less likely to stay in the default fund in Treatment 3 than in Treatment 1. As time passes they are more likely to stay in, while the older participants are more likely to stay in, as well as those with higher incomes (in real life); those that are more risk-averse are more likely to move out of the default fund. It is not clear what implications these findings have for policy.

4.2.10. We should note that our comments on the effect of risk-aversion on behaviour in the paragraphs above are concerned with relative risk-aversion - the more risk-averse a participant the smaller his or her pension allocation. Our questionnaire-based index of risk attitude does not tell us anything about absolute risk attitude. So, if we have two participants, one more risk-averse than the other we cannot know whether both are risk-averse, both are risk-loving or one risk-averse and one risk-loving.

4.3. Conclusions Which Might Assist the DWP in Understanding Pension Behaviour and the Possible Impact of Policy Measures

4.3.1. In trying to draw some conclusions about the success or otherwise of the experiment, we should begin by deciding our criteria for judging success. It was clear from the outset that the experiment was ambitious both in its design and in its basic intentions: those of reproducing in the laboratory the ‘pension problem’, and expressing it in such a way that it was accessible to the target population. The experiment could then be judged as to how successful it was in doing that; this may be considered the first of the criteria for judging its success or otherwise. On the other hand it could be judged on the basis of how well it answered the key questions of the research (concerning opting-out behaviour, fund choice, the effect of the different treatments and so on).

4.3.2. As far as the first of these criteria are concerned we showed that we could reproduce in the laboratory a stylised version of the pension problem and organise the

experiment in such a way that it was accessible to the target population. We are not claiming that it was 100% successful as there were some participants who were unable to understand the instructions and others who only understood the instructions as they proceeded through the experiment, but the majority of the participants reported little difficulty and responded to the control questions correctly.

4.3.3. As far as the second criterion is concerned, we should carefully distinguish between whether the experiment produced results (relating to the key questions of the research) and what those results were. Given that between Treatment 1 and 2 there was only one difference - that of the existence of the savings scheme - a comparison of the results of these two treatments tells us immediately about the likely impact of the introduction of the savings scheme (relative to Treatment 1). Treatment 2 also tells us about likely opt-out rates. Again, given that between Treatment 3 and Treatment 2 there was only one difference - that of the existence of a risk-free fund - a comparison of the results of the two treatments tells us immediately about the likely impact of the introduction of such a fund. So, in this respect, the experiment was bound to be a success. If it had an alternative research agenda, it would have had a different set of treatments.

4.3.4. The experiment also had a further aspect - the collection of some self-reported demographic data, and its use in analysing the data from the experiment. However, it is difficult to know what might be meant by judging the success of so doing. It is clearly of interest to know whether these demographics affect behaviour (though generally they did not) but whether they do or not does not help us to judge whether collecting the data was a successful exercise.

4.3.5. We should note at this stage that the experiment, by its very nature, was unable to start from the true *status quo* of the target population. This was a consequence of two features: first, that the experiment was conducted in the laboratory and lasted only between two and three hours; second, that we provided information in a form that would not normally be available to the target population.

4.3.6. The first of these features inevitably implied that participants were forced to take active decisions – they could not simply sit there and do nothing. We have already discussed how we might get round this problem – by conducting the experiment over an extended period of time with the participants having to take some positive action (like logging onto the internet) in order to take decisions. This is a possibility that might be considered for the future but which would require serious design considerations to be resolved.

4.3.7. The second feature – that of **information provision** – seems to us to be really crucial. We decided to give as much relevant information as possible about the likely future pension, and in a form that was relatively easily digestible. This was not an easy task as the provision of information in a ‘relatively easily digestible’ form is not straightforward. Indeed, we cannot claim that the information was easily digestible to all the participants. However, we may be able to able to claim that the way that information was provided in the experiment was perhaps easier to digest than information that the participants would obtain in real life. So, for example, we told them the distribution of their likely pension for any given planned level of pension allocation. We doubt that any pension salesman would give potential clients information in such a precise and unambiguous form.

4.3.8. We suspect that providing information in such a direct and as-transparent-as-possible format influenced behaviour. It perhaps made the 'pension problem' more obvious and transparent. So perhaps the way that we were providing information was driving our results. There are ways to discover whether this was in fact the case, which would have to be discussed before any future experiment. We could organise the software so that participants had to ask for information and had to specify the form in which it was provided. There are technical problems with writing the software in such a way that this would be possible, and there are also conceptual issues.

4.3.9. We should also make a comment about the context of the experiment. In this experiment, we consciously used the word 'pensions'. In addition, the participants were aware that the project was funded by the Department for Work and Pensions. Inevitably, therefore, participants were invited to think about and concentrate on pensions. We could have made the experiment totally context-free – by telling them that they received an income every period and could allocate this either to A or B. Income allocated to A would earn them money directly, while income allocated to B (or some fund within B) would yield them money only from periods 16 onwards (if they survived that long). In other words, we could have completely stripped the context from the experiment – and never mentioned the word 'pensions'. **This may well have led to different behaviour** (though, of course, we cannot test this proposition without carrying out such a further experiment)**.** Indeed, one of the major modifications at the design stage was to reduce the emphasis on pensions, though not removing it completely – and we noticed that this led to a reduction in the pension allocations (though we are not claiming that this was the reason for the reduction as there were other changes between the pilot and the main experiment). The idea of having such a context-free design might be considered for future experiments - though, of course, whether such a design would be a good idea or not does depend upon the objectives of the experiment.

4.3.10. There is another important issue. We, as experimental economists designing this experiment, presumed that we were presenting the participants with a problem to solve – and one that was the same, within any given treatment, for all the participants in that treatment. However, it is clear from the results that we present in the Report, and particularly in our analysis of the effects of demographic variables on behaviour, that **participants were importing their own personal circumstances into the laboratory**. This may have partly been for 'solving' the decision problem and also partly for understanding it in the first place[8]. For some of these variables, for example risk aversion, this is clearly something that will affect someone's decisions about pensions – though it is not clear how this could be incorporated into any pension design. Others are more subtle – like age. If age has an affect on pension allocation, in that the older think more carefully than the younger, how can this be incorporated into a pension policy? Basically it seems to be implying that the older may be less myopic (in either of the two senses that we have used this term) than the younger. But this implies that a given individual varies in his or her myopia over the lifetime. This, in turn, implies that the individual is dynamically inconsistent – wanting to do one thing when young and a different thing when old. How might the government (be able to) react to such an inconsistency?

4.3.11. **In conclusion, it may be useful to summarise what have we learned from the experiment about the behaviour of people with respect to their pensions.**

[8] With lengthy and potentially confusing Instructions, participants may well have interpreted the instructions based on their own experience.

4.3.11.1. First, we observe some kind of myopia (especially amongst the younger participants - this myopia seems to decrease with the real age of the participant). We infer this from the relative under-allocation of income to pensions early in the working life and the relative over-allocation in later working life. This may be due to an incomplete understanding of the effect of compound rates (though we cannot claim that this is an unambiguous result of this experiment). If this supposition is true, it suggests (all other things being equal, of course) that policy-makers should carefully consider or tailor messages to different age segments and communicate better to the target population the effect of compound interest rates.

4.3.11.2. Second we note that the real-life pension problem is very complex. Even in the simplifying world of the laboratory, the problem is difficult for the experimenter to present and explain, and is difficult for the participants to understand and 'solve'. The experiment showed that, even when presented in a more 'simplified' setting with ready-to-hand information, solving the pension problem remained complex for participants of the target population. We note that some evidence suggested that participants found it difficult to abstract from their personal characteristics - partly in terms of understanding the lengthy and apparently complex instructions, but also in terms of responding to and 'solving' the pension problem posed to them. Nevertheless, the behaviour of the participants in all treatments of the experiment suggested that they were tackling the problem in a reasonably adequate way, and certainly better than the target population in real life. This suggests that the provision of information in the experiment, although complicated, was significantly less complicated than in real life. Moreover, the information was all provided in one place and at one time. This suggests that the provision of simple, accessible (ideally found in one place) information might be important in supporting pension reforms. This conclusion would agree with some of the results coming from behavioural economics which suggest that avoiding information overload may subsequently lead to avoiding procrastination.

4.3.11.3. Third, we note that the experiment, by its very nature, involved the participants in engaging with the pension decision: except for the three people who said that they could not understand the decision task, all the other participants completed the experiment. Clearly the experiment was an artificial way to engage participants, but they were actively prompted to make decisions and were rewarded with a financial incentive and a participation fee. Proving information in a relatively simple and accessible format may be part of the solution, but 'forcing', in some sense, people to actually engage in the pension problem seems equally important.

4.3.11.4. The issue of risk and of risk-aversion is a difficult one - and the message emerging from the experiment is mixed. Risk-aversion appeared to have a strong effect on behaviour, with the more risk-averse participants allocating less to pensions (presumably because they saw investment in pensions a risky activity while the returns from consumption/conversion were immediate and safe). Moreover the introduction of a risk-free asset in Treatment 3 caused a reduction in the amount of money allocated to pensions. In addition, and perhaps rather obviously, the more risk-averse participants allocated relatively more money to the safer pension funds.

4.3.11.5. We should note that the information on risk attitude was obtained from the answers to the questionnaire – which was self-reported data. We used that data to construct a risk-aversion index – measuring the risk attitude of the participant (details of the construction are given in the Appendix). It is clear, from the analyses that we have carried out, that the risk aversion index so obtained contains useful information – in the sense that the effects of this risk-aversion on behaviour accord not only with intuition but also with previous results. Nevertheless, it is not clear

how this information may be used in any policy measures. It would be difficult in practice to measure the risk attitude of particular people in the population and it would be even more difficult to tailor policies for them. Perhaps the lessons from the economic literature on adverse selection may be exploited: by offering a package of pension policies, some targeted at, and tailored for, the more risk-averse, and others targeted at, and tailored for, the less risk-averse.

4.3.11.6. Penultimately, we should comment on inertia in behaviour and whether this could be useful in building pensions policies. Unfortunately, no clear pattern emerges from the experimental data. While there are participants who chose roughly the same allocation to pensions each period and stuck largely with the same pension fund, there were many participants whose pension allocations changed markedly from period to period, as did the fund they chose in which to invest their allocations. Moreover the evidence on the tendency of participants to stay in the default fund does not appear to be particularly strong and does not give any clear direction for policy. Given the design of the experiment, no general finding on inertia emerges. This does not preclude the possibility that other designs may have encouraged inertia.

4.3.11.7. Finally, we should note that the other experimental implementations may have shed more light on certain aspects of behaviour. While the experiment has shown how people ‘solve’ the pension problem while being provide with directly relevant information and actively engaged with it, the experiment sheds no light on what people do in other circumstances. This suggests that other experiments – with less prompted involvement and with information less readily available –may help us understand how the target population are currently responding to their pension problem.

VOLUME 2: THE TECHNICAL APPENDIX

5. The Actual Implementation of the Main Experiment

5.1. Challenges with Implementing the Experiment

5.1.1. Volume 1 of this report described the basic structure of the experiment and the various treatments that we implemented. We now turn to some practical challenges involved with its implementation. Later, in section 6.2, we will describe the practicalities of its execution.

5.1.2. One immediate problem is the complexity of the experiment, not necessarily in terms of the practicalities of taking decisions[9], but rather in conveying information to the participants about the experiment. **Indeed, describing the experiment turned out to be a difficult matter: the instructions went through numerous iterations before we finally converged onto the Instructions which we used in the experiment itself** – those for Treatment 3 are reproduced in Appendix 3 to this Report. The difficulty was in explaining quite complicated concepts to the participants in a language that would be understandable to them. In addition, considerable efforts were made to reduce the sheer volume of the information contained in the Instructions - though it may well have been the case that some participants found the Instructions too lengthy. (There is a serious problem here in getting the appropriate trade-off between giving all the information that participants needed and their presentation in an appropriately straightforward manner - a problem not aided by the fact that the information was complicated.)

5.1.3. Particularly difficult was explaining to the participants the implications – for their future pension – of any proposed allocation to pensions. The main problem was that interest rates were, in general, risky and therefore that the implied pension was not certain. How does one describe a risky pension? Of course, given that we knew the possible rates of interest on each of the pension funds, we could work out mathematically the probability distribution of the future pension. But we still had to convey this information to the participants. There are different ways that one can do this: (1) in words, for example, "there is a 0.125% chance that your pension may be 100 per period, a 0.25% chance that it may be 110 per period…."; (2) graphically in the form of a histogram; or (3) in various summary forms – giving the mean, median, mode, standard deviation and so on. Because of the large number of possible values it was decided that form (1) was simply impractical; form (2) was more practical but we realised that some participants may not know how to interpret a histogram; form (3) was perhaps the simplest but we still had to decide which summary statistics to provide. In the end we provided both forms (2) and (3) and restricted the summary statistics to be the quartiles[10] of the distribution. An example can be seen in the "Main Decision Screen in the Experiment" which can be found in the Instructions in Appendix 3.

[9] In essence all the participants had to decide in each of the periods before retirement were: (1) how much of the income to allocate to pensions; (2) whether to opt out or stay in the savings scheme (in Treatments 2 and 3); (3) and in which fund to make the allocation.

[10] The quartiles are as follows: (1) the lowest possible pension; (2) the value below which the pension will be with a 25% chance; (3) the value below which the pension will be with a 50% chance; (4) the value below which the pension will be with a 75% chance; (5) the highest possible pension.

5.1.4. An additional problem was the fact that the future pension would not only depend on the allocations that the participant had made in the past, but would also depend on the allocations that were made in the future. We decided, therefore, to give information about the distribution of the future pension in two scenarios: Scenario 1, in which the participant made no contributions to pensions in the future; and Scenario 2, in which it was supposed that the participant made exactly the same allocation to exactly the same fund in all remaining periods before retirement.

5.1.5. We should note that, while we had to provide this information for the benefit of the participants in the experiment, the provision of the information in this form is more than the target population would normally receive. Indeed, in real life the relevant information is often (1) difficult to obtain; (2) possibly uncoordinated and conflicting (in the sense that different documents may contain different assumptions) and (3) difficult to assimilate. Thus, participants in the experiment were at a considerable advantage in this respect. This should be taken into account when interpreting the results of the experiment. Nevertheless it remains true that some participants found the Instructions sufficiently lengthy to cause confusion by their sheer length.

5.1.6. We also had to decide on the conversion scale – from consumption to real money – that we would use in the experiment. We have already explained that it has to be increasing, but at a decreasing rate. The reason for this has been explained elsewhere. In essence the conversion scale (from experimental units to money) mimics the utility function (from consumption to utility) of the economic theory. It is usually assumed that this utility function is increasing at a decreasing rate (that is, marginal utility is declining) - hence the reason for this particular shape for the conversion scale. The actual scale used in the experiment is shown in Figure 1.

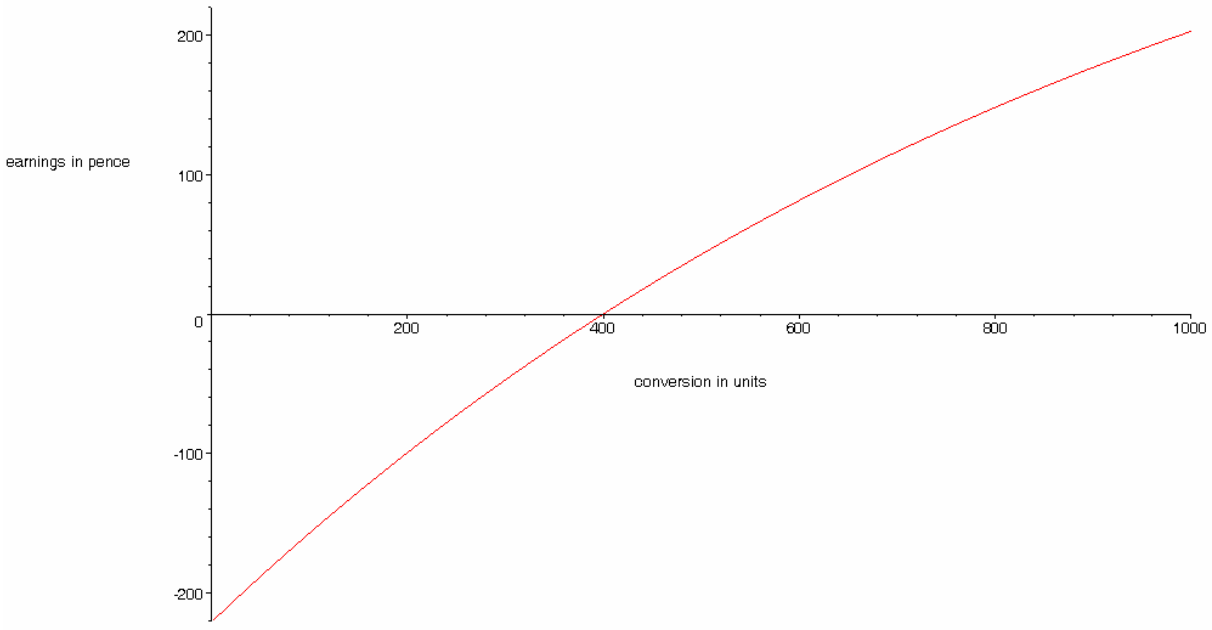

Figure 1: The Conversion Scale

5.1.7. **Before we implemented the main experiment, we carried out a pilot experiment with student participants**. In the light of the results from this pilot experiment a number of

modifications were adopted. Most of these involved simplifying the instructions and the presentation of the experiment. A further pilot was carried out with 30 participants selected from the target population. Further minor modifications were implemented as a consequence. Then the Main Experiment was carried out. This section concentrates on the Main Experiment.

5.2. The Main Experiment

5.2.1. **The Main Experiment involved a total of 174 active participants**. They were recruited by Ecotec, and constituted a quota sample from the target population – namely those members of the UK population believed to be under-saving and subsequently at risk of having insufficient private pension provision. Details are provided in Appendices 1 and 2. It should be noted that typically participants in experiments are normally recruited from the student population, so that the use of participants from the target population represents a major innovation in the implementation of an experiment. In principle it means that the results can be generalised to the target population. However, the use of these participants added further problems (which we had anticipated during the design of the experiment), not least the problem of ensuring that the participants understood the nature of the problem and what they were meant to be doing (if not how they should be doing it). As a consequence we put great efforts to ensure that the experimental task was explained in the clearest possible way. To this end, we proceeded as follows.

5.2.2. When the participants were gathered in the laboratory one of the experimenters made a brief statement, telling the participants about who was financing the experiment and what the participants should get out of it. We made a point of reminding the participants that they were guaranteed a £30 participation fee and that they could earn considerably more as a consequence of their decisions during the experiment. We emphasised that their objective should be maximising the amount of money that they took away from the experiment. The participants were then taken into the laboratory and seated at separate screened terminals, with the males on one side of the room and the females on the other. The reason for this is that the males and the females had different life tables. They were asked to read the written Instructions and then call over an experimenter when they had finished.

5.3. The Written Instructions

5.3.1. The Instructions for Treatment 3 are attached in Appendix 3. Despite considerable efforts to make these as simple as possible, the nature of the experiment itself means that the experiment is difficult for anyone. We note that, while many participants reported difficulties with fully understanding the written Instructions, most admitted that things became clearer once the experiment started. However, it should be borne in mind in interpreting the results, that some of the behaviour may have been caused by confusion in the minds of the participants, not necessarily in terms of what their objective was (making as much money as possible from participating in the experiment) but rather in how they should achieve this objective. **One could argue that this confusion reflects their real-life confusion concerning what they should do concerning their pensions**. In taking such decisions they not only need to decide how much to allocate to pensions but also how they

should allocate it. The experiment captures the essence (though not the detail) of these real-life decisions.

5.4. The PowerPoint Presentation

5.4.1. To help participants further with the Instructions we prepared a PowerPoint presentation which the participants viewed after reading the written Instructions. The PowerPoint presentation was played at a pre-determined speed. Both before and after the presentation, experimenters were available to answer any questions that the participants may have had. Rather obviously, the response of the participants to the Instructions and to the presentation varied from participant to participant: some said that they felt that they understood the Instructions well, some said that they had some doubts which they expected to be resolved once they started the experiment, and some expressed considerable doubts and confusion. Three of the participants were so confused by the Instructions that they decided not to participate in the experiment.

5.5. The Control Questions

5.5.1. **To satisfy ourselves that the participants truly understood the Instructions we asked the participants to respond to a set of control questions.** Those for Treatment 3 are attached in Appendix 4. We have kept a record of their answers to these control questions[11]. A summary is presented in Appendix 5. After the participant had answered the control questions, one of the experimenters went through, with the participant, all the questions answered incorrectly until the participant said that he or she was happy with the correct answer. At this stage, after answering any further questions that the participant had, he or she was allowed to start the experiment. During the experiment, the participants could ask clarification from an experimenter at any stage.

5.5.2. The PowerPoint presentation lasted about 10 minutes. The time for answering and discussing the Control Questions varied from participant to participant. Some participants answered very quickly and had all their answers correct. Others took longer, got answers wrong and we had to spend time explaining to them what the correct answers were. We did not monitor the time taken.

5.6. The Experimental Software

5.6.1. This was developed and written in Visual Basic. The core of the program as seen by the participants is the 'Main Decision Screen' – which for Treatment 3 can be seen in Appendix 3 (the Main Decision Screen for the other treatments was almost identical – just omitting features that were not present in these other treatments). Via this screen, participants could enter their decisions regarding their (proposed) allocation to pensions and the fund in which they wished to allocate it. In Treatments 2 and 3 they could also indicate their decision as to

[11] Preliminary analyses suggest that there is no relationship between the number of correct answers to the control questions and the performance of the subjects.

whether to stay in the savings scheme or opt out of it. The screen also told them how much real money that they earned up to that point in the experiment and how much they would earn that period from the amount of experimental units that they were considering converting into real money. We should reiterate that the income of the participants in the experiment was denominated in experimental units - which had no value in and of themselves. Experimental units became valuable only when they were converted into real money. The reason, as we have explained before, is that 'experimental units' in the experiment are the equivalent of 'consumption' in the economic theory while 'real money' in the experiment is the equivalent of the 'utility' of the theory.

5.6.2. Also the screen provided information about their likely future pension. As we have already discussed, this future pension is uncertain (due to the riskiness of the rates of interest) and therefore we had to find ways of presenting the information to the participants about the *distribution* of their pension in a form that was easy to understand. We did this in two ways: (a) by providing a histogram displaying visually the distribution of the pension; and (b) by giving them a table providing information on: (1) the lowest possible pension; (2) the 25% quartile: (3) the median; (4) the 75% quartile; and (5) the highest possible pension. We tried to ensure that the information was phrased in a way that would be understandable to the participants. Again as we have already noted, because their likely future pension not only depends on their past allocations but also on their future allocations, we gave this information (on the distribution of their pension) under two different scenarios, which we called 'Scenario 1' and 'Scenario 2', and which we have described earlier.

5.6.3. We should note that participants were allowed between 3 and 5 minutes to take decisions in each of the 15 periods prior to retirement - they could not move on to the next period until 3 minutes had elapsed and they had to move on after 5 minutes had elapsed. After retirement, when no further decisions were required, each period lasted just a few seconds. Overall, including the reading of the Instructions, the watching of the PowerPoint presentation, answering and correcting the Control Questions, and doing the experiment lasted between 2 and 3 hours.

5.7. The Questionnaire

5.7.1. After the participants completed the experiment, they were asked to fill in a brief questionnaire. This is reported in Appendix 6. This provides us with useful demographic information as well as information on the participants' attitudes to risk – though we should note that all this information is self-reported.

6. Possible Reference Points?

6.1. The Model-Optimal Strategy for a Risk-Neutral Participant

6.1.1. As we noted in section 4.2, experimental economists usually like to have a reference point for behaviour – against which to compare the actual behaviour of the subjects. If there is a uniquely 'best' way to respond to the experiment, this is taken as the reference point. In this experiment, however, there is no unique 'best' strategy, and we therefore have to make further assumptions to get some reference point.

6.1.2. The objective of the participants in this experiment was 'to make as much money as possible' from the experiment. However, their actual payment was not certain: it was inevitably risky. There are two sources of risk in the experiment – the interest rates and the length of life. Participants could influence to a certain extent the amount of risk from the first source (through their fund choices) but had no control over the length of their life (which was determined by the computer using the Government Actuary's Department Life Tables). Hence, whatever strategy participants adopted, their actual payment was not certain – it was risky. Thus, the idea of 'maximising their payment from the experiment' has to be more precisely defined. One way to do this is to interpret it as maximising their *expected* payment from the experiment. **This implicitly assumes that the participant is *risk-neutral* – that is, does not take into account any risk involved in the strategy.** Crucially, however, 'maximising the expected payment from the experiment' is a well-defined objective and we can find the strategy which achieves this objective. We do so in Appendix 7. We call this strategy *the risk-neutral model-optimal*[12] strategy in the experiment. It provides a reference point against which to compare the actual behaviour of the participants. This strategy is illustrated in Figure 2. This shows what a risk-neutral participant should do (in terms of pension allocations) in the various periods before retirement, in order to maximise their expected payment from participating in the experiment. In the experiment there were 15 such periods. It will be noted that this risk-neutral model optimal strategy is different for males and females (because females live longer than males) and is different in Treatment 1 compared to Treatments 2 and 3.

6.1.3. We note that participants in practice had to decide *three* things: (1) how much to allocate to pensions; (2) (in Treatments 2 and 3) whether to stay in or opt out of the savings scheme; (3) in which fund to place their pension allocations. It can be shown that it is never optimal to opt out of the scheme (basically because the optimal allocation is always more than the minimum 10% required). Moreover, a risk-neutral individual will always invest in the fund with the highest expected return. Hence this section considers only the *allocation decision* - for someone who is concerned only with maximising the expected earnings from the experiment.)

6.1.4. Some comments may be useful. It will be seen that this strategy involves allocating most to pensions early in the working life (a little under 400 of the income of 1000 for females and somewhat over 300 for males), and then decreasing the allocation through the working life until close to retirement. The generally downward sloping shape of this strategy results from the fact that early in the working life survival probabilities are high but they fall as retirement approaches. (The actual survival probabilities are in the Life Tables in the Instructions.) Of course, it is also a consequence of the fact that interest rates are positive. Money saved in early life earns interest at a compound rate: so a certain amount of money saved early leads to a bigger increase in the pension than the same amount of money saved later in life. As this point is important we should repeat it: for a risk-neutral participant, that

[12] 'Model-optimal' because it assumes that the objective of the participant is the same as in the model which we are using to describe behaviour.

is one who wants to maximise his or her expected earnings from the experiment, he or she should have a pension allocation profile which generally declines through time[13].

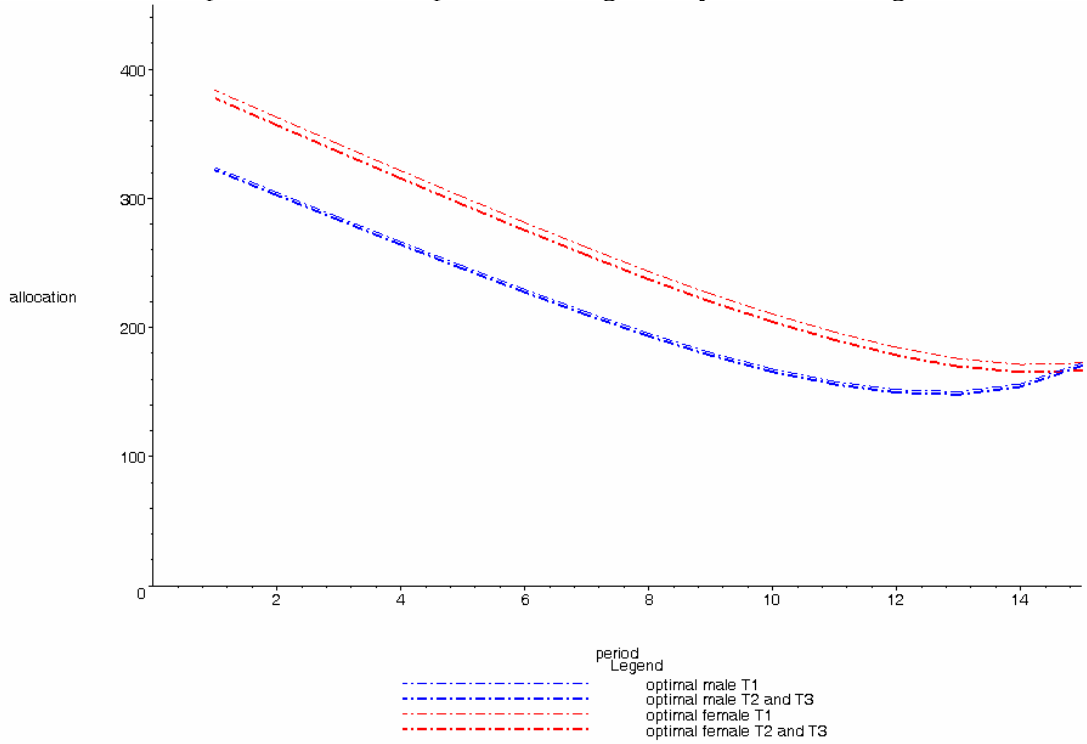

Figure 2: Risk-neutral model-optimal pension allocations

6.1.4. **We also note that the profiles for males are below the profiles for females; this is simply a consequence of the fact that females live longer.**

6.1.5. Finally, we note that the risk-neutral model-optimal strategy for both males and females is slightly lower in Treatments 2 and 3 than in Treatment 1. This is a consequence of the fact that, in Treatments 2 and 3, pension allocations are augmented by 75%.

6.1.6. This risk-neutral model-optimal strategy is one possible point of reference against which we could compare the actual behaviour of the participants. However, we should be careful in so doing. If a particular participant is not risk-neutral, then his or her model-optimal behaviour is not that illustrated in Figure 2. For example, if a participant is averse to risk (does not like risk) then he or she would prefer a strategy which may have a lower expected payment but implies less risk. The precise implications depend on the participant's actual aversion to risk and we cannot calculate the appropriate model-optimal strategy without knowing more about the participant's risk aversion. The difficulty is that 'risk-aversion' is not uni-dimensional and that attitude to risk may depend on the source of that risk, For example, it may be argued that a risk-averse participant would want to allocate less to pensions than a risk-neutral participant – because consumption is a safe activity while pension allocation is risky. However, a counter-argument might say that a risk-averse subject might want to allocate more to pensions so as to ensure that their future pension is sufficiently high. Indeed, the precise implications are not clear – in principle we would have to compute a model-optimal strategy for each participant – given knowledge of their attitude to risk.

[13] We should note that income in each period before retirement for all participants was 1000 experimental units. It was the same for all participants to remove individual differences and was the same in each period to simplify the decision problem of the participants.

Unfortunately we do not know their precise risk attitudes[14]. So we must simply conclude: if a particular participant is not risk-neutral then his or her model-optimal behaviour is *not* that illustrated in Figure 2. However, we might be tempted to argue that the *shape* of the profile should be similar – that is, generally downward sloping through time. This is a conclusion that depends on positive interest rates and high survival probabilities and appears to have nothing to do with the individual's attitude to risk. Indeed, it is difficult to see how one could argue that the profile should be flat or upward-sloping. However, we should stress that that this is a tentative conclusion.

6.1.7. We also note that this proposed reference point not only assumes that the participant is risk-neutral participant, but it also assumes that the participant's objective is that of maximising his or her (expected) payment from the experiment. If, however, the participant has some other objective, then this strategy is not the strategy that is optimal for them – *given their actual objective*. Of course, we as the experimenters do not know what the true objective of the participants in our experiment was. This is always the case in experiments: while the experimenter may assume that the participants want to earn as much money as they can from participating, this may well not be the case. Psychologists who do experiments are familiar with this and can provide many examples of different objectives: to please the experimenter; to finish the experiment as soon as possible; to do what they think that they ought to do; and so on. Economists argue that one should make the monetary incentive sufficiently high to outweigh all these other possible incentives – but one can never be sure that one has done so. The problem, however, is simple: if the participants have some other objective, then, unless the experimenter knows what that is, there is simply no reference point against which one can compare the behaviour of the participant.

6.2. The *Status quo*?

6.2.1. **A second possible reference point is what we might call the *status quo* of the target population.** The reason for the concern of the DWP for this target population is that they allocate little or nothing to their pension provision. While in no sense can this be considered optimal behaviour it does provide a useful point of reference. Accordingly, when we refer to the *status quo* strategy, we will be referring to participants who allocate nothing to pensions and hence end up in retirement with just a basic pension of 100 units.

7. Findings from the Experiment

7.1. Preamble

7.1.1. In this section, we provide some descriptive statistics of behaviour in the experiment; some more formal analysis of behaviour, including an analysis of demographic effects; a comparison of behaviour across treatments; a comparison of behaviour with the two 'reference points' introduced and defined in section 3. The material is divided by the

[14] We do have some information on risk-attitude *via* the questionnaire but this is not sufficient to identify accurately the risk-attitude of the participants.

variables of interest. We start with earnings; then we consider pension allocations; then the decision concerning opting out of the saving scheme (the stylised version of the personal accounts scheme); then opting out of the default fund; and finally the fund choice. The main items of interest should be kept in mind, namely: did participants maximise their earnings from the experiment?; how much did they allocate to pensions?; how close was their behaviour to the the risk-neutral model-optimal strategy?; how close was it to the *status quo* strategy?; in which funds did they make their allocations?; (in Treatments 2 and 3) did they opt out of the savings scheme?; were there any demographic effects?; was there evidence of inertia?

7.2. Earnings

7.2.1. Recall that participants were instructed to try and maximise their earnings from the experiment. Thus, achieved earnings are a crude measure of the extent to which they did this. Table 2 shows the mean and standard deviation of earnings over all the (relevant[15]) participants. These earnings do not include the participation fee of £30 which all participants received.

Table 2: Mean (standard deviation) of earnings from the experiment.

	Males	Females
Treatment 1	26.28 (9.89)	27.89 (10.77)
Treatment 2	31.95 (13.25)	34.41 (9.28)
Treatment 3	28.29 (13.62)	34.03 (9.72)

7.2.2. It can be seen that earnings were highest in Treatment 2 for both males and females; earnings were second highest in Treatment 2 and lowest in Treatment 1. However, it should be remembered that in Treatments 2 and 3 the participants could choose not to opt out of the savings scheme, and staying in meant that their pension allocations were increased by 75% by the experimenter. As a consequence, potential earnings were higher. Additionally, females could earn more than males because of their longer life expectancy. One way to standardise these figures is to compare them to earnings under our two reference points: the risk-neutral model-optimal strategy and the *status quo*. The maximal expected earnings under the first of these are: Treatment 1: £28.10 (males) and £28.93 (females; Treatments 2 and 3: £33.63 (males) and £35.88 (females). The earnings under the second of these are £14.93 (males) and £14.03 (females).The results of these comparisons are shown in Table 3.

Table 3: Earnings as a percentage of earnings in the two reference point strategies

	Males		Females	
Treatment	actual earnings as a percentage of those under the risk-neutral model optimal strategy	actual earnings as a percentage of those under the *status quo* strategy	actual earnings as a percentage of those under the risk-neutral model optimal strategy	actual earnings as a percentage of those under the *status quo* strategy
1	93.5%	176%	96.4%	199%
2	95.0%	214%	95.9%	245%

[15] Note that of the 207 participants who completed the experiment, 27 have been excluded from all the analyses because they did not fit with the recruitment profile given to Ecotec in the DWP guidelines. Details are given in Appendices 1 and 2. The analysis that we present is therefore based on the responses of 174 participants.

3	84.1%	190%	94.8%	242%

7.2.3. **This is an informative Table**. We note a number of things, most notably that generally participants did well compared to the risk-neutral model-optimal strategy. Recall that under this strategy expected earnings are maximised – and we see that in all except one case participants achieve more than 93.5% of this maximum. We also note that they do very well compared to the *status quo* – often achieving earnings of twice those under the *status quo*[16]. As to the treatment effects we note that males do best (relative to both reference points) in Treatment 2, but rather worse in Treatment 3 as compared with both Treatments 1 and 2. Females, in contrast, show modest decreases in earnings relative to the risk-neutral model-optimal strategy from Treatment 1 to Treatment 2 and again from Treatment 2 to Treatment 3, but relative to the *status quo* strategy earn most in Treatments 2 and 3. This latter result, of course, stems from the fact that the *status quo* strategy does not take advantage of the savings scheme (the personal accounts scheme). It is not clear why we observe these treatment effects. The experiment can really only shed light on what happens and not on why it happens. We would need a differently designed experiment to shed light on the latter.

7.2.4. It should, however, be noted that there is considerable variation in individual earnings. As is clear from Table 2, the standard deviations are large. A histogram of earnings is given in Figure 3: along the horizontal axis are the observed values of earnings (going from a small negative earning to almost £60) over and above the participation fee of £30; on the vertical axis is the frequency density of the observations. The histogram shows the distribution of actual earnings. Table 4 gives summary statistics. The variable *es* denotes earnings.

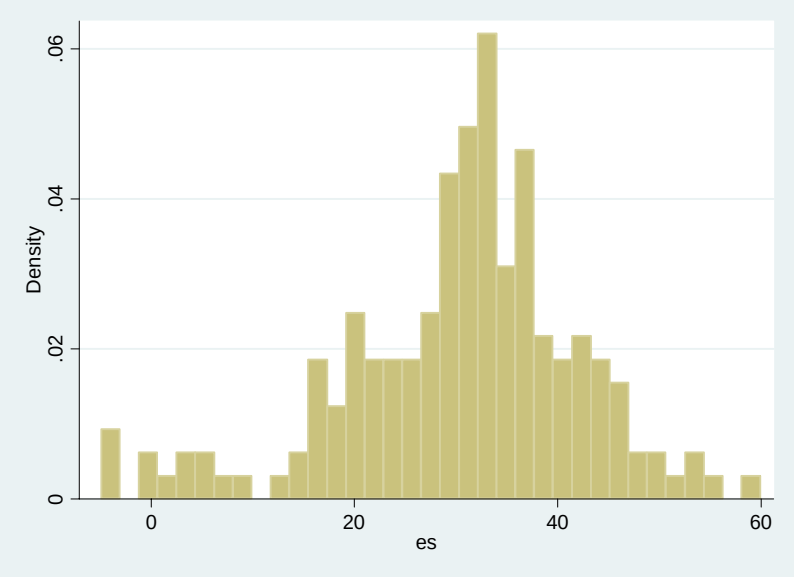

Figure 3: The distributions of earnings from the experiment

Table 4: Minimum/maximum earnings in the experiment

	Males		Females	
	Minimum	Maximum	Minimum	Maximum
Treatment 1	-0.07	42.87	-4.92	42.20
Treatment 2	-0.98	54.75	15.09	53.65
Treatment 3	-4.24	50.57	16.95	58.08

[16] We note that the *status quo* strategy - implying effectively no involvement with the pension decision - is almost bound to lead to very low earnings compared with any strategy of active involvement.

7.2.5. A comment on these numbers is relevant here, particularly in the light of the statements about the risk-neutral model-optimal earnings in the experiment. These earnings are the expected earnings following the risk-neutral model-optimal strategy as viewed from the beginning of the experiment. Actual earnings will deviate from this because of the randomness in the experiment. There are two sources of this: first, the randomness in the realised interest rates and secondly, randomness in the date of death. If we look at those participants who received very high earnings, they were either those who were lucky in their realised interest rates, or were those who built up higher pensions and survived quite far in the experiment. For example, those participants who had built up a pension of more than 2,000 experimental units would receive over £3.59 for every period that they survived in retirement. If they lived for the maximum 12 periods after retirement they would thus earn over £43.08 whilst in retirement – pushing up their earnings considerably. In Treatments 2 and 3 achieving a pension of over 2,000 was not difficult to the extent that the participants received a 75% augmentation from the experimenter[17]. On the other hand, those participants who received very low earnings either were those who were unlucky with interest rates or those who died very young. To achieve negative earnings they would have had to save on average more than 600 each period while alive and experimentally die before retirement. There were some who did[18].

7.2.6. The above descriptive statistics are confirmed by a more formal analysis – in which we try and explain earnings in terms of potential explanatory variables. These include treatment effects and demographics. To incorporate the former we introduce dummy variables *dt1, dt2* and *dt3*, which are such that *dti* takes the value 1 in Treatment *i* and the value 0 elsewhere. The demographic information is obtained from the questionnaire. The key variables are *gender* and *rac* – a risk aversion index calculated from the answers to questions 10, 11 and 12. It takes values between 0 and 1, with 0 indicating the lowest risk aversion and 1 the highest. Details of the construction of this index are given in Appendix 9.

7.2.7. We carried out a number of (ordinary least-squares) regressions to explain earnings. None of the demographics were statistically significant[19] except for gender which was on the borderline of statistical significance at 5%. The risk aversion coefficient was also not statistically significant. However the treatment dummies were statistically significant. We report the results of a regression explaining earnings in Table 5. In this and subsequent tables reporting the results of regression analyses, we indicate significance[20] at 1% with ** and significance at 5% with * attached to the t-statistic.

Table 5: The determinants of earnings

	Coefficient	standard error	t-statistic
dt2 (T1=0,T2=1,T3=0)	6.04	2.15	2.81**
dt3 (T1=0,T2=0,T3=1)	3.92	2.10	1.87
gender (male=1 female=2)	3.30	1.74	1.90
constant	22.16	2.98	7.43**

[17] The allocations necessary to achieve a pension of over 2,000 depend on when the allocations were made and on the rates of interest that were realised - as well as on the gender of the participant. To give an example, if the rate of interest averaged 2.5% then a male would have to allocate 302 of his income of 1,000 every period to achieve a pension of 2,000 (including the basic pension of 100).

[18] There were 7 participants who died before retirement who earned less than £5 from the experiment, and 5 participants who died before retirement who earned a negative amount (excluding the participation fee). We should note that we did not enforce these losses on these 5 participants and hence no-one left the experiment with less than the £30 participation fee.

[19] See Appendix 10.

[20] See Appendix 10.

7.2.8. These simply corroborate the results from Table 2 (though note that we have excluded any interaction terms involving sex and the two treatment dummies – they were not significant). **So earnings are higher in Treatment 2 than in Treatment 3 and in Treatment 3 than in Treatment 1. Gender has a moderate effect – implying increased earnings by £3.30 for females compared with males.** We also note that earnings generally are not too far below those implied by the risk-neutral model optimal strategy. This could be taken as an indication that participants were at least trying to maximise their earnings from their participation. This might be considered a possible indication of one aspect of the success of the experiment (though this latter is multidimensional).

7.3. Pension Allocations

7.3.1. We begin with some descriptions of the pension allocations[21]. Then we do some more formal analyses. The main items of interest in this section are: (1) what were the pension allocations of the participants? (2) how close were they to the risk-neutral model-optimal allocations? (3) how close were they to the *status quo* allocations? (4) how did they appear to be determined? (5) were there treatment effects? (6) were there demographic effects?

7.3.2. One of the more obvious features of our data is that there is a lot of noise in it. Some of the noise is *between* subjects and some *within* subjects. By the former we mean that different subjects behaved differently. By the latter we mean that, for some participants, pension allocations fluctuated considerably during the experiment. However, for some subjects their behaviour was smoother. Some idea of this noise can be gleaned from Figure 4 and Figure 5 below, in which we present scatter graphs (Figure 4 with males on the left and females on the right, and Figure 5 by treatment). Plotted on the vertical axis is the individual mean pension allocations and, on the horizontal axis, the standard deviation of their pension allocations. The graphs for males and females are similar but there are clear treatment effects. There is generally a tendency for higher means to be associated with higher standard deviations. To give some point of reference we present in Table 6 the means and standard deviations of pension allocations under the risk-neutral model-optimal strategy – though these figures obscure the fact that the risk-neutral model-optimal strategy is downward sloping through time.

Table 6: Means and standard deviations of allocations under the risk-neutral model-optimal strategy

	Males		Females	
	mean	standard deviation	mean	standard deviation
Treatment 1	213.5	57.2	255.6	70.4
Treatment 2	211.3	57.2	249.6	70.4
Treatment 3	211.3	57.2	249.6	70.4

[21] We note that the only constraint on pension allocations was that they could not be negative and had to be less than the experimental income of 1,000 units.

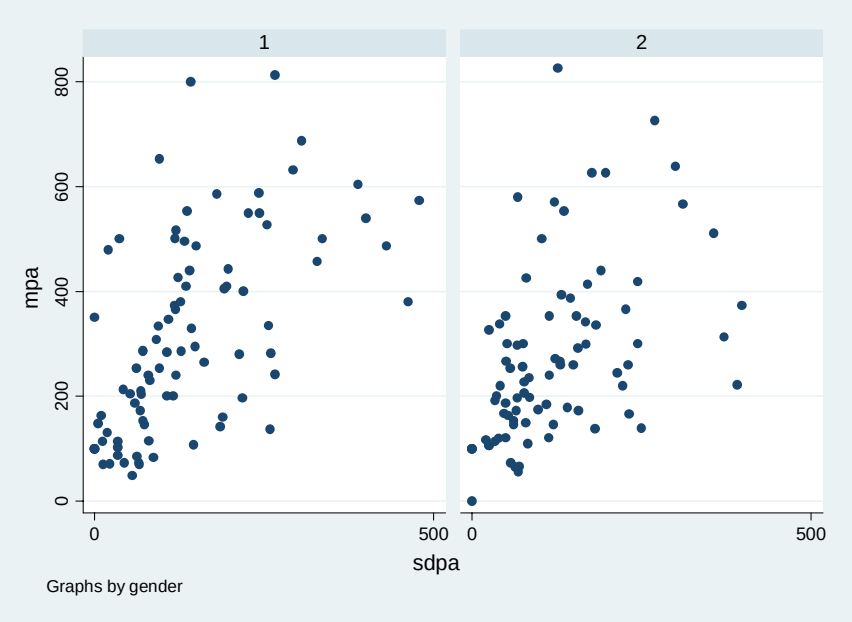

Figure 4: Mean and standard deviation of pension allocations by gender

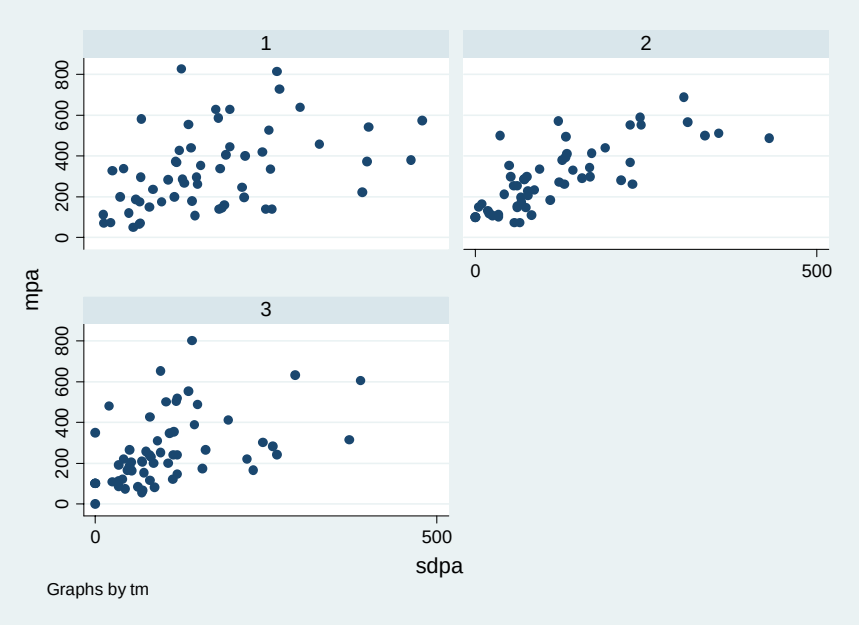

Figure 5: Mean and standard deviation of pension allocations by treatment

7.3.3. In addition to the considerable variability *between* participants, there is also considerable variation *within* participants – particularly in terms of their pension allocations through time. As can be seen from Figure 2 above, the risk-neutral model-optimal strategy involves a decline through time in the allocations (except in the final one or two periods). Indeed we have argued that this is necessarily the case even for participants who were not risk-neutral, though we should emphasise that we have not proved that this is so. There were some participants who realised that this declining allocation was a necessary feature of what appears to be a good strategy, but there were many who did not realise this. Some had a roughly constant strategy while some increased their allocations through time. Others oscillated wildly. Figure 6 shows an example of a participant who had a strategy with a profile similar to that of the risk-neutral model-optimal strategy. This shows the pension

allocations that subject 1101 made in the 15 pre-retirement periods in the experiment – starting off from around 600 in period 1 and declining to around 200 in period 15.

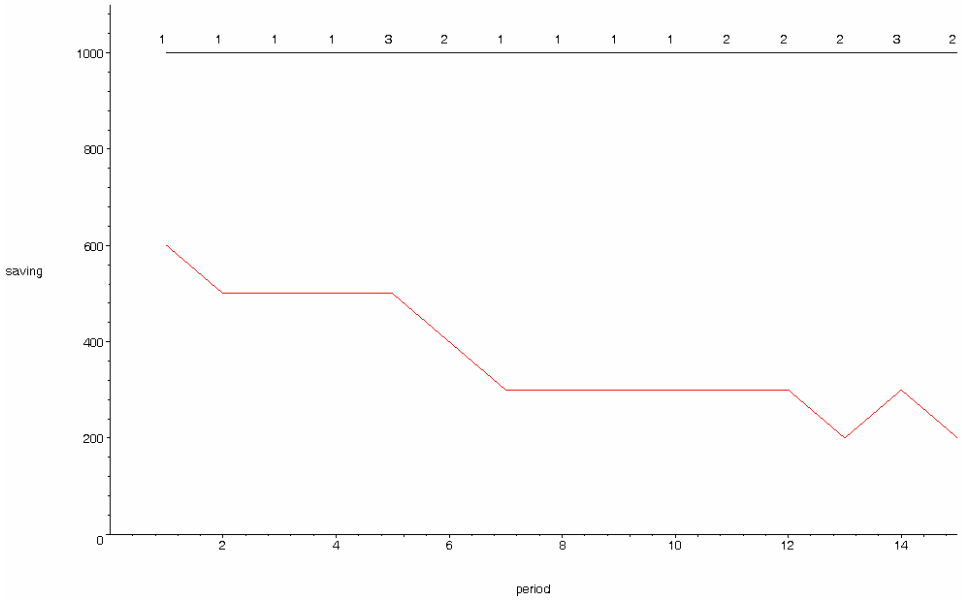

Figure 6: Treatment 1 Subject 1101

7.3.4. At the top of the graph is indicated the fund in which the savings were made. Although the savings are generally in excess of the risk-neutral model-optimal savings, the general shape is in accordance with that strategy.

7.3.5. We do not have the space to include graphs of the behaviour of all the participants in the experiment, but we think it important to give some idea of the types of behaviour exhibited by the participants. A preliminary classification is as follows:

(a) participants who broadly had a downward sloping profile
(b) participants who broadly had a constant profile
(c) participants who broadly had a upward sloping profile
(d) participants who broadly had an oscillating profile
(e) participants who had a wildly fluctuating profile and/or who appeared to be totally confused.

7.3.6. In Appendix 8, Figures 1 through 6, we give examples of each of these types. A formal separation of the participants into these various categories is, however, not easy - and the above list is given merely to indicate the range of different kinds of behaviour. It might be interesting, if any further analyses were to be conducted on this data, to try and separate the participants into different categories and examine separately the behaviour within each category.

7.3.7. Figure 7 (males) and Figure 8 (females) show the average pension allocations through time. These Figures also show the risk-neutral model-optimal allocations.

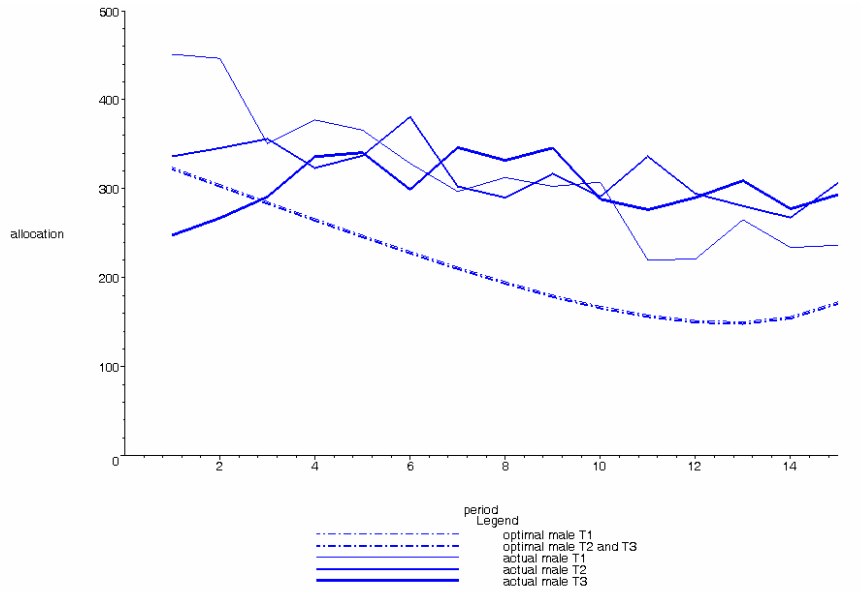

Figure 7: Pension allocations males

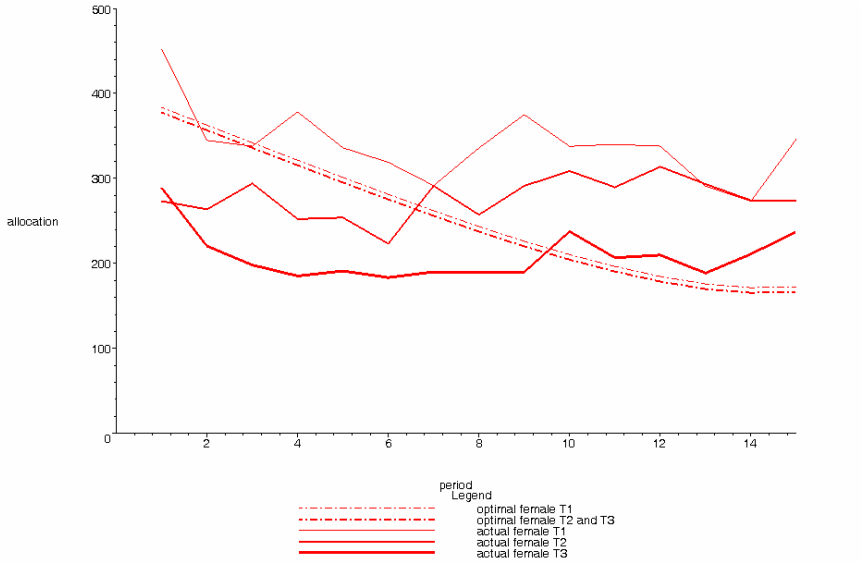

Figure 8: Pension allocations females

7.3.8. **Some of the noisiness in the data may be due to learning effects**. If this is true then we would expect that noisiness would be reduced in the later periods of the experiment relative to the earlier periods. To investigate this possibility we carry out two comparisons of allocations in the first and second halves of the pre-retirement period. First, in Table 7 we show the standard deviation of allocations in the first and second halves. Second, in Table 8 we show the *mean absolute deviations* of the actual allocations relative to the risk-neutral model-optimal ones and in the first and second halves. These tables shows some modest reduction in the noise relative to the risk-neutral model-optimal strategy in the later periods of the experiment, but there is no strong evidence of convergence - either to a constant strategy or to the risk-neutral model-optimal strategy.

Table 7: Standard deviations of actual allocations

	males		Females		

	t < 8	*t>7*	*t < 8*	*t>7*	
Treament 1	313	249	285	260	
Treatment 2	246	239	206	193	
Treatment 3	239	232	183	167	

Table 8: The mean absolute deviation of actual from risk-neutral model-optimal allocations

	males		Females	
	t < 8	*t>7*	*t < 8*	*t>7*
Treament 1	258	197	228	215
Treatment 2	194	191	179	168
Treatment 3	199	189	180	130

7.3.9. Apart from the noisiness in the data, there are a number of features that emerge from this descriptive analysis. **First, the males tend generally to allocate more to pensions than they would be doing if they were following the risk-neutral model-optimal strategy**. This could be due to all sorts of reasons: risk aversion (though interestingly the females, who are more risk-averse, do something different); myopia; learning and so on, though we can only speculate at this stage. We do note that in Treatment 1 there is some sign of a decline in contributions through time. **On the other hand, the females have a flatter strategy but are somewhat closer to the allocations that would be optimal if they were following the risk-neutral model-optimal strategy** (we have already seen this from Table 2 above). There is also clear evidence of a treatment effect on the pension allocation of the females though there is very little obvious effect in the males' allocations. This is shown more clearly in Table 9.

Table 9: Mean (standard deviation) of pension allocations

	Males	Females
Treatment 1	315 (286)	340 (273)
Treatment 2	318 (243)	275 (200)
Treatment 3	303 (235)	208 (174)

7.3.10. The mean male allocation is modestly highest in Treatment 2 and lowest in Treatment 3, while for females much less is allocated in Treatment 3 relative to Treatment 2 and much less in that relative to Treatment 1.

7.3.11. Table 9 refers to the means (and standard deviations) of the participants' allocations. In Treatments 2 and 3 these are augmented by the 75% added by the experimenter. From Table 5 we can calculate the means (and standard deviations) of the augmented allocations. These are given in Table 10.

Table 10: mean (standard deviation) of augmented pension allocations

	Males	Females
Treatment 1	315 (286)	340 (273)
Treatment 2	556 (425)	481 (350)
Treatment 3	530 (411)	364 (304)

7.3.12. **With the augmentation of the 75% from the experimenter, both male and female allocations are highest in Treatment 2; they are next highest in Treatment 3.** If the objective was to maximise (augmented) pension allocations then participants in Treatment 2 did better to achieve that objective than the participants in the other treatments. The lower

allocations in Treatment 3 might result from the feeling of security generated by the existence of a risk-free fund.

7.3.13. The implied pensions and the pensions under the risk-neutral model-optimal strategy are shown in Table 11.

Table 11: Mean actual pensions and pensions under the risk-neutral model-optimal strategy

	Males		Females	
	mean actual (% of model-optimal)	risk-neutral model-optimal	mean actual (% of model-optimal)	risk-neutral model-optimal
Treatment 1	1635 (186)	867	1467 (164)	896
Treatment 2	2139 (150)	1429	1745 (119)	1462
Treatment 3	1967 (139)	1429	1338 (92)	1462

7.3.14. It is clear that participants in all treatments had higher pensions than under the risk-neutral model-optimal strategy, except for females in Treatment 3. Indeed, generally, females were closer (on average) to those implied by the risk-neutral model-optimal strategy, while the ratio between the mean actual and that under the risk-neutral model-optimal fell from Treatment 1 to Treatment 2 and again to Treatment 3. All things being equal participants' pension allocations in Treatment 3 were closer to the pension allocations in the risk-neutral model-optimal strategy than the participants' allocations in Treatments 1 and 2.

7.3.15. Table 8 above summarises the total allocations to pensions made by the participants. We now explore the impact of the various demographic variables on allocations, by carrying out tobit[22] regressions of the total pension allocations on various possible explanatory variables. Table 12 was the result after following the usual procedure of eliminating statistically insignificant variables. **With the exception of income (in real life) and risk aversion (*rac*), none of the demographics were statistically significant.** The treatment effects were also not statistically significant (probably because of the noise in the data) but we include them in Table 12 to avoid distortions.

Table 12: The determinants of total pension allocations

variable	coefficient	standard error	t-statistic
dt2	-326.7	476.7	-0.69
dt3	-796.4	477.1	-1.67
income	315.6	103.5	3.05**
rac	-2783.5	1019.4	-2.73**
constant	5134.2	867.1	5.92**

7.3.16. **The coefficient on the risk aversion variable is interesting** – indicating that the more risk averse participants allocate significantly less to pensions than those less risk averse – but it could be explained by the fact that putting money into pensions is a risky activity, partly because of the risk on the rates of interest and partly because of risk concerning the date of death, while converting/consuming is a certain activity. Thus, the more risk-averse participants put relatively more experimental money into the safe activity (conversion/consumption) and relatively less money into the risky activity (pensions). The income variable is also of interest – suggesting that participants with a high income in real life allocate more than those with a low income. This may be because participants could not

[22] We note that the dependent variable is truncated between 0 and 15000 - hence making an ordinary least-squares regression inappropriate.

abstract from their personal characteristics. Of course, the level of participants' real life income should be irrelevant to their tackling of the problem posed to them in the experiment.

7.3.17. The effect of risk aversion, if not of income, also shows up in an tobit[23] regression analysis of pension allocations through time. Here the important variables are *t* (the period number), *rac, gender* and the treatment dummies, as Table 13 shows.

Table 13: Pension allocations through time

variable	coefficient	standard error	t-statistic
gender	-71.4	22.5	-3.17**
dt2	-111.8	27.2	-4.10**
dt3	-156.0	27.3	-5.70**
t	-21.3	4.27	-4.98**
*dt2*t*	9.84	3.03	3.24**
*dt3*t*	11.4	3.04	3.73**
*gender*t*	6.23	2.50	2.49*
rac	-210.7	28.1	-7.49**
cons	656.9	41.1	16.0**

7.3.18. Here we get a strong influence of risk aversion on savings: the most risk averse allocate on average over 210 less to savings than the least risk averse. This effect is not only strongly significantly statistically but it also has a strong economic significance – given the magnitude of the coefficient, moving from the least risk averse to the most risk-averse decreases allocations by over 210 each period. Set against the risk-neutral model-optimal values of pension allocations (see Figure 2), 210 is a large decrease. The Table also shows the gender effect (as is clear from Figure 4 and Figure 5): at the start, females allocate less but the decay in their contributions through time is also less than the males. We also note the strong treatment effects: in Treatment 2 allocations start out below those in Treatment 1 but decay less rapidly with time; in Treatment 3 allocations start out below those in Treatment 2 (and hence below those in Treatment 1) but they decay less rapidly than those in Treatment 2 (and hence less rapidly than those in Treatment 1). The fact that females are generally more risk-averse than males should be kept in mind when interpreting these effects.

7.3.19. The magnitude of the time coefficients should also be noted. Table 14 gives the details. In general allocations fall much less than under the risk-neutral model-optimal strategy. Moreover, the allocations made by females go more strongly against the latter strategy than do those for males.

Table 14: Actual and risk-neutral model-optimal average per period changes in allocations

	males		females	
	actual	risk-neutral model-optimal	actual	risk-neutral model-optimal
Treatment 1	-13.56	-10.8	-8.10	-15.0
Treatment 2	-4.16	-10.8	+1.20	-15.1
Treatment 3	-2.69	-10.8	+2.67	-15.1

7.3.20. **This all provides confirmation of the fact that participants allocate relatively little at the start of their lifetime and relatively a lot at the end. This could be referred to as evidence of the existence of a form of myopia**, but we should be careful about the use of

[23] This type of regression is appropriate when the dependent variable is truncated - in this case between 0 and 1,000.

this word. One could argue that a totally myopic person would not allocate anything to pensions. Clearly this is not the case for our participants. Indeed they were making substantial allocations to pensions – and much more than under the *status quo*. So we may be observing a form of myopia where participants did not appear to realise that making allocations early in life (early in the experiment) were more effective than making allocations later in life (later in the experiment). One may possibly conclude that participants did not understand the effect of compound interest rates.

7.3.21. If this misunderstanding of the effects of compound interest is true, then it seems to be the case that this lack of understanding is constant across treatments.

7.4. Opting Out of the Savings Scheme

7.4.1. **We now explore the determinants of the decision as to whether to stay in or opt out of the savings scheme.** This is one of the features of interest in this experiment. What this section is addressed to is whether, given that participants are automatically enrolled into the savings scheme, they then stay in the scheme. Note that we cannot check, given the experimental design, whether they are more likely to stay in with the automatic enrolment than without it, but we can still analyse the possible reasons for staying in and opting out.

7.4.2. Opting-out is relevant only to Treatments 2 and 3. In the risk-neutral model-optimal solution participants should never opt out – as the risk-neutral model-optimal pension allocation is always in excess of the 10% of salary required for staying in the scheme. In the experiment, only on very few occasions – just on 56 out of 840 occasions in Treatment 2 and just on 77 out of 900 occasions in Treatment 3 – did a participant opt out. Moreover, 82 out of the 116 participants in Treatments 2 and 3 *never* opted out, while 111 opted out at most in 5 periods. Note that it follows that the higher is the rate of opt out the lower will be the earnings in the experiment.

7.4.3. We begin with a simple descriptive table showing the numbers staying in the savings scheme and opting out of it, by period, treatment and gender. This is given in Table 15. There are some slight fluctuations, but, given the small number of observations in each cell, there seems to be nothing of statistical significance out of which inferences might be made.

Table 15: Opting out of savings scheme by treatment, period and gender

	Treatment 2						Treatment 3					
	males			females			males			females		
	out	in	all	out	in	all	out	in	all	out	in	all
1	2	24	26	0	29	29	1	32	33	1	26	27
2	2	24	26	0	29	29	3	30	33	1	26	27
3	3	23	26	1	28	29	2	31	33	4	23	27
4	1	25	26	2	27	29	3	30	33	3	24	27
5	4	22	26	2	27	29	3	30	33	2	25	27
6	1	25	26	4	25	29	4	29	33	4	23	27
7	1	25	26	2	27	29	2	31	33	4	23	27
8	3	23	26	1	28	29	1	32	33	5	22	27
9	2	24	26	0	29	29	0	33	33	5	22	27
10	3	23	26	0	29	29	2	31	33	2	24	26
11	2	24	26	1	27	28	0	33	33	2	24	26

12	4	22	26	1	26	27	2	29	31	4	22	26
13	3	23	26	1	26	27	3	27	30	4	21	25
14	2	21	23	2	25	27	3	26	29	2	23	25
15	6	17	23	0	27	27	3	25	28	2	22	24

7.4.4. Despite the fact that very few participants decided to opt out, it is possible, in principle, to carry out probit[24] analyses of the decision to opt out, though, in practice, nothing of any statistical significance seems to emerge. A process of elimination of the least statistically significant variable leads to the estimates in Table 16.

Table 16: The determinants of the decision to stay in (1) or opt out (0) of the savings scheme

Variable	coefficient	standard error	't-statistic'
dt3	-0.14	0.09	-1.55
t	-0.01	0.01	-1.4
constant	1.61	0.11	14.74**

7.4.5. Neither the treatment dummy *dt3* or the period variable *t* are significant, though there is slight evidence of more opting out in Treatment 3 and a slight increase in opting out through the periods of the experiment. These results are not surprising. **It might be inferred that those opting out were those who did not understand the implications of the saving scheme and hence their behaviour might be rather unpredictable.** Additionally one possible explanation of the slight fall in opting out as the experiment progressed is that participants were learning during the experiment.

7.4.6. Finally we note a consequence of the savings scheme on contributions: in Treatments 2 and 3 there is a noticeable spike at 10% of income. See Figure 9.

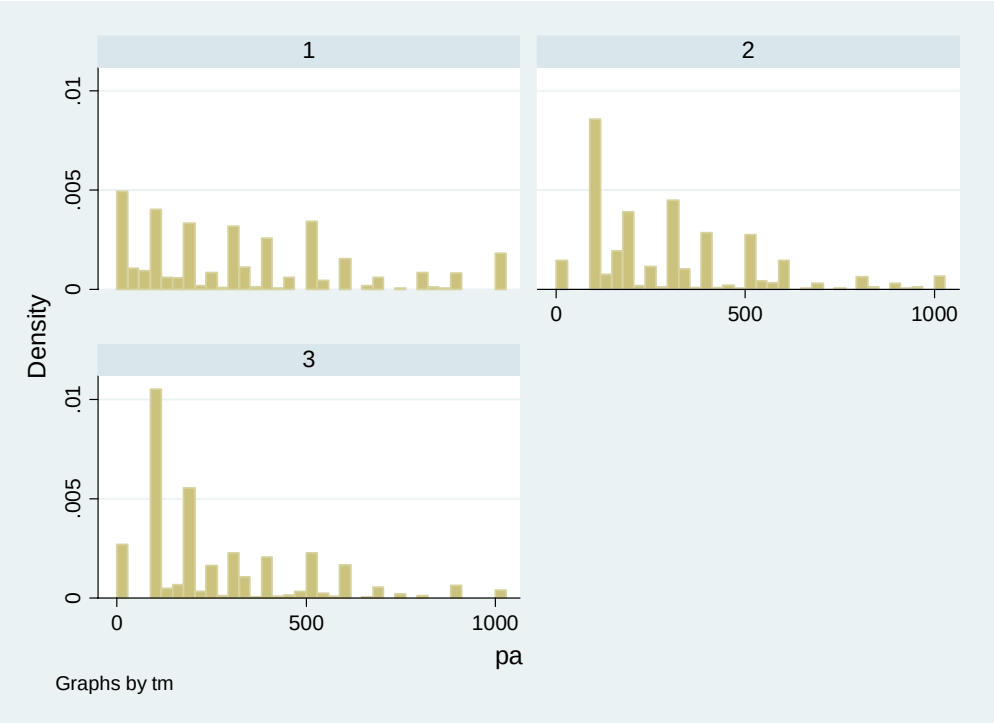

Figure 9: Histogram of pension allocations by treatments

[24] A probit analysis is appropriate when the dependent variable just takes the values 0 and 1.

7.5. Opting Out of the Default Fund

7.5.1. This section analyses the decision to opt out of the default fund, thus looking for signs of inertia in participants' behaviour.

7.5.2. **The default fund is Fund 2 or the fund in which the participant made the allocation in the preceding period.** If participants stay in the default fund (whether it is the default fund as provided by the experiment, or the default provided by their past behaviour) then this is some sign of inertia in behaviour. Of course, we cannot tell whether the decision to stay with the default fund **is** an indicator of inertia or an indicator that the participant is choosing what he or she thinks is the best fund; a change to the design of the experiment would be needed to distinguish between these two possibilities.

7.5.3. We analyse the determinants of the variable *idf* which takes the values 1 if the participant is in the default fund and the value 0 if the participant is in some other fund (including no fund). This variable is an indicator of inertia in the behaviour of the participant. We begin with a simple description of the data; this is presented in Table 17. If participants were changing funds at random, we would see entries of 33% in Treatments 1 and 2,and of 25% in Treatment 3.

Table 17: Percentage of time participants stay in the default fund

	males	females
Treatment 1	44%	46%
Treatment 2	60%	52%
Treatment 3	56%	47%

7.5.4. There is some modest evidence of inertia here - participants are more likely to stay in the default fund than choose at random amongst the available funds. There is some evidence here of more inertia for males in Treatment 3 than in Treatment 2 and more there than in Treatment 1. There is a similar pattern for females but the evidence is not as strong. If we do a more formal analysis, we get similar results. We carry out a time series analysis, using *idf* as the dependent variable. This just takes values 0 or 1, so, once again, a probit analysis is appropriate. We get the results in Table 18 after the usual process of eliminating the least significant variables[25].

Table 18: The determinants of the decision to stay in (1) or opt out (0) of the default fund

variable	coefficient	standard error	t-statistic
dt2	0.29	0.06	4.62**
dt3	0.18	0.06	3.01**
t	0.01	0.01	2.26*
gender	-0.13	0.05	-2.55*
constant	-0.27	0.10	-0.50

7.5.5. Here we find that participants are more likely to stay in the default fund in Treatment 2 as compared with in Treatment 1, and that they are also more likely to stay in the default fund

[25] Starting with all potential explanatory variables included and then successively eliminating the least significant (those with the smallest t-statistic) until just the most significant variables are left.

in Treatment 3 as compared with in Treatment 1, but here the effect is less strong. Males seem to show somewhat more inertia than females. Interestingly, we see here no effect of the other demographic variables. As time passes they are more likely to stay in - though the magnitude of the effect is small.

7.6. **The Fund Choice**

7.6.1. Here we present results on the fund choice. The questions of particular interest to us are: which funds do participants choose?; is there a time effect?; are there treatment effects (particularly in Treatment 3)?; are there demographic effects? We want to understand what determines the fund choice.

7.6.2. Figure 10 represents a histogram (with the frequency density on the vertical axis) of the fund choice by gender (all treatments together) while Figure 11**Error! Reference source not found.** represents a histogram (with the frequency density on the vertical axis) of the fund choice by treatment (both genders together). Note that 'Fund 0' refers to a situation in which the pension allocation is zero, and hence in which there is no fund choice to be made. Figure 10 shows some minor differences between males (the left hand graph) and females (the right hand graph): the modal fund for males is Fund 1 while that for females is Fund 2 – which is less risky. This seems to be consistent with the fact that females tend to be more risk-averse than men.

7.6.3. The differences between Treatments 1 and 2 seem to be minor (shown in Figure 11), while naturally the distribution in Treatment 3 is different – with the inclusion of Fund 4. Somewhat over 14% of the allocations in the latter treatment were in Fund 4, and its existence seems also to have changed the distribution over the other three funds - in that Fund 2 was the modal choice in Treatments 1 and 2 while Fund 1 is the modal choice in Treatment 3. The introduction of the risk-free fund has moved some people into it and others into Fund 1.

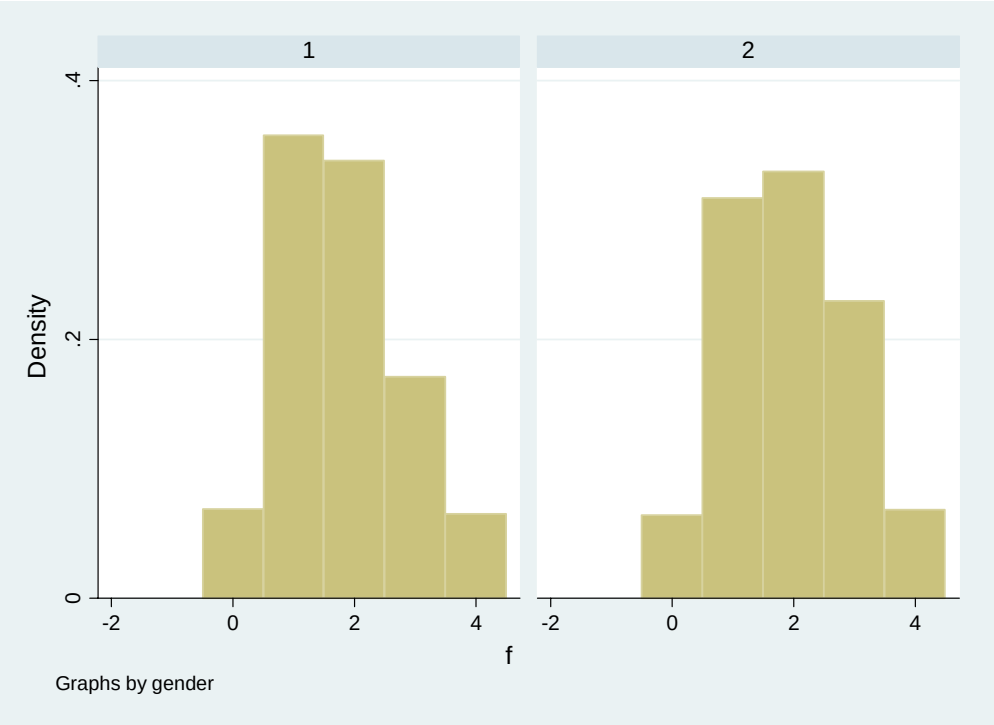

Figure 10: Histograms of the fund choice by gender

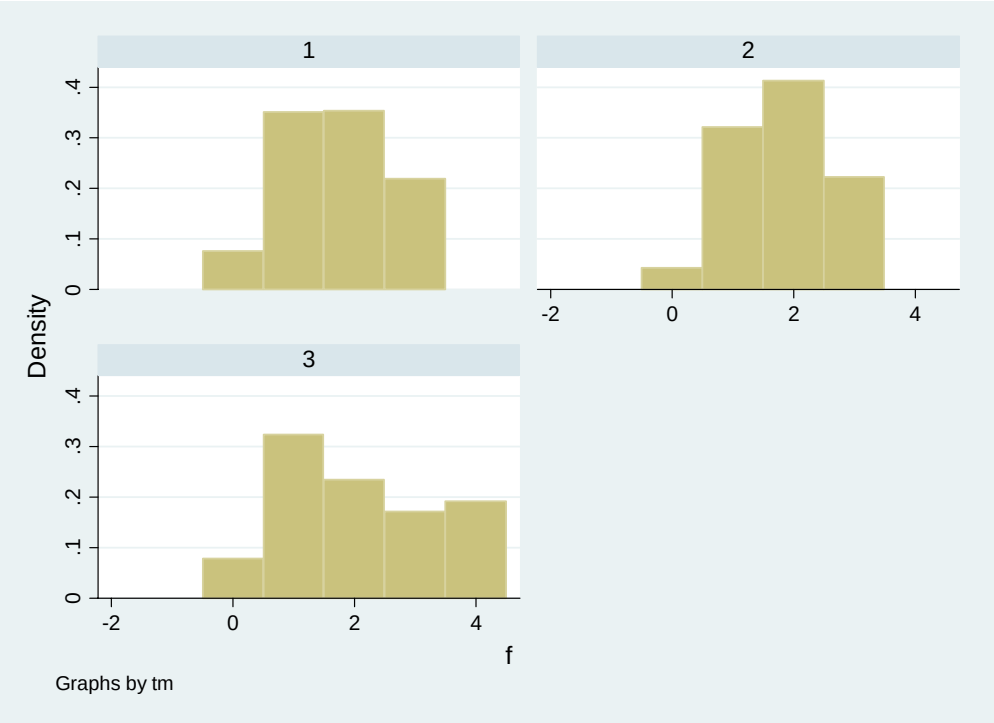

Figure 11: Histograms of the fund choice by treatment

7.6.3. A more formal analysis shows some interesting demographic effects. After following the usual procedure of eliminating the least statistically significant variables, we arrive at Table 19 (note that we have left *t* and *dt2* in the equation, despite their statistical insignificance, as they are economically of interest). The dependent variable in this table is the number indicating the fund choice.

Table 19: The determinants of fund choice

variable	coefficient	standard error	t-statistic
dt2	0.06	0.05	1.23
dt3	0.31	0.05	6.41**
t	0.00	0.00	0.69
age	0.09	0.02	5.10**
income	-0.04	0.01	-3.47**
rac	0.61	0.11	5.7
constant	1.23	0.10	12.18**

7.6.4. Treatment 3 induces a less risky fund choice (recall that going from Fund 1 through to Fund 4 the riskiness of the find is decreasing), and there is a modest (but statistically insignificant) increase in Treatment 2. **The passage of time in the experiment has no effect while *age* increases the fund choice – the older participants choose the safer funds.** Participants with higher income in real life choose riskier funds. Finally, the risk aversion variable has a statistically significantly positive effect: the more risk-averse, the less risky the fund chosen.

7.6.5. If we restrict the above analysis to Treatments 2 and 3 (in which there was the savings scheme) we get the results in Table 20.

Table 20: The determinants of fund choice in Treatments 2 and 3

variable	coefficient	standard error	t-statistic
dt3	0.25	0.05	4.94**
t	0.00	0.01	0.47
age	0.08	0.02	3.63**
income	-0.03	0.01	-2.30*
rac	0.76	0.14	5.31**
cons	1.20	0.13	9.03**

7.6.6. **The results are similar: time (in the experiment) is not significant; older people are more likely to choose the less risky funds; the people with higher income (in real life are more likely to choose the riskier funds; more risk-averse people are more likely to choose the safer funds; in Treatment 3 there is the obvious move into Fund 4.** Note that we are talking about the participants' real age and real income and not their age or income in the experiment. The fact that demographic variables such as age and income appear (as important explanatory variables in explaining behaviour) reinforces the point made earlier on several occasions: these variables are irrelevant to the solution of the problem in the laboratory and hence the fact that they are significant has two implications: first that these variables actually influence the ability of the participants to tackle the experiment; second, that these variables somehow activate personal heuristics. We cannot distinguish between these two possibilities given the information at our disposal.

7.6.7. An analysis using dummy variables reinforces the above analysis. If we define the variable *df1* as taking the value 1 if the allocation is made in Fund 1 and 0 otherwise, and we define *df2* and *df3* similarly, then we get the results in Table 21. (Note that in this table we are just reporting the estimated coefficients and whether they are significant (** at 1%. * at 5%). Note also that we are including in the reported equation the demographic variables *gender* and *rac* - because we found them important elsewhere).

Table 21: The determinants of individual fund choice

Dependent →				

variable Independent variables ↓	*df1*	*df2*	*dt3*	*df4*
dt2	-0.09	0.17**	0.01	-
dt3	-0.08	-0.35**	-0.14*	-
gender	-0.07	-0.36	0.11*	-0.03
age	-0.08**	0.07**	-0.01	0.27**
rac	-0.02	0.02	0.25**	-0.12
constant	-0.04	-0.54**	-1.06*	-1.62**

7.6.8. What is particularly interesting about this table, in addition to the obvious move out of Funds 2 and 3 in Treatment 3, is the strong influence of age on fund choice: older people are much more likely to choose Fund 4.

8. References

Ballinger, T.P., Palumbo, M.G., and Wilcox, N.T. (2003) 'Precautionary saving and social learning across generations: an experiment', *Economic Journal*, 113, 920-947

Carbone, E. (2005) 'Demographics and behaviour', *Experimental Economics*, 8, 217-232

Carbone, E. and Hey, J.D. (2004) 'The effect of unemployment on consumption: an experimental analysis', *Economic Journal*, 114, 660-683.

Cox, J. C. and Harrison, G. W. (2007), eds, *Risk Aversion in Experiments*, Greenwich, CT: JAI Press, Research in Experimental Economics, Volume 12.

Deaton, A. (1992), *Understanding Consumption*, Oxford University Press.

Dohmen, T., Falk, A., Huffman, D., et al. (2005) 'Individual risk attitudes: new evidence from a large, representative, experimentally validated survey', IZA Discussion Paper No. 511

Eichberger, J., Guth, W., and Müller, W. (2003) 'Attitudes towards risk: an experiment', *Metroeconomica,* 54, 89-124

Fellner, G. and Maciejovsky, B. (2002) 'Risk attitude and market behavior: evidence from experimental asset markets', Max Planck Institute for Research into Economic Systems, Papers on Strategic Interaction No. 34-2002.

Hey, J.D. and Dardanoni, V. (1988) 'Optimal consumption under uncertainty: an experimental investigation', *Economic Journal*, 98, 105-116.

Rohrmann, B. (1994) 'Risk perception of different societal groups: Australian findings and cross-national comparisons', *Australian Journal of Psychology*, 46, 150-163

Wärneryd, K.-E., (1996) 'Risk attitudes and risky behavior', *Journal of Economic Psychology,* 17, 749-770

Appendix 1: Sampling

A1.1: Methodology

This experiment used quota sampling. We used this method in order to select as broad a range of the target population as possible. Accordingly we used the following sampling variables: income, age, gender, and employment pattern (i.e., full-time or part-time). Appendix 2 gives our target and achieved quotas. The remaining variables, household composition, marital status, employment sector (e.g., public, private), and pension contribution, were monitored and considered in the analysis.

Participants were recruited by the consultancy Ecotec.

A1.2: The target group

The target population for this study is the target group for personal accounts. This group is discussed in the DWP White Paper (December 2006) 'Personal accounts: a new way to save'. Broadly the group consists of employees from 22 to state pension age (currently 65 for men and 60 for women) earning between £5,000 and £33,500 a year and not in a private pension scheme.

We did not specify pension provision in the recruitment since employees who already had a private pension might move into the personal accounts scheme, for example, by a change of employer. We did, however, specify that participants should not be self-employed: although it is intended to allow the self-employed to opt in to personal accounts, the design of the employer's contribution in the experiment is not appropriate for this group. Similarly we excluded from the sample the economically inactive, for example, those who are caring for children and not doing paid work.

Appendix 2: Target and Achieved Quotas

Some of those recruited had to be excluded from the analysis because they fell outside the target group. Three of those recruited attended an experiment session but withdrew from the experiment. The following table gives details of the participation rates.

Table A2.1 Participation rates

	N	**%**
Number sampled	204	
Ineligible (not population of interest)	27	
Opt-outs (attended session but declined to participate)	3	
In scope of experiment:	201	100
Ineligible cases		
Participant self-employed	17	
Participant outside earnings band	10	
	27	
In scope of analysis:	174	87
Participation rate		**99**

Table A2.2 summarises our specified and achieved sample. The variables were monitored using the questionnaire in Appendix 6.

Table A2.2 Target and achieved quotas

	Target	**Achieved**	**Ineligible**	**In scope of study**	**In scope of study %**
Gender					
Male	105	102	13	89	85
Female	105	98	14	84	80
Age					
22 to 34	70	83	9	74	106
35 to 44	70	46	4	42	60
45 to 64	70	70	13	57	81
Income (£ 000)					
5 to 20	108	114	9	105	97
20 to 33	102	72	8	64	63
Work pattern					
Part time	at least 54	52	15	37	-
Full time	at least 54	142	11	131	-

Note that subtotals may not add up to the sample size in the scope of the experiment because not every participant completed all parts of the questionnaire.
The final Table in this section shows how the participants were distributed across the three treatments.

Table A2.3 Distribution of participants across treatments

	Treatment 1		Treatment 2		Treatment 3	
	Achieved in sample	**In scope of study**	**Achieved in sample**	**In scope of study**	**Achieved in sample**	**In scope of study**
Gender						
Male	33	30	34	26	35	33
Female	33	28	31	29	34	27
Age						
22 to 34	27	27	30	25	26	22
35 to 44	13	11	12	12	21	19
45 to 64	26	20	22	18	22	19
Income (£ 000)						
5 to 20	37	34	39	37	38	34
20 to 33	23	21	24	18	25	25
Work pattern						
Part time	13	8	22	18	17	11
Full time	50	47	43	38	49	46

Appendix 3: Instructions for Treatment 3

Note: those for Treatments 1 and 2 are the same except that neither contain a reference to Fund 4 and Treatment 1 does not contain any reference to the personal accounts scheme.

EXEC Centre for Experimental Economics at the University of York

3

Welcome to this experiment. The Department for Work and Pensions (DWP) of the UK Government has provided the funds to finance this research. Depending on your decisions you may earn a considerable amount of money which will be paid to you in cash immediately after the end of the experiment. This sum will be in addition to the £30 that you have been promised.

There are no right or wrong ways to complete the experiment, but the decisions that you take will have implications for what you are paid at the end of the experiment. This depends partly on the decisions that you take during the experiment and partly on chance. So you will need to read these instructions carefully.

At the end of the experiment you will be asked to complete a brief questionnaire and to sign a receipt for the payment that you received, and to acknowledge that you participated voluntarily in the experiment. The results of the experiment will be used for the purpose of academic and governmental research and will be published and used in such a way that your anonymity will be preserved.

How these Instructions are organised

We begin with a brief overview and then we give you all the details that you need.

OVERVIEW

- You should try and maximise your payment from participating in this experiment.
- This you can do by the decisions that you take.
- The experiment lasts a random number of periods.
- In each period you get an income denominated in units of experimental money.
- You have to decide each period how much of this income you want to convert into real money.
- Your payment for participating (in addition to the £30 that you have been promised) is the sum of the amounts of real money that you have converted from experimental money over all the periods of the experiment.

DETAIL

Periods

The experiment simulates a life from the age of 22 to the age of death. This time is divided up into 3-year blocks, which we call periods. The experiment starts with you at age 22, in period 1, and continues until you experimentally die. At this point the experiment is over for you. The number of such periods is random, as is the length of life in practice. We will use life tables to simulate *your* experimental date of death; we will describe this process shortly.

Retirement

If you reach age 67, that is, if you reach period 16, you retire. Thus, periods 1 through 15 are before retirement and periods 16 onwards are after retirement.

Your income

In each period before retirement, you will receive an income of 1000 units of experimental money. After you retire you receive a pension, also denominated in units of experimental money. This pension will consist of a guaranteed basic pension of 100 units of experimental money per period, plus an additional pension that depends on any savings that you have made before retirement.

Your decisions

In each period before retirement, you have to decide how much of your income of 1000 units of experimental units you want to convert into real money and how much you want to save. If you decide to save some of your income, you will have a choice of different *Funds* in which you can put these savings. After retirement you have no further decisions to make – your pension is automatically converted from units of experimental money into real money.

Converting experimental money into real money

The Conversion Scale (that is, the rate at which experimental money is converted into real money) is shown in the Tables in the other document that you have in front of you. You might like to refer to this now. If the amount consumed in any individual period is below 400 you actually earn a *negative* amount – that is, you lose money. So, for example 100 units consumed become -157p in real money (that is, a loss of £1.57); 400 units become 0p; 700 units become 117p (that is, a gain of £1.17); and so on.

An example

Suppose that convert your entire income every period before retirement and thus you save nothing. Then your pension would just be the basic pension of 100 units per period. Suppose the experiment lasts for you a total of 23 periods (15 before retirement and 8 after retirement) then your conversions in these 23 periods would be:

period	1	2	3	4	5	6	7	8	9	10	11	12
income	**1000**	**1000**	**1000**	**1000**	**1000**	**1000**	**1000**	**1000**	**1000**	**1000**	**1000**	**1000**
period	13	14	15	16	17	18	19	20	21	22	23	
income	**1000**	**1000**	**1000**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	

and hence your corresponding earnings would be, all in pence:

period	1	2	3	4	5	6	7	8	9	10	11	12
earnings	**203**	**203**	**203**	**203**	**203**	**203**	**203**	**203**	**203**	**203**	**203**	**203**
period	13	14	15	16	17	18	19	20	21	22	23	
earnings	**203**	**203**	**203**	**-157**	**-157**	**-157**	**-157**	**-157**	**-157**	**-157**	**-157**	

Thus, your earnings from the experiment itself would be £17.89 (that is 203 pence for 15 periods minus 157 pence for 8 periods).You would be paid this in addition to the £30 that you have been

promised – implying a total payment of £47.89. Note carefully that it follows that what you are paid (over and above the promised £30) depends on the decisions that you take during the experiment.

The Number of Periods in the Experiment

The number of periods that the experiment lasts for you depends upon chance. We will use life tables to simulate a lifetime. The life table in the other document that you have in front of you shows, for men and women separately, the chances of surviving to the various ages in the table from the age of 22. You might like to refer to this now. The older a person gets the less likely they are to survive.

We have assumed that no-one survives to age 103 or more. So the maximum number of periods in retirement is 12. The number before retirement is 15. We will tell you throughout the experiment the chance of you surviving to future periods.

Surviving to the next Period in the Experiment

Survival to the next period will be determined by the computer using the life table.

Your Savings

In each period before retirement you will be asked to decide how much of your income of 1000 units of experimental money you want to convert into real money and how much you want to save. You will also be asked if you want to stay in or opt out of a Saving Scheme. This will be described below. Once allocated to pensions, you will not be able to access any units allocated until you retire. If you decide to save some of your income, then you will also be asked in which of three different *Funds* you want to make the saving. The default fund is either Fund 2 or the Fund in which you put your savings in the preceding period. If you do not want to make a decision about the Fund your saving will be put into the default fund. These Funds differ in the *rates of interest* on money saved in them. We shall describe these differences below.

The Saving Scheme

Unless you choose otherwise you will be opted into a Savings Scheme. This obliges you to save a minimum of 10% of your income. If you stay in the Savings Scheme then any savings that you make will be increased by 75% with a contribution paid by the experimenter. You can, however, choose to opt out of the Saving Scheme, and make lower savings. If you do opt out, your savings are no longer increased by 75%. You will be asked each period if you want to stay in the Saving Scheme or opt out of it. Membership of the Saving Scheme makes no difference to the rate of interest you get on funds invested.

The Rate of Interest

This specifies how money saved grows in value between periods. If, for example, the rate of interest is 2%, then the value of savings grows by 2% between periods: 100 this period becomes 102 next period, 200 this period becomes 204 next period, and so on; if the rate of interest is 4%, then the value of savings grows by 4% between periods: 100 this period becomes 104 next period, 200 this period becomes 208 next period, and so on.

The Different Funds

If you save some of your income of 1000 units of experimental money, then these savings will earn interest through time. You will be asked to state into which of four Funds you want to put these savings. Three of these Funds are risky – in the sense that the rate of interest on them is not certain – and they differ in terms of their riskiness. Fund 1 is the most risky, Fund 2 the next, and Fund 3 the least risky. Each of them has two possible rates of interest – each equally likely – with possible values are as given below. In contrast, Fund 4 is certain and has a known rate of interest. Each of the risky Funds has two possible rates of interest – each equally likely – with possible values as follows; in contrast the rate of interest on Fund 4 is certain:

Fund 1: either 1.0% or 5.2% (average 3.1%)
Fund 2: either 1.5% or 4.3% (average 2.9%)
Fund 3: either 2.0% or 3.4% (average 2.7%)
Fund 4: certain to be 2.2%

So with Fund 1, units saved in the Fund either grow in value by 1.0% or grow by 5.2% in a period. With Fund 2, units saved in the Fund either grow in value by 1.5% or grow by 4.3% in a period. With Fund 3, units saved in the Fund grow in value by either 2.0% or grow by 3.4% in a period. With Fund 4, units saved in the Fund are certain to grow by 2.2%. So, if you save 100 units in Fund 1, they either grow in value to 101 or grow to 105.2 by the following period; if you save 100 units in Fund 2, they either grow to 101.5 or grow to 104.3 by the following period; if you save 100 units in Fund 3, they either grow to 102.0 or grow to 103.4 by the following period. If you save 100 units in Fund 4, they are certain to grow to 102.2 by the following period. The rate of interest on any one Fund in any one period is independent of its value in any other period and of its value in any other Fund. So whether a fund grows in value by a higher or a lower amount is a matter of chance, and it is *not* affected by the fund's growth in the previous period, or by growth in the other funds at any time.

It is important to note that you will not be able to change your choice of fund for past savings decisions. For example, if you decide to save 100 units in period 1 in Fund 2, those 100 units saved in period 1 will remain invested in Fund 2 in periods 2, 3, 4, and so on, until retirement. If you decide to save 200 units in period 2 in Fund 3, those 200 units saved in period 2 will remain invested in Fund 3 in periods 3, 4, 5, and so on, until retirement.

Your Pension

When you retire at the age of 67, you will be given a basic guaranteed pension of 100 units of experimental money per period, to which will be added an extra pension determined by your savings. The savings that you have made before retirement, plus the accumulated interest, will determine the amount of your pension over and above the basic pension of 100 units per period. Every period in retirement you will get this pension, and, as we have noted above, it will be automatically converted into real money. This will form part of your earnings from this experiment – if you survive past the age of 67.

As the formula used to convert your savings during your working life into the pension you receive each period during retirement is rather complicated, we will simplify things for you by providing, every period before retirement, a prediction of the likely pension that you will get. Because the rate of interest on savings is not certain the prediction will not be precise. However, we can tell you the *distribution* of your likely pension, and, in particular, the lowest and highest pension that you might receive.

We do this for two different scenarios, Scenario 1 and Scenario 2.

Scenario 1: we tell you the distribution of your future pension given the savings that you have already made and that which you are considering making for this period, under the assumption that you will make no further savings in the future.

Scenario 2: we tell you the distribution of your future pension given the savings that you have already made and that which you are considering making for this period, *under the assumption that you will make exactly the same saving in exactly the same Fund in all the remaining periods until retirement*.

We will provide this information in the form of a graph (showing the distribution) and in the form of a table. Both the table and the graph contain the same information, though they show the information in different formats. An example is shown in the other document that you have in front of you which contains a screen shot of the main decision screen in the experiment. You might like to refer to this now. At the right of the screen can be seen the two Scenarios. In Scenario 1 (in this case when you save 100 in Fund 2 only in this period), the pension is certainly not going to be less than 143, and is certain to be less than 164. In Scenario 2 (in this case when you save 100 in Fund 2 in *all* periods before retirement), the pension is certainly not going to be less than 680 and is certain to be less than 831. The graphs show the same information in a different format.

On retirement, you will be told the precise value of your pension. You will receive this amount every period on and after retirement until you experimentally die. This will be automatically converted into real money each period using the Conversion Scale that we have already described.

Summary

In essence, the experiment is simple. Each period from the age of 22 until retirement you receive an income of 1000 units of experimental money. You have to decide whether to stay in or opt out of the Saving Scheme, and whether you want to save some of your income, and, if so, in which Fund you want to save it. If you stay in the Saving Scheme any savings that you make will be increased by 75% by a contribution from the experimenter. Any units that you do not save will be converted into real money, which will contribute to your payment for taking part in this experiment. Every period in retirement, assuming you have survived to retirement, you will receive, until you experimentally die, a pension which will be automatically converted into real money, again part of your payment for participating in this experiment.

More Instructions

As your payment for participating in this experiment depends upon the decisions that you take, it is vitally important that you fully understand these Instructions. Therefore, after you have read these Instructions, we will show you a PowerPoint Presentation which repeats these Instructions and gives you more detail. After this Presentation, if you have any questions, there will be an opportunity for you to ask these questions to one of the Experimenters.

Timing

As we want you to think carefully about the various decisions, we will give you a predetermined time to take the decisions in each period; you will see that you cannot speed the process up. Each period before retirement, we have programmed the software to allow between three and five minutes for the decision-making; thus the 15 periods before retirement will take up to one hour and 15 minutes to complete. The periods after retirement, involving no further decisions, will be

completed relatively rapidly. The PowerPoint presentation, which plays at a predetermined speed, is programmed to take 10 minutes – so expect the experiment to last around one hour and 30 minutes from the time that the PowerPoint presentation starts.

Control Questions

As we want to satisfy ourselves that you have understood these Instructions, we will ask you to respond to some *Control Questions.* When you have correctly responded to these Control Questions you will be free to start the experiment.

In Conclusion

At the end of the experiment, when you experimentally die, we will ask you to complete a short questionnaire and to sign a receipt for the payment which you have earned. We will then pay you this amount in cash and you will be free to leave.

*We thank you on behalf of **EXEC** and the **DWP***
for your participation in this experiment

Life Table

Age	Period	Men		Women	
		Number surviving	% Chance of surviving from age 22	Number surviving	%Chance of surviving from age 22
22	1	100000	100.00	100000	100.00
25	2	99760	99.76	99912	99.91
28	3	99508	99.51	99810	99.81
31	4	99235	99.23	99689	99.69
34	5	98924	98.92	99541	99.54
37	6	98573	98.57	99353	99.35
40	7	98155	98.16	99106	99.11
43	8	97639	97.64	98794	98.79
46	9	96972	96.97	98357	98.36
49	10	96097	96.10	97759	97.76
52	11	94945	94.95	97001	97.00
55	12	93483	93.48	96041	96.04
58	13	91609	91.61	94807	94.81
61	14	89096	89.10	93174	93.17
64	15	85753	85.75	91055	91.05
67	16	81537	81.54	88297	88.30
70	17	76178	76.18	84694	84.69
73	18	69453	69.45	79987	79.99
76	19	61086	61.09	73676	73.68
79	20	51219	51.22	65673	65.67
82	21	40256	40.26	55956	55.96
85	22	29398	29.40	44781	44.78
88	23	19014	19.01	32401	32.40
91	24	10666	10.67	20586	20.59
94	25	5044	5.04	11021	11.02
97	26	1913	1.91	4781	4.78
100	27	571	0.57	1645	1.64
103	28	0	0.00	0	0.00

Conversion Scale

Please note that a negative amount (for example, -170) means that you lose money (in this example 170 pence) while a positive amount (for example, 43) means that you earn money (in this example 43 pence). Please note that the computer will tell you at any stage how much you will lose or earn for any conversion that you are contemplating.

Amounts between 0 and 100 experimental units

Units	0	10	20	30	40	50	60	70	80	90	100
Pence	-221	-215	-208	-201	-195	-189	-182	-176	-170	-164	-157

Amounts between 0 and 1000 experimental units

Units	0	100	200	300	400	500	600	700	800	900	1000
Pence	-221	-157	-100	-47	0	43	82	117	148	177	203

Amounts between 0 and 10000 experimental units

Units	0	1000	2000	3000	4000	5000	6000	7000	8000	9000	10000
Pence	-221	203	359	417	438	445	448	449	450	450	450

You may choose to use or not use the graph below which presents the conversion scale in a different format.

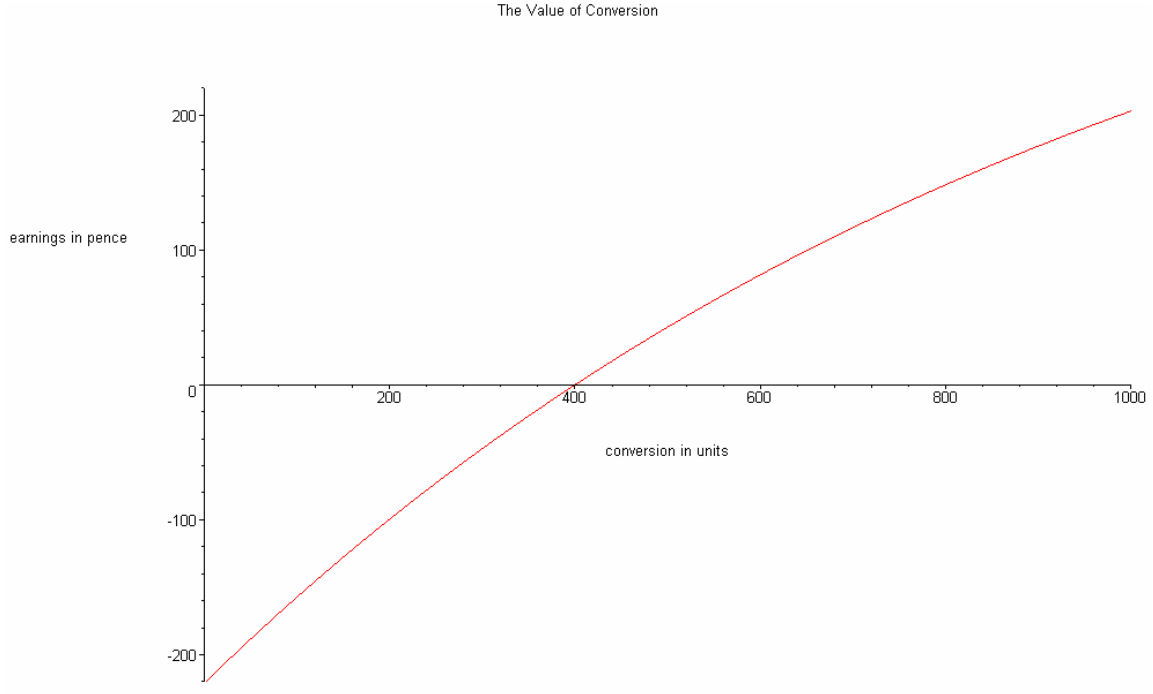

The Main Decision Screen in the Experiment

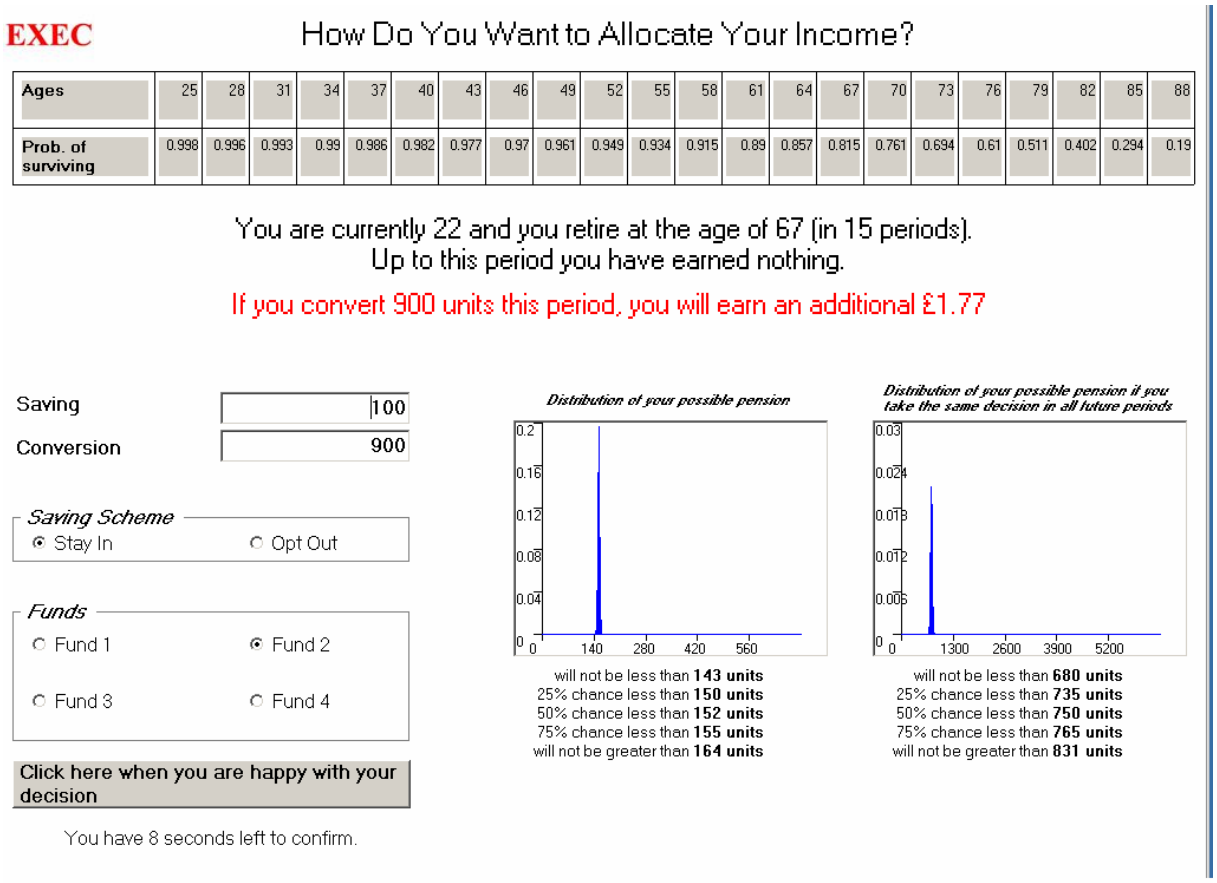

Appendix 4: Control Questions for Treatment 3

Note: those Treatments 1 and 2 are the same except that neither contain a reference to Fund 4 and Treatment 1 does not contain any reference to the personal accounts scheme.

EXEC Centre for Experimental Economics at the University of York

Control Questions

In order to satisfy ourselves that you have fully understood these Instructions, we would like you to answer some simple questions. Just tick the answer that you think is correct. Call over an experimenter when you have answered all the questions. When the experimenter agrees that you have answered these correctly, you may start the experiment.

1) If you convert 400 units in a particular period, how much real money have you earned for that period?
- You lose 221 pence
- You earn nothing
- You earn 203 pence

2) If you convert 0 (zero) units in a particular period, how much real money have you earned for that period?
- You lose 221 pence
- You earn nothing
- You earn 203 pence

3) If you convert 1000 units in a particular period, how much real money have you earned for that period?
- You lose 221 pence
- You earn nothing
- You earn 203 pence

4) If you do **NOT** opt out of the Saving Scheme, what is the minimum savings that you can make?
- 5% of your income
- 10% of your income
- 20% of your income

5) What happens to your saving if you do **NOT** opt out of the Saving Scheme?
- The experimenter adds 10 units to your saving
- Nothing – your saving remains unchanged
- The experimenter adds 75% of your saving to your saving

6) What happens to your saving if you **DO** opt out of the Saving Scheme?
- The experimenter adds 10 units to your saving
- Nothing – your saving remains unchanged
- The experimenter adds 75% of your saving to your saving

7) If you same some of your units, how many Funds can you chose between?
- 1
- 3

- 4

8) Which is the least risky Fund?
- Fund 1
- Fund 2
- Fund 4

9) If you DEFINITELY want a rate of interest of at least 2.2% which Fund should you chose?
- Fund 1
- Fund 3
- Fund 4

10) If you want the POSSIBILITY of a rate of interest of at least 5.0% which Fund should you chose?
- Fund 1
- Fund 2
- Fund 3

11) What happens to the rates of interest on the various Funds if you do **NOT** opt out of the Saving Scheme?
- Nothing
- The rates of interest are all doubled
- The rates of interest are all halved

12) What happens to the rates of interest on the various Funds if you **DO** opt out of the Saving Scheme?
- Nothing
- The rates of interest are all doubled
- The rates of interest are all halved

13) What is the basic guaranteed pension that you get whether or not you made any savings?
- 0 units
- 100 units
- 200 units

14) If your pension amounts to 400 units, how much real money do you earn in each period in retirement?
- You lose 221 pence
- You earn nothing
- You earn 203 pence

15) If your pension is the basic pension of 100 units, how much real money do you earn in each period in retirement?
- You lose 157 pence
- You earn nothing
- You earn 203 pence

Appendix 5: Summary of Responses to the Control Questions
(Note that different questions appear in the different treatments)

This Table reports the percentage of correct responses.

	Treatment 1	Treatment 2	Treatment 3
Control Question 1	91.4	92.9	93.3
Control Question 2	65.5	75.0	76.7
Control Question 3	94.8	89.3	93.3
Control Question 4	94.8	85.7	91.7
Control Question 5	84.5	91.1	85.0
Control Question 6	87.9	76.8	83.3
Control Question 7	91.4	100.0	96.7
Control Question 8	94.8	92.9	91.7
Control Question 9	77.6	78.6	93.3
Control Question 10	75.9	96.4	96.7
Control Question 11	no such question	87.5	88.3
Control Question 12	no such question	85.7	76.7
Control Question 13	no such question	98.2	98.3
Control Question 14	no such question	89.3	81.7
Control Question 15	no such question	73.2	70.0

Appendix 6: The Questionnaire

EXEC Centre for Experimental Economics at the University of York

Participant Number……

Questionnaire

We would be grateful if you could answer a few questions about yourself. Your answers will remain anonymous but will be useful to us when analysing the results of the experiment.

1. What is your gender?

[ ] Male
[ ] Female

2. Which age group do you belong to?

[ ] 16-24
[ ] 25-34
[ ] 35-44
[ ] 45-54
[ ] 55-64
[ ] 65+

3. What is your marital status?

-1- [ ] Never married
-2- [ ] Married / Civil Partnership
-3- [ ] Cohabiting
-4- [ ] Separated
-5- [ ] Divorced
-6- [ ] Widowed

4. Do any of the following live with you in your household?
(Please tick all that apply)

-7- [ ] Child under the age of 18
-8- [ ] Child over the age of 18
-9- [ ] Parent or grandparent
-10- [ ] Other relative
-11- [ ] Other non-relative (dependent or otherwise)
-12- [ ] None of the above

5. Which of these applies to you?

-13- [ ] Working full time (30 or more hours per week)
-14- [ ] Working part time (8 to 29 hours per week)
-15- [ ] Working part time (Less than 8 hours a week)
-16- [ ] Full time student

-17- [ ] Retired
-18- [ ] Unemployed

-19- [ ] Other not working

6. If you are in work, please tell us which type of organisation you currently work for. Otherwise please go to Question 7.

-20- [ ] Self-employed
-21- [ ] Private sector firm or company (e.g. limited companies, and PLCs)
-22- [ ] Nationalised industry or public corporation (e.g. post office, BBC)
-23- [ ] Other public sector employer (e.g. local government, Civil Service, LEA, NHS, Police, armed forces)
-24- [ ] Charity/voluntary sector (e.g. charities, churches, trade unions)
-25- [ ] Other

-26- [ ] Have never worked

7. What would you say best describes your current pension arrangements (pensions to which you are currently paying into)? Please tick all that apply.

-27- [ ] Employer pension scheme
-28- [ ] Personal pension scheme
-29- [ ] Do not currently belong to a pension scheme
-30- [ ] Don't know

9. What is your personal gross basic wage / salary? (gross wage / salary is before tax deductions)

-31-	[ ]	Less than £5,000 a year	(less than £100 a week)
-32-	[ ]	£5,000 to £9,999 a year	(£100 to £199 a week)
-33-	[ ]	£10,000 to £14,999 a year	(£200 to £299 a week)
-34-	[ ]	£15,000 to £19,999 a year	(£300 to £399 a week)
-35-	[ ]	£20,000 to £24,999 a year	(£400 to £499 a week)
-36-	[ ]	£25,000 to £29,999 a year	(£500 to £599 a week)
-37-	[ ]	£30,000 to £39,999 a year	(£600 to £799 a week)
-38-	[ ]	£40,000 to £49,999 a year	(£800 to £999 a week)
-39-	[ ]	£50,000 a year or more	(£1,000 a week or more)
-40-	[ ]	Don't know	

10. How do you see yourself:

Are you generally a person who is fully prepared to take risks or do you try to avoid taking risks?

Please tick a box on the scale, where the value 1 means: "not prepared to take risks" and the value 5 means: "fully prepared to take risks".

Not prepared to take risks				**Fully prepared to take risks**
1	2	3	4	5
[ ]	[ ]	[ ]	[ ]	[ ]

11. People can behave differently in different situations.

How would you rate your willingness to take risks in the following areas? Please tick a box in each line of the scale.

Not prepared to take risks				**Fully prepared to take risks**
while driving?				
1	2	3	4	5
[ ]	[ ]	[ ]	[ ]	[ ]
in financial matters?				
1	2	3	4	5
[ ]	[ ]	[ ]	[ ]	[ ]
during leisure and sport?				
1	2	3	4	5
[ ]	[ ]	[ ]	[ ]	[ ]
in your occupation?				
1	2	3	4	5
[ ]	[ ]	[ ]	[ ]	[ ]
with your health?				
1	2	3	4	5
[ ]	[ ]	[ ]	[ ]	[ ]
– **with your faith in other people?**				
1	2	3	4	5
[ ]	[ ]	[ ]	[ ]	[ ]

12. Please consider what you would do in the following situation:

Imagine that you had won £100,000 in the lottery. Almost immediately after you collect the winnings, you receive the following financial offer from a bank, the conditions of which are as follows:
There is the chance to double the money within two years.
It is equally possible that you could lose half of the amount invested.
You have the opportunity to invest the full amount, part of the amount or reject the offer.

What share of your lottery winnings would you be prepared to invest in this investment?

[ ] £100,000
[ ] £80,000
[ ] £60,000
[ ] £40,000
[ ] £20,000
[ ] Nothing, I would decline the offer.

Many thanks for your participation.
Before you leave we would like you to complete and sign this receipt.

EXEC The Centre for Experimental Economics at the University of York

---

This is to certify that I received a payment of £…….… for my participation in this experiment. I participated voluntarily. I understand that the results of this experiment will be used anonymously for the purposes of academic and governmental research.

Number……..…….

Name……...…….....

Signature……...……..

Date…….

Appendix 7: Analytical Solution of the Risk Free Case

Let us begin by denoting by k the treatment effect. In Treatment 1, $k=1$, and in Treatments 2 and 3, $k = 1.75$. The parameter k indicates by how much allocations are increased – in Treatment 1 they are not increased while in Treatments 2 and 3 they are increased by 75% (assuming that they are above 10% of income – which they always are in the optimal solution). The conversion scale is denoted by $u(.)$ – so that c converted becomes $u(c)$ in money (payment to the participant).

Denoting by S_t the saving in the t'th period, Y the income (always equal to 1000 in the experiment), p_t the probability of surviving from period t to period $t+1$ and by P the pension, the problem of the individual is to maximise the following expression:

$$u(Y-S_1)+\sum_{t=2}^{t=15}[\prod_{s=1}^{s=t-1}p_s]u(Y-S_t)+\sum_{t=16}^{t=27}[\prod_{s=1}^{s=t-1}p_s]u(P) \quad (1)$$

In this equation the first term is the amount of money earned in the first period; the second term is the expected earnings from the remaining periods before retirement and the final term the expected earnings from the post-retirement periods.

We know that P is such that:

$$F=(P-P_0)A$$
$$\text{where} \quad A=\{1+\sum_{t=16}^{27}[\prod_{s=16}^{t}p_s/R]\} \quad (2)$$

where R is the rate of return after retirement (equal to 1.025 in the experiment), F is the pension fund on retirement and P_0 the guaranteed basic pension (always equal to 100 in the experiment). This equation simply states that the pension is such that its expected discounted value (over and above the guaranteed pension) is equal to the pension fund, From equation (2) we have that:

$$P=F/A+P_0 \quad (3)$$

Of course, the pension F is given by the savings. We have:

$$F=k\sum_{t=1}^{15}S_t r^{16-t} \quad (4)$$

Here we assume that all allocations are greater than the minimum required by the saving scheme. In this equation r is the rate of return before retirement (which in the experiment varies according to the fund, and which we put equal to its Fund 4 rate (1.002) in this risk-free computation).
From the equation (1) we can find the first order conditions. We have simply for S_t the condition:

$$u'(Y-S_t)=k[\prod_{s=t}^{15}p_s]\frac{a}{A}r^{16-t}u'(P)$$
$$\text{where} \quad a=1+\sum_{t=16}^{27}\prod_{s=16}^{t}p_s \quad (5)$$

These equations are true for all $t = 1,\ldots,15$.
From equation (5) we get immediately that for any $t=1,\ldots,14$ we have:

$$u'(Y-S_t)=p_t r u'(Y-S_{t+1}) \quad (6)$$

Now if we take the utility function:[26]

$$u(x)=A-B\exp(-\alpha(x-y)) \quad (7)$$

where y is subsistence consumption, it follows that:

[26] This utility function is inique only up to a linear transformation, so that the coefficients A and B can be chosen arbitrarily. In the experiment *A and B both took the value 459 and* α the value 0.001.

$$u'(x) = B\alpha \exp(-\alpha(x-y)) \tag{8}$$

The term $B\alpha$ outside the exp clearly is going to cancel out in what follows.
We thus have:

$$\exp(-\alpha(Y - S_t - y)) = p_t r \exp(-\alpha(Y - S_{t+1} - y)) \tag{9}$$

From this it follows, taking logs, that:

$$-\alpha(Y - S_y - y) = \ln(p_t r) - \alpha(Y - S_{t+1} - y) \tag{10}$$

and hence:

$$S_t = S_{t+1} + \frac{\ln(p_t r)}{\alpha} \tag{11}$$

This gives us the differences between the savings in the various periods.

To simplify what follows let me denote by $\gamma_t = \frac{\ln(p_t r)}{\alpha}$ for *t = 1,..,14*. Then we have that

$S_t = \gamma_t + S_{t+1}$. Further, denoting by $\beta_t = \sum_{s=t}^{14} \gamma_s$ *(t = 1,.., 14)* and $\beta_{15} = 0$, it follows that:

$$S_t = \beta_t + S_{15} \quad \text{for } t \text{ from 1 to 15} \tag{12}$$

Equation (12) tells us the relative levels of the optimal savings in the various periods. It remains to find the absolute level.
If we take the first order condition for the optimal level of saving in period 15, S_{15}, we have:

$$u'(Y - S_{15}) = \frac{kp_{15}ar}{A} u'(P) \tag{13}$$

From this we get:

$$\begin{aligned} \exp(-\alpha(Y - S_{15} - y)) &= \frac{kp_{15}ar}{A} \exp(-\alpha(P - y)) \\ &= \frac{kp_{15}ar}{A} \exp(-\alpha(\frac{F}{A} + P_0 - y)) \end{aligned} \tag{14}$$

Hence, taking logs, we get:

$$-\alpha(Y - S_{15} - y) = -\alpha(\frac{F}{A} + P_0 - y) + \ln(\frac{kp_{15}ar}{A}) \tag{15}$$

From this, after a bit of rearranging, we get:

$$F + AS_{15} + AP_0 = AY + \frac{A}{\alpha}\ln(\frac{kp_{15}ar}{A}) \tag{16}$$

Now recall the equation for *F.* From this we have:

$$\begin{aligned} F = k\sum_{t=1}^{15} S_t r^{16-t} &= k\sum_{t=1}^{15} (\beta_t + S_{15}) r^{16-t} \\ &= k(\sum_{t=1}^{15} \beta_t r^{16-t} + S_{15}\sum_{t=1}^{15} r^{16-t}) \\ &= k(\beta_s + S_{15} r_s) \end{aligned} \tag{17}$$

using an obvious notation, namely $\beta_s = \sum_{t=1}^{15} \beta_t r^{16-t}$ *and* $r_s = \sum_{t=1}^{15} r^{16-t}$

Hence we have:

$$k\beta_s + kS_{15}r_s + AS_{15} + AP_0 = AY + \frac{A}{\alpha}\ln(\frac{kp_{15}ar}{A}) \tag{18}$$

and so:

$$S_{15}(kr_s + A) = AY + \frac{A}{\alpha}\ln(\frac{kp_{15}ar}{A}) - k\beta_s - AP_0 \tag{19}$$

from which we can obtain the *model-optimal* value for S_{15}.

In order to calculate the earnings under the *model-optimal* strategy we need to evaluate the following expression:

$$\max u = u(Y - S_1) + \sum_{t=2}^{t=15} [\prod_{s=1}^{s=t-1} p_s] u(Y - S_t) + \sum_{t=16}^{t=27} [\prod_{s=1}^{s=t-1} p_s] u(P) \quad (20)$$

This is used to calculate the expected earnings of a participant following the model-optimal strategy and is used as the point of reference in section 7.1.

Appendix 8. Examples of the Different Types of Behaviour

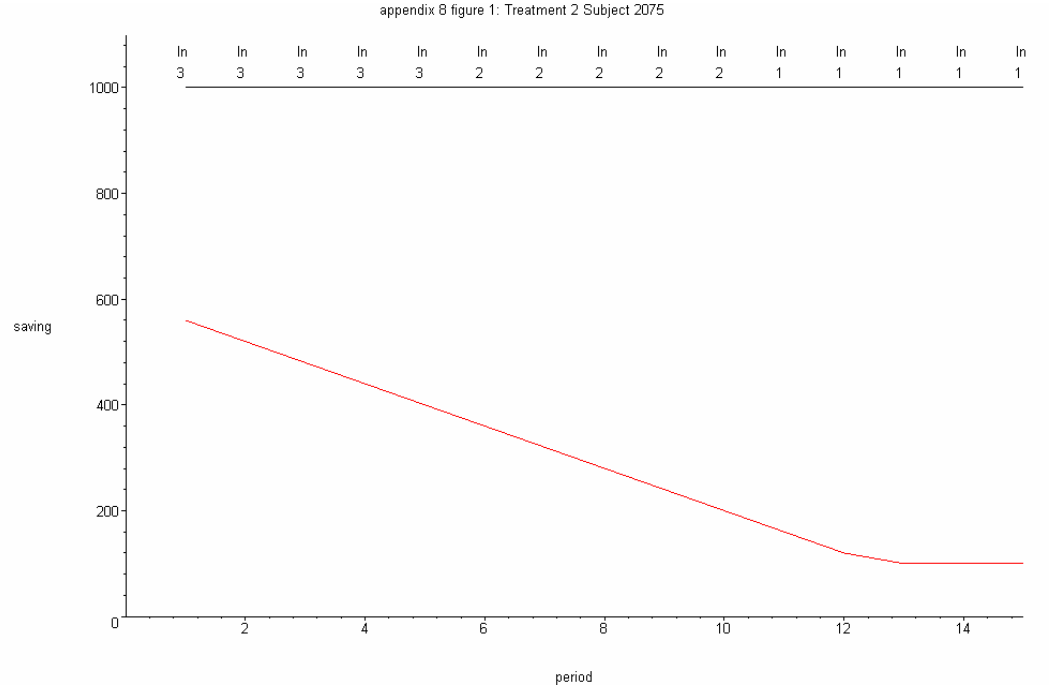

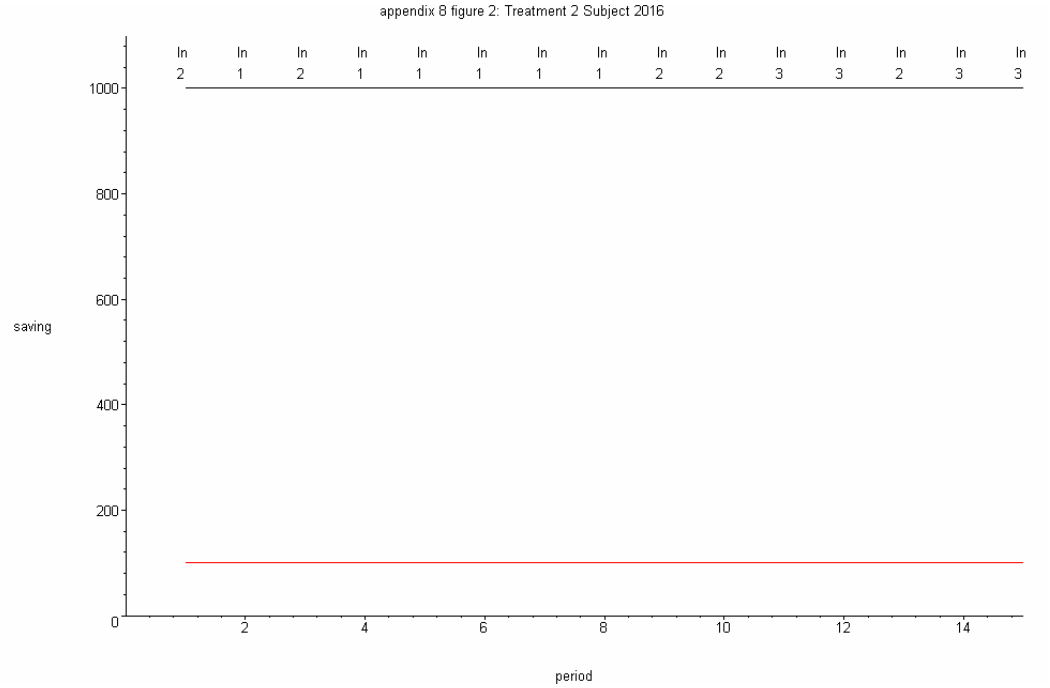

appendix 8 figure 3: Treatment 2 Subject 2065

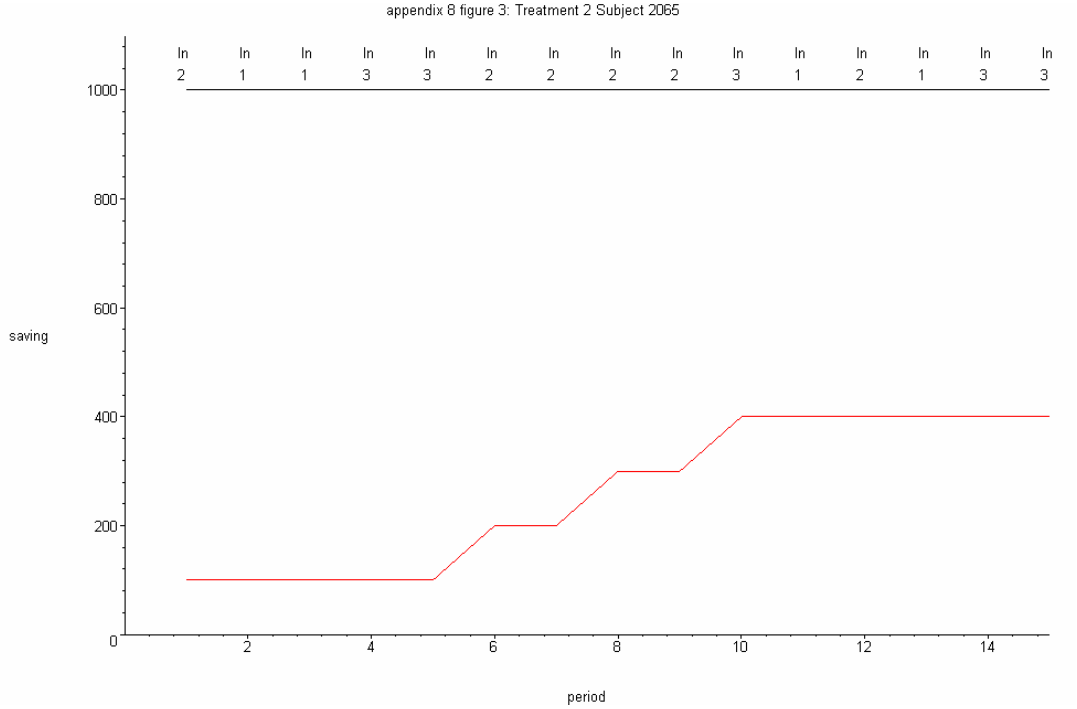

appendix 8 figure 4: Treatment 1 Subject 1108

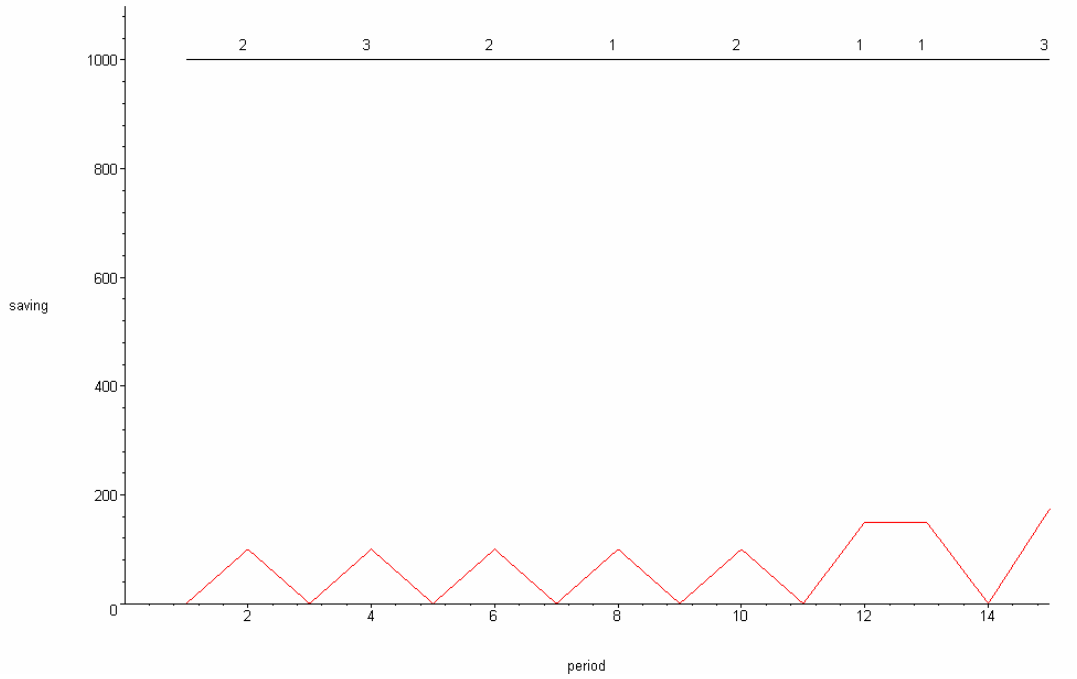

appendix 8 figure 5: Treatment 1 Subject 1140

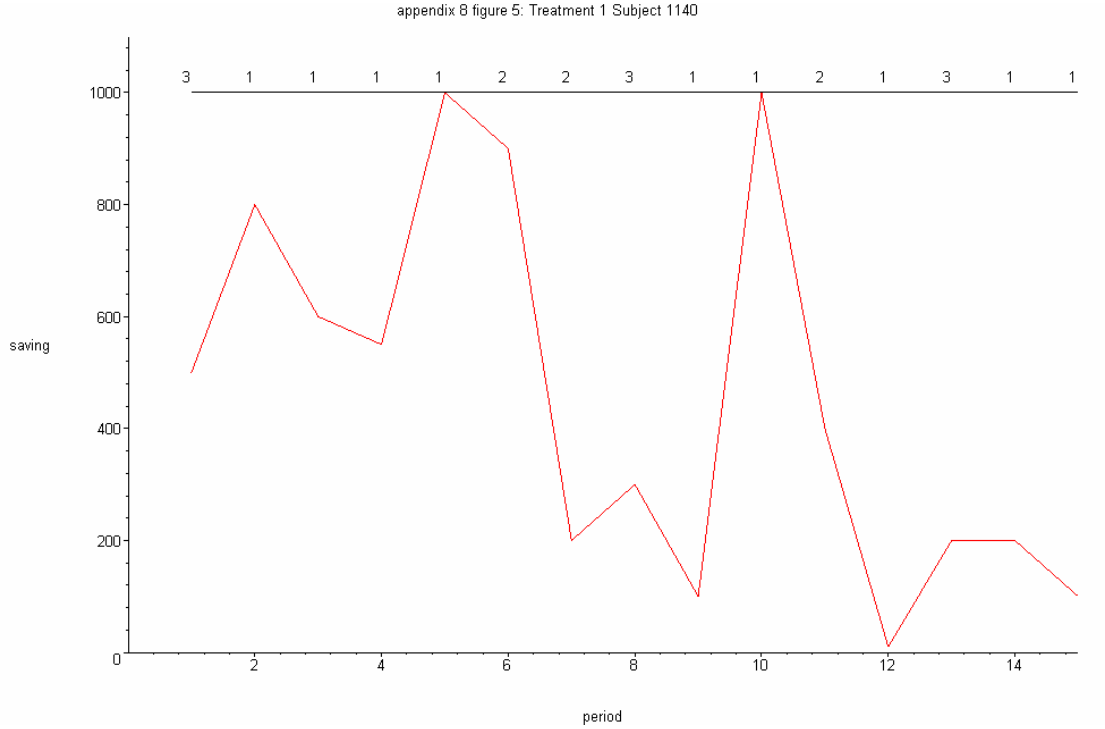

appendix 8 figure 6: Treatment 2 Subject 2097

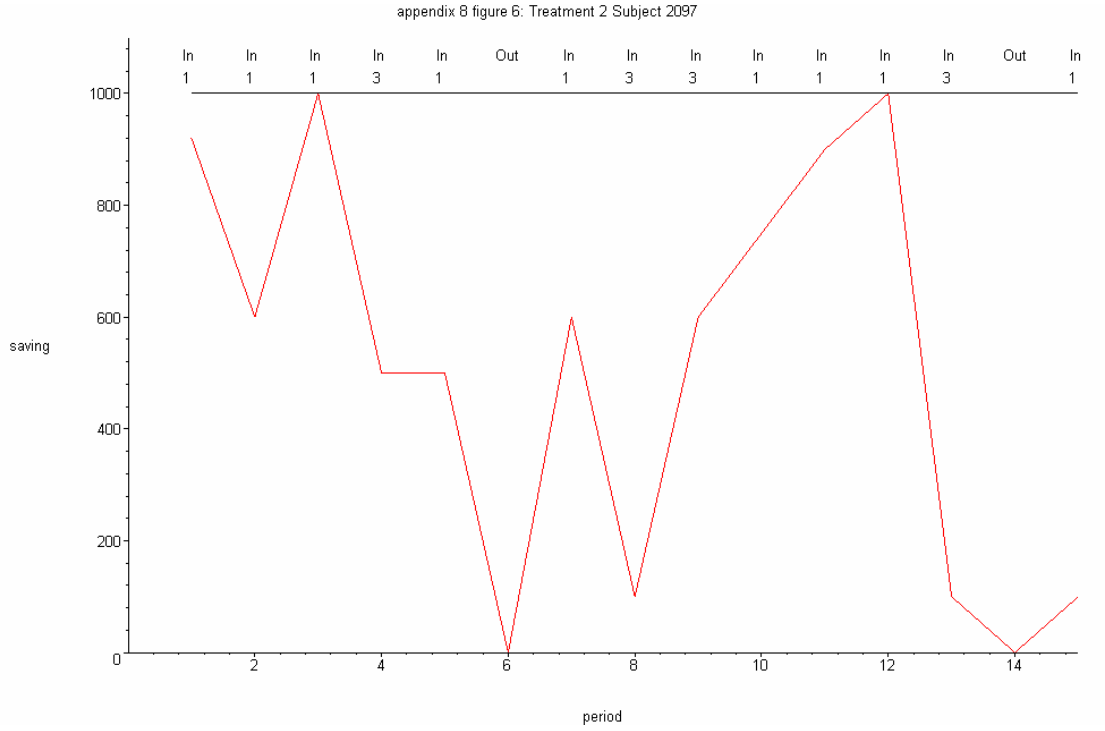

Appendix 9: Description of the variable on risk aversion

9.1 Introduction

'Risk' is a situation in which the outcome is not known with certainty. Risk attitude may then be defined as an individual's propensity or preference for engaging in risky behaviours. According to psychology (Rohrmann, 1994), attitude to risk is a stable personality trait which can predict risk-taking behaviour across a variety of domains, from social and recreational behaviour to financial and occupational decisions.

Risk attitudes can be illustrated with respect to financial decision-making. For example, suppose that an individual is given the choice of a bet with equal chances of winning £100 and nothing or, alternatively, receiving some fixed payment with complete certainty. The expected reward for choosing to bet is £50 in this case. We say that the individual is *risk-averse* if they would rather accept a certain payment of less that £50, *risk-neutral* if they would have no preference between the bet or a certain payment of £50, and *risk-loving* if they would require the certain payment to be more than £50 before declining to bet. We refer the reader to the standard economics literature on risk aversion for formal definitions, e.g., of the 'Arrow-Pratt measure of relative risk-aversion'.

There are many measures of risk attitudes in the economics literature (see, for example, Eichberger et al (2003), Fellner and Maciejovsky (2002), and Wärneryd (1996)). In this report, we establish a numerical measure of an individual's risk aversion and investigate its usefulness, for example, in helping to predict which people are likely to opt-out of personal accounts and how risk aversion affects the choice of fund.

9.2 Calculation of the risk-aversion coefficient (RAC)

We used participants' responses to questions about their risk preferences to calculate the RAC. Please see Section 5.6, Questions 10 to 12. The questions were taken from the literature (Dohmen et al, 2005).

The respondent is asked to rate their likeliness to take risks on a scale of 1 to 5 for Question 10 to Question 11 (part vi). In Question 12 the respondent is asked a hypothetical question about a lottery and the responses are numbers from 1 to 6. For this question lower numbers indicate a higher preference to take risks.

To compute the RAC, the responses to these eight questions were first normalised so that they took values between 0 and 1 (with 0 indicating the most risk-averse response and 1 indicating the most risk-loving response). The RAC is then defined to be 1 minus a weighted sum of these normalised responses. Thus the RAC takes values between 0 and 1 and is 0 if and only if all the participant's responses were the least risk averse and the value is 1 if and only if all the responses were the most risk averse.

The weights chosen for our analysis of the data were as follows:

Table 9.1 Weights used to calculate the RAC in the analysis

Weight	Question	Description
5%	10	General attitude to risk.
5%	11 i	Prepared to take risks while driving.
25%	11 ii	Prepared to take risks in financial matters.
5%	11 iii	Prepared to take risks during leisure and sport.
10%	11 iv	Prepared to take risks in occupation.
5%	11 v	Prepared to take risks with their health.
5%	11 vi	Prepared to take risks with their faith in other people.
40%	12	Amount prepared to invest if given £100,000 lottery win and equal chances of doubling and halving the stake.

9.3 Further remarks on the RAC

The weights were chosen to reflect the importance of financial decision making in the experiment. Although it may be a criticism that these weights are arbitrary, sensitivity analysis shows that the resulting versions of the RAC do not vary much for reasonable choices of weights.

We should note that the RAC is not the same as the 'coefficient of relative risk aversion' (CRRA) or the 'coefficient of absolute risk aversion' both of which appear in the economics literature. Some other measures depend on an individual's wealth while our statistic ignores this. On the other hand, the values in the lottery question we used are all high relative to the incomes of the participants in our sample. Furthermore the following factors suggest that it is not appropriate to use the information on income available to us in the questionnaire responses:

- Individuals' debts are not known;
- individuals' household composition may affect their income; and
- understanding of the lottery question was not monitored.

Having said this, our RAC may be interpreted as a proxy for (a rescaled version of) CRRA. The RAC agrees with results in the literature about what variables affect CRRA (see, for example, Dohmen et al, 2005). For example, the RAC is significantly higher for women than it is for men in our sample. The RAC also varies in a similar way to the CRRA with respect to other variables, for example, age and income, but the results here and in the literature are less clear-cut.

An important validation of the coefficient is that it has a significant effect (cf. Section 8.5) on fund choice: participants with high values of RAC were more likely to have chosen less risky funds.

Appendix 10: Glossary of technical terms used in the Report

Annuity: This a financial instrument that allows for a seller (issuer), typically a financial institution such as a life insurance company, to provide a series of future payments to a buyer (annuitant) for a known sum with a net present value; the payment stream has an unknown duration based principally upon the life expectancy of the annuitant. (Definition taken from Wikipedia).

Consumption Smoothing: If income is not constant through time the individual should try and smooth consumption through time in some way – usually by borrowing or saving. For example, if income is low early in life and after retirement and is high in the mid-life, then the individual is better off borrowing in early life, and then, in mid-life, saving for retirement. What is the 'best' way of doing so depends upon the available interest rates, the individual's discount factor and his or her utility function.

Discounting: This refers to the fact that people consider consumption in different periods differently. More precisely, a given consumption received at some point in the future is usually considered less important to the individual than the same consumption received earlier. The technical term is that people *discount* future consumption relative to present consumption.

Exponential and Hyperbolic Discounting: Usually in economics it is assumed that people discount future consumption at a constant rate. So, for example, consumption received 1 year hence is considered as being worth a fraction $1/(1+\rho)$ of consumption received today, consumption received 2 years hence is considered as being worth a fraction $1/(1+\rho)^2$ of consumption received today, and so on, where ρ is the individual's rate of discount. Such constant discounting is a necessary condition for dynamic consistency – that people do not revise tomorrow plans made today (unless something relevant has changed in the meantime). This is referred to as *exponential* discounting. In practice, however, it seems that people have a present bias and discount the immediate future at a higher rate than the distant future. Such practice is called *hyperbolic discounting* and can lead to inconsistencies in behaviour, so that plans made today are not implemented tomorrow even though nothing relevant has changed in the meantime.

Life-Cycle Theory: This refers to a theory in economics in which people make consumption decisions on the basis of considering their whole life – their life-cycle. In other words, people are not myopic, but consider all their future life when making a decision at some time.

Myopia: the tendency not to consider (or to consider with too little weight) the implications of present behaviour for the future utility/happiness of the individual.

Risk-aversion: an individual is said to be risk-averse if he or she is willing to pay money to get rid of risk. The more that the individual is willing to pay the more risk-averse he or she is.

Risk-neutral: an individual is said to be risk-neutral if he or she is only concerned with maximising the expected return of some project and is not concerned with the amount of risk associated with it.

Significance: loosely we state that an estimated coefficient is significant (more precisely is significantly different from zero) if its t-statistic is greater than 2 in magnitude. This means that it is almost certain that the true value of the coefficient is not zero.

Standard error: when the various estimated equations are reported (both ordinary regression equations and probit equations) we also report the standard error. This gives an indication of the

accuracy of the reported estimated parameters. Very loosely, if the estimated coefficient is *a* and its standard error is *s*, then the true value of the coefficient lies in the interval *(a – 2s. a +2s)*. Hence the smaller is *s* the more accurate is the estimate of the underlying true parameter.

Tobit regression: this is the appropriate estimation technique when the dependent variable is truncated.

Treatments: this is a term used by experimenters to refer to a particular set of sessions of the experiment – in which all parameters and other design features are kept constant. In different treatments parameters take different values. Thus, a comparison of the behaviour of participants between different treatments enables an analysis of the effect of parameter changes on behaviour.

t-statistic: is simply the ratio of the estimated coefficient to its standard error. As a rough rule of thumb, we can say that if this t-statistic is larger than 2 in magnitude, then the true coefficient is almost certainly not zero.

utility: In economics, utility is a measure of the relative satisfaction or desired-ness from consumption of goods. Given this measure, one may speak meaningfully of increasing or decreasing utility, and thereby explain economic behaviour in terms of attempts to increase one's utility (Definition taken from *Wikipedia.*)