Inequality Aversion and Self-Interest:

An Experimental Approach

Matthew Robson

PhD
University of York
Economics
January 2018

Abstract

Preferences relating to *inequality aversion*, the trade-off between equality and efficiency, and *self-interest*, the degree to which the 'self' is weighted in relation to 'others', are incorporated within structural models to explain *prosocial behaviour*, the act of giving to others. To observe such behaviour, incentivised laboratory and lab-in-the-field experiments are run. Structural preferences parameters are then estimated, at the individual, cluster and sample level, within the utility functions proposed. *Noise* in decision making is formally modelled with the Beta and Dirichlet distributions, which are formulated as random behavioural models. In the first chapter, distributional decision problems amongst groups of three are presented to participants within a laboratory experiment. Using *multiple experimental designs* and *alternative perspectives*, within-subject treatment effects are tested. The second chapter incorporates *oneness*, the closeness of connection to others, within a structural model to better explain the differential effects that social distance can have on distributional decision making. In a lab-in-the-field experiment in Mbale, Uganda, modified three-person dictator games are presented to participants to enable the observation of such behaviour, alongside an extensive survey. Finally, the third chapter focuses on N-person giving. Five alternative utility functions are formulated, which incorporate differing behavioural preference parameters; accounting for the distinction of *self-other* and *between-other* inequality aversion, *congestion* and *minimum threshold* levels. Both the goodness-of-fit and predictive accuracy of each model are compared, to identify the 'best' model for each individual. Within each of the three chapters, results show extensive heterogeneity in prosocial behaviour, which is accounted for through the estimated preference parameters. On average, participants have a substantial regard for others and a preference for reducing inequality, rather than increasing efficiency. The experimental design, perspective, *oneness* levels and number of recipients are shown to have significant, but differential, effects on prosocial behaviour.

Contents

List of Figures

List of Tables

Acknowledgements

Undertaking a doctorate makes you realise just how little you know. Spending years specialising on one subject, in one area, of one discipline leaves you with a sense of the infinitely vast array of thoughts, opinions and approaches that exist; of which you know the smallest fraction. It is a humbling thought, and one I am grateful to the PhD for. The little I do know would be a fraction of itself were it not for my friends, family and colleagues, to whom I am indebted.

To my parents, Nigel and Steph, I am ever grateful for all you have done for me; encouraging me in every endeavour and teaching me that with hard work anything can be achieved. I am especially thankful to Harriet. The epic adventures, incredible food and never ending support (including ever so patient proof-reading) have ensured I have loved life throughout the PhD. Charlotte, thanks for always being there (especially with emergency fish and chips). To the rest of my family, especially my grandparents, I also give my thanks. I am particularly grateful to Nancy and Lewis Owen, you are the reason I love reading, without which I would not be here today. To my friends from Bangor and home you are partially to blame for who I am today, so for that I give my thanks. To Bernard, Amy, Josh and Yanos I thank you for your continued friendship, random meet-ups and brilliant conversation. To everyone I have played squash with over years I am very grateful, as through squash the stresses of the PhD have always been much reduced.

I would like to thank Arjan Vershoor, Joshua Balungira and everyone at the Field Lab for the incredible support and advice you gave, helping me fully embrace my time in Uganda. To everyone I met while in Bergen, thank you for making twenty-seven days of consecutive rain in June thoroughly enjoyable! Further thanks goes to Sabina Alkire and the OPHI team. More generally, I would like to thank everyone I have met through various conferences, spring schools and workshops; who have offered great friendship, feedback and advice.

For the financial support from the ESRC White Rose DTC (ES/J500215/1) I also give my thanks. The funding provided allowed me to embrace each opportunity the PhD threw at me. Further thanks goes to all the wonderful people who enabled me to CrowdFund my first experiment.

For those at York, thank you for sharing in the pain of the PhD and making it thoroughly enjoyable. Especially to Gloria, for the copious amounts of tea, to Vigyan, for letting me win at chess occasionally, as well as to André, Ieva and Mashrur for many brunch excursions. To the support and IT staff in the department all would be lost without you, so many thanks for all your help. I am further indebted to all the academics in the department, especially to those with whom I have been taught by or have taught with. I am particularly thankful to John Hey and Richard Cookson, for offering exceptional advice and support throughout.

Finally, I wish to thank my supervisor John Bone. It is through your direction, and our frequent enlightening meetings, that the ideas within this thesis have emerged and been developed. Without you the thesis would simply not be as it is today. I am incredibly grateful for all you have taught me, for being the most supportive supervisor I could have wished for and for keeping my rambling thoughts in check. I consider myself very fortunate to have had you as a supervisor.

Declaration

I, Matthew Robson, declare that this thesis entitled, "Inequality Aversion and Self-Interest: An Experimental Approach" is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as references.

“ *How selfish soever man may be supposed, there are evidently some principles in his nature, which interest him in the fortune of others, and render their happiness necessary to him, though he derives nothing from it except the pleasure of seeing it.*”

- Adam Smith, *The Theory of Moral Sentiments*, (1790)

Introduction

A regard for others is something inherent in human nature. Acts of kindness, generosity and compassion are frequently observed throughout society. Yet, much research in the field of economics assumes that all individuals are purely selfish. The intricacies of the varying degrees of altruism are often ignored, and through it a substantial understanding of human nature is missed.

The constructs of *inequality aversion*, the trade-off between equality and efficiency, and *self-interest*, the degree to which the 'self' is weighted in relation to the 'other', form the foundations of this thesis. These preferences account for the regard for others that particular individuals feel, allowing for heterogeneity in *prosocial behaviour*, the act of giving to others, to be explained. Within each of the three chapters of this thesis incentivised experiments are run to enable the observation of individual-level prosocial behaviour and subsequent estimation of preference parameters, which strive to explain such behaviour.

This introduction will outline the general approach taken, before introducing the three chapters. Following the introduction will be the three chapters. Each are written as self-contained papers, with extensive appendices included. The conclusion will, then, briefly summarise the whole thesis, describing its contributions to the literature and limitations, before discussing possible extensions and proposed future research.

***

Throughout the thesis, the methodological approach taken falls under the umbrella of *experimental economics*. Experimental methods allow for the collection of data not easily observed in the real world. It allows for the abstraction and isolation of particular phenomena enabling the testing of economic theory, estimation of preferences and evaluation of treatment effects. Through utilising experiments, and analysing the data produced, light can be shed on particular aspects of human nature, thus increasing our knowledge and understanding.

To utilise such methods, incentivised experiments were run in the EXEC laboratory, at the University of York, and as a lab-in-the-field experiment, in Mbale, Uganda. Within the experiments participants were presented with incentivised decision problems. Participants made individual choices, using interactive on-screen interfaces, which had real monetary consequences for themselves and others in the experiment. It is through this that a *revealed preference* approach can be taken, which is distinct from the *stated preference* approach. In the latter, participants are asked to state what they would do in a particular scenario, while the former incentivises each decision problem to *reveal* how participants actually behave. Decisions in the laboratory experiments were made using computers, while a temporary laboratory was set-up in the lab-in-the-field experiment, with touch-screen tablets. The controlled environment within both approaches allows for complex experiments to be run, allowing for extensive datasets to be constructed.

The samples within the laboratory experiments predominantly consist of university students, while in the lab-in-the-field experiment participants are drawn from both general population and university student pools. Across both types of experiment rich demographic data are obtained alongside the experiments, through running parallel questionnaire-based surveys. Utmost care was taken when running each experiment, to ensure that the methods used were transparent, replicable and void of deception, with participants decisions remaining anonymous.[1]

To analyse the data obtained from the experiments descriptive, econometric and structural modelling analytical techniques are used. The first is primarily used to illustrate the observed behaviour, while the second is used to identify significant treatment effects, test hypothesis and examine demographic trends. Within the third, utility functions are proposed to explain behaviour through incorporating structural preference parameters. The preference parameters both have an intuitive meaning behind them and strive to account for differences in behaviour. Before going into detail on the preference parameters, a divergence into the notion of utility will be taken. This discussion will clarify the philosophical standpoint of the research.

Utility is a construct which is often thought of as synonymous with happiness, felicity or pleasure. To explain behaviour, economists formulate utility functions; where agents are assumed to behave in a manner which maximises their utility, given their constraints. As individuals prefer one alternative over another, and therefore receive higher utility from it, that alternative is chosen by that individual. Here, utility is considered in the 'modern' context (see Binmore (n.d.) for discussion). Individuals are not assumed to behave in the way they do *because* they are actually maximising the particular utility function proposed. Rather we imagine individuals behave *as if* they are maximising utility, according to the particular functional form proposed and preferences estimated. The utility function is *not the reason why* individuals exhibit such behaviour, but merely allows for a *description* of their behaviour.

A utility function, ultimately, is only a model. A model simplifies something which is complex; in doing so it allows for interesting observations to be drawn, but still falls short of encompassing the truth. A paper plane models how aerodynamic forces act on the paper plane while in flight, while a model aircraft shows a to-scale replica of a real aircraft. We begin with the knowledge that the models are false (aeroplanes are neither made of paper, nor are they miniature) but use them to simplify matters allowing for a greater understanding to emerge. Our objective is to find the model, and preference parameters within the model, which best explain the behaviour observed and provide intuitive insights into the reasoning behind such behaviour.

[1] Indeed, this research focuses upon altruism, rather than reciprocity. While in the latter generosity is often considered as a calculated self-interested act for future return, the former concerns the 'pleasure' derived purely from an increase in the welfare of others. By ensuring anonymity the concerns of reciprocity are removed.

Preference parameters are at the core of the models. Different parameter values within the utility functions allow for alternative preference orderings, which determine which behaviour is optimal (i.e. that which gives the most utility). The flexibility of these parameters allows for heterogeneity in behaviour to be accounted for, and this is the crux of the utility functions proposed. By explicitly modelling heterogeneity the differences in behaviour can be accounted for. In our setting, the interest is in prosocial behaviour, and through it other-regarding preferences. Rather than assuming *all* individuals are purely self-interested an array of preferences can be accounted for, which allows for the variation in behaviour.

In this emerges an interesting paradox. The assumption that all individuals are *purely self-interested* is questioned, yet at the same time the assumption that individuals (selfishly) seek to maximise their own utility is relied upon. This apparent paradox points to some of the main critiques of mainstream economics, however, if a clear distinction is made this paradox is removed. *Egoism*, or *self-love*, refers to those who are purely self-interested in terms of monetary payoffs (income, wealth, etc.). *Self-interest*, can instead be used to refer to agents maximising their own utility. By distinguishing between the two, and relaxing the assumption of *egoism* to that of *self-interest* then the other-regarding preferences (heterogeneous though they are) can be explicitly accounted for in mainstream economic thought.

A further point of discussion relates to assumptions of deterministic choices, that agents always act optimally. If a *deterministic* approach is taken, then individuals are assumed to behave optimally and make no error in their decisions, but if behaviour is assumed to be *stochastic* then decisions are assumed to be *noisy*, meaning that error models need to be proposed to rationalise observed behaviour. If the former is assumed then *elicitation* methods can be used to establish parameters, while preferences need to be *estimated* within the latter. Both *elicitation* and *estimation* techniques are used within the first chapter, with the focus shifting to estimation techniques in the latter two chapters. The type of error model predominantly used is a *random behavioural* model; where the error is made when calculating (or choosing) the optimal decision. The Beta distribution and Dirichlet distribution are formulated as random behavioural models, offering a novel flexible approach to accounting for noise in allocation problems.

Once the utility functions and error models are formulated, and preference parameters estimated, the goodness-of-fit of the model can be evaluated. If the models fit well, they achieve their aims, if not, perhaps better models are needed. However, it is not just fit, but prediction which is important in economics models; as the following illustrates. The quadrant may well be a useful tool, enabling a sailor to pinpoint where in the seas they are at that precise moment, and indeed allowing them to map where they have been, but as Captain Ahab laments: “Thou canst not tell where one drop of water or one grain of sand will be to-morrow noon” (Melville, 1851). Unlike a quadrant, an economic model should be able to predict behaviour, as well as fit previously observed behaviour. Often experimental papers estimate preference parameters which well fit the data, but neglect an assessment of their predictive accuracy. The focus of the third chapter is shifted to address this, enabling the analysis of both the fit and predictive accuracy of the models proposed.

In general, the thesis relies upon experimental methods to gather data on prosocial behaviour. To analyse such data, descriptive, econometric and structural modelling analytical techniques are used. A revealed preference approach is taken, where utility functions are proposed and preference parameters estimated (or elicited) to account for the heterogeneity in the behaviour observed. The extent to which the models explain the data, is then be analysed for both goodness-of-fit and predictive accuracy. The following delves into the specifics of each chapter.

***

In the first chapter *multiple experimental designs* and *alternative perspectives* were presented in an incentivised laboratory experiment (n=30). Within the experiment participants made distributional decisions amongst groups of three. In order to explain the behaviour observed, preferences relating to *inequality aversion* and *self-interest* are estimated at the individual, cluster and sample level, within the utility functions proposed. Within-subject treatment effects are tested between the multiple designs - Pairwise Choice, Discrete Slider and Continuous Allocation - and alternative perspectives - Place-in-Society, Veil of Ignorance and Impartial Spectator. Furthermore, methods relating to both *deterministic* and *stochastic* behaviour are explored; with the Beta distribution proposed as a random behavioural model, to account for *noise* in decision making. The goodness-of-fit of different utility functions and error models is further analysed. Within the chapter is an extended literature review and an extensive discussion, encompassing: differences in design and perspective, individual level prediction and risk under the veil of ignorance, in addition to comparisons with the heath-related social welfare function and distributional preferences literatures. Results reveal extensive heterogeneity exists between-subjects, but also significant within-subject treatment differences. The majority of the sample exhibit behaviour which is inconsistent with a purely individualistic model; with substantial regard for others and a preference for reducing inequality over reducing inefficiencies.

The second chapter incorporates *oneness*, the closeness of connection to others, within a structural model to better explain the differential effects that *social distance* has on distributional decision making. A CES utility function is formulated which builds upon previous models that incorporate *inequality aversion* and *self-interest*. Preferences parameters are introduced which reflect behavioural responses to changes in *oneness*. These parameters distinguish between how elastic *self-other* and *between-other* trade-offs are, to better explain the distributional effects that differential *oneness* can have. Further to this, in order to rationalise *noise* in decision making, the Dirichlet distribution is proposed as a random behavioural model. To observe behaviour an incentivised lab-in-the-field experiment was run in Mbale, Uganda (n=156). The experiment was in the form of a modified three-person dictator game, where two within-subject treatments varied if the identity of the recipients of giving was *anonymous* or *known*. Decision problems were repeated (54 rounds) to ensure individual-level preferences could be estimated. The experiment was run on touch-screen tablets, and recruited both general population and student samples. Alongside the experiment an in-depth survey was conducted to establish an extensive list of demographic characteristics, including indices of multidimensional poverty and asset wealth which are particularly important in the context of development. Results show that entwining both distributional preferences and social distance are crucial in understanding giving behaviour, when the identity of the recipient is known. On average, oneness is shown to have large and significant effects on giving; with a greater willingness to trade-off between the self and others, than between-others, as relative oneness levels change. The defining feature of the results is the extensive heterogeneity in both preferences and prosocial behaviour.

Within the third chapter a modified N-person dictator game was run, in an incentivised laboratory experiment (n=83), to observe how giving behaviour changes in response to a varying number of 'recipients'. Using a *within-subject* design two treatments are presented to participants across 45 rounds; the *multiple* slider treatment, allowing precise allocations to each player, and the *single* slider treatment, only allowing for the proportion of allocations to the self to be chosen. *Individual-level* preference parameters are estimated within *five* alternative utility functions; each incorporating inequality aversion and self-interest. Additional preference parameters are formulated within *extended* models, to account for alternative behavioural responses to changes in N. The first models the distinction between *self-other* and *between-other* inequality aversion. The second incorporations *congestion*, the trade-off between *average* and *total* payoffs to others. While the third accounts for *minimum threshold* levels. The relative goodness-of-fit and predictive power of each model is tested, allowing for the identification of 'types' of individuals. This approach allows the flexibility to explain heterogeneity in individual behaviour not only through preference parameters within a particular model, but between different behavioural assumptions made in alternative models. Additional error parameters are further incorporating in the random behavioural formulation of the Dirichlet distribution, allowing for differential error as the *complexity* in decision making increases. Results show that increasing the number of recipients changes the behaviour of individuals in different ways. On average, participants are willing to give more *total* payoffs to others as the number of players increase, but not to maintain *average* payoffs to others. However, extensive heterogeneity is found in individual preferences, with no model 'best' fitting all individuals.

There are, of course, differences between the chapters. The experiments of the first and third chapters are run in a laboratory, with student samples from the University of York, while in the second chapter both student and general population samples are recruited in a lab-in-the-field experiment in Mbale, Uganda. The first chapter varies both perspective and design, while the latter two keep both constant. Within the second chapter the anonymity of the others is removed, while throughout the other chapters anonymity is maintained. The Beta distribution is formulated as a random behaviour model in the first chapter, while the Dirichlet distribution is used in the second and third. Only one model is formulated in the second chapter, while several utility functions and error models are analysed in the first and third. However, it is this due to the variation in the three chapters which allows for a more complete analysis to emerge.

The three chapters of the thesis strive to account for a host of different factors which surround *inequality aversion* and *self-interest*. Each the experimental design, perspective of the participant, closeness to others and the number of recipients, alongside the prices of giving and size of the budget, are incorporated in the thesis, and their effects on prosocial behaviour and other-regarding preferences are tested. This broad inquiry into a specific component of human nature will hopefully be of interest to the reader, shed light on the subject matter and provoke further thoughts, questions and debate.

Chapter 1

Multiple Designs and Alternative Perspectives

Preferences relating to *inequality aversion* and *self-interest* are estimated, at an individual, cluster and sample level, using an incentivised laboratory experiment (n=30). The experiment presents participants with decision problems concerning distributional decisions amongst groups of three. Using *multiple experimental designs* and *alternative perspectives* within-subject treatment effects are tested. To model behaviour utility functions are proposed which question the integrity of a purely individualistic models, by incorporating preference parameters relating to *self-interest* and *inequality aversion.* Results reveal that while preferences are heterogeneous between-subjects, within-subject differences caused by treatment effects are significantly larger. Furthermore, the majority of the sample exhibit other-regarding behaviour which is inconsistent with a purely individualistic model.

***

1 Introduction

References to inequality aversion and other-regarding preferences are becoming ever more present in the field of economics. Stemming from experimental, behavioural and welfare economics, these notions question the integrity of purely individualistic utility functions. Models which offer greater explanation of individual behaviour are being formulated, by accounting for heterogeneous preferences over distributional decisions. Yet, different experimental designs are used to observe individual behaviour and participants are asked to consider alternative perspectives when making distributional decisions. Do such design decisions have a significant effect on the behaviour observed and preferences estimated?

In this paper, distributional decision problems over groups of three are presented to participants within an incentivised laboratory experiment. Individual, cluster and sample level preference parameters are estimated within utility functions which incorporate concerns of *inequality aversion* and *self-interest*. By using multiple experimental designs and alternative perspectives within-subject treatment effects are analysed. Methods relating to both *deterministic* and *stochastic* behaviour will be explored; where the Beta distribution is proposed to rationalise *noise* in individual decision making. The goodness-of-fit of different utility functions and error models is compared, and the importance of distinguishing between risk aversion and inequality aversion under the 'veil of ignorance' is discussed.

The following section will review the relevant literature. The experiment will then be described, with details of the three experimental designs - Pairwise Choice, Discrete Slider and Continuous Allocation - and the three alternative perspectives - Place-in-Society, Veil of Ignorance and Impartial Spectator. The theoretical framework will follow, where the functional form of the utility functions will be explained alongside the random behavioural error model. The results will then be presented, followed by the discussion and conclusion.

1.1 Literature

1.1.1 Distributional Preferences and Social Welfare Functions

In mainstream economics one fundamental assumption is that individuals are innately selfish. They aim to maximise an objective function where their only concern is their own monetary payoffs, not the payoffs of others. Yet a body of literature has emerged, primarily in the field of experimental economics, which reveals that utility functions often better explain individual behaviour by incorporating other-regarding preferences. Two main streams of thought, within the *distributional preferences* literature, have been proposed; those which incorporate *social preferences* and those which incorporate *reciprocity* (E. Fehr and Schmidt, 2006). It is the former which will be concentrated upon here. Models of social preferences not only assume an individual's utility is based upon their own income, but also upon the distribution of resources amongst others. The degree of self-interest can, however, vary between individuals. Three major models are those concerning altruism (Andreoni and Miller, 2002), relative income and envy (Bolton, 1991), and inequity aversion (E. Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). Various hybrid models have also been established (Charness and Rabin, 2002; Erlei, 2008; Cox, D. Friedman, and Gjerstad, 2007). These models primarily seek to model individual behaviour, which is observed in economic experiments.

Another strand of literature, focusing upon Social Welfare Functions (SWF), is concerned with establishing a single numerical welfare index from "aggregating all information about an income distribution" (Amiel and F. Cowell, 1999). Implicit within a SWF is a value judgement concerning the trade-off society is willing to make between equality and efficiency. It can be used to rank different societal distributions. There are a wide range of forms a SWFs; the Atkinson Index (Atkinson, 1970), Kolm Index (Kolm, 1976), Health-Related Social Welfare Function (Dolan and Tsuchiya, 2009) and the Parabolic Health-Related Social Welfare Function (Abasolo and Tsuchiya, 2004), for example. This strand of literature, rather than focusing upon explaining individual behaviour seeks to establish societal moral value judgements, which enable a ranking of different distributions amongst the population. The parameter of interest in such models is the *inequality aversion* parameter. Predominantly through questionnaire-based surveys techniques, in the realm of health, several studies have sought to elicit a societal inequality aversion parameter (Dolan and Tsuchiya, 2011; Abásolo and Tsuchiya, 2013; Robson et al., 2017).

Deeply rooted in both strands of literature is the concept of inequality, and the degree of aversion to it. Yet, the two strands seem to vary somewhat on the exact meaning of this aversion. The concept of *inequality aversion* in the distributional preferences setting is determined from a self-regarding perspective, while within the SWF context it is a philosophical concern (or lack of concern) for unequal distributions amongst the entire society. What would, therefore, seem to be a plausible assumption is that it is both of these elements which are important; in both individual behaviour and moral value judgements. It is the combination of these two elements, the degree of inequality aversion for distributions amongst all present and a level of self-interest, which shapes the decisions of individuals, which will be the focus of this research.

1.1.2 Experiments or Questionnaire-Based Surveys

While these two separate strands of literature often intertwine their methodological standpoints frequently conflict. Those concerned with distributional preferences mostly use incentivised laboratory experiments, while questionnaire-based surveys are predominantly used in the SWF literature. Experimental economists often argue that questionnaire type data is second best; that not only do they not truly *reveal* what the populace actually would do, as there are no monetary consequences, but they are prone to various biases and errors. The power to explain individual behaviour is, therefore, argued to be significantly decreased. On the other hand, those who utilise questionnaire-based surveys argue that when eliciting values about social norms "the subjects should be genuinely interested in the underlying issues which the experimenter wishes to study, and not primarily in the sums of money they can bring home after the experiment" (Gaertner and Schokkaert, 2012). Indeed, strongly-defined views may not translate directly into concrete actions (Amiel and F. Cowell, 1999). What is clear is that, although these arguments could point to issues with the other method, the main issue is a lack of cohesion between the aims of the two; while the experimental economists strive to ascertain individual behaviour, those conducting questionnaire-based surveys are more often concerned with social values. Yet, when considering individual behaviour which is influenced by social norms an area between the two strands emerges. While having to remain wary of the two arguments, a synergy can be established by combining necessary elements.

1.1.3 Philosophical Standpoints

Often there is a lack of integration between economics and philosophy. However, the concept of *inequality aversion* as proposed in Atkinson (1970) carefully intertwines the two. Inequality aversion is defined as the trade-off between equality and efficiency. The concept is encapsulated as a parameter in an economic model, but the standpoints the parameters represents are grounded in philosophy. The Atkinson Index in Equation (1.1) is a social welfare function, which incorporates ε as the inequality aversion parameter, where $\varepsilon \in [-1, \infty]$. The income to each agent i is denoted by x_i, μ is the mean level of income and $f(x_i)$ is the proportion of agents with the same level of income as i:

$$I = 1 - \left[\sum \left(\frac{x_i}{\mu}\right)^{-\varepsilon} f(x_i)\right]^{\frac{1}{-\varepsilon}} \tag{1.1}$$

To visualise what different values of ε represent, iso-welfare curves denoting the indifference between distributions between individual i and j are shown in Figure 1.1. By following the line of each curve the different distributions of income between which an individual with that level of inequality aversion is indifferent can be observed. When $\varepsilon = -1$ the curve is linear, representing the philosophical standpoint of 'Utilitarianism' where total income, regardless of its distribution, is all important. As ε increases more weight is attached to income transfers lower in the distribution; representing a 'Weighted Prioritarian' standpoint, where a higher weight is given to the worse off. The standpoint of 'Maximin' means that welfare is solely dependent upon the income of the worst-off in society, and is represented by $\varepsilon = \infty$. To extend the Atkinson inequality aversion parameter; if ε was unbounded below -1 then it could incorporate those who actually preferred inequality, those who are 'Inequality Seeking'. Another extension, is to consider the standpoint of an 'Egalitarian'. 'Egalitarians' are so averse to inequality that they are willing to sacrifice the income to the worst-off in order to reduce it. To account for this preference an alternative functional from, which allows for such violations of monotonicity, needs to be proposed as in Abasolo and Tsuchiya (2004).

1.1.4 Perspective

When answering the question of how averse to inequality is an individual? The perspective from which the individual considers the question needs to be accounted for. Several studies have been conducted assessing the extent to which there are significant differences in value judgements, given an alternative perspective. Core to the literature are four different perspec-

Figure 1.1: Iso-Welfare Curves for Alternate Inequality Aversion Parameters

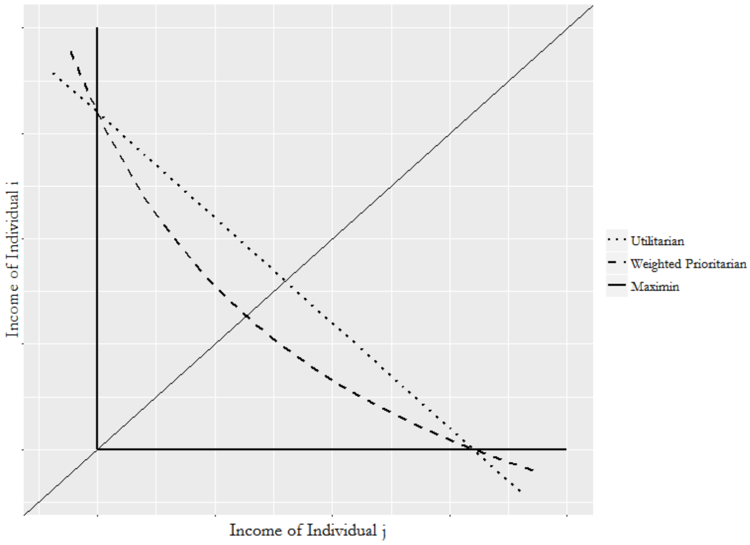

tives where the individual: knows their Place-in-Society, is under a 'thin' Veil of Ignorance (Harsanyi, 1953), is under a 'thick' Veil of Ignorance (Rawls, 1999) or is an Impartial Spectator (Smith, 1790). To be specific, when considering the Place-in-Society the individual knows exactly who they are in society. A 'thin' Veil of Ignorance is where they know that, with a certain probability, they will be one of the members of society, but do not know which. While a 'thick' Veil of Ignorance is where individuals are completely shrouded by the veil, they know not who they are, nor the probability of being anyone in the distribution. An Impartial Spectator, is someone who decides upon the distribution of a society when they are not involved; often referred to as a social planner.

Traub et al. (2005) focus upon two of these perspectives: the Impartial Spectator and the Veil of Ignorance. They also have two levels of information, where the probabilities are known, 'ignorance', and when they are not, 'risk'; which conform to the 'thick' and 'thin' veil's, respectively. They found that in the 'risk' treatment the Impartial Spectator is more inequality averse than when under the veil, yet in the 'ignorance' treatment the reverse was true. Further papers have also conducted research on the difference of perspective (Bosmans and Schokkaert, 2004; Herne and Suojanen, 2004; Amiel, F. A. Cowell, and Gaertner, 2009). The choice of this perspective is therefore important when establishing parameter values of inequality aversion.

1.1.5 Is Behaviour Deterministic or Stochastic?

When establishing individual-level preferences there are two alternative streams of thought; that behaviour is either deterministic or stochastic. Behaviour is *deterministic* when individuals make no error in their decision making. If this is the case then experimenters can present individuals with decision problems, and through observing their choices *reveal* 'true' preference parameters through direct *elicitation*. The alternative is to assume that behaviour is *stochastic*. Rather than individuals always behaving 'optimally' their decisions are *noisy* and prone to *error*.[1] Here, rather than direct elicitation, a particular stochastic error structure has to be assumed and preferences *estimated* from observed behaviour in multiple decision problems.

Three families of models have been proposed to account for stochastic error: Random Utility, Random Behavioural and Random Preference (Conte and Moffatt, 2014). Random Utility (RU) models, or Fechner models (Fechner, 1965 [1860]), assume the utility valuation of each alternative choice is subject to random error. When the participant then ranks alternatives according to the level of utility there may be errors, due to this miscalculation prior to ranking. The stochastic term is applied additively to the utility function. Random Behavioural (RB) models assume participants aim to calculate the optimal decision $y*$ directly, but then make an error when calculating or choosing $y*$. The actual behaviour is then assumed to be distributed according to some stochastic error distribution. Random Preference (RP) models assume that it is the preferences of the individual which are stochastic, meaning that different rankings of alternatives will be made and a different optimal decision chosen, from that new ranking.

The stylised graph, in Figure 1.2, highlights the three forms of error. The solid curve represents the level of utility for each alternative choice y, where the optimal choice is y^*. The vertical arrows represent the errors which are encompassed in the RU model, while the horizontal arrows the RB model. The grey dotted curve represents the alternative utility rankings, from the RP model, which has y_P as the 'optimal' decision. The model used throughout the main analysis is the RB model, intuitively for the continuous nature of the decision problem used it is more suitable than RU and, unlike RP, can 'rationalise' all possible decisions made.

Figure 1.2: Stochastic Error

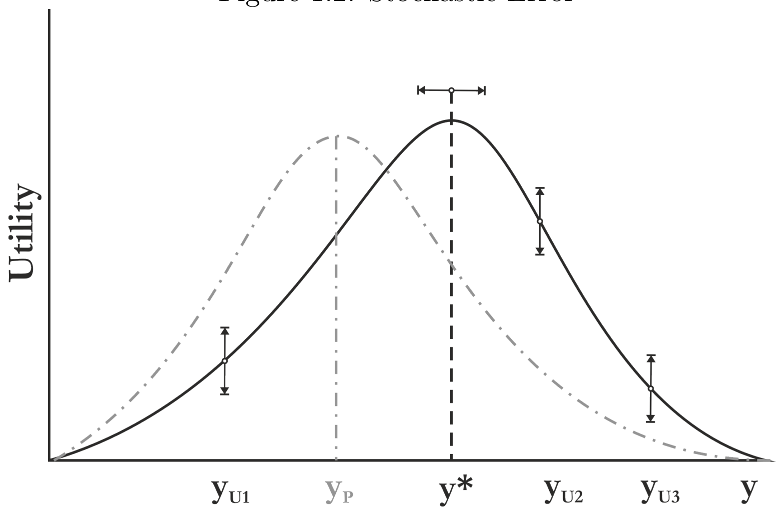

[1]Of course, it is not necessarily the case that individuals are making error, but that the predictions of the proposed model are wrong. Indeed this is the viewpoint taken here. That models by their very nature are wrong; they cannot possibly capture every element of human decision making. But by simplifying such complex decision making into a mathematical model, we can better understand behaviour. This enables us to predict differential choices by accounting for preferences. The aim is then to find the most simple model which is least wrong.

While *deterministic* models may be seen as naïve, and *stochastic* models the more sophisticated of the two, there is a trade-off between the two methods. In practical terms, participants have to be paid to be in the laboratory, which implies a time constraint. A greater range of hypotheses can be questioned and parameters elicited if deterministic preferences are assumed. Indeed, if too many repeat questions are asked it may be argued that, in trying to reduce the issues associated with error, one ensures error in the response; as participants become tired and bored of a repetitive task. A trade-off between the two has to therefore be established. Not enough questions leave conclusions prone to potential error in responses, while too many may prove inefficient and actually create error.

***

This paper contributes to the literature with an experimental design which allows for an inclusive analysis of the concerns of the above literature. The three-person design of the experiment enables both *self-other* and *between-other* distributional decisions to be observed. Through this, preferences parameters can be estimated (and elicited), across multiple experimental designs, when individuals consider both *partial* and *impartial* perspectives. The preference parameters established both strive to explain individual behaviour and establish moral value judgements over distributional decisions.

2 Experimental Design

The experiment is broken up into designs, games and rounds. There are a total of 3 designs, 10 games and 68 rounds. Throughout the experiment participants made incentivised decisions about the distribution of monetary payoffs between group of three players. The three designs are different methods of establishing the parameter values of interest, they are: *Pairwise Choice* (PC), *Discrete Slider* (DS) and *Continuous Allocation* (CA). Within each of these designs there are three different perspectives which the participants were asked to assume, there were: *Place-in-Society* (PS), *Veil of Ignorance* (VOI) and *Impartial Spectator* (IS). There were multiple rounds within each of these individual perspectives. More specific details about the experiment design and the sample can be found in Appendix A.1, alongside the paper instructions in Appendix A.2.

2.1 Three Designs

In the PC design participants were offered the choice between two different distributions of *payoffs*; Choice A and Choice B. The payoffs were iteratively reduced, for Choice B, each round, creating equality-efficiency and self-other trade-offs. The DS design is in the form of a modified dictator game; participants were given a *budget* (£30) to allocate amongst three players. For each player the allocations given were multiplied by a *multiplication factor* to give the final payoff. The choices allowed were discrete distributions and they were made by selecting a notch along a slider. The CA design was, again, a modified dictator game, however, the decisions allowed were (almost) continuous, to distribute to the nearest pence.

While there are three different designs in place there are several aspects which remain common throughout. Firstly, each round is supplemented by *instructions*. There are the paper instructions to begin with, which participants keep beside them, then on-screen tutorials to explain the layout of each design. Before each game there was a prompt screen explaining the particulars of that game, and each round there are shortened instructions to the side. Within each round there were the possible payoffs for each player in the group, shown by the height of the orange bars, with numbers showing the precise payoffs. There were accompanying individual payoffs, alongside the 'total sum of payoffs' (the sum of the payoffs in the group) and the 'gap between payoffs' (the difference between the best-off and worst-off player). Common to each round was a timer, the timer counted down to a minimum time (before which they could not advance to the next round) and then to a red maximum time (after which they were forced to the next round). Participants had to interact in some way, particular to the design, to indicate their preferred distribution of payoffs amongst the group. The specifics of each design are shown below.

2.1.1 Pairwise Choice

The first design presents respondents with a set of pairwise choices, which offer the choice between two different distributions of payoffs: Choice A and Choice B. Participants made their choice by clicking either distribution. Within each of the 4 games (2 PS, 1 VOI and 1 IS) there were 8 rounds. Throughout each of the rounds Choice A remained the same, while Choice B started at some initial point and was incrementally reduced. Figure 1.3 shows the screen of one such round.

Figure 1.3: Pairwise Choice: Z-Tree

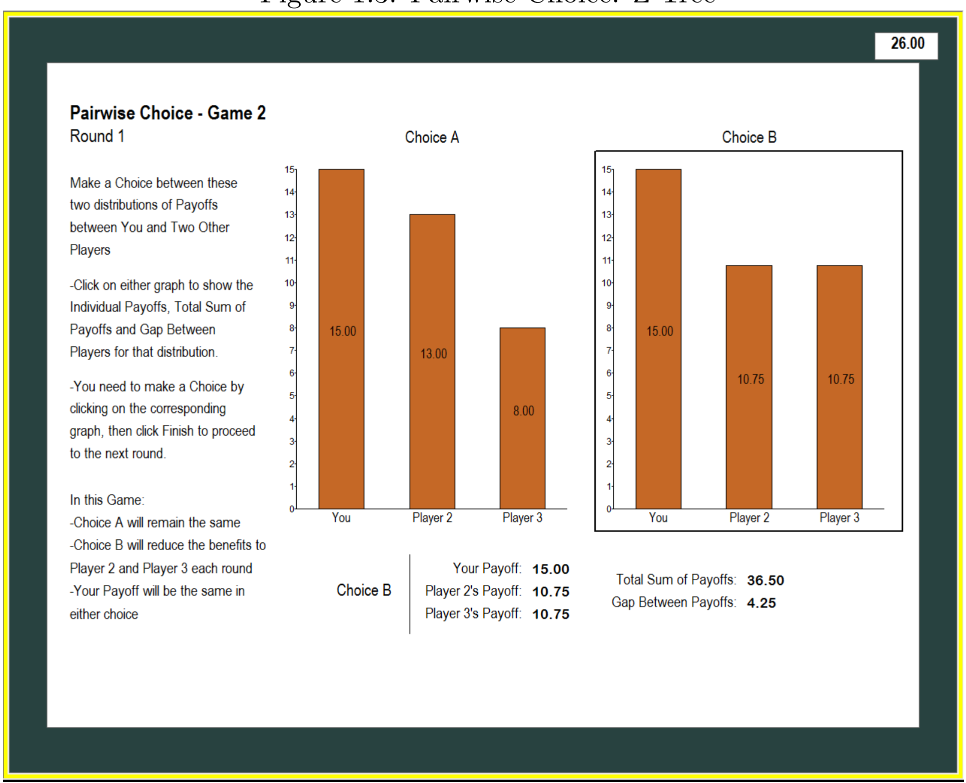

To *elicit* individual parameter values the properties of their indifference curves need to be established. For inequality aversion, r, it is the curvature which is required, while it is the gradient which is needed for self-interest, α. Figure 1.4 shows a simplified two person case, where each curve represents a different degree of inequality aversion. At each point along a given curve the individual is indifferent between the distributions of payoffs. To establish the curvature, and the corresponding r, two distributions of payoffs between which the participant is indifferent needs to be found. As previously mentioned Choice A remains the same, but Choice B changes each round, as shown by the horizontal dotted line in Figure 1.4. By finding the point at which the participant 'switches' from B to A, using the midpoint between the two, the point of indifference can be found. This 'switch' is expected as there is a trade-off between equality and efficiency. Choice B is the more equitable choice of the two, but in each round the total payoffs are reduced. A similar method is used for α, considering the gradient of the curve.

Figure 1.4: Indifference Curves for Varying Level of Inequality Aversion

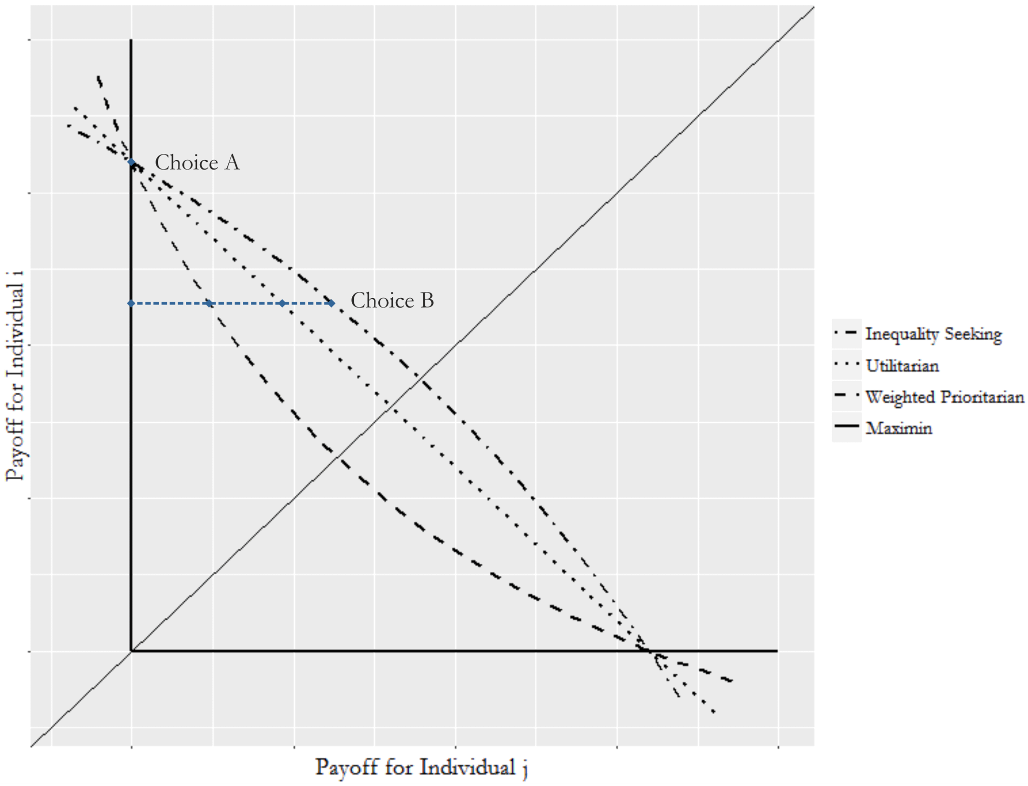

For both the VOI and IS treatments only the r parameter needed to be established, as α is assumed to be 1/3. This meant that only one set of 8 rounds needed to be presented to elicit an r parameter. Choice A remained constant at [15, 13, 8], for Player 1, 2 and 3 respectively. Choice B on the other hand kept the payoff to each player equal; these were [12.4, 11.6, 10.8, 10, 9.2, 8.4, 7.6, 6.8] for the corresponding eight rounds. To illustrate this Table 1.1 shows the categorised choices, with a clear 'switching' point, alongside their r parameter and category for the VOI and IS treatments. Take the 'Utilitarian' category as an example. In order to elicit an r value of -1, the participant must have chosen BAAAAAAA, that is B in Round 1, then A for the seven rounds after. As they 'switched' between Round 1 and 2 the indifference point is assumed to be 12, halfway between 12.4 and 11.6. From this we infer that the participant is indifferent between the points [15, 13, 8] and [12, 12, 12]. By numerically solving equivalent utility functions an r parameter of -1 is then elicited.

Table 1.1: Pairwise Choices: VOI and IS

Choice	Indifference Point	r	Category
AAAAAAAA	12.8	-3.5445	Inequality Seeking
BAAAAAAA	12	-1	Utilitarian
BBAAAAAA	11.2	0.9293	Weighted Prioritarian 1
BBBAAAAA	10.4	2.8702	Weighted Prioritarian 2
BBBBAAAA	9.6	5.4989	Weighted Prioritarian 3
BBBBBAAA	8.8	11.4793	Weighted Prioritarian 4
BBBBBBAA	8	∞	Maximin
BBBBBBBA	7.2	N/A	Egalitarian 1
BBBBBBBB	6.4	N/A	Egalitarian 2

For the PS perspective, both r and α needed to be elicited. This was done by ensuring the choice set could elicit r irrespective of α first and then, contingent on this r, elicit α. This was achieved by keeping the payoff to P1 identical in both choices, while changing the distribution of payoffs between P2 and P3 throughout the rounds. Choice A remained as [15, 13, 8]. Choice B kept the payoff to P1 at 15 in each round but had P2 and P3's payoff equal at [10.75, 10.25, 9.75, 9.25, 8.75, 8.25, 7.75, 7.25] for the corresponding round. Once r had been established α could then be elicited. A set of choices was established for each possible r, they adhered to the following equation, where x_a,x_b and x_c are the income for P1, P2 and P3 in Choice A, respectively, and x_1,x_2 and x_3 is income for P1, P2 and P3 in Choice B, respectively:

$$\alpha = \frac{x_b^{-r} + x_c^{-r} - x_2^{-r} - x_3^{-r}}{2x_1^{-r} - 2x_a^{-r} + x_b^{-r} + x_c^{-r} - x_2^{-r} - x_3^{-r}} \tag{1.2}$$

2.1.2 Discrete Slider

The *discrete slider* is a discretised modified dictator game. Participants are given a budget, which they must allocate amongst the three players. They are, however, constrained to a discrete set of distributions from which they must choose. These choices will be represented with a slider, which at each notch provides different levels of payoffs for each of the three players. Participants must move the slider to each possible notch, to see each available option, then move the slider to the notch of their preferred distribution. The distributions used are the optimal allocations for given levels of r and α (see x_i^* in Section 3). A screenshot of the design is shown in Figure 1.5.

In order to elicit parameter values of r and α two stages are necessary. In the first stage, the game is simplified to the specialised case where and $\pi = [1, 1, 1]$. Due to this specification, for any given value of x_1^* it must be that $x_2^* = x_3^* = \frac{m - x_1^*}{2}$ in order to ensure that such allocations are optimal values for a set of r and α values. Eight notches of discrete choices

Figure 1.5: Discrete Slider: Z-Tree

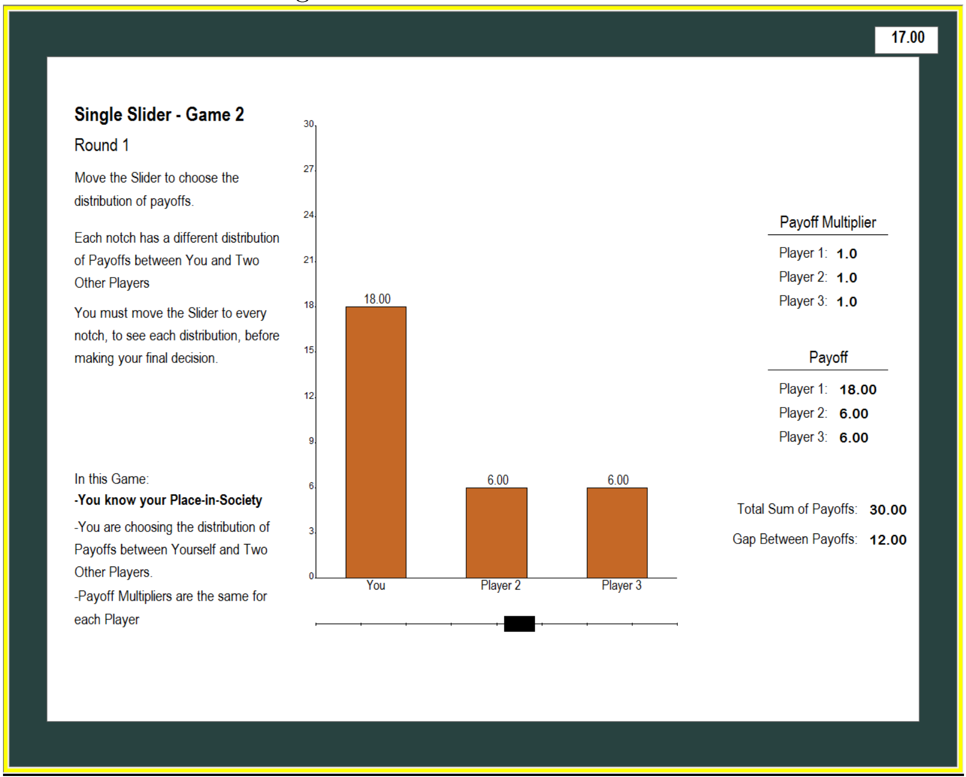

are provided, whereby $x_1^* \in [10, 9, 8, 7, 6, 5, 4, 3.33]$. Given these discrete choices, and the specialised case of the game, there exist a set of corresponding α and r parameters for which that allocation is optimal; i.e. there is no unique solution. After the first stage the second slider was presented, the choice set given was dependent upon the response to the first slider. Importantly, the productivity factor was changed to $\pi = [1, 0.6, 1.4]$, in order to find a unique solution. It is dependent upon this second choice that a unique categorised α and r parameter can be elicited for a participant.

2.1.3 Continuous Allocation

The final design allows participants to precisely distribute allocations amongst the three players. Participants were given a budget, £30, and had to distribute the entire budget. The *allocation* for each player was multiplied by their respective *payoff multiplier*; the resulting amounts were denoted by the *payoff*. Within the design there were the three perspectives, each of these had multiple rounds; VOI and IS had 5 rounds each, while PS had 20. Within these rounds the payoff multipliers were varied, to provide different decision problems to the participants. A screenshot of the design is shown in Figure 1.6.

To make their allocations participants had multiple options. Firstly, they could move the *sliders*. These sliders are accurate to £0.01, are between an allocation of £0 and an allocation of the Total Budget, £30. The second option was to click the *arrows*, either side of the sliders, these made the small changes of £0.01 easier, so that participants could more accurately adjust their allocations, once near with the slider. The final option was to use

Figure 1.6: Continuous Allocation: Z-Tree

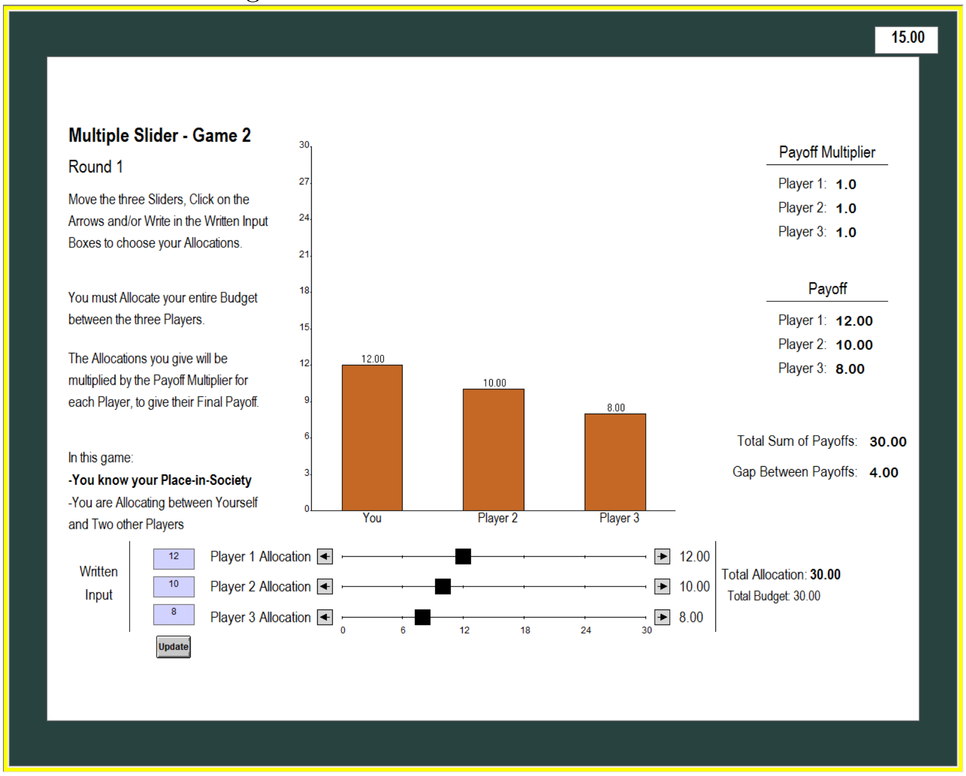

the *written input* boxes; by inputting numerical values and clicking update the participants could enter exact allocations for each player. The slider, arrows and written input boxes are available for each of the three players. When the allocation was input, the payoffs to each player updated immediately.

2.2 Three Perspectives

One question that may be asked is to what extent does the perspective of the individual affect their preferences and behaviour? The three perspectives which are included within each design are the: Place-in-Society, 'thin' Veil of Ignorance and Impartial Spectator. For the first perspective, Place-in-Society, the participant knows which of the three players they are in the distribution; making distributional decisions concerning themselves and two others. The Veil of Ignorance stems from work by Rawls (1999) and Harsanyi (1953), where participants are placed under a 'veil'; they know they are distributing between themselves and two others, but do not know which player is which (see Section 6.5 for further discussion). Our approach is perhaps closer to Harsanyi's, in that participants know the probabilities of being either player is equal. Finally, the last perspective is where participants are outside of the distribution, and are making distributional decisions about three other players. This stems from the concept of the impartial spectator (Smith, 1790), an outside figure looking down upon a society in which they are not involved; often referred to as a social planner.

2.3 Questionnaire

Following the experiment was a questionnaire; asking questions on demographics, opinions and testing cognitive abilities. The demographic questions are on: gender, age, nationality, ethnicity, religion, highest level of education and household income. The opinions questions concern: political persuasion, opinions on reducing inequality, the importance of equality vs incentives from income differences and the trade-off between total income vs income equality. While the cognitive questions ask three cognitively difficult questions to reveal if participants are willing to exert the effort, and are capable of answering such questions.

3 Utility Function

In order to intertwine the distributional preferences and SWF literature, a utility function must be proposed which incorporates the essential elements of individual behaviour and moral value judgements. The functional form is derived from the Atkinson Index (Atkinson, 1970); a Constant Elasticity of Substitution (CES) function with additive separability and convexity to the origin. Importantly, the 'utility' value corresponds to the *Equally Distributed Equivalent* (EDE), which is defined as the mean level of payoffs "which if equally distributed would give the same level of social welfare as the present distribution" (Atkinson, 1970). The index has been reformulated to incorporate the α parameter, which relaxes anonymity and allows for varying degrees of self-interest, by Andreoni and Miller (2002), and is similar in form to that used by Dolan and Tsuchiya (2009) and Fisman, Kariv, and Markovits (2007). Specific to our case are the Multiplication Factors, π_i; the inclusion of which means participants are concerned with the outcome, or payoff, rather than the allocations, x_i. This approach allows a SWF to be translated into an individual utility function, where the individual has some other-regarding preference and aversion to inequality. The *outcome* model is as follows:

$$U_i = \left(\sum_{i=1}^{N} \alpha_i (\pi_i x_i)^{-r} \right)^{-\frac{1}{r}} \tag{1.3}$$

The parameters of interest are *inequality aversion*, r, where $-1 \leq r \leq \infty$ and $r \neq 0$, and the degree of *self-interest*, α, where $0 \leq \alpha \leq 1$ and $\alpha_1 = \alpha$. The regard for others $\alpha_j = (1-\alpha)/n$, denotes the weight given to each other, where $j > 1$ and n is the number of others. The level of inequality aversion corresponds to the convexity of the indifference curves. When $r = -1$ the curve becomes linear, representing the 'Utilitarian' standpoint. As r increases the curves become more convex, implying a higher weight to those who are worse off, which represents a 'Weighted Prioritarian'. When $r = \infty$ 'Maximin' preferences are represented. The gradient of the indifference curve is determined by the degree of self-interest. When $\alpha \to 1$ individuals become more self-interested, as $\alpha_j \to 0$, while when $\alpha \to 0$ they become less self-interested, as $\alpha_j \to 1/n$. The other parameters are x_i, π_i and N where: x_i is the allocation given to Player i, π_i is the *multiplication factor* for player i and N is the number of players in the group.

Given the above utility function and a budget constraint $m = \sum_{i=1}^{N} x_i$, where m is the budget, the following optimal allocations, which maximise utility, can be obtained:

$$x_i^* = \frac{m}{1 + \sum_{j \neq i}^{N} \left(\frac{\pi_i}{\pi_j} \left(\frac{\alpha_j \pi_j}{\alpha_i \pi_i} \right)^{\frac{1}{1+r}} \right)} \tag{1.4}$$

In the specific case of three individuals, we have $\alpha_2 = \alpha_3 = (1-\alpha)/2$. Further to this, given that the allocations are optimal we can obtain r and the self-interest parameter α contingent on that r, where $j, k \neq 1$ and $j \neq k$:

$$r = \frac{-ln\left(\frac{\pi_j}{\pi_k}\right)}{ln(\pi_k x_k) - ln(\pi_j x_j)} - 1, \qquad \alpha = \left(\frac{n}{\left(\frac{\pi_j x_i}{\pi_1 x_1}\right)^{-(1+r)} \frac{\pi_j}{\pi_1}} + 1 \right)^{-1} \tag{1.5}$$

Through observing any allocation the participant chooses, then r and α can be established; assuming that individuals are behaving optimally, so to maximise the objective function.

3.0.1 Graphical Intuition

To provide some intuition behind the interaction of the parameters in the proposed utility function Figure 1.7 shows the optimal allocation to Player 1, x_1^*, for given inequality aversion, r, and self-interest, α, parameters. The scenario is normalised to where $m = 1$ and each of the multiplication factors are unity. When individuals are entirely self-interested, where $\alpha = 1$, r becomes irrelevant and the allocation goes solely to Player 1. Similarly, when $\alpha = 0$ and the individual is entirely selfless, x_1^* goes to zero. When $r \to \infty$, meaning that individuals are 'Maximin', given that $\alpha \neq 0, 1$ the level of 'self-interest' have little effect and the payoffs are split equally. When an individual is a 'Utilitarian', when $r = -1$, they allocate all the Player 1, if $\alpha > \frac{1}{3}$, when $\alpha < \frac{1}{3}$ they allocate zero to Player 1, and if $\alpha = \frac{1}{3}$ then they are indifferent between any allocation. When $r \to 0$ 'Cobb-Douglas' preferences are represented, where x_i^* is proportionate to α.

3.1 Allocation Model

The above *outcome* model proposes that individuals are concerned with the distribution of payoffs, meaning that decisions are made by accounting for the multiplication factors. An alternative functional form, the *allocation* model, assumes that individuals ignore the multiplication factors; and instead focus only upon the allocations they are giving. This is a simplified version, where the optimal allocations would be somewhat easier to calculate. It is hypothesised that for certain individuals different models will best explain their behaviour. Therefore, the parameter values for each model and the goodness-of-fit will be estimated for each model; showing the specification and suitability of each. The *allocation* model is as follows:

Figure 1.7: Optimal Allocation, x_1^*: Varying r and α

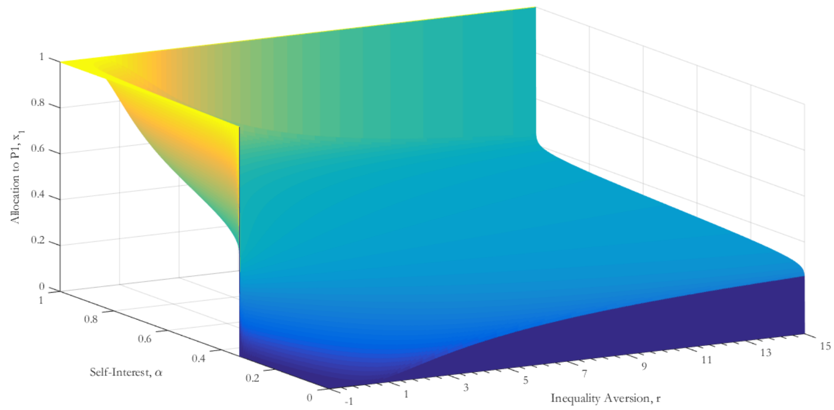

$$U_{2i} = \left(\sum_{i=1}^{N} \alpha_i (x_i)^{-r} \right)^{-\frac{1}{r}} \tag{1.6}$$

The corresponding optimal allocations are:

$$x_i^* = \frac{m}{1 + \sum_{j \neq i}^{N} \left(\left(\frac{\alpha_j}{\alpha_i} \right)^{\frac{1}{1+r}} \right)} \tag{1.7}$$

4 Error Modelling

The two alternative methods which can be used to establish preferences of individuals are: *elicitation* and *estimation*. The former assumes that individuals make deterministic choices, while the latter assumes they are stochastic. Deterministic behaviour means that participants make no error in their decisions, therefore any observed response can be used to *elicit* parameter values within a utility function. Stochastic behaviour, on the other hand, assume that individuals make some error in the decisions they are making. As a result, multiple decision problems need to be presented and a particular error model needs to be assumed in order to *estimate* an individual's preferences. The error model assumed here is a random behavioural model; where participants have some 'true' preferences which lead to an 'optimal' decision (which maximises utility). This 'optimal' decision, however, may be difficult to calculate and therefore error is made in its calculation. Estimation is used to establish the 'most likely' parameter values, given the observed behaviour of the participant.

Both Pairwise Choice and Discrete Slider designs assume participants make deterministic choices, but the Continuous Allocation design has been expanded to consider stochastic choices. To model decisions within the CA treatment we assume individual's actual allocations x_i are drawn from a Beta distribution, a flexible distribution constrained between 0 and 1, and specified by two parameters; (a_1, β_1). Individuals are assumed to have some optimal allocation, x_1^*, x_2^* and x_3^*, which correspond to the expected value of the Beta distribution. In this way parameters for the utility function and a precision term, denoting the extent of the error, can be estimated at both individual and sample level. This work is an extension of work done by Hey and Panaccione (2011) and Hey and Pace (2014) where the Beta distribution is used for similarly bounded decision problems, for ambiguous and risky decisions.

4.1 Continuous Allocation Stochastic Error

In order to model the decisions made for the Continuous Allocation design three assumptions are made, (1) $x_1 \sim B(a_1, \beta_1)$. Due to the nature of the continuous allocation game this is an appropriate distribution to assume. Like the experiment where allocations are constrained between two points (£0 and £30) the beta distribution is constrained between 0 and 1, therefore by normalising the allocations to between 0 and 1, the decisions made are likened to the beta distribution. The second assumption is that (2) $E(x_1) = x_1^*$, the expected value of x_1 is some optimal allocation x_1^* which corresponds to the optimal values established above. Furthermore, we assume (3) $var(x_1) = \frac{x_1^*(1-x_1^*)}{s_1}$ where more variance is observed further from the bounds, as there are more possibilities for error, and s_1 is a precision parameter; the higher the precision the lower the variance, and vice versa. Given these assumptions, and the form of the beta distribution:

$$E(x_1) = \frac{a_1}{a_1 + \beta_1} = x_1^* \tag{1.8}$$

$$Var(x_1) = \frac{a_1 \beta_1}{(a_1 + \beta_1)^2 (a_1 + \beta_1 + 1)} = \frac{x_1^*(1 - x_1^*)}{s_1} \tag{1.9}$$

Solving these equations, it follows that:

$$a_1 = x_1^*(s_1 - 1), \qquad \beta_1 = (1 - x_1^*)(s_1 - 1) \tag{1.10}$$

The above has established the model for x_1^*, however, the participants are making allocation decisions over three players. As a result we also need to model x_2^*. Rather than assuming that x_2 is distributed according to the Beta distribution we must, however, assume that it is $x_2/(1-x_1)$ which is distributed according to the Beta distribution $B(a_2, \beta_2)$. What this represents is the remainder that x_2 is, in comparison to x_3, from $m - x_1$. The allocation to Player 3, x_3^*, does not need to be modelled explicitly as it is simply the residual; $x_3 = m - x_1 - x_2$. This ensures the allocations are constrained to the budget, and therefore $\sum_{i=1}^{N} x_i = m$.

For x_2 we therefore have:

$$E\left(\frac{x_2}{1 - x_1}\right) = \frac{a_2}{a_2 + \beta_2} = \frac{x_2^*}{1 - x_1^*} \tag{1.11}$$

$$Var\left(\frac{x_2}{1 - x_1}\right) = \frac{a_2 \beta_2}{(a_2 + \beta_2)^2 (a_2 + \beta_2 + 1)} = \frac{x_2^*(1 - x_1^* - x_2^*)}{(1 - x_1^*)^2 s_2} \tag{1.12}$$

It follows that:

$$a_2 = \frac{x_2^*(s_2 - 1)}{1 - x_1^*}, \qquad \beta_2 = \left(1 - \frac{x_2^*}{1 - x_1^*}\right)(s_2 - 1) \tag{1.13}$$

Now that we have established the form for x_1 and x_2, the form of the log-likelihood function that needs to be maximised, with respect to the parameters α, r, s_1 and s_2, needs to be derived. As the decisions made by the participants are not truly continuous we need to evaluate the likelihood that the optimal value is between the possible options the participant had. The slider the participants could use was accurate to £0.01, as a result there is £0.005 either side of the allocations where their true decision could lie. As we have normalised the parameters to between 0 and 1, we also need to divide by m to ensure that it is consistent. Therefore, for x_1 the following needs to be established:

$$prob\left(x_1 - \frac{0.005}{m} \leq x_1^* \leq x_1 + \frac{0.005}{m}\right) \tag{1.14}$$

The log-likelihood function to be maximised for x_1, over each $t \in T$ rounds, with respect to the parameters is, therefore:

$$\sum_{t=1}^{T} log \int_{x_{1t}-\frac{0.005}{m_t}}^{x_{1t}+\frac{0.005}{m_t}} \left(\frac{\Gamma(a_{1t}+\beta_{1t})}{\Gamma(a_{1t})\Gamma(\beta_{1t})}\left(z^{a_{1t}-1}(1-z)^{\beta_{1t}-1}\right)\right) dz \tag{1.15}$$

For x_2:

$$prob\left(\frac{x_2 - \frac{0.005}{m}}{1 - x_1} \leq \frac{x_2^*}{1 - x_1^*} \leq \frac{x_2 + \frac{0.005}{m}}{1 - x_1}\right) \tag{1.16}$$

So our log-likelihood function becomes:

$$\sum_{t=1}^{T} log \int_{\frac{x_{2t}-\frac{0.005}{m_t}}{1-x_{2t}}}^{\frac{x_{2t}+\frac{0.005}{m_t}}{1-x_{2t}}} \left(\frac{\Gamma(a_{2t}+\beta_{2t})}{\Gamma(a_{2t})\Gamma(\beta_{2t})}\left(z^{a_{2t}-1}(1-z)^{\beta_{2t}-1}\right)\right) dz \tag{1.17}$$

'True' parameter values are estimated by searching for those parameter values which maximise the sum of the log-likelihood functions. This analysis will be conducted with MatLab, using *fmincon* (constrained non-linear minimisation). The parameter values estimated will be the 'most likely' parameter values that individual has, from the data gathered, given they made error according to the assumed distribution.

Intuitively, the log-likelihood function evaluates the probability that, with given parameter values, the *optimal* allocation lies between the feasible options the participant had each round of the game. By finding those parameter values which give optimal allocations which maximise the probability of the actual allocations being observed the model provides a 'best' fit. Goodness-of-fit measures can then be calculated to identify how well the model fits actual behaviour.

Figure 1.8 shows an example of the *actual* and *optimal* allocations, alongside *estimated* and *optimal* preference parameters for one individual over 18 rounds. We can observe from the bottom three panels the actual x allocations for each round. From this, optimal parameter values can be found for r and α through the maximum likelihood-estimation. Using these estimated parameter values the optimal allocations, found in the lower three panels, can be established. As a goodness-of-fit measure, by calculating the absolute difference between the optimal and actual allocations, it is possible to see how close the estimated parameter values are to the actual allocations. For x_1 allocations, for example, the average absolute difference between estimated and actual allocations is 0.0332, this corresponds to an average error in estimation 99.6p per round. Alongside the *estimated* r and α, the *optimal* r and α (see Equation 1.5) are shown in the top two panels. This highlights the importance of *estimating* parameters, rather than *eliciting* them; there is extreme volatility in the *optimal* parameters, if *deterministic* choices are assumed.

Figure 1.8: Goodness-of-Fit Example

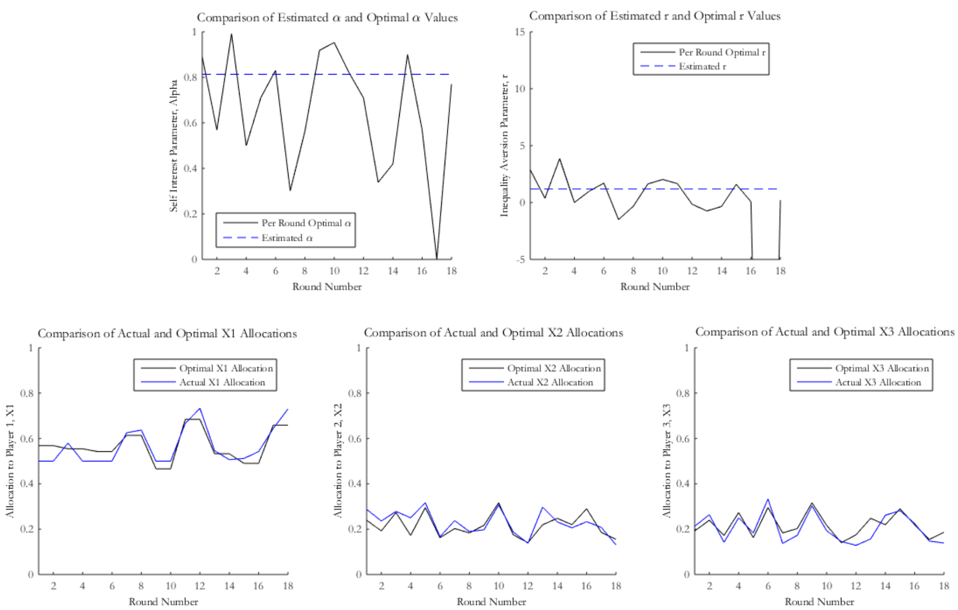

4.1.1 Two-Beta

In the extreme case where $\alpha = 1$, then $x_1^* = 1$, the above models will not suffice. As a result a more specific form of the above is necessary, whereby $\beta_1 = 1$. Due to the nature of the problem, the usual assumptions relating to the expected value and the variance are no longer correct. What is, therefore, done is to estimate a_1, this not only provides the shape of the pdf, but also a degree of precision for these estimates. For the sake of continuity $a_1 = s_3$. The higher a_1 the higher the probability that x_1 will be close to 1. Likewise we know the opposite is true as $x_2^* = 0$, for which we assume $a_1 = 1$. As before the aim is now to find the optimal β_2, or s_4, the measure of precision in actual allocations.

These extreme case models will be run parallel to the standard model, then for each individual the most suitable model will be used. The maximum likelihood estimates will be used to determine which model, normal or extreme the individual is most likely to conform to.

4.1.2 Beta-with-Bias

The second family of error models proposed are the Beta-with-Bias models. These models revert back to the standard error story, but a bias is introduced. For the first model, Beta-with-Equality-Bias we define: $x_i' = \frac{bm}{3} + (1-b)x_i^*$ where b is the bias parameter. This variable x_i' replaces x_i^* in the standard error model above. There is no bias if $b = 0$, but if $b > 0$ then bias exists and the optimal allocation is biased towards an equal distribution of allocations. For second model, Beta-with-Effort-Bias, we define: $x_i' = b\ddot{x}_i^* + (1-b)x_i^*$, where again b is the bias parameter, however we now have $\ddot{x}_i^*$ which corresponds to the optimal x allocation in the allocation model, specified above. Again there is no bias if $b = 0$, however if $b > 0$ then individuals make some decisions based on the simpler Allocation model, they do not consider the final payoffs but the allocations themselves. Intuitively the Beta-with-Equality-Bias model combines the complex optimal allocation decision the individual could make, with respect to their 'true' parameter values, with a basic heuristic to share their allocation equally between the three players. For the Beta-with-Effort-Bias the higher the bias the less 'effort' participants are inputting to calculate their desired distribution of payoffs, relying on the distribution of allocations instead. Both models will be used to estimate parameter values. The results will be compared and the connotations of the assumptions behind each discussed.

4.1.3 Local Maximums

One issue with the estimation procedures necessary for the error modelling described above is that starting values need to be specified in order to find optimal values within MatLab. As a result depending on the starting value used a local maximum may be found, rather than a global maximum. To overcome this issue multiple starting values will be used, each individual maximisation problem ran, then the result with the highest fval will be used as the global optimal result.

5 Results

The results shown below will be in the following order. First, the sampling and data will be explained. Then, the aggregated distributions of individual-level behaviour and preferences will be shown. Sample level treatment effects will then be analysed, alongside an in-depth analysis of the treatment effects on the Equally Distributed Equivalent (EDE). Cluster analysis will be presented, with preferences estimated for representative agents of each cluster. Throughout this analysis preferences are estimated within the *outcome* utility function, using the two-beta error story. The final section, tests between the *outcome* and *allocation* utility functions and the alternative error models; assessing the goodness-of-fit of each model.

For further analysis see the appendices for: disaggregated within-design comparisons (A.3), the pilot study, simulated comparisons of experimental designs and precision of estimated parameters (A.4) and the use of the beta-with-effort-bias model to detect effort (A.5). Alternative analytical techniques can be found to the EDE random-effects model (A.6) and finite mixture model (A.7). Sensitivity analysis is also conducted in A.8, on the assumption that $\alpha = 1/3$ and in A.9 on the distribution of payoffs assumed for the calculation of the EDE.

5.1 Data

The experiment took place in the EXEC laboratory, at the University of York. Using hroot (Hamburg Registration and Organisation Online Tool), randomised invites were sent out amongst a pool of 1866 users to acquire a sample of 30 participants. Three sessions were run and 15 users were invited to each session. In order to ensure correct group sizes, only samples which were multiples of 3 could be used. For the first session there were 10 users who showed up, so the experiment was run with 9 participants. For the second session 11 showed up and 9 participated, while there were 13 who showed up and 12 who participated in the third session. This resulted in a total sample size of 30. Details on the demographic characteristics of the sample are found in A.1.

5.2 Aggregated Individual Distributions

The distribution of the Proportional Payoff to Player 1 (PPP1), inequality aversion (r), self-interest (α) and Equally Distributed Equivalent (EDE) values, derived from individual participants decisions, are shown in Figure 1.9. While the r parameters and EDE distributions are from all treatment groups, PPP1 and α are restricted to particular sets. The PPP1 distribution contains only DS and CA data, from every round, as PC decisions are not comparable. The α distribtution is only from the PS perspective, as for VOI and IS α is assumed to be 1/3.

Perhaps the most intuitive of the results, the PPP1 shows the share of the total payoffs the participant distributed to themselves. When rounding to 0.01, the majority of decisions show participants are willing to share some of the budget, 81.97%, yet a substantial minority of responses, 18.03%, do exhibit purely individualistic behaviour. Those decisions which give Player 1 a third are 9.55%, with 32.42% taking less than 40% and 53.94% with taking less than 50%. With a mean and mode of 0.5562 and 0.4737, respectively, the decisions the participants have taken show two clear findings: participants are heterogeneous and the vast majority violate assumptions of pure self-interest.

For preferences in relation to inequality aversion this heterogeneity of behaviour is reflected. A median parameter of 2.5552 reveals an 'average' population preference reflecting 'Weighted Prioritarianism'. In terms of 'defined' preferences there are 6.67% of parameters established which reflect 'Utilitarianism', 4.81% 'Cobb-Douglas', 50.37% 'Weighted Prioritarian' and 30.74% who are extremely inequality averse, or 'Maximin'.[2] When considering self-interest, for the PS perspective, a mean of 0.7526 and median of 0.85, show an aggregate preference which weights the self higher than others. Disaggregating this 9.3% weight themselves equally and 23.26% have a weight of less than 0.5, however,58.14% weight themselves higher than 0.75, while 26.74% have an α close to 1.

Figure 1.9: Aggregate Distribution of PPP1, r, α and EDE

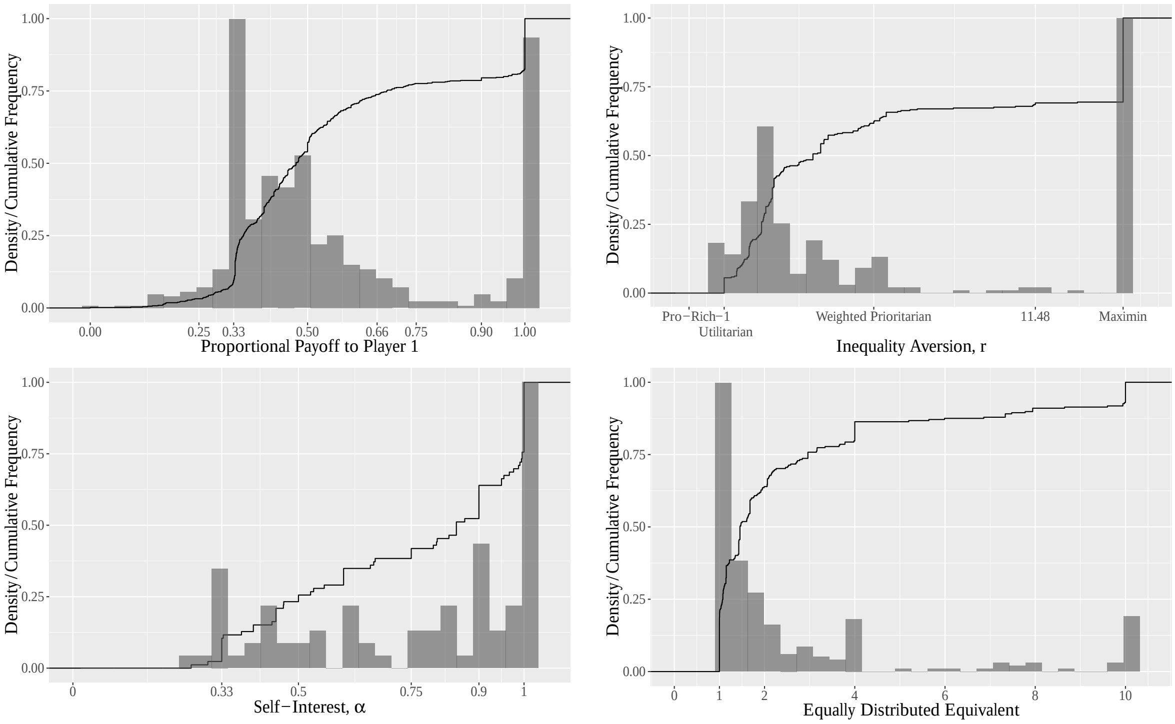

The distribution of EDE is shown in the bottom right graph. The EDE is useful as it is comparable across all treatments and avoids the issues that emerge from viewing parameter estimates individually, [3] This variable is equivalent to the utility values that each individual gets from the distribution of payoffs in a certain society. The society in question is arbitrary, so we have used an unequal society where the distributions are [10, 1, 1] to P1, P2 and P3 respectively (see A.9 for sensitivity). The EDE is a continuous variable from 1, where individuals are 'Maximin', to 10, where individuals are entirely self-interested. 22.96% of participants have an EDE valuation of 1, 63.04% have a value lower than 2, with 7% at 4 and 7.39% at 10. This distribution needs to be view with caution, however, as a large proportion of lower EDE's are due to the inclusion of the VOI and IS data, which has an assumed α, which skews the data. A more in-depth analysis, and resulting cumulative frequency plots, will be shown later.

[2]Those with an r greater than 15 were grouped.

[3]Estimations which show maximin preference often have misleading α parameters, as do α values of 1 for the inequality aversion parameters.

5.2.1 Self-Interest and Inequality Aversion Scatter

Within the Place-in-Society perspective a cross design comparison can also be mapped for the two parameters, r and α, together. Figure 1.10 plots both parameters, for different designs. For visual ease, if $r \geq 15$, it is simply plotted at 15. Similar patterns from the graphs above emerge, however, the pattern of responses and clustering become more apparent. For DS there appears to be clustering around lower self-interest and lower inequality aversion, while PC appears slightly more self-interested and more inequality averse. CA seems to be distributed more evenly, but has more points within the extremely inequality averse region. For the pattern of responses, it is clear that PC and DS were *elicited* values, while CA has been *estimated*, due to the rigid grid pattern of the former and non regularity of the latter.

Figure 1.10: Parameter Value Comparison Between Designs, PS

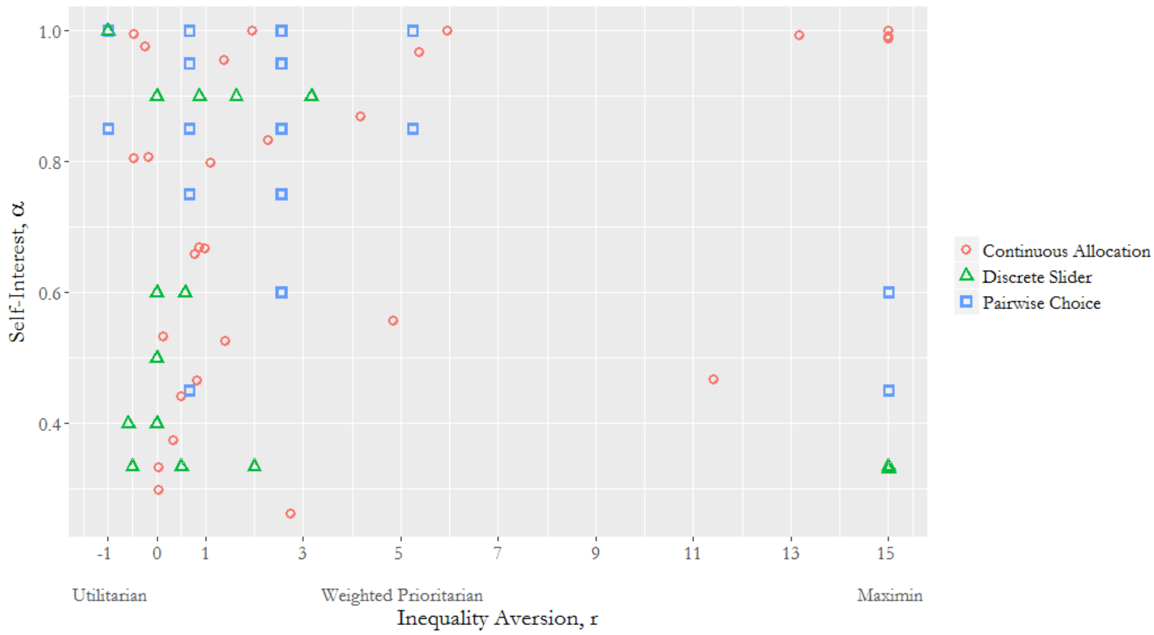

To summarise the elicited and estimated preference parameters, the median responses are as follows. Between alternative designs and perspectives the range of median estimates for *inequality aversion*, r, is from 0.5 to 3.590. This range is encompassed by the standpoint of a 'Weighted Prioritarian', where a higher weight is given to the worst-off. The range of median *self-interest*, α, is from 0.792 to 0.9. Showing that participants have other-regarding preferences; weighing their own welfare more, but willing to sacrifice some of their own payoffs to benefit others.

5.3 Treatment Effects

To delve deeper into the results, Table 1.2 shows multiple models which assess the treatment effects on the PPP1, α, r and EDE. The dependent variables are as before, with the exception of PPP1, which (to ensure comparability between designs) consists only of the first decisions taken, when $\pi = [1, 1, 1]$. With the exception of (3) each model is a Random Effects model (with robust standard errors). Model (3), with r as the dependent variable, is a Random

Effects Ordered Probit, as r had to be categorised due to issues of large and infinite values representing Maximin preferences.

The first model seeks to establish differences in pure responses, due to either perspective or design. The constant shows that, given the treatment is CA-PS, the average PPP1 is 0.5536. Both perspectives, VOI and IS, show a large and significant reduction of 0.2392 and 0.2112, respectively. There is no significant difference between the VOI and IS perspectives, however. Behaviour is significantly different when the participant knows who they are, but when they decide upon distributions when they are either in the society, but do not know who they are, or are not in the society, decisions made are similar. When changing from continuous to discrete, there is a significant positive effect ($p < 0.05$). A 5.07*ppt* increase, shows that when presented with discrete options participants appear less generous than the continuous case.

Moving to model (2) it appears that the self-interest is relatively robust with regard to experimental design. Self-interest is not significantly different between PC and CA, and only significant at the 10% level between PC and DS. Inequality aversion does, on the other hand, have significant treatment effects. Moving from PC to either DS or CA have significant positive effects at the 5% and 1% level receptively, meaning that participants are more averse to inequality. There are, however, no significant effects between DS and CA. In terms of the *perspective*, in the VOI participants show significantly more aversion to inequality, however, IS is only significant to the 10% level. There are also no significant differences between VOI and IS. In line with the higher magnitude for PPP1, this perhaps shows that when participants are within the society and know who they are they become less averse to inequality and more willing to trade-off equality for efficiency. The magnitudes shown here are not directly interpretable, however, as the model is an Ordered Probit, though it does give an idea of the significance and direction of effect.

Table 1.2: Random Effects Model: PP to P1, α, r and EDE

	(1) **PP to P1** Coef./Std. err.	(2) **Self-Interest** Coef./Std. err.	(3) **Ineq. Aversion** Coef./Std. err.	(4) **EDE** Coef./Std. err.
Perspective				
Veil of Ignorance	-0.2392^{***}		0.3897^{***}	-3.0730^{***}
	(0.044)		(0.151)	(0.512)
Impartial Spectator	-0.2112^{***}		0.2935^{*}	-3.0072^{***}
	(0.041)		(0.162)	(0.479)
Design				
Discrete Slider	0.0507^{**}	-0.0995^{*}	0.4296^{**}	-0.7717^{**}
	(0.022)	(0.054)	(0.200)	(0.333)
Continuous Allocation		-0.0694	0.5508^{***}	-1.2203^{***}
		(0.061)	(0.174)	(0.292)
Constant	0.5536^{***}	0.8106^{***}		5.4361^{***}
	(0.045)	(0.041)		(0.594)
N	30	30	30	30
Observations	180	86	246	257
R-Squared	0.2981	0.0284		0.3171
Model	RE	RE	RE-OP	RE

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Finally, to consider the EDE brings a more complete, and comparable, analysis to the decisions made. A lower EDE represents a lower level of self-interest and/or a higher aversion to inequality, and vice versa. The results from model (4) show the significant negative effects for perspective, with a slightly higher effect for VOI rather than IS. While this effect is apparent in this model, there are no significant differences between VOI and IS. Both DS and CA are significantly different from PC, 5% and 1% respectively, with CA having a greater negative effect (-1.2203) than DS (-0.7717).

These results show that both design and perspective are crucial considerations when conducting similar experiments. Place-in-Society treatments reveal much lower giving and lower levels of inequality aversion. So when estimating parameters for SWF's, or Atkinson Indices, this difference needs to be considered. While the results from the perspective were expected, the differences between design are somewhat more worrying. Although we assume participants enter an experiment with specific preferences, it appears that the design influences the parameters elicited/estimated. The PC design elicits lower aversion to inequality than dictator games. Even within the category of dictator game, whether the decisions made are continuous or discrete affects the decisions made and preferences established.

5.3.1 Equally Distributed Equivalent

To further demonstrate the treatment effects Figure 1.11 shows cumulative frequency plots of the EDE values for the sample. The three left-hand panels show differences in perspective, within each design, while the three right-hand panels show differences in design, within each perspective. There is a clear difference between the perspectives; while VOI and IS appear very close, within each design, Place-in-Society is consistently to the right. This is not surprising, and is partially driven by the assumption that $\alpha = 1/3$ within VOI and IS. What is of interest is that VOI and IS are very close together. Even though participants are within the former society, and not within the latter, participants appear to have similar preferences. When considering the differences between designs, given a perspective, there are apparent differences. While bias could be down to either design, the PC design appears to bias preferences towards selfishness, while CA appears to bias preferences towards 'Maximin'. In the PS perspective, for instance, those who are 'Maximin' in CA are slightly over 30% and just under 15% for PC. At the opposite end of the scale, less than 20% are totally self-interest in CA, but this reaches almost 40% in PC. This pattern also holds true, though perhaps to a lesser extent, for both VOI and IS.

One observation of note, is that within the VOI and IS treatments, the difference appears to be mainly driven by the *coarseness* of the PC design; indeed simulations in Appendix A.4 show that the PC treatment underestimates r (in the PS perspective), which would lead to a higher EDE. Had a more *fine* set of choices been made available perhaps the differences may not have been as large.

The observations in this subsection, and that above, are further confirmed by the Skillings-Mack results, in Appendix A.6. The results more closely dissect the EDE results, without needing to make the assumptions of normality the random effects model requires. As a complete model, for design effects within each perspective and for perspective effects within each design the Skillings-Mack results generally show that the PC design and the PS perspective elicits/estimates higher EDE values. With results significant in all but the between designs in the IS perspective.

Figure 1.11: EDE Differences in Perspective and Design

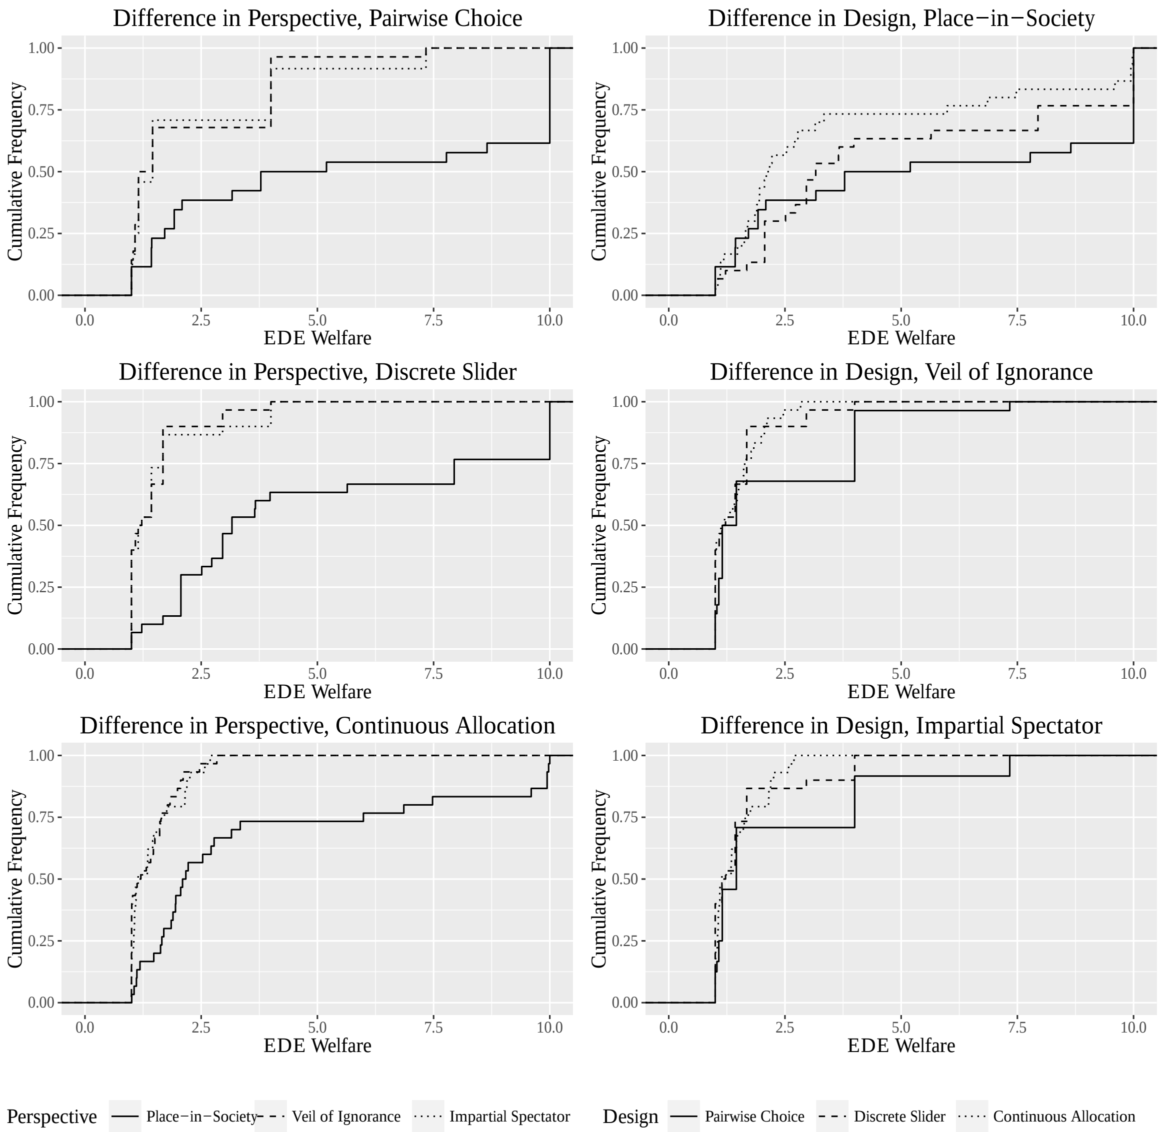

5.3.2 Questionnaire Responses

Further analysis was conducted to establish an individual effects from the Demographic characteristics, Opinions and Cognition, from the questionnaire. For both PPP1 and EDE there were no significant effects for any characteristic. For α and r there were no significant effects at the 5% level, however two variables were significant at the 10% level. For α if the participant was female they were, on average, more self-interested (0.1860). For r females and those who valued reducing income inequality rather than increasing national income were more averse to inequality. Age, nationality, education, parental income, political preference, views on reducing inequality and cognitive error were all insignificant at the 10% level. These results are, however, presented with caution, as the sample size is relatively small.

5.4 Cluster Analysis and Finite Mixture Model

While data can be analysed on a sample level, taking account of averages, in this context a cluster based approach may be appropriate. When considering that individuals have heterogeneous preferences in relation to inequality aversion and self-interest, one may also propose that there are certain *clusters*, or types, of individuals who hold more similar preferences to those within their cluster in comparison with those outside their cluster. If this were to be the case, then exactly who belongs to which cluster, and how many clusters are in the sample needs to be established.

Proportional payoffs to players 1, 2 and 3 will be used from the CA, PS, treatment in order to identify any clustering of payoff allocations between subjects. Three methods have been utilised, Hierarchical Clustering, K-means and Finite Mixture Model Clustering; the latter is presented here, the former two in Appendix A.7. These methods ascertain how these clusters may be composed and propose an optimal number of clusters within the sample.[4] The analysis will be run using *r*, where the commands used are predefined (and relatively straightforward) commands to use.

A Finite Mixture Model (FMM) proposes a 'mixture' of individual density functions which accurately fit the data. Here the standard form will be used; where a set of normal densities, with varying variance and mean, will be chosen to most accurately fit the data. The form of the finite mixture model, assuming we have a mixture of Gaussian densities is:

$$f(x) = \sum_{g=1}^{G} \pi_g f(x|\theta_g) \tag{1.18}$$

The mixing proportions of each of the normal densities is denoted by π_g, it shows that a particular proportion of the population fits within that corresponding density, or cluster. Each of these individual clusters are defined by the function representing that density. In order to establish the parameters of the model a maximum likelihood estimation is run, using the Expectation-Maximisation Algorithm . The *r* code 'mclust' was used to establish the parameters of the model; enabling a clustering of observations which was the 'most likely'.

The results find that a two cluster solution is optimal, that the split was a 22:8 split between the two clusters. For those separated into either cluster the probability that they should belong to their cluster was 1.000 for every individual, with the exception of individual 9 for whom the probability was 0.985. Such high individual probabilities show that the clusters are very well defined and that individuals 'belong' to that cluster. The mixing proportions, that show the probability that the individual belongs to either cluster, are 0.7329 and 0.2671 respectively. The mean proportional payoff that each cluster can be shown to be [0.4256,

[4]I am very grateful to Nema Dean who gave a workshop on Cluster Analysis Using R, for the Applied Quantitative Methods Network at the University of Sheffield. The methods used are closely intertwined with what she taught.

0.2871, 0.2873] for P1, P2 and P3, respectively, within the first cluster and [0.8905, 0.0557, 0.0538] for P1, P2 and P3, respectively, within the second. The split populations could be seen as a more altruistic group, 73.29% of the population who gave on average 57.44% to others, and a more self-interested group, 26.71% of the population who only gave on average 10.95% to others.

Table 1.3: Finite Mixture Model Clustering Results

		Mean Proportional Payoff		
	Mixing Proportions	Payoff to P1	Payoff to P2	Payoff to P3
Cluster 1	0.7329	0.4256	0.2871	0.2873
Cluster 2	0.2671	0.8905	0.0557	0.0538

In order to visualise the clusters which have been established Figure 1.12 is a scatter plot and accompanying kernel density plots which show the Proportional Payoffs between P1 and P2 for each individual. Those observations highlighted in blue represent Cluster 1, while those in red represent Cluster 2. A clear separation between the observations appears within the scatter plot. Note that the Proportional Payoff to P3 is simply the residual of the two from one. There is a very distinct trend which emerges from this data, that while individuals may differ on how much they allocate to themselves, invariably they aim (on average) to share the remaining payoffs equally between the other two, which is in line with the theoretical prediction of the utility model.

Figure 1.12: Clustered Proportional Payoff: Scatter and Kernel Density Plots

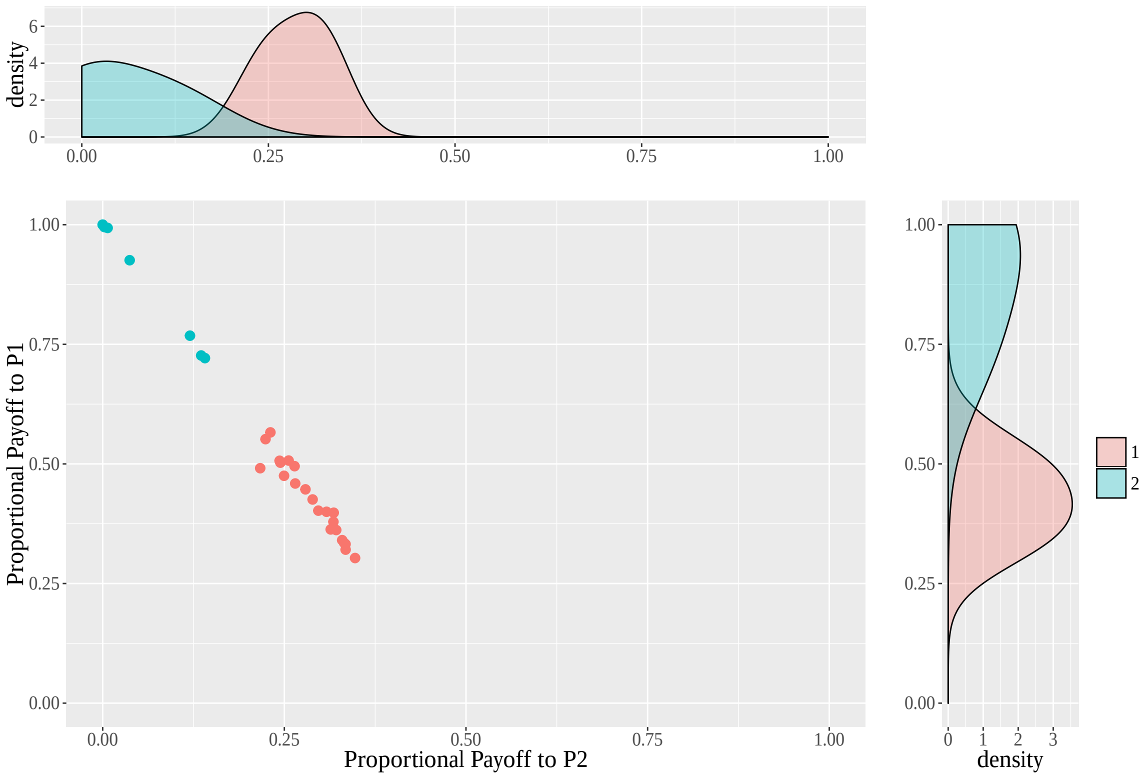

If this clear separation into two clusters is to be observed then, given this analysis has been done, 'social preferences' can be established for each of the clusters. Optimal inequality aversion and self-interest parameters can be established for the 'representative agent' of each cluster, then instead of using median preferences of the entire sample (which may be sub-optimal for the majority of the population) these two clustered social preferences could be used, with the mixing proportions to advocate policy implementations.

Table 1.4: Cluster Estimates of Parameter Values and Goodness of Fit

			Average Goodness of Fit		
	r	α	P1	P2	P3
Cluster 1	1.5445	0.5818	2.5427	1.6736	1.7453
Cluster 2	-0.1262	0.8753	4.2576	2.1474	2.1918
Sample	0.613	0.8327	7.7383	3.9391	3.9555

Table 1.4 shows the parameter estimated and average absolute goodness of fit statistics for each cluster separately, and the entire sample. The analysis is run by using each decision, for each individual, (i.e. 18 rounds for all individuals within a cluster), and the resulting r and α parameters are those estimated to be the most likely. The goodness of fit measures show the absolute difference between the actual allocations and the optimal allocations (according to the estimated parameter values) averaged for each observation. When considering the entire sample there is an inequality aversion parameter of 0.6130 and a self-interest parameter of 0.8327, meaning there is some weighting on the income of the worst-off, but significant levels of self-interest. When moving to Cluster 1 there is an increase in the r and a decrease in α, which is somewhat intuitive; those who are allocating on average 0.4256 to themselves care more about inequality and less about their own payoffs. Cluster 2 shows the opposite pattern, less aversion to inequality (hence more emphasis on efficiency) and higher self-interest. The estimation procedure is most accurate for Cluster 1, where the average GOF for Player 1 is 2.5427; meaning that on average the difference between the actual and optimal x allocations is £2.54 out of the budget of £30. This figure falls for both P2 and P3, and is relatively similar for both. The GOF worsens when considering Cluster 2 across each of the allocations. This is perhaps due to the switching behaviour of the two individuals who have a certain proportion of their allocations to P1 at 1, and the rest around 0.33. As expected when looking at the full sample, the GOF worsens again.

When these estimated parameters are used to establish the 'optimal' allocations (when $\pi = [1, 1, 1]$) these allocations can be compared with the actual average allocations from each cluster. For Cluster 1 the 'optimal' allocations are [0.4278, 0.2861, 0.2861] to Players 1, 2 and 3, respectively, while for Cluster 2 they are [0.9113, 0.0443, 0.0443]. When this is compared to the mean response we find that for Cluster 1 the allocations are very close, while for Cluster 2 the 'optimal' allocations are slightly lower.

5.5 Individual Level – Continuous Allocation

While the previous section compares specific treatment effects on a sample wide level, this section seeks to determine to what extent the *stochastic error models* specified explain the behaviour of individuals throughout the experiment. The CA design was created specifically to be able to test between different utility and error models. Due to the repeated nature of the questions, *estimation* procedures, rather than *elicitation* methods can be used. To simplify this comparative task only the PS perspective will be used; this was one major aim of the study and has the most rounds making the estimation procedure more accurate. Firstly, comparisons can be made between the two utility models, *outcome* utility and *allocation* utility. Secondly, comparisons can be made between the two error stories, the *Two-Beta* model and the *Beta-with-Bias* models.

Figure 1.13 shows a comparison of the parameter values between the different models. The allocation model shows the most divergence from the other models. For the three models which incorporate the outcome, many have identical values (usually those with zero bias in their responses) while a trend appears that those not identical have less aversion to inequality. This latter trend can perhaps be explained, as a bias towards equality would be consistent with a higher inequality aversion.

Figure 1.13: Comparison of Estimated Parameters; Utility Models and Error Stories

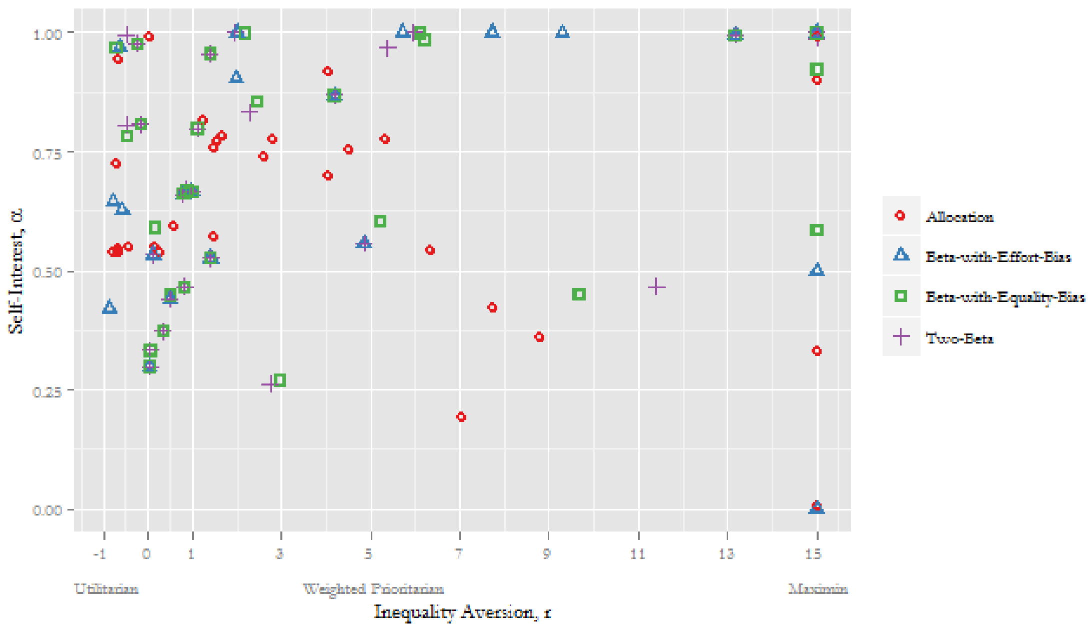

5.5.1 Goodness of Fit

In order to establish which of the models best fits the responses made the goodness of fit measures from each individual can be compared. Several measure will be shown here: the Euclidean Distance (ED), Maximum Log-Likelihood (MLL) and the corrected Akaike information criterion (AICc).[5] The ED established how close the optimal allocations, x_i^*, are to the observed allocations, x_i, while the MLL establishes how likely the observed behaviour is, given the estimated preference parameters and error model proposed. The AICc is a statistic used for model selection; using the MLL but penalising models which use more parameters. Each measure is calculated for each individual, for each utility function and error model. Comparing the measures allows an analysis of how well the models fit the data, relative to one another.

Figure 1.14 shows the distribution of GOF measures for the entire sample (excluding MLL as visually it is very similar to AICc). The left panel show the mean ED, for each individual, where a lower value shows a better fit. The distribution shows that each the Two-Beta, Beta-with-Effort-Bias and Beta-with-Equality-Bias, are very close. The allocation model does worse for the majority of the sample, meaning that most participants, appear to consider the problem in its entirety, they consider the total payoff to each of the individuals, and distribute accordingly. The Two-Beta egoist model, however, fits much worse, for the majority of the sample, with the exception of the few individuals who are distributing payoffs primarily to themselves. The right panel shows the distribution of the AICc, which shows a similar trend. Notice here, however, that the Two-Beta egoist model dominates for individuals at very low levels of AICc, these two individuals are those who allocated everything to themselves each round. Note also, that the Two-Beta model appears to be doing slightly better than the Beta-with-Bias models, this is due to the additional parameter the Beta-with-Bias models include.

Figure 1.14: Comparison of Utility Functions and Error Models; Goodness of Fit

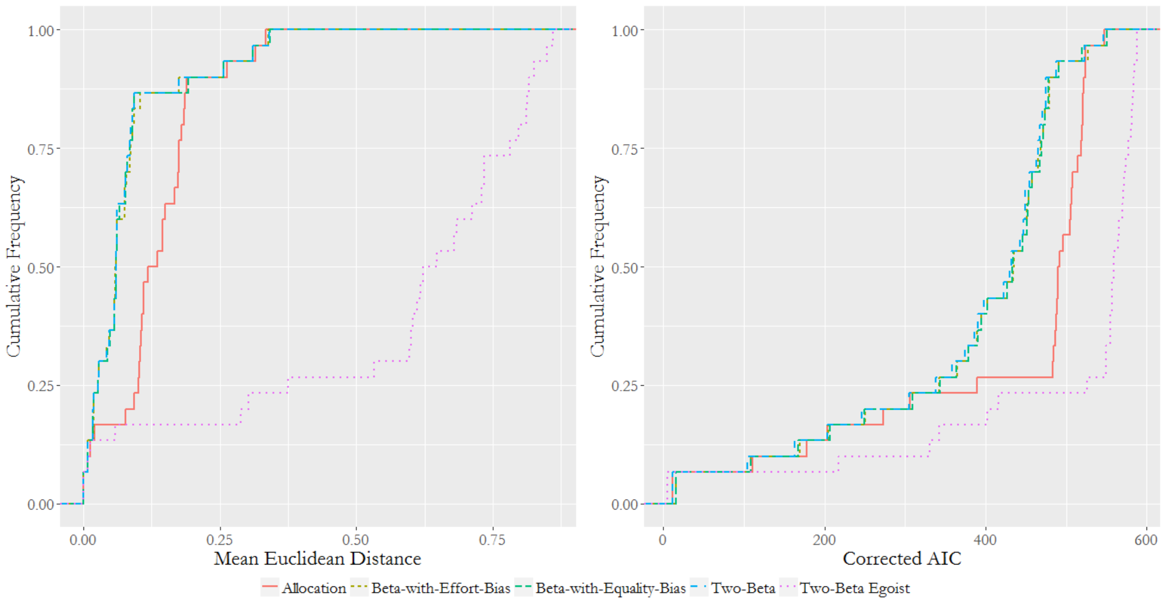

Table 1.5 summarises the three GOF measures, showing the sample mean values for each model and the number of 'types' within each model, where the GOF measure of that model

[5]The Euclidean Distance is calculated as: $ED = \sqrt{\sum_i^N (x_i^* - x_i)^2}$. The Maximum Log-Likelihood is the sum of the maximised log-likelihood values in Section 4.1. The corrected Akaike information criterion is calculated as: $AICc = 2k - 2(MLL) + 2k(k+1)/(n-k-1)$ (Sugiura, 1978).

is dominant. The sample mean ED is similar for the Two-Beta and Beta-with-Bias models, with larger values for the allocation model and even larger for the Two-Beta egoist model. Similar patterns emerge for the MLL (with a higher value denoting a better fit), but results from the AICc show that the Two-Beta model seems to do better than the Beta-with-Bias models. Indeed, when considering 'types' for both ED and MLL there are a mix of individuals classified into each model. With the Beta-with-Bias models performing similarly to the Two-Beta models. However, when the AICc is considered the small additional benefit that the Beta-with-Bias models offer is not large enough to offset the additional parameter needed. The majority of the sample are 'best' characterised by the preference estimates within the Two-Beta model, with only one participant classified as a Beta-with-Equality-Bias type.

Table 1.5: GOF Summary of Alternative Utility Functions and Error Models

		Two-Beta		**Beta-with-Bias**		**Allocation**
		Standard	Egoist	Effort	Equality	
Mean	ED	0.0788	0.5570	0.0800	0.0798	0.1328
	MLL	-180.94	-243.69	-180.90	-180.74	-204.71
	AICc	372.95	492.18	376.81	376.48	420.95
Type	ED	9	6	9	4	2
	MLL	8	2	8	12	0
	AICc	27	2	0	1	0
	N	30	30	30	30	30

6 Discussion

6.1 Difference in Design

An important question to ask when conducting laboratory experiments is: to what extent does the specific design of the experiment affect behaviour? The majority of experiments focus upon one design, and therefore risk not knowing if that design has any inherent biases. By posing this question and presenting respondents with three different designs - Pairwise Choice (PC), Discrete Slider (DS) and Continuous Allocation (CA) - this paper aims to identify potential biases, strengths and weaknesses of each design.

From our results differences between designs are apparent. Preference parameters and EDE values can be compared between all treatments, while raw decisions (PP to P1) can only be compared between DS and CA (see Table 1.2). On average participants allocate more to themselves in the DS, than in the CA. Self-interest, α does not vary much between designs; only parameters in DS are significantly lower than in PC, and this only at the 10% level. When considering inequality aversion, r, and EDE significant differences emerge. Inequality aversion is significantly higher, and EDE significantly lower, in both DS and CA in comparison to PC. This reveals that participants are choosing distributions which favour the worst-off to a greater extent in the DS and CA treatments. The Skillings-Mack results show similar patterns, for the EDE. Showing that the rank of EDE is lower in the DS and CA, compared to PC, for each perspective. Though differences are not significant in the impartial spectator perspective.

Further to experimental results, simulations were conducted, assessing if the biases between the designs were due to factors inherent in the design. This was conducted for the Place-in-Society perspective, as it was the most susceptible to differences in design and allows for differences in self-interest to be compared. First, *precision* was analysed, where CA was shown to allow a much *finer* parameter estimation, while PC and DS were more *coarse* elicitation methods. Indeed this distinction proved to be important when considering 'deterministic' choices simulated individuals would take. When no error was made by the individuals, their 'true' parameter values were more precisely captured, in CA, compared to the other designs. However, when a 'random preference' error was added to the individual's preference parameters, CA gave more biased results. The discrete nature of the choices in the PC and DS designs mitigate the effects of such an error.

A further difference between the designs is that of *elicitation* (PC and DS) vs *estimation* (CA). The choice between the two primarily depends on the trade-offs which the experimenter needs to make between *time*, *precision* and concerns of *noise*. *Elicitation* techniques assume that participants make no error, and therefore only need one set of decision problems, while *estimation* assumes decisions are *noisy* and usually require a large battery of decision problems.

Assume all designs were *elicitation* designs. In terms of *time*, PC requires multiple pairwise choices to be presented (eight or sixteen in our case), while DS and CA need only one or two. As discussed above the discrete nature of PC and DS leads to less *precision* in parameter estimates, in comparison to CA. Yet, *noise* can manifest itself in several ways. If *noise* exists and is (as in Appendix A.4) due to 'random preference', then the PC and DS are preferred to CA. However, participants could also make 'random behavioural' error, which in the case of PC could lead to incoherent 'switching', meaning that no parameters can be elicited (as was the case for 8 of 30 for r in PS). So the issues for PS are *time*, *precision* and 'random behavioural' *error*, for DS *precision* and CA 'random preference' *error*. If, alternatively, *estimation* techniques were used then all models would require a longer *time*, due to large batteries of problems being needed, but they would reduce issues of *precision* and allow *noise* to be effectively modelled, and its bias reduced.

It appears, therefore, that the appropriateness of the design depends on the context. Alternative designs can be used depending upon the relative importance the experimenter places on *time*, *precision* and *noise*. Both *elicitation* and *estimation* methods can also be used for each design shown here, allowing for additional flexibility. While experimenters have to be wary that differences in design could have significant effects on behaviour and estimation (or elicitation) of preferences, simulations can be ran for alternative designs to identify such issues without using valuable experimental funding.

6.2 Alternative Perspectives

By incorporating alternative perspectives: Place-in-Society (PS), Veil of Ignorance (VOI) and Impartial Spectator (IS); the following questions were sought to be answered. To what extent does the perspective of the individual affect their value judgements? Are preference parameters significantly different when the participant is asked to consider different perspectives? This question is intertwined with the more normative judgement of which perspective should be used for such distributional decision making. While this analysis is a positive one, the results from this chapter could inform such a normative discussion.

Unlike when comparing experimental designs, the raw responses can be used to identify differences between perspectives, within designs. When considering PC, individual's decisions directly map onto inequality aversion groups, see Figure 1.18. While differences between the median groups are not very apparent, the median response is either 'Weighted Prioritarian 1' or '2', the distribution of responses changes somewhat between perspectives. While neither perspective have any individuals willing to violate monotonicity, and a similar proportion of individuals are 'Maximin', differences appear at lower levels of inequality aversion. While only 4.5% are 'Utilitarian' in PS, there are 20.8% and 28.6% 'Utilitarian's' and even 8.3% and 3.6% 'Inequality Seeking' individuals in IS and VOI, respectively. Interestingly, some individuals appear to prioritise efficiency, rather than equality, when the perspective is more impartial.

When moving into the DS comparisons are somewhat different, here there is a stark contrast between PS and the other perspective, see Figure 1.19. More equal allocations emerge from IS and VOI, with 40% of subjects distributing payoffs equally in both perspectives, in comparison to the 0% in PS. In CA the trend is similar, distributions in VOI and IS are far more equal (see Figure 1.20); with a median response of 0.46 in PS, and median responses of 0.33 in both VOI and IS.

The results of the random effects model in Table 1.2 results show that inequality aversion is significantly higher in both the VOI and IS (10% level), while EDE is significantly lower for both VOI and IS. The Skillings-Mack results (see Table 1.12) confirm this for EDE, for both aggregate and within design. While the difference in EDE is not surprising, due to the assumption that $\alpha = 0.33$ in VOI and IS, the difference is partially driven by those significant differences in r. The hypothesis that $\alpha = 0.33$ is not rejected for VOI and IS (see Appendix A.8), and indeed is significantly different from the estimated α within the PS, with a mean of 0.74.

Perspective is therefore important, it has large effects on the distributional decisions that individuals make. Individuals allocate far more equally when *under the veil* or as an *impartial spectator*. Across all designs participants appear to be less averse to inequality when they know who they are. Furthermore, the estimates of self-interest are not significantly different from 1/3 for the VOI and IS perspectives, but are significantly higher in the PS perspective.

6.3 Individual Level Prediction

Within the structural approach to experimental economics, assessing how well the proposed model, or models, fit to the data is important. Here there are two proposed utility functions, *outcome* utility and *allocation* utility, and two families of error model, *two-beta* and *beta-with-bias*. The CA PS treatment was chosen to test between utility functions and error models. Results show that the *outcome* utility model strictly dominated the *allocation* utility model, for AICc statistics. Individuals appeared to fully consider the final payoffs, rather than simply distributing the allocations. In terms of error models the *two-beta* model appeared somewhat better; however, the *beta-with-bias* models perform similarly, but are penalised by having an additional preference parameter. Between the utility functions the *outcome* model dominates, however, with regard to the error models the choice is debatable.

Due to analysis in Section 5.5.1 the Two-Beta model was chosen for the main analysis. Primarily this was due to higher AICc values, which selected the model as best for all but one participant. Although the Beta-with-Bias models are more general, they were not able to explain behaviour better than the Two-Beta model. An additional concern is that the bias parameter may alter the other estimated value judgements. While the Two-Beta model does little to affect the actual parameter values, r and α, by design the Beta-with-Bias models affects the parameters. Individuals with true α values of 1, will be shown to have an α of less than 1 for instance. As these values are meaningful in terms of the philosophical standpoint which they represent the Beta-with-Bias models are, perhaps, less desirable. For the Beta-with-Equality-Bias model specifically, the assumption is that individuals are biased towards an equal allocation, the issue is this equal allocation may correspond to a set of parameter values; which, with the inclusion of the bias parameter, will not be estimated.

There are, however, several benefits of using the Beta-with-Bias models. The first of which is related to the Beta-with-Effort bias model. Behaviourally it is appealing, in the sense that it can be proposed that by linking the *outcome* and *allocation* utility models a proxy of effort is received. Individuals who put in more effort consider the payoff outcomes each round, even though it is more demanding in terms of cognition and effort, while those who only consider the allocations put in less effort. The higher the bias towards this allocation utility function, the less effort is put in. Indeed, with the regression in Appendix A.5 there is shown to be a significant relationship. The models also smooth the effects of those individuals who switch between entirely selfish allocations and sharing allocations. While the Two-Beta model may be able to estimate error for them, it does have to assume they are entirely self-interested to begin with, while the Beta-with-Bias models can begin the estimations without imposing that assumption.

6.4 Comparisons

6.4.1 Social Welfare Functions in Health

There are a number of studies within the health economics literature which utilise a Social Welfare Function, from which an inequality aversion parameter is elicited. A. Williams, Tsuchiya, and Dolan (2004), Tsuchiya and Dolan (2007), Dolan and Tsuchiya (2009), Dolan and Tsuchiya (2011) and Robson et al. (2017) conducted questionairre-based surveys, using a similar Pairwise Choice design, in England, while Abasolo and Tsuchiya (2004), Abasolo and Tsuchiya (2008) and Abásolo and Tsuchiya (2013), conducted such surveys in Spain. Each of these studies revealed that the majority of the population were averse to inequality. In most studies a comparable median inequality aversion parameter was not found, either as it was not reported, or was not comparable. Median values of 27.9 and 9.95 were reported in Dolan and Tsuchiya (2011) and Robson et al. (2017), respectively. From two of the studies in Spain, Abasolo and Tsuchiya (2004) and Abásolo and Tsuchiya (2013), the median response was found to violate monotonicity, meaning that an r parameter would not be identifiable as they were more inequality averse than Maximin preferences. The most comparable median response estimated here, from the Pairwise Choice design and Impartial Spectator perspective, was 1.9. This parameter value is significantly less than the elicited value from the other studies, meaning that our sample was significantly less averse to inequality.

A more in-depth comparison, concerning the entire distribution of responses can be conducted using data from Robson et al. (2017). The format of the survey is comparable to that used here, however, there are key differences. As mentioned above, the dimension is health, rather than income. Participants were not given monetary incentives for each question, but paid a set turn-up fee. While there were similarly 8 rounds of Pairwise Choices, those in the questionnaire had the option of 'Programme A and Programme B are Equally Good' in addition to the choice of Programme A or Programme B.

For the purpose of comparison the *health questionnaire* results have been split into various samples. The first is the 'paper' based questionnaire, where members of the public attended a day-long session at the University of York and invigilators were present to explain, and form discussion groups, in relation to the questionnaire. The bottom two bars were both done in the form of an *online survey*, hosted by Smart Survey. They have been split into a set of *student* only responses and the larger *general population* responses.

When comparing the results there are several differences. The median responses show higher inequality aversion in each of the health questionnaire results ($r = 9.95$) compared to the experimental results ($2.5552 <= r <= 2.8702$). The distribution of responses at the extremes is also significantly different. While not one individual was classed as an 'Egalitarian' (and hence violated monotonicity) within the experiment between 14.3% and 34.2% of the health questionnaire responses did. At the opposite extreme 20.8% and 28.6% of participants were 'Utilitarian's' in the Impartial Spectator and Veil of Ignorance perspectives, respectively, compared to 3% of in the general population, and 0% in the other samples. Arguments may emerge that the population sample may be that which effected the differing distributions. However, as Figure 1.15 highlights when considering only students within the health sample the difference is even more stark.

There are two potential reasons for this; domain and method. The first, domain, refers to the difference between health and income. Individuals could be more averse to inequality within the dimension of health. It could be seen as a necessity in life, and something which

Figure 1.15: Distribution of Inequality Aversion for Alternative α Values

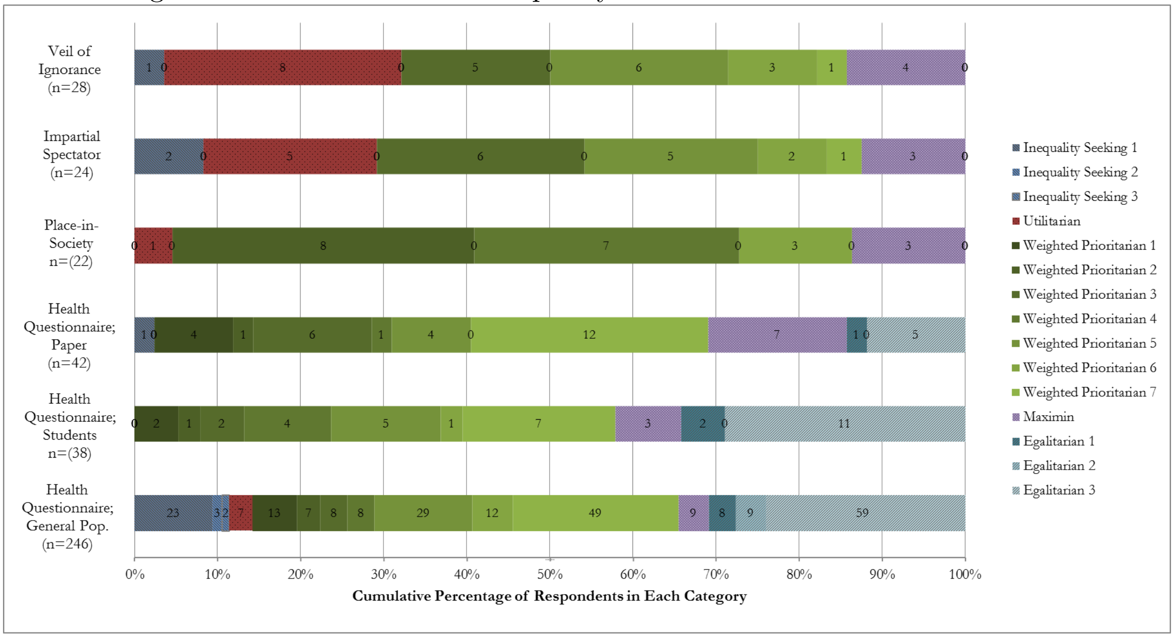

should not be unequally distributed. The second factor refers to the difference between incentivised laboratory experiments and questionnaire-based surveys. As discussed in Section 1.1, the two strands of literature, Distributional Preferences and Social Welfare Functions, use these different methods and each debate the other's method. Without straying into the debate, it is clear that these differences could also be responsible for the disparity in estimates. However, one aim of this experiment was to mitigate the arguments against the experimental method. By utilising the Impartial Spectator perspective the concern that individuals are primarily concerned with 'sums of money they can bring home after the experiment' is mitigated. The underlying issues of inequality aversion are hopefully assessed, while the incentive compatible experiment hopefully mitigated the biases and errors that questionnaire responses may contain. If this is the case then the differences emerge from either the potential biases within the questionnaire method or the differing aversions regarding the domain. Indeed, when formulating the Atkinson index (in the context of income) Atkinson (1970) took values of between -1 and 1.5 when proposing the index. Whilst not based on experimental evidence, this intuitive range is not too dissimilar to the elicited value of 1.9

6.4.2 Dictator Games

Engel (2011) performs a meta-analysis of dictator games, which allows for a comparison of giving in other experiments. From 616 treatments, the 'grand mean' showed that dictators were willing to give 28.35% of their share on average. From our accumulated responses the mean PP to P1 was 0.56, giving a share of 44.02%, which is a significantly higher share than in absolute terms. Yet this share is 22.01% to each of the other players. Indeed, Engel finds that giving is significantly higher in multiple recipient dictator games, however, his results suggest that giving per person actually increases.

Consistent with our results is a large degree of heterogeneity. From the 328 treatments that gave full distributions, 36.1% were entirely self-interested, 16.74% shared equally, 34.09% gave less than equal but more than none, and 13.07% gave more than half their share. From our results, in relative terms to the three players, 17.67% of the responses were entirely self-interested, 9% shared exactly equally, 65.32% gave less than equal but more than none, and 8.02% gave more than half their share. There are roughly half the number of self-interested responses in our sample, less *equal sharers* and fewer *more than equal sharers*.

Cappelen et al. (2007) propose that individual's fairness ideals are heterogeneous, proposing three such ideals: Strict Egalitarianism, Libertarianism, and Liberal Egalitarianism. Where strict egalitarians dislike unequal outcomes, libertarians believe individuals should give exactly what they produce and liberal egalitarians should be held responsible only for their choices. A structural model is estimated and they find that the proportions of the above three fairness ideals are 0.435, 0.381 and 0.184, respectively. Their experiment was a one-shot dictator game, with factors that were both within and outside the individual's control. The mean share the dictator offered to the other player was 27.1%, and the median share 29.2%.

Andreoni and Miller (2002) have an approach somewhat similar to ours. They consider a CES utility function and vary prices, while ensuring individuals are subject to a budget constraint. Again, they conclude that individuals are heterogenous and, as in Cappelen et al. (2007), group individuals into specific groups. While no aggregate parameter values are found for their CES function for the sample as a whole, the distribution of respondents who fit into each specific group is comparable. They find that 43% of their subjects adhere strongly to either, Purely Selfish, 22.7%, Maximin, 14.2%, or Utilitarian, 6.25%, preferences. The remaining subjects weakly adhere to these utility functions, and they utilise the CES function to predict similar self-interest and inequality aversion parameters. Converting our results into comparable groups we find approximately between 12% and 26.67% of respondents adhere to Maximin, 13.33% to Purely Selfish and 0% Utilitarians. As with other comparisons the number of Purely Selfish individuals seem somewhat lower, while our Maximin responses vary depending upon the r cut-off the range does span that of Andreoni and Miller's value. There are, however, zero Utilitarians within our sample.

While there are a whole host of experiments which consider similar dictator games, there are several conclusions which appear throughout the literature. Firstly, that the self-interest hypothesis, that all individuals are solely motivated by their own payoffs, is consistently refuted E. Fehr and Schmidt (2006). Indeed, many individuals are motivated by other-regarding preferences, of some form. Second, preferences are heterogeneous, far from conforming to a specific social norm, or distinct preference. The implications of this affect the majority of economics, not only do these models better explain behaviour but point to alternative methods of incentivising individuals in the real world.

6.4.3 Ugandan General Population and British University Students

One contentious issue with traditional laboratory experiments is the demographic from which the sample is taken. The samples are usually made up mostly, or entirely, of students who themselves are a specific sub-population and are not representative of society as a whole. The reason for this sample is more often than not due to cost and convenience; students are readily available and are perhaps more willing to be rewarded with less money than non-students may be. When considering preferences over elements which may have policy implications a representative sample is often desirable. One way in which to do so, which will be explained below, is to use samples from developing countries. These samples can consist of representative samples and the compensation needed is far less in relative terms.

A pilot study (which was a precursor to Chapter 2) was ran in the Mbale region of Uganda with 24 respondents. The design was similar to that of the Continuous Allocation, for Place-in-Society. The numbers were made slightly simpler, to deal with poor numeracy skills, but the principle remained similar. Rather than a computerised laboratory it was run with paper and pen. Participants moved physical tokens into boxes to denotes their allocations to P1, P2 and P3 respectively. Each box had different multiplication factors each round. Facilitators were present to record responses, and to move to the next round of the game. Apart from these aspects the experiment was kept relatively close to the experiment in the UK, to allow for comparison.

Figure 1.16 shows the cumulative frequency plot of EDE for both samples. It is clear that, while for three quarters of the population the EDE values are very similar, there is a substantial difference for the higher EDE values. The University of York sample appears to have a greater proportion of more self-interested individuals that the Ugandan sample.

Figure 1.16: Distribution of Inequality Aversion for Alternative α Values

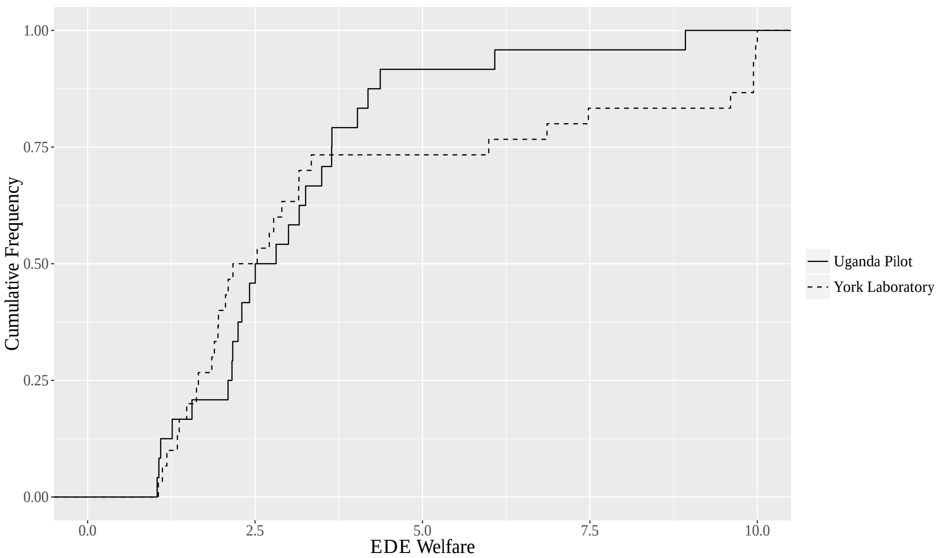

To further test this difference a mixture model was ran to identify clustering within the sample. Within both samples the optimal number of clusters was two and these clusters, as in our previous analysis for the whole sample, was split between a more egalitarian cluster and a more self-interested cluster. The difference, however, is the proportion of those within each sample. As Table 1.6 shows 92% of the Uganda sample fall within the first cluster, while the proportion is only 73% of the York sample. While these results are just illustrative, they pave the way for a full comparison between student and representative sample which will take place in Chapter 2.

Table 1.6: Mixture Model Cluster Comparison: Uganda and York Samples

	Uganda		York	
Cluster	Mixing Proportions	Mean EDE	Mixing Proportions	Mean EDE
1	0.92	2.64	0.73	2.06
2	0.08	7.63	0.27	8.72

6.5 Risk Under the Veil of Ignorance

The Veil of Ignorance perspective, within the experimental context, creates two normative dilemmas. The first concerns the distinction between aversion to inequality and aversion to risk, the second the assumptions behind self-interest; both are intertwined. Two alternate models will be proposed which reflect these two concerns.

In the VOI treatment participants are presented with a decision problem where they do not know which of the three players they are, but know that the decisions they make could determine the payoffs for themselves and two others. Participants could view the situation purely as a risky choice, or as a distributive decision which will end in a degree of inequality. If an individual had risk averse preferences they may act *as if* they had inequality aversion, and vice versa. The CES utility functions used are common to both literatures: the Constant Relative Risk Aversion and Atkinson Index (when reformulated as is here) are all but identical. Indeed, as Rawls notes "of course the two principles are not the same [...] but there is this similarity", both views weight "more heavily the advantages of those whose situation is less fortunate" (Rawls, 1999). So the interpretation of the γ parameter, specified below, is left to the reader.

The matter is further complicated by assumptions concerning self-interest. If an individual is assumed to be purely self-interested, adhering to egoism, then the problem becomes a risky decision. Rawls, although not ruling it out, does not specify that individuals behind the veil follow egoism, but that they must be "mutually disinterested". Meaning they are "conceived as not taking an interest in one another's interests"; they are "not willing to have their interests sacrificed to others" (Rawls, 1999). In our context, this is still open to interpretation. If an individual's *interests* are only the payoff to themselves then "surely his conception of the good is egoistic". If an individual does have other-regarding preferences their *interests* could be considered to be 'utility', leading them to share the payoffs (to some extent). By maximising their *own* 'utility' they are still 'mutually disinterested' and not sacrificing their own interests.

To delve into these issues two models are presented. The first, which has been assumed throughout, has taken the EDE utility function as previously defined, but by setting $\alpha = 1/3$ we can assume that participants are purely self-interested individuals, with their own payoffs as a proxy for individual utility and establish a degree of risk aversion. This method assumes individuals are making decisions about their possible payoffs in each of the three states. An alternate assumption could be that individuals are considering how to distribute resources amongst the 'society' which they will be in, meaning inequality aversion would be the preference of interest. The model is shown below:

$$U_{VOI1} = \left(\sum_{i=1}^{N} \left(\frac{1}{3} (\pi_i x_i)^{-\gamma} \right) \right)^{-\frac{1}{\gamma}} \tag{1.19}$$

The second model assumes that individuals do have some other-regarding preferences, as a result the decision becomes more complex. Now participants must consider a decision where they have a probability of 1/3 of being in the position of Player 1, 2 or 3, and must consider the utility they would receive from allocations to all players in each state. Participants have a degree of self-interest, α, so must consider each possible state when their *self* is Player i. More formally:

$$U_{VOI2} = \left(\sum_{i=1}^{N} \left(\frac{1}{3} (z_i)^{-\gamma} \right) \right)^{-\frac{1}{\gamma}} \tag{1.20}$$

Where:

$$z_i = \left(\alpha(\pi_i x_i)^{-r} + \sum_{j \neq i}^{N} \frac{1-\alpha}{2} (\pi_j x_j)^{-r} \right)^{-\frac{1}{r}} \tag{1.21}$$

Note the difference between r and γ, while r denotes inequality aversion in the case where an individual knows they are Player i, γ is denoted to highlight the parallels between risk and inequality aversion, when the participant does not know which Player they are. To establish γ, for Equation 1.20, the parameter values estimated in the Place-in-Society treatment are used in Equation 1.21, then given those parameters γ is estimated.[6]

The distribution of estimated parameter values, from these alternate models, are shown in Figure 1.17. A divergence appears between the two models. For γ 43.33% of the population are classed as Maximin/Extremely Risk Averse, in Model 1, with 28.57% for Model 2. Further more, while there are 21.43% Utilitarian/Risk Neutral in Model 2, there are 0% in Model 1. While both models shows that Maximin/Extreme Risk Aversion are the model group, Model 2 shows a significant proportion of Utilitarian preferences.

6.5.1 Rawls and Harsanyi

Integral to the concept of the Veil of Ignorance is the debate between Rawls and Harsanyi. Both adopt a definition of the Veil and propose 'rules' that individuals would adhere to

[6]Results shown below are estimated from the Continuous Allocation treatment. Unlike previous analysis a analytical solution for x_i^* could not be obtained, so results are estimated using numerical solvers to ensure the set of corresponding first order conditions equal zero. Similar analysis can be conducted in the Pairwise Choice treatment, using PS data to elicit α and r, and VOI data to elicit γ, given α and r.

Figure 1.17: Comparison Between Inequality Aversion and Risk Aversion

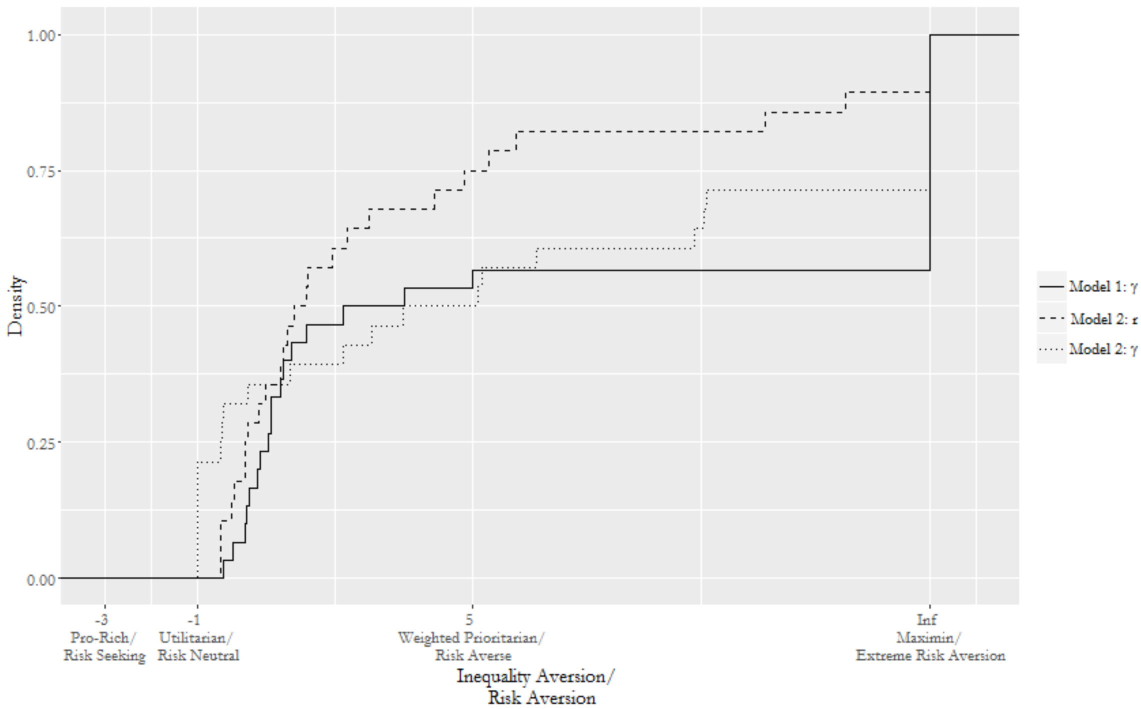

once under the Veil. These 'rules' are, however, contradictory; with Rawls pointing towards Maximin and Harsanyi to Expected-Utility Maximisation.

Harsanyi points to a necessary impartiality when making a "value judgement on the distribution of income." Individuals making such a choice must be "ignorant of what his own relative position [. . .] would be within the system." That this would be the case if they had "exactly the same chance" of being in either position in society (Harsanyi, 1953). This being the case then, for an individual, expected-utility theory would be used as their decision rule, moreover they would choose the society with the highest average utility level (Harsanyi, 1975).

Rawls postulates that the Veil of Ignorance is where "no one knows his place in society, his class position or social status; nor does he know his fortune in the distribution of natural assets and abilities, his intelligence and strength, and the like" (Rawls, 1999). That when under the veil 'rational' and 'mutually disinterested' individuals should adhere to the Maximin Rule. To conform to the Maximin Rule one must "identify the worst outcome of each available alternative and then adopt the alternative whose worst outcome is better than the worst outcomes of all the other alternatives" (Rawls, 2001).

However, Rawls does specify that it is only under certain "conditions in which it is rational to be guided by the Maximin rule". Perhaps the condition most appropriate refers to probabilities; the "maximin rule takes no account of probabilities" and that parties involved have no "reliable basis for estimating the probabilities" (Rawls, 2001). Harsanyi argues for Laplace's principle of indifference, that if we are in a situation of complete ignorance "then it is reasonable to assign equal probabilities to all the possibilities". Moreover, for moral reasoning, the same "a priori weight should be given to all members of society" leading to these equal probabilities (Harsanyi, 1975). Out of a situation of ambiguity a situation emerges where an equal weighting is reasonable; the ambiguous problem emerges as an uncertain one. Indeed, this is the crux of the argument; while Harsanyi proposes a feasible assumption to enable a rule to shape decisions, Rawls argues that Maximin would only be adhered to if such probabilities did not exist at all. Perhaps this is the reason the proposed rules are opposed.

Both rules are justified from a normative standpoint, yet they lead to opposite conclusions. One maximises total income, while the other increases the income of the worst-off. Rawls states "in working out the conception of justice as fairness one main task clearly is to determine which principle of justice would be chosen in the original position" (Rawls, 1999). So, rather than adopting a normative standpoint, the work done here attempts to answer this question from a positive standpoint.

Results shows that if participants are assumed to be 'egoistic' then the modal group adhere to Maximin (43.33%) with no participants adhering to Utilitarianism. Suggesting that under those assumptions more individuals adhere to Rawls' 'rule'. If we assume individuals have some other-regarding preferences, then Maximin still remains as the modal group (28.57%), however a significant minority (21.43%) adhere to Utilitarianism. Although the Maximin rule consistently has a majority of 'votes' it is clear that, even under the Veil, preferences are still heterogeneous, and although the same behaviour is observed a more simple model (the first) may over estimate the proportion of Maximin preferences.

By using these models to estimate heterogeneous preferences, the 'just' principles estimated can be extended to policy decision making. When policy makers are faced with decisions (relating to redistribution of income) Social Welfare Functions, incorporating estimated preferences, could be used to analyse the welfare effects and rank the decisions; allowing decision makers to choose that with the highest welfare gain. While a single 'rule' could be chosen, perhaps from a modal response, or an aggregate preference could be established, the median or 'most likely' societal preference, individual level preference data could actually be used, and then aggregated. In this approach more in-depth analysis could be conducted, and if larger experiments were conducted, relating to the welfare gains and loses for particular groups. Indeed, it would point to an ideal distribution in society, according to potentially 'just', evidence based, principles. While only meagre steps towards this aim have been made within this experiment, it hopefully provides a way in which to push this agenda forward.

7 Conclusion

To conclude, an incentivised laboratory experiment ($n = 30$) was run, where participants were required to make distributional decisions amongst groups of three. Multiple experimental designs and alternative perspectives were presented to allow the identification of within-subject treatment effects. Utility functions have been proposed which incorporate parameters relating to self-interest and inequality aversion; combining behavioural modelling with moral value judgements. Preferences are then estimated on an individual, cluster and sample level to explain the behaviour observed in the experiment. The goodness-of-fit of alternative utility functions and stochastic error models is further analysed, alongside the disentanglement of risk and inequality aversion while under a veil of ignorance.

Economic models which assume *all* individuals are purely self-interested have been shown to be lacking in explanatory power, as extensive heterogeneity in behaviour is found between-subjects. Through formulating models which incorporate other-regarding preferences, however, the differences in individual-level behaviour can be explained. The majority of the sample is shown to possess some regard for others and an aversion to inequality. Yet, altering the design of the experiment, or perspective which the individual considers, has been found to have significant effects on the behaviour of the individual. This leads to the need to estimate different preference parameters. Furthermore, results have shown the importance of accounting for *noise* in behaviour. These factors are important, and in future research the effect that the design and perspective can have on the behaviour observed and preferences estimated should be accounted for.

A Appendices

A.1 Running the Experiment

A.1.1 Sample Characteristics

Table 1.7 shows the demographic composition of the sample, which was established through the questionnaire. The sample contains significantly more female participants (70%), with 93.3% between the ages of 18 and 30. In general participants are highly educated, 33.3% at undergraduate level and 46.7% at postgraduate level, with 28 out of 30 participants being students. Income is low, as to be expected with students, with 66.6% earning below £12,600, while parental income is slightly more representative there is still a high proportion, 36.7%, in the lowest income bracket. In terms of nationality and ethnicity the sample is relatively international with 10 different nationalities being present; the majority groups are British, 46.7%, and white, 53.3%, with a sizeable proportion of Chinese nationals, 26.7%. Social scientists make up the largest majority group of students, 40%, (half of which are in the general area of economics) but arts and humanities and sciences are also well represented. In terms of religion, there is an eclectic mix, and in terms of political preference the distribution seems bimodal, with peaks to the left and centre. Participants were given the option of not responding to any of the questions they felt uncomfortable with answering, so there are many N/A responses in questions of ethnicity, religion and income which should be considered while looking at the demographics.

Table 1.7: Demographic Composition of Sample; N=30

Demographic	Category	Number	Percentage
Gender	Male	9	30.0%
	Female	21	70.0%
Age	18-21	10	33.3%
	22-25	11	36.7%
	26-30	7	23.3%
	31+	1	3.3%
	N/A	1	3.3%
Nationality	British	14	46.7%
	Chinese	8	26.7%
	Canadian	1	3.3%
	Greek	1	3.3%
	Hungarian	1	3.3%
	Indian	1	3.3%
	Japanese	1	3.3%
	Latvian	1	3.3%
	Lithuanian	1	3.3%
	Turkish	1	3.3%
Ethnicity	White	16	53.3%
	British Asian	1	3.3%
	Greek	1	3.3%
	Japanese	1	3.3%
	Mixed Race	1	3.3%
	Turk	1	3.3%
	Chinese	2	6.7%
	N/A	7	23.3%
Religion	Agnostic	1	3.3%
	Atheist	5	16.7%
	Buddist	3	10.0%
	Christian	4	13.3%
	Hindu	1	3.3%
	Orthadox	1	3.3%
	Other	1	3.3%
	None	6	20.0%
	N/A	8	26.7%
Highest Education	A Levels	5	16.7%
	Other	1	3.3%
	Postgraduate	14	46.7%
	Undergraduate	10	33.3%
Subject/ Occupation	Arts and Humanities	4	13.3%
	Science	6	20.0%
	Social Science	12	40.0%
	Misc. Student	6	20.0%
	Administrator	1	3.3%
	N/A	1	3.3%
Income	£0	10	33.3%
	£1 - £12,600	10	33.3%
	£12,601-£20,600	1	3.3%
	£20,601-£32,100	1	3.3%
	£32,101-£49,900	0	0.0%
	£49,901 or more	0	0.0%
	N/A	8	26.7%
Parent's Income	Less than £12,600	11	36.7%
	£12,601-£20,600	4	13.3%
	£20,601-£32,100	4	13.3%
	£32,101-£49,900	4	13.3%
	£49,901 or more	6	20.0%
	N/A	1	3.3%
Political Preference	Left: 1	1	3.3%
	2	0	0.0%
	3	8	26.7%
	4	4	13.3%
	5	6	20.0%
	6	8	26.7%
	7	1	3.3%
	8	1	3.3%
	9	0	0.0%
	Right :10	0	0.0%
	N/A	1	3.3%

A.1.2 Instructions

In order to increase the understanding of the experiment, participants were given instructions throughout the experiment. The instructions came in four forms: paper instructions, on-screen tutorials, reminders and passive on-screen instructions. The *paper instructions* were given at the beginning of the session, before the on-screen experiment started. This consisted of general instructions, which explained the experiment as a whole, and design specific instructions, which explained the specifics of each design (see AppendixA.2). Participants were requested to read only the general and pairwise choice instructions before they began, to reduce information overload and to ensure they were focused on one design at a time. Following the first paper instructions there was an *on-screen tutorial* which showed the participants the specifics of the next design. The layout, and meaning of each section, of the screen was explained, showing exactly how participants were to interact with the screen and make their decisions. The tutorial itself was interactive, allowing the participant to click forward and backward to understand the each design. They were made to correctly fulfil each requirement of each design in order to proceed, and were given a minimum time to ensure they did not rush through. Once the tutorial was completed the *reminder* screen would show, as it would before each different game, to show the specifics of each game, including the perspective they would take, the number and nature of rounds and other specific information. Finally, within each round there were very brief *passive on-screen* instructions, giving just enough information incase participants forgot some information. When the first design was complete participants were asked to refer to the next section of the written instructions, and then would repeat each form of instructions for that design, and the next. The aim of providing information in small sections repeated in different, and hopefully interesting ways, was to ensure participants remained focused and absorbed the necessary information to really understand each element of the experiment.

A.1.3 Randomisation

Randomisation is a vital element in the experimental design. Firstly, as each participant entered the room they were asked to take a number, which corresponded to certain computer, from a bag to randomise who was allocated to each seat. Secondly, subjects were randomly allocated into groups of three for each round of the experiment. While participants made decisions each round, ‘as if’ they were the dictator, in each round one dictator was randomly (and anonymously) chosen to determine the payoffs to the group. For the Place-in-Society and Veil of Ignorance perspectives the groupings were straight forward, the ‘dictator’ was P1, and their decisions for P2 and P3 are played out for the participants in particular seats. For the Impartial Spectator perspective the dictator in Group 1 makes decisions for Group 2, Group 2’s dictators decides for Group 3, and so forth, until the final dictator in the final group decides for Group 1. These random group orderings remained anonymous, but ensured that the design was incentive compatible.

A.1.4 Ordering

Due to the small sample size randomisation of the treatment order was not possible. If the ordering was randomised in each session there would have only been 9 to 12 participants to observe any effects, these effects could, therefore, be because of random sampling rather than due to the order effect. As a result focus was made on an ordering which aided understanding of each design. Firstly, the designs were ordered in terms of how cognitively demanding they were. Pairwise Choice, was deemed the least demanding as participants were required to make a decision between only two choices, clicking on the distribution they preferred. Discrete Slider, was second, as they had to move a single slider to state a preference amongst multiple distributions. Last was Continuous Allocation, as not only did participants have to calculate an exact preference in terms of allocation, but had three, rather than one, mediums to interact with; the multiple sliders, arrows and input boxes. Secondly, the order of perspective was Impartial Spectator, Place-in-Society and then Veil of Ignorance, as the complexity of the perspective was deemed to be in that order. Decisions for the Impartial Spectator are about decisions purely about others. When considering the Place-in-Society the matter is more complex as the participants payoffs are within this, alongside others. But the Veil of Ignorance is more complex as decisions about both individual and others need to be made, when who each player is, is not certain.

A.1.5 Incentives

In order to incentivise participants the reward mechanisms used need to be 'incentive compatible'. What this, in essence, means is that the respondent provides "truthful responses to the specific questions that the experimenter wants to observe responses to" (Bardsley et al., 2010). One such mechanism, which will be used here, is the random-lottery incentive scheme. This mechanism entails giving participants specific tasks which each has a well-defined reward structure. Once all of the tasks have been finished one of the tasks will be selected at random, and it is from this the participant will receive their payoff (Bardsley et al., 2010). Theoretically the random-lottery incentive scheme is unbiased, if one were to assume expected utility preferences; the participant should give the exactly the same response to each task as they would have done if they were only faced with that single task.

An important element of the design is the group dimension. As a result the payments to those within the group and themselves are a vital when considering incentives. As three players will be in each game, each receiving payoffs, it is necessary to ensure these payoffs are considered to be real. As a result each participant made their decisions 'as if' they were the dictator. One dictator was randomly (and anonymously) chosen for each round, and their decisions played out for the participants within that group. In this way each participant received one 'set' of payoffs, either that which they gave to themselves (when they were chosen as the dictator) or that with another participant in their group gave them (when that other participant was the dictator). This payoff mechanism, should, ensure that all players behaviour optimally, as each of their allocation decisions had an equal weight of playing out.

A.1.6 Software

The experiment was run in a computerised laboratory, with participants interacting with an on-screen programme. This was programmed in z-Tree (Fischbacher, 2007), a toolbox designed for experimental economists which has the additional benefit that it saves data routinely; which reduces the issues associated with server crashes.

A.2 Instructions

Instructions

Please **Read These Carefully**.

Everyone Will Receive the **Same** Instructions.

General Instructions

In this experiment you will be making **decisions** about the **distribution of payoffs** between **yourself** and **other participants** in this room. These **payoffs** are in addition to your **turn-up fee of £2.50**.

There will be **three** different **experiments**, each made up of multiple **games**, which in turn consist of multiple **rounds**. Your **actual payoff** will be determined from **one randomly selected round**. It is from this **one round** that **all participants** will receive their **payoff**. This means that **every round** has an **equal chance** of determining your **final payoff**, so consider each choice you make **carefully**. **Everyone** will finish at the **same time**, as you need to wait for **every participant** to finish **each round** before you can move onto the next.

The choices you make will determine the payoff for **three players**. You will make choices which concern the **distribution of payoffs** between each player. For **each round**, you will be **randomly** linked with **other** players. The three of you will make your choices **independently** of one another, but only **one** of the player's **choices** will be selected, randomly, to provide the payoffs for **all three players**. You **will not** know whose decision has been chosen and will receive you payoff, individually, at the end.

A **simplified version** of the experiment is shown on the opposite page. It shows a possible **distribution** of payoffs. The **height** of the orange bars shows the **payoff to each individual**; with the payoff level shown on the vertical axis and the corresponding player on the horizontal axis. **Payoffs** are **always shown in pounds**, in this example, at the top of the orange bars and to right hand side next to the corresponding player. The **'Total Sum of Payoffs'** shows the payoffs of all the players added together, while the **'Gap Between Payoffs'** shows the difference between the best-off and worst-off player. The Finish button in the bottom right corner allows you to move to the **next round**. You must **make a decision** in **every round** and **then click Finish** to confirm your decision.

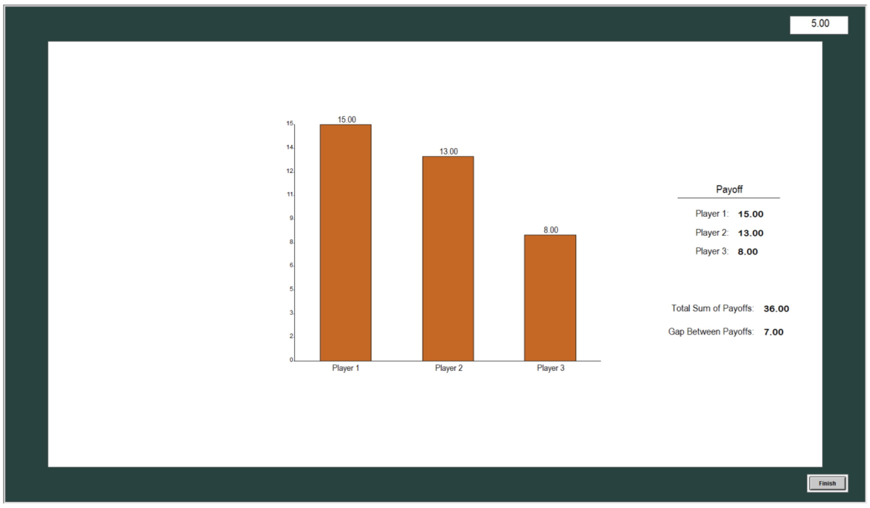

A **minimum time** will be displayed in the **top right corner** in every round, in black; this time **must have elapsed** before you progress to the next round. There will also be a **maximum time**, in red, which will be **double the minimum time**, if you **do not make** a decision in this time and **click Finish** you will receive a **payoff of zero** for that round, and your decision **will not count** for the **other two players**.

Throughout the three experiments you will face **three different perspectives**; they are **Outside Observer**, **Place-in-Society** and **Under the Veil**. For the **Outside Observer** your decisions will concern **three other players**. **You** will **not** be making a decision about your payoff; that will be done by some **other player**. In **Place-in-Society** the three players in the group are **yourself** and **two other players**. You **will know** which player you are when you make your decision. **Under the Veil** will be the same; however, you **do not know** who you are in the distribution, this will be **randomly decided**.

After the experiment you will be required to fill out a **questionnaire**. Your responses from the **questionnaire**, and from the **entire experiment**, will be treated **anonymously**. You will be given **instructions** before **each** of the **three experiments**, explaining the precise situation. You should **read these carefully**, as each experiment will be different. After the **paper instructions** you will be given a **tutorial**, so that you can understand each experiment fully. If you require help at any time, **please raise your hand**.

Pairwise Choice Instructions

The first experiment will be a set of **Pairwise Choice Questions**; where you have to **choose** between **two** different **distributions of payoffs** between **three players**. The decision you make will be between **Choice A** and **Choice B**, which each contain a different distributions of Payoffs. You will be shown two **graphs** side-by-side, which represent these choices.

You are required to make a **choice** between the two. You can make your choice by simply **clicking anywhere** on either **graph**. You can change back and forth, to see in greater detail the payoffs. Once selected, your choice will be **highlighted with a box**. Once you have made your choice you may click the **Finish Button**, in the bottom right, to proceed. A screenshot of the experiment is shown below.

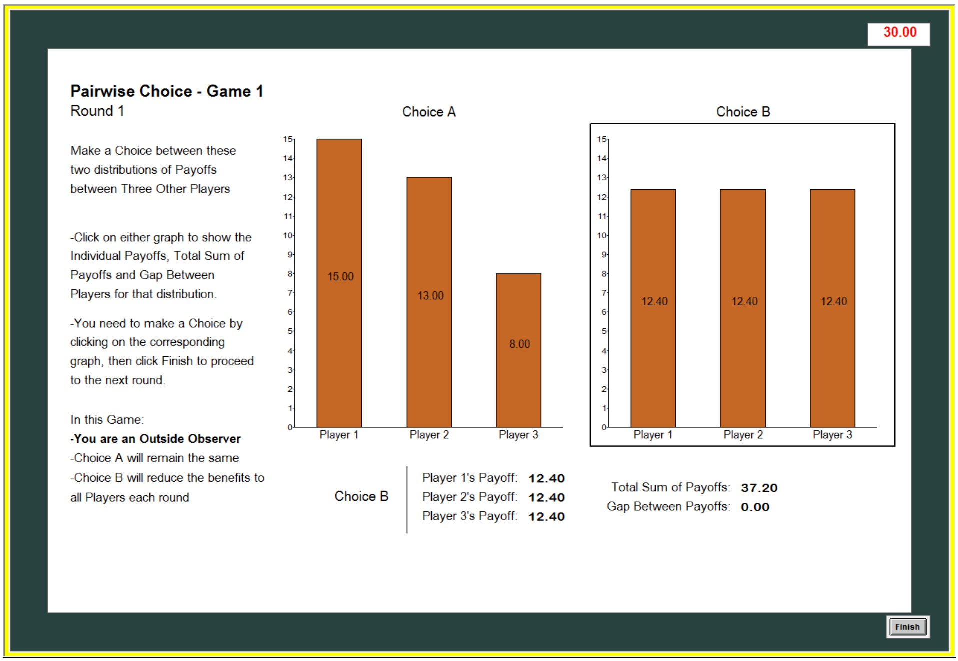

You will firstly be shown an **onscreen tutorial** which, with **blue information boxes**, will explain the interface and show you how to make your decisions. Then you will do **4 games**, each made up of **8 rounds**. In the **first** game you will be an **Outside Observer**, in the **second** and **third** games you will know your **Place-In-Society** and in the **fourth** you will be **Under the Veil**.

Throughout each game there will be **trade-offs** between the **Total Sum of Payoffs** and the **Gap Between Payoffs**. When making your decisions be sure to **check the differences** between the two in Choice A and Choice B, **as well as** the distribution of Payoffs.

In **every** game **Choice A** will **remain the same** throughout the **8 rounds**, while **payoffs** in **Choice B** will be gradually reduced in some way.

End of Pairwise Choice Paper Instructions

Please Proceed to the On Screen Tutorial

Single Slider Instructions

The second type of experiment will be in the form of a **single slider**; there are a **set range of options** which you can choose from, **using the slider**. Unlike the first experiment you will only see **one graph**, shown below. It will have different distributions of payoffs between **three players**. These distributions are all chosen from a **total budget of £30**. What you must do is to **move the slider** along the axis to find the **distribution** that you prefer. There are **multiple distributions** to choose from, so be sure to **check each of them**.

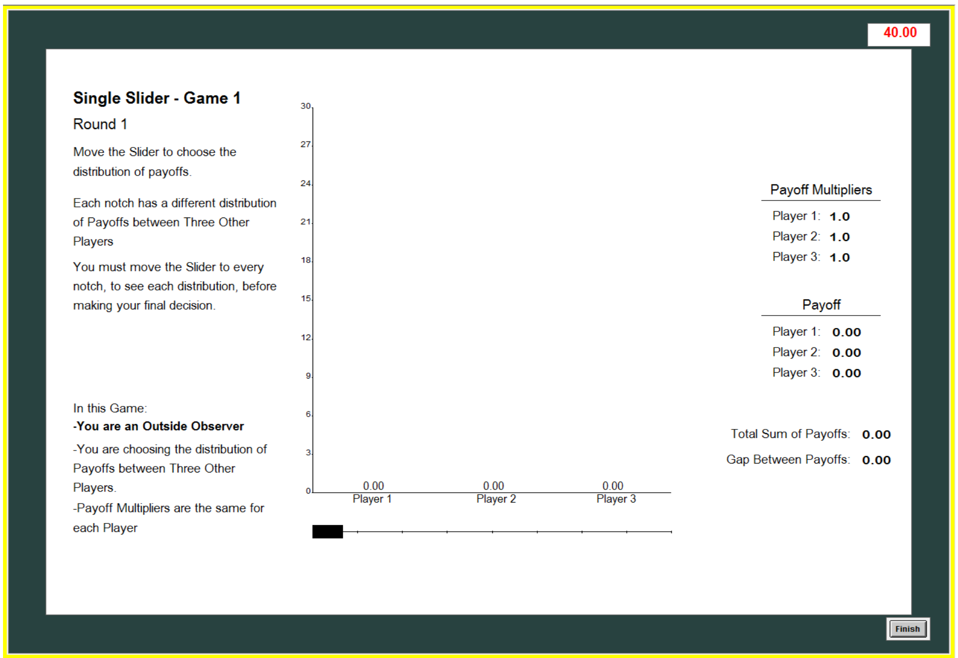

There will be **3 games** for this experiment, each made of **two rounds**. In the first game you will be an **Outside Observer**, in the second you will know your **Place-In-Society** and in the third you will be **Under the Veil**. For the **second game**, the options you see in the **second round** will depend on your choice in the **first round**. For the **first and third games** each round will be independent.

Within each round there will be **Payoff Multipliers** implicit in the choices available. In the **first round** of **every game** they will be [**1, 1, 1**] for Player 1, Player 2 and Player 3, respectively. In the **second round** of **every game**, however, they will be [**1, 0.6, 1.4**] for Player 1, Player 2 and Player 3, respectively. What this means is that the **allocations** from the budget of £30 will be multiplied by the **Payoff Multiplier** when given to each **Player**. You will see the effect of this in the **second round**, where the **Total Sum of Payoffs** will change depending on which choice is made. An example of how this works is shown below:

	Player 1	Player 2	Player 3
Allocation	10	10	10
Payoff Multiplier	1	0.6	1.4
Payoff	10	6	14

The addition of these **different Payoff Multipliers**, in the **second round** of each game, means that there emerges a **trade-off** between the **Total Sum of Payoffs** and the **Gap Between Payoffs**. As the **Total Sum of Payoffs** increases, so does the **Gap Between Payoffs**.

Whichever **choice** you make will be the **distribution** from which you may receive your **payoff**. You will firstly be shown an **onscreen tutorial**, which, with **blue information boxes**, will explain the interface and show you how to make your decisions. Then you will do the **experiment for real**.

End of Single Slider Paper Instructions

Please Proceed to the On Screen Tutorial

Multiple Slider Instructions

The final type of experiment is the most complex. An extension of the single slider, **the multiple slider choice** allows you to make **very precise allocations** between **all three players**. There are **no predefined distributions** to choose from, **you** must therefore **choose allocations to each player** on a continuous level. The allocations you make will determine the potential **payoff** of the **three players** in that round.

In order to do this there are **three options**. Firstly, by **moving the slider**, for **either Player**, you can give them the **allocation you desire**. Secondly, to enable you to **more precisely** change the allocations the **arrow buttons** either side can increase or decrease the allocations by a **small amount**. Thirdly, by using the **Written Input** boxes you can precisely **type** the Allocations you want. A screenshot of the experiment is shown below.

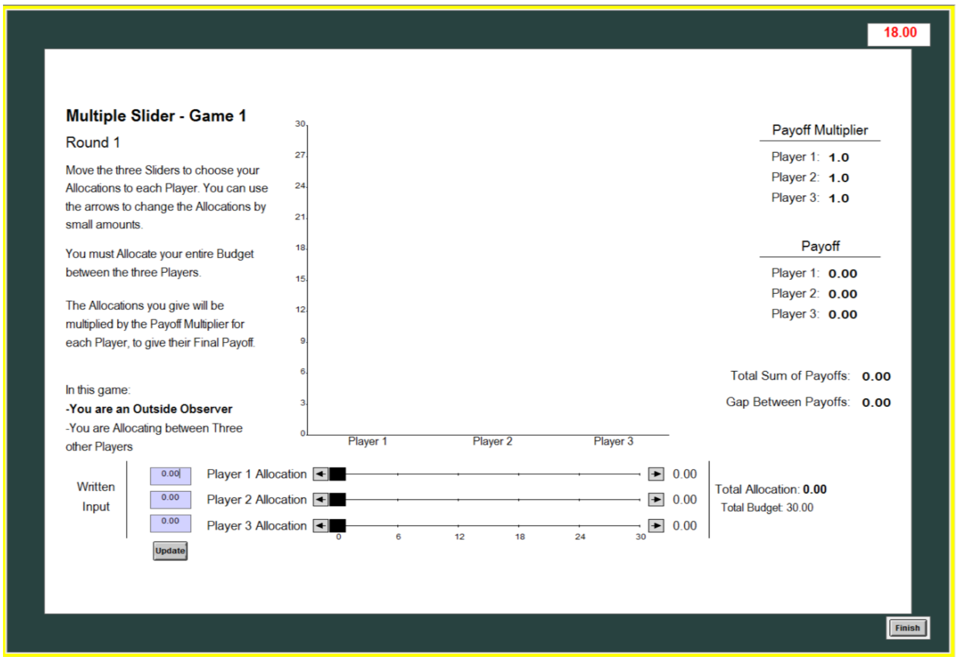

The **Allocations** you give to **each Player** are shown **beside the slider**, with the **Payoff Multipliers** shown in the top right. The **Payoffs** are represented by the **orange bars**, the **value next to them** and the number to the **right hand-side**, the **'Total Sum of Payoffs'** and **'Gap Between Payoffs'** is also shown.

One thing which will **alter between rounds** of this experiment is the **Payoff Multipliers**. For each round each player will be given a different **Payoff Multiplier**. The **payoff** that player gets will be the **Allocation**, you give, **multiplied by the Payoff Multiplier**. An example of how this works is shown below:

	Player 1	Player 2	Player 3
Allocation	10	10	10
Payoff Multiplier	1	0.8	1.2
Payoff	10	8	12

These **Payoff Multipliers** are **important**, if you choose to allocate to the Player with a **higher** Payoff Multiplier, then the Payoffs they get are **greater** than if you choose to allocate to someone who has a **lower** Payoff Multiplier. You will be able to see the effect of this in the **Total Sum of Payoffs**. There is, however, a **trade-off** as by allocating more to the Player with the higher Payoff Multiplier you are likely to increase the **Gap Between Payoffs**.

First, you will be shown an onscreen tutorial, which **doesn't count** for your payoff. Then you will do the **experiment for real**. It will be made of **3 games**; the **first game** has **5 rounds**, while the **second** has **20** and the **third** has **5 rounds**. In the first game you will be an **Outside Observer**, in the second you will know your **Place-In-Society** and in the third you will be **Under the Veil**. In each round, you must ensure that the entire **Total Budget** is **allocated** amongst **the three players** before you continue.

End of Multiple Slider Paper Instructions

Please Proceed to the On Screen Tutorial

A.3 Within-Design Comparisons

The results here identify the differences in responses due to changes in perspective, within each experimental design. Figure 1.18 shows the distribution of responses between perspectives in the Pairwise choice design. The nine categories denote different 'switching' points (see Table 1.1), which correspond to varying degrees of inequality aversion. The median response is either Weighted Prioritarian 1 or 2, for each of the three perspectives. There are no 'Egalitarian' responses within any of the designs, implying that no subject is willing to violate monotonicity. While the distribution within the IS and VOI treatments are very similar, there are less 'Utilitarians' and no 'Inequality Seeking' individuals within the PS perspective.

Figure 1.18: Pairwise Choice, Distribution of Responses

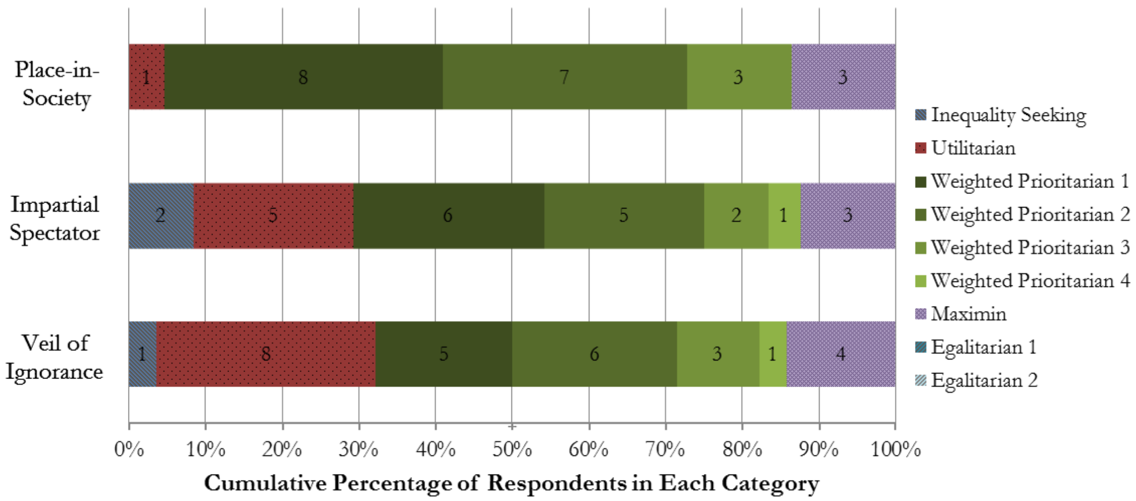

Figure 1.19 shows the comparison of decisions made by participants in the discrete slider design. The numbers denote the notches that participants could choose; where Notch 1 denotes the choice which gave everything to Player 1 in the PS perspective, and to the player with the highest multiplication factor in VOI and IS. Notch 8 represents the choice which allocates most equally amongst the group of three. Results show very similar behaviour in the IS and VOI perspectives, with a large proportion of the sample (40%) choosing the most equal distribution. The PS perspective is somewhat different, with no participants choosing Notch 8, and over half choosing Notch 1.

Figure 1.20 shows that a similar pattern emerges for the continuous allocation design. The VOI and IS perspectives are very similar, while the PS perspective reveals individuals allocating more to P1. Here each individual response, from each of the available rounds, is shown. The median responses are 0.46, 0.33 and 0.33 for PS, VOI and IS, respectively, with clear modal spikes at 0.33 for VOI and IS, and 0.33 and 1 for PS. There is a clear shift to allocations to P1 when participants know who they are in society.

Figure 1.19: Discrete Slider Comparison between Perspectives

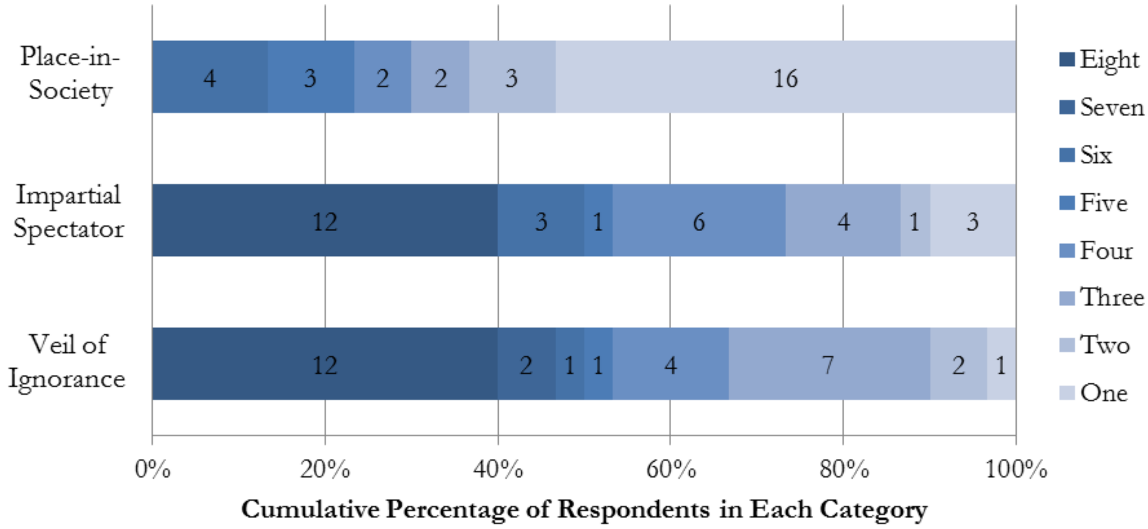

Figure 1.20: Continuous Allocation; Proportional Payoffs to P1, between Perspectives

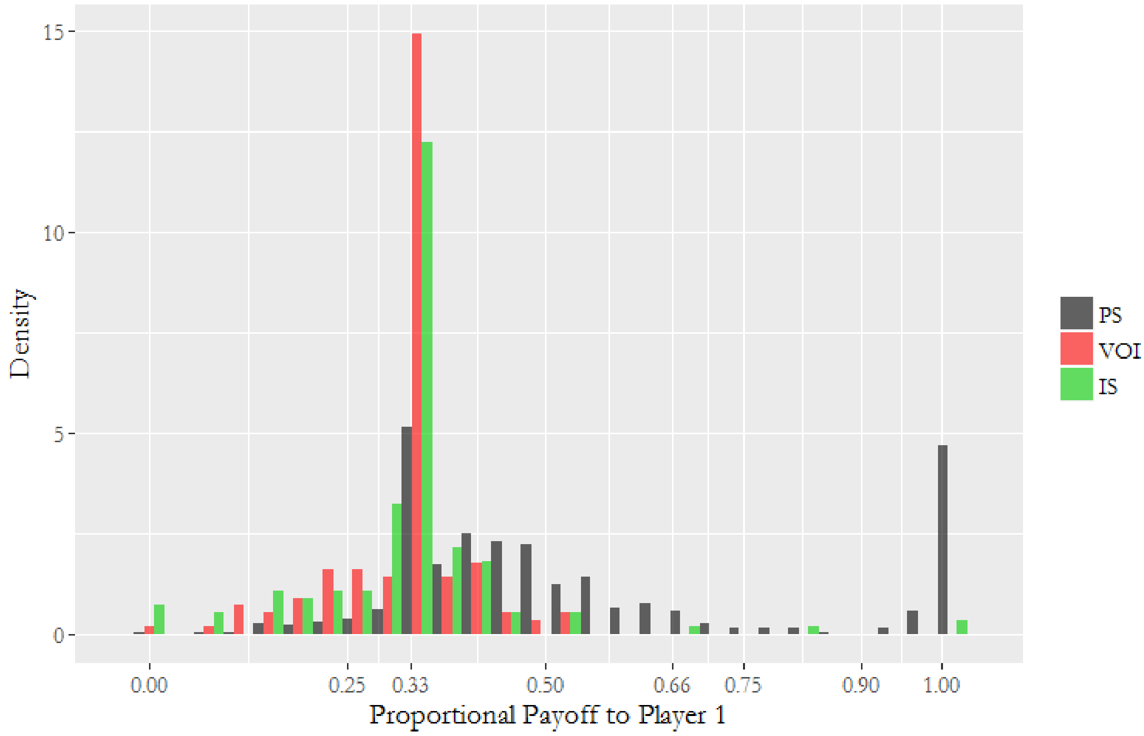

A.3.1 Pure Gains vs Efficiency

One design element of the experiment was a set of control questions, within the discrete slider design, to answer the question: does behaviour change when faced with decisions of pure gains, compared to decisions which are subject to some efficiency trade-offs? The pure gains problems posed had a multiplication factor of one for each P1, P2 and P3, therefore there are no competing objectives of equity and efficiency as no trade-off is required. The efficiency questions had different multiplication factors, meaning participants needed to trade-off between the competing aims of maximising total income and distributing equally. The results are shown below in Figure 1.21. The results are shown in terms of the notch which the participant chose as optimal. Where 'Eight' denotes an equal distribution and 'One' denotes an allocation entirely to P1, in the PS perspective, and to the player with the highest multiplication factor in VOI and IS. The lower of the bars denotes Pure Gains, while the higher denotes Efficiency.

From these descriptive results there is a clear shift towards a more equal distribution when moving into the pure gains questions. This is most stark in the VOI and IS treatments, where responses within the most equal notch change from 40%, in both cases, to 80% and 73.33%, respectively. Within PS the change is less, but there is still a clear shift; the median response, for instance, shifts from One to Six. While these results are perhaps not shocking, they do show that individuals are thinking about the decision task. There appears to be an understanding of the multiplication factors, and as a whole do trade-off these competing objectives when necessary. It appears that in a perfect world, where no trade-offs are necessary, preferences shift towards a more equal society. Even when individuals know who they, there is a shift away from more selfish allocations.

Figure 1.21: Pure Gains vs Efficiency Comparison; Discrete Slider

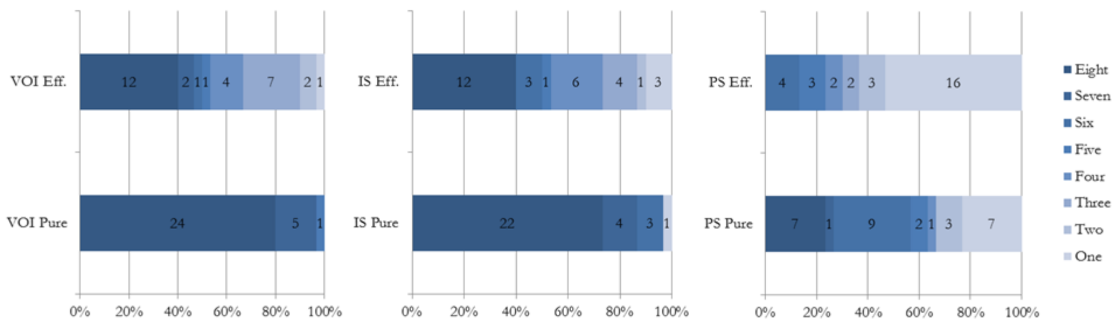

A.4 Simulations and Pilot Study

In order to ensure the experimental design was optimal two preliminary stages were be carried out; simulations and a pilot study. In order to assess if the experiment can tease out the parameters required preliminary simulations were ran. Randomly generated 'participants' with given parameters 'played through' the experiment. Through utilising the generated choices they make the test is to assess if the parameter values of each 'participant' can be estimated correctly. If this is the case the experimental design achieves its aims, if not reworking of the parameters needs to occur. Once an 'optimal' experimental design has been ensured a pilot study can be ran. The aims of the pilot study are to ensure that every part of the experiment works smoothly in the laboratory, and to provide a more probable distribution of responses. Once this distribution has been established further simulations can be conducted, to see if the experiment holds up to more realistic participant interactions.

A.4.1 Pilot Study

Before the main experiment was ran a Pilot Study was conducted to test the design. Four participants, each of them postgraduate students at the University of York, participated in the experiment. The experiment was run in full, but incentives were divided by ten to reduce costs. While the experiment ran smoothly several design features were altered to ensure the design was optimal for the real experiment. The interaction between subject and interface was observed and questions were asked at the end, in order to gain an understanding of what needed to be changed and what worked well.

In order to increase understanding of the Pairwise Choice design an additional page was added to the Tutorial section. This section showed exactly how each round would progress, with Choice A remaining static, with Choice B reducing each round. Through doing this it was made clearer that a single switch at some point was a logical and consistent decision to make. Alterations were made to the Continuous Allocation in two forms; firstly, the Input Boxes were added to enable quicker and more precise allocations to be made and secondly, the wording was changed from Productivity Factor to Multiplication Factor to avoid implications that Productivity may have. The ordering of perceptions was changed from PS – VOI – IS to IS – PS – VOI, as respondents felt that was the easiest and most logical ordering. Alterations were made to the questionnaire, allowing participants to refrain from answering any question they preferred not to. Timings were altered slightly to allow more time when it was needed, and less when participants were waiting for long periods of time. Further changes were made to make both paper and on-screen instructions more easily understandable; by clarifying wording and ensuring repetition of vital instructions were made. With the changes in place, the main experiment could be conducted, fully incentivised with a full and random sample.

A.4.2 Simulating Comparisons of Experimental Designs

In order to establish any potential effect that the experimental design may have on the estimated parameter values simulations need to be ran to discover exactly which parameter values can be estimated, if there is any potential bias in either design and how robust they are to error. Below a comparison is made between the three experimental designs, the method is drawn from (Crosetto and Filippin, 2015). What is shown, in Figure 1.22, is the corresponding parameter values an individual could be observed to have for each possible choice within each experimental design. The upper-right and lower-left graphs show individual ticks for each possible self-interest and inequality parameter respectively. The top-left graph represents the possible choices for both parameter values for each design. It is clear that, the Continuous Allocation (CA) design offers a much greater range of parameters, while PC and DS offer significantly less. Although more spread out, DS and PC are evenly spread when considering the self-interest parameter. However, when considering the distribution of possible inequality aversion parameters DS is significantly right skewed and PC is slightly right skewed. When considering the top-left graph it is clear that any individual with true parameter values that are inequality averse and relatively selfless would find few options in the DS which would capture this. While PC has more options than DS it is clear that large jumps to the nearest parameter value would need to be made. It is clear that the CA method captures a far greater range of possible values. Moreover, while the CA generates this range from just one decision problem (in this case where $\pi = [1, 0.8, 1.2]$) there are twice as many needed for DS and fifteen as many in PC.

Figure 1.22: Comparisons of Possible Parameter Values for each Design

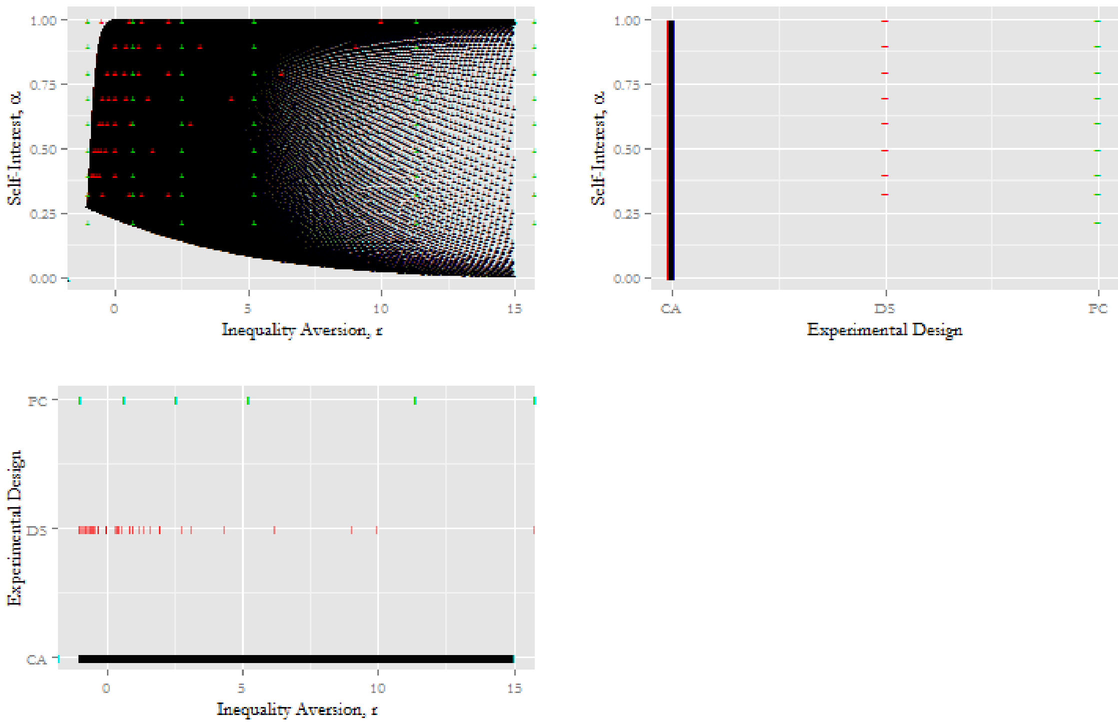

The density plots below, in Figure 1.23, reveal more information about the distribution of possible parameter values. While we know from the previous graph that CA does have a far greater number of choices, what is not clear is how they are distributed. We can see that, although relatively low on numbers PC has a very even distribution across the potential parameter values. DS is, as before, equally spread for α but skewed towards parameters which reflect less inequality aversion. CA is now revealed to have a very unequal distribution amongst the parameters; with a greater density at 0, 0.33 and 1 for α, and towards -1 for r.

Figure 1.23: Density Plots of Options for each Parameter Value

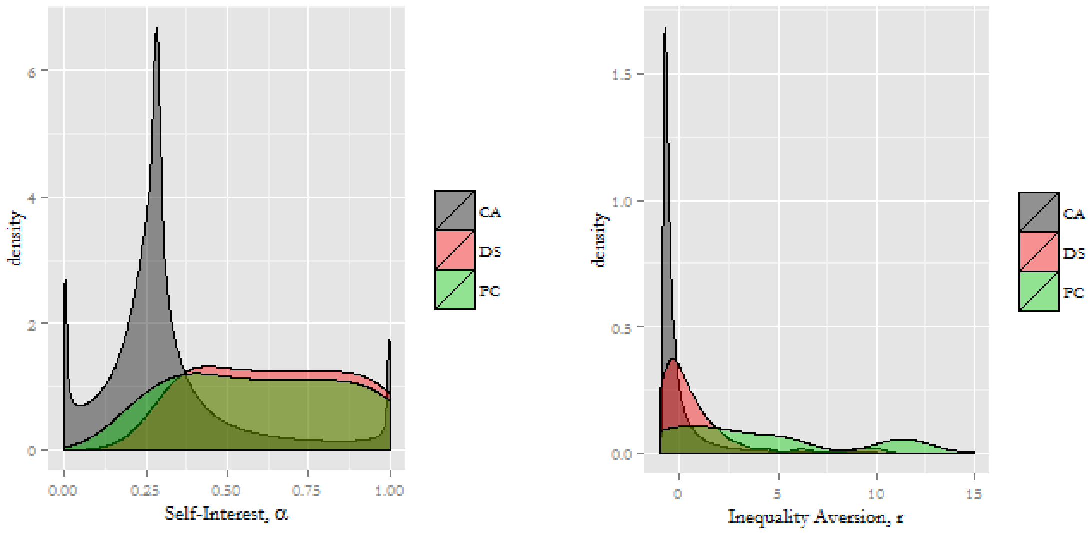

In order to test what the effect these elements of each design could have upon the elicitation and estimation of parameter values a simulation was ran. Parameter values were was simulated for 10,000 hypothetical individuals, characterised by the Outcome Utility function. The parameter values were drawn from distributions which reflected those observed within the experiment, perhaps not perfect but with similar characteristics. To ensure a similar distribution $r \sim B'(5, 1.8)$, while usually characterised as being bounded between 0 and ∞, this was readjusted to accommodate our -1 parameter values. The median value is similar to that established from the results of the experiment, at 2.1262. For the Self-Interest parameter, again a similar distribution was mapped to that found from the results of the experiment. This time $\alpha \sim B(4, 1.1)$, with a median value of 0.8216. The distributions of parameter values are shown in Figure 1.24.

Figure 1.24: Distribution of Randomly Generated Parameter Values

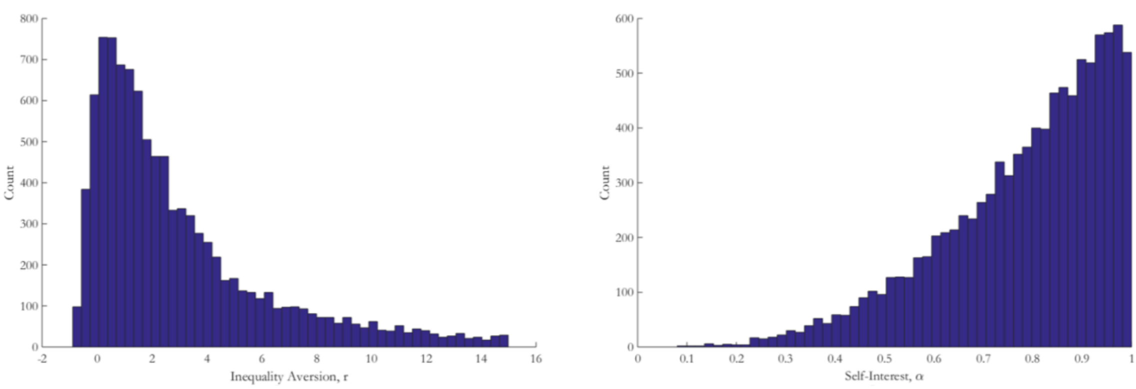

Once a distribution of 'true' parameter values was established then a comparison of those parameter values which could be elicited/estimated from each design could be compared with those 'true' values. There were several proposed simulations to compare, but each consisted of matching the closest possible parameter value from the design for that simulation. The first was Deterministic, where there was no error in responses, and individuals were able to choose the option which was closest to their 'true' parameter. The second two simulations incorporate stochastic preferences, where error is directly added to the 'true' preferences.[7] A simple normal error is added, with zero mean with varying variance. The low error variance parameter is 0.3 and 2 for the self-interest and inequality aversion simulations, respectively. While for the high error the values are 0.6 and 4, respectively. The final simulation aims to highlight a proportion of individuals within the experiment who exhibit confused behaviour. Rather than adhering to the utility function these individuals choose any parameter value with equal probability. A significant proportion, 10% of these responses are mixed with other deterministic individuals, to provide our final simulation.

Figure 1.25 shows cumulative frequency plots which show the distribution of each of these simulations for each design; for self-interest, the left column, and inequality aversion, the right. Each of the simulations, are plotted against the 'true' parameter values. The graphs highlight two key aspects for each design; the precision of the simulations and their susceptibility to error. First if deterministic preferences are considered CA is by far the more accurate, with the deterministic curve almost exactly mapping the 'true' curve. For both PC and DS, whole closely following the curvature of the line, both stair-step, due to the nature of the discrete intervals. While both perform equally for Self-Interest, DS is more precise in inequality aversion, especially at lower levels.

When considering the individuals could make error when thinking of their true preferences, the accuracy of the estimates diminished in all cases. However, the effects that error has on each design vary. While both PC and DS have predefined notches, which correspond to theoretically appealing values, the effect seems to be less extreme than in the case of CA. For Self-Interest, in DS for example, due to a minimum cap on α that error which has led to a parameter value below 0.33 (which has a very low probability of happening) does not have much effect, and the estimates are somewhat closer. Compare this to CA where, as any parameter value is attainable, the effects of the specific error make a large difference to the curves. Indeed, if the error proposed had been different, the effects of that error on the simulations would have been as acute. For the random error, an interesting pattern which reflects the histograms in Figure 1.23 emerges. While the effects on PC and DS for self-interest are minimal, for CA there is a clear bias towards the 0.33 parameter value. While for inequality aversion, CA appears to be slightly biased towards -1, while PC has very little change.

To provide a more concise summary of these figures Table 1.8 and Table 1.9, show mean, median and standard deviation statistics for Self-Interest and Inequality Aversion, respectively. Those numbers shown in bold, for each simulation, shows the closest value to the 'true' values. In both tables, CA is incredibly precise within the deterministic simulation. PC and DS are relatively accurate for α, but much less so for r. When error is added a different

[7]This approach draws from the Random Preference approach; used to enable analysis across all three designs.

Figure 1.25: Cumulative Frequency Plots; Comparing 'True' Parameter values to Established Parameter Values

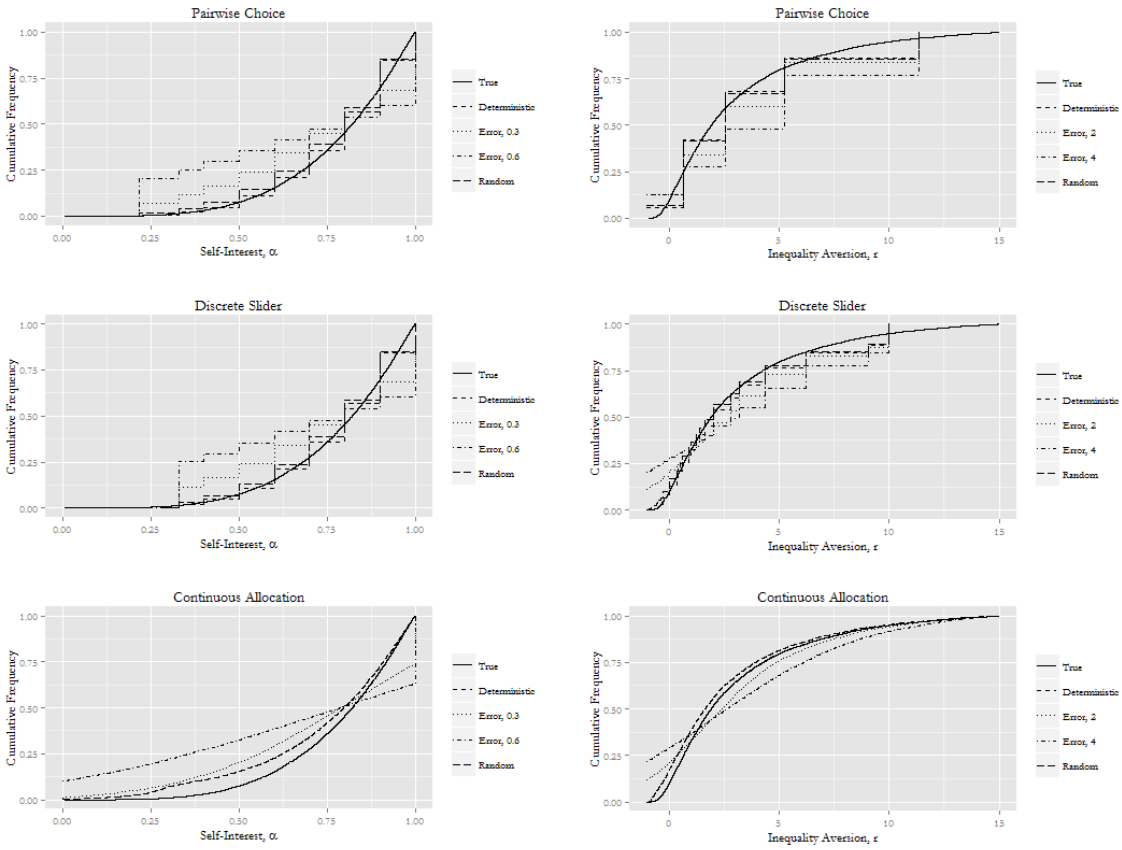

picture emerges. CA becomes less accurate the more error is added in. For α, DS becomes the most accurate for the mean and standard deviation statistics, and jointly dominant with PC for the median, in each of the two error and random simulations. For r, however, CA has the closest mean, while PC has the closest median and standard deviation in the first error simulation. But in the second simulation DS becomes dominant in the mean and standard deviation. For the random simulation CA remains superior for the mean, while DS is dominant in the other two statistics. Throughout each of the simulations, for both parameters, the median response for PC remains constant, always slightly imprecise but robust to error, the opposite of CA.

Table 1.8: Summary Table for Self-Interest: Comparing Simulations

	Deterministic			Error = 0.3			Error = 0.6			Random		
	Mean	Mdn	SD	Mean	Mdn	SD	Mean	Mdn	SD	Mean	Mdn	SD
True	*0.782*	*0.822*	*0.168*	*0.782*	*0.822*	*0.168*	*0.782*	*0.822*	*0.168*	*0.782*	*0.822*	*0.168*
PC	0.784	0.8	0.172	0.738	**0.8**	0.253	0.691	**0.8**	0.317	0.765	**0.8**	0.189
DS	0.784	0.8	0.17	**0.746**	**0.8**	**0.237**	**0.714**	**0.8**	**0.283**	**0.772**	**0.8**	**0.179**
CA	**0.782**	**0.822**	**0.168**	0.732	0.79	0.266	0.66	0.788	0.365	0.736	0.798	0.222

From these simulations we can pose several hypotheses. First, that assuming individuals make no error, the CA method is superior in terms of how accurately we can establish individual parameter values. As there are very few possible parameter values that cannot be arrived

Table 1.9: Summary Table for Inequality Aversion: Comparing Simulations

	Deterministic			Error = 2			Error = 4			Random		
	Mean	Mdn	SD	Mean	Mdn	SD	Mean	Mdn	SD	Mean	Mdn	SD
True	*4.593*	*2.126*	*8.028*	*4.593*	*2.126*	*8.028*	*4.593*	*2.126*	*8.028*	*4.593*	*2.126*	*8.028*
PC	3.862	2.555	7.848	3.898	**2.555**	**8.053**	4.016	**2.555**	8.331	5.286	2.555	14.13
DS	3.779	2.000	7.697	3.906	2.820	7.805	**4.181**	3.170	**8.043**	3.621	**2.000**	**8.067**
CA	**4.592**	**2.128**	**8.024**	**4.750**	2.709	8.160	5.155	3.135	8.490	**4.166**	1.846	7.723

at through the CA method, in comparison to PC and DS there is less potential for bias in the design of the design. However, if we pose that individuals are likely to be less precise in their estimates, and have more error, then CA may indeed bias their responses towards whichever error they are making. If we observe differences within the results from the experiments, then these are potential aspects that could ensure they are accounted for.

A.4.3 Simulated Data Estimates – Comparing Levels of Precision

While for both Pairwise Choice and Discrete Slider designs elicitation procedures can be used, for the Continuous Allocation design estimation is a necessity. While more complicated an additional benefit is that we can simulate data, according the model and error story constructed. This can be done before the experiment to identify potential pitfalls in the experiment and to reveal necessary considerations. By plugging in 'true' parameter values and running the model it can be observed how accurate the estimation will be for different assumed levels of error. For this estimation the 'true' parameter values of $\alpha = 0.7$ and $r = 0.5$ have been used with varying levels of precision, s, where; $s = s_1 = s_2$. By generating random 'actual' values for x_1, x_2 and x_3 for each of the 18 rounds, from the beta distribution, the estimation process for an 'estimated' α and r value could begin. In order to establish a large sample, 2000 separate random allocations were used; the resulting 'estimated' parameter values are shown in Figure 1.26.

What is clear is that they higher the s, and therefore lower the error, the more tightly clustered is the 'estimation' around the 'true' parameter values. Table 1.10 shows a summary of the parameter values from the simulations that were ran. To put these parameter values into context Table 1.11 shows the resulting X allocations from each of the corresponding mean and confidence interval values. There are several notable differences between each precision value. First is the size of the confidence intervals, perhaps obviously, the higher the degree of precision the smaller the confidence interval. Importantly the number of correct observations also differs, with a higher degree of precision the estimation procedure generates estimates 87.6% of the time, in comparison to 40.9% of the time in the lowest degree of precision. The number of observations could have been increased by specifying a larger number of initial starting values, therefore if there is a lower degree of precision in estimates then this number may need to be increased. While it is clear that the highest precision value has the closest mean parameter values, this is not the case for the second highest in comparison to the lowest. Indeed, it appear that a jump from 20 to 50, in terms of s, gives mean parameter values further away from the true values. Yet, when considering the optimal X allocations, it is clear that

Figure 1.26: Comparison of Simulated Data Estimates; $\alpha = 0.7$, $r = 0.5$

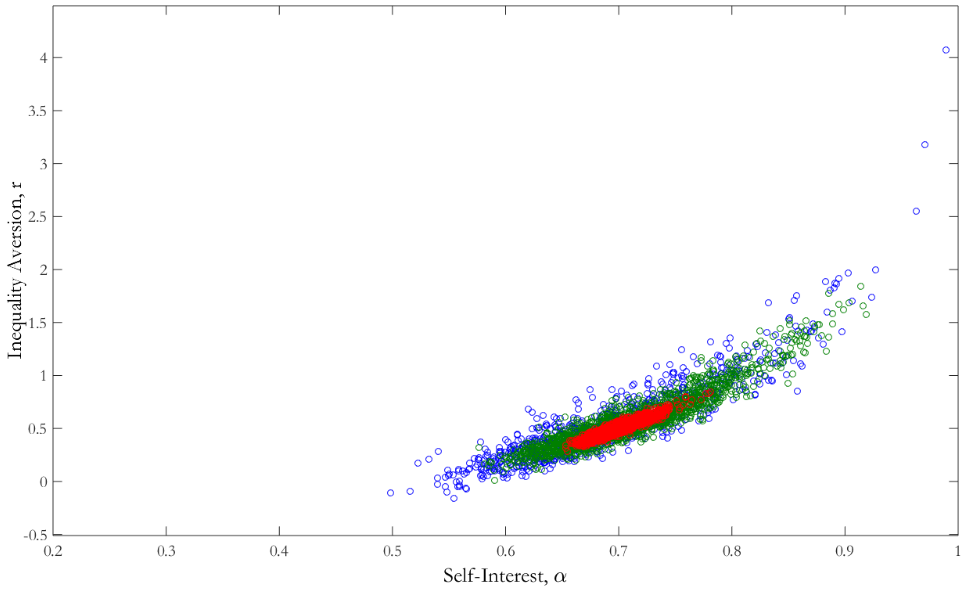

this holds. For the estimation procedure what is important is not the precise values of α and r, but what they mean in the context of the utility function, and the resulting optimal allocations it prescribes.

Table 1.10: Simulation Summary

Precision, s	Variable	Obs.	Mean	Std. Dev.	Lower 95% CI	Upper 95% CI
20	α	818	0.691	0.08	0.535	0.847
	r	818	0.555	0.398	-0.224	1.335
50	α	1548	0.713	0.059	0.597	0.829
	r	1548	0.584	0.264	0.066	1.102
1000	α	1751	0.701	0.018	0.665	0.737
	r	1751	0.508	0.077	0.356	0.659

The lessons to be drawn from these simulations is that the higher s the greater confidence can be put in the estimates. That if individuals are prone to making larger errors, then estimating their parameters is more difficult. But also that the parameter values on their own may have less importance than the implications they have, in terms of the optimal allocations they may give.

Table 1.11: Optimal X Allocations: Simulation Estimates

	x_1			x_2			x_3		
	Lower 95% CI	Mean	Upper 95% CI	Lower 95% CI	Mean	Upper 95% CI	Lower 95% CI	Mean	Upper 95% CI
20	17.83	16.94	17.39	3.7	5.26	5.35	8.47	7.8	7.26
50	17.42	17.3	17.75	4.57	5.17	5.19	8.02	7.54	7.06
1000	17.37	17.4	17.47	4.97	5.09	5.15	7.66	7.52	7.38
Actual		17.41			5.08			7.51	

A.5 Using the Beta-with-Effort-Bias Model to Detect Effort

The *bias* parameter within the *beta-with-effort-bias* model can arguably be used as a proxy for effort. There are two functional forms within the beta-with-effort-bias model, that associated with b is the *allocation* functional form, while $(1-b)$ the *outcome* functional form. The former can be seen as low effort, as individuals only consider a basic proportion of allocations to give to themselves, and the remainder to the others, not taking into account the *multiplication factors*. While the latter demands more cognitive effort to establish the final payoffs which will be consistent with their 'true' preferences. Figure 1.27 shows a scatter plot of the individual-level estimates of b against the average time taken to click Finish once the minimum time has ran out.

Figure 1.27: Scatter Plot and Fitted-Line; Average Time Taken vs Bias

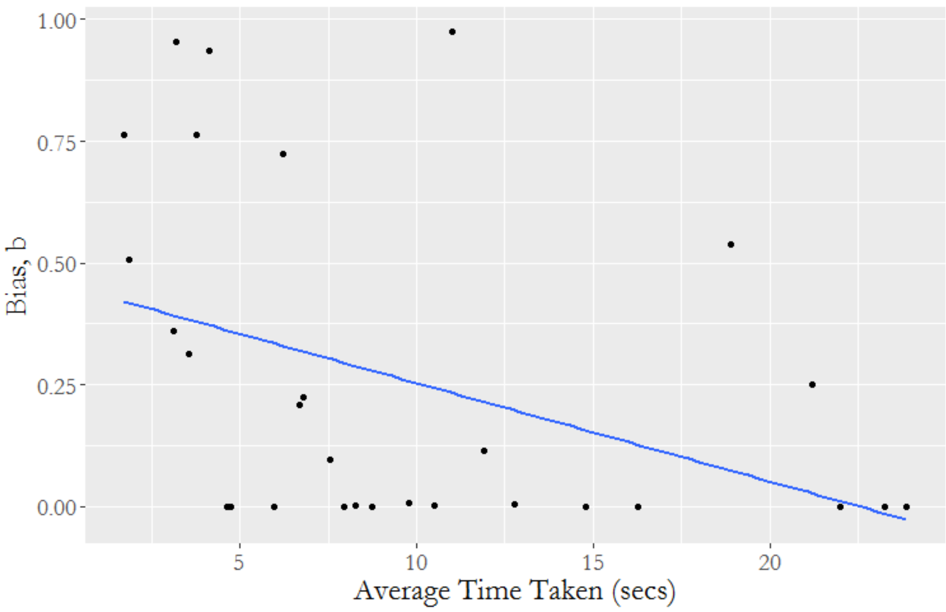

From a bootstrapped and clustered OLS regression we find a constant of 0.4554(0.1051) and an observed coefficient of -0.0202 (0.0079), which have p-values of 0.000 and 0.010 respectively. Meaning that for every extra second, after the minimum time had expired, the participant took their bias parameter was reduced by 0.0202. An individual who took an average of 10 seconds longer would have an expected bias parameter of 0.2534, and would take an average of 22.54 seconds extra to have a beta parameter of 0. The results from here are quite intuitive. Individuals who were willing to put more effort, as proxied by time, were more likely to have a lower bias parameter. The lower the bias parameter the closer the individual's responses fit to the *outcome* utility function, rather than the *allocation* function. In order to calculate the optimal response for an *outcome* individual more effort, and time, would be required.

A.6 Skillings-Mack Test

An alternative method to identifying significant differences between treatment effects is the Skiilings-Mack test statistic. Stemming from the Friedman test (M. Friedman, 1937), the Skillings and Mack test (Skillings and Mack, 1981) is a test which can be applied to data with an *unknown distribution*. Importantly, as the statistic is dependent upon ranks the usual assumption of normally distributed errors in not needed. Moreover it can be used in any block design with arbitrarily missing data and is appropriate for small samples, which is why here it is preferable to the Friedman test (Chatfield and Mander, 2009). The general form of the randomised block design is:

$$Y_{it} = \mu + \beta_i + \tau_j + \epsilon_{ij} \tag{1.22}$$

Where Y_{it} denotes the response for the jth treatment in the ith block. The overall mean is denoted by μ, the jth treatment effect by τ_j and the independent and identically distributed errors by ϵ_{ij}. The null hypothesis states that the treatment effects τ_j are all identical (Chatfield and Mander, 2009).

The SM test requires data to be arranged long into blocks, or treatments, then removes any block with only one observation. For those blocks left the observations are ranked, where ties are computed as the average rank of those tied. When there is missing data the ranks are centred, then the ranks are weighted; from this the Weighted Sum of Centred Ranks, A_j, is calculated. A simple covariance structure is produced, and the test statistic is calculated from the covariance of the treatment sums. This test statistic is then tested against the null hypothesis that all the treatment effects are equal, and is rejected if the SM statistic is greater than, or equal, to some critical value (Chatfield and Mander, 2009).

Within the SM tests conducted there are several reported numbers. Firstly, the Weighted Sum of Centred Ranks, denoted as W. Sum, shows which treatments have the higher or lower ranked observations. The Standard Errors are reported, alongside the Weighted Sum of Centred Ranks/Standard Errors which allows for an informal examination of these differences. The Skillings-Mack statistic will be reported, alongside two P-Values. The first P-Value is calculated for no-ties, and is a conservative estimate. The second is a simulated P-Value when ties occur, which is necessary (and less conservative) as when there are ties there is less variation amongst A_j (Chatfield and Mander, 2009).

A.6.1 Results

Table 1.12 shows the results from the Skillings-Mack test. The Empirical P-Values show that there are significant differences, to the 5% level, for all of the models, with the exception of (6) to the 10% level and (7) which is not significant. The WS/SE numbers are again intuitive, a positive sign implies a higher weight, which implies a higher EDE, which corresponds to greater self-interest and less inequality aversion, while a negative sign represents the opposite.

When considering (1) a clear pattern, similar to that of the above regressions emerges, each of the PS treatments are ranked significantly higher than either VOI or IS in all of the designs. However, the pattern for within perspective is less clear, indeed from the other models these conditional differences become apparent. Models (2), (3) and (4) do not reveal much more, confirming what was already known, and further showing that the effect of the assumption that $\alpha = 1/3$ makes a significant difference. The final models (5), (6) and (7) do, however show more detail. When considering IS there is no significant difference between designs, the EDE established appears robust regardless of the experiment design. However, when PS or VOI are assumed the differences between designs are mixed. For PS, PC gives a higher EDE, while CA gives a lower EDE with DS in the middle, as found from the previous regressions. Yet, for VOI it is DS which gives the lower EDE, PC the highest and CA lies in the middle.

It appears that regardless of the design the perspective significantly affects the EDE, but when considering the effect of difference of design it depends on which perspective the individual sees the problem from as to which biases the EDE in which way.

Table 1.12: Skillings-Mack Results: EDE Welfare

		All	Differences in Perspective			Differences in Design		
		(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**
	N	WS/SW	WS/SW	WS/SW	WS/SW	WS/SW	WS/SW	WS/SW
PC								
PS	26	5.01	3.25	.	.	1.68	.	.
VOI	28	-0.07	-1.73	.	.	.	2.31	.
IS	24	-0.44	-1.52	.	.	.	.	1.64
DS								
PlS	30	4.90	.	4.92	.	1.00	.	.
VOI	30	-2.99	.	-2.57	.	.	-0.49	.
IS	30	-2.97	.	-2.35	.	.	.	-1.00
CA								
PS	30	3.38	.	.	4.87	-2.62	.	.
VOI	30	-3.89	.	.	-2.95	.	-1.78	.
IS	29	-2.74	.	.	-1.93	.	.	-0.55
Skillings-Mack		89.67	10.57	24.22	23.98	7.08	5.91	2.753
P-value		0.000	0.005	0.000	0.000	0.029	0.052	0.253
E. P-value		0.000	0.002	0.000	0.000	0.018	0.056	0.254

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.7 Alternative Cluster Analysis Methods

A preliminary method for establishing cluster solutions is that of *agglomerative hierarchical clustering*. This establishes a hierarchy of different cluster solutions from which one solution can be chosen. The method begins by beginning with each observations as singleton clusters, calculates the distance between these clusters, then merges the two clusters with the smallest dissimilarities into a new single cluster. This is iterated until there is only one cluster remaining. There are two variable components herein which need to be chosen; the *distance* measure and the *linkage* method.

The *distance* measure defines the distances between the clusters, but are defined in several ways. Measures include: the Euclidean distance, weighted Euclidean distance, Mahalanobis distance, Manhattan distance and Maximum distance amongst many others. Similarly the different methods define how the clusters are to be linked; Complete Linkage (Sørensen, 1948), Single Linkage (Florek et al., 1951), Centroid Linkage (Sokal, 1958), Average Linkage (Sokal, 1958) and Ward's Linkage (Ward, 1963) are all examples of *linkage* methods. The nature of the clusters created is dependent upon these two components. The results shown below use the Euclidean Distance, as this is standard in the literature, and the Average Linkage method. Each of the alternative linkage methods above were conducted but results found identical clusters being formed. The Average Linkage method is defined to be the "average of all the distances between all pairs of points with one point from cluster A and one point from cluster B".[8]

One intuitive way in which to conduct hierarchical clustering is through a *dendrogram*. This graphical plot shows each observation (or cluster) as a leaf, the vertical line, which are joined at certain heights to form larger clusters, at the horizontal lines. The different clusters which emerge due to close proximity are show as the height is increase until only one cluster remains at the top. The larger the gap between clusters the further away they are, and less suitable they are to be considered a cluster. While the number of clusters left is subjective the Dendogram provides an intuitive way in which to visualise the potential clusters within the observations.

Figure 1.28: Cluster Dendogram; Proportional Payoffs to P1, P2 and P3

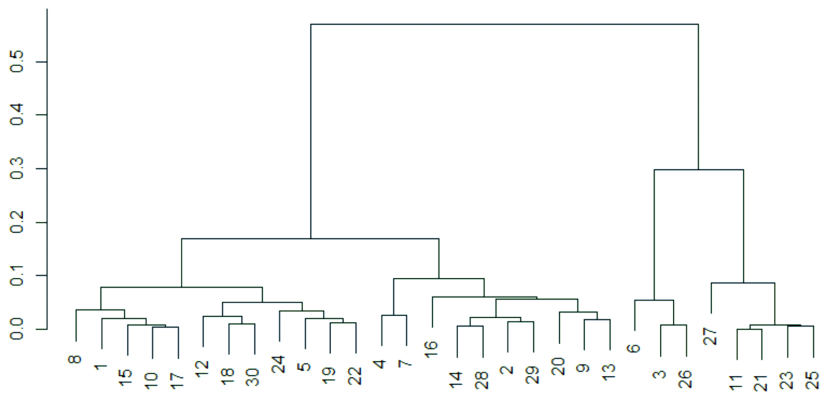

[8]Quoted from Nema Dean's workshop.

The figure shows a range of possible clusters which are apparent in the data. There is a definite two cluster scenario which appears with 22 individuals in the first cluster and 8 individuals in the second. A further 'cut' could be made where the second cluster is split into a group of 3 and 5; giving three clusters. Alternatively four cluster could be made, with the first cluster being split into 12 and 10. As mentioned previously this 'cut' is subjective, therefore further methods are required to ascertain the 'best' clusters.

The second method to be used relies upon the K-means approach (MacQueen, 1967). Unlike the previous method there is no hierarchical component. K-means tries to find an assignment of the observations into a predefined K number of clusters. It does this by "minimising the sum over all clusters of the sum of squares within clusters".[9] While again this can be done easily with r-code the predefined number K is still a subjective decision to make; K-means can find the 'best' clustering for a predefined number of clusters (given certain properties) but choosing which number K should be is more difficult. One method which allows a circumvention of this issue is by using *silhouette widths*. A silhouette width gives a criterion which allows comparison and ranking for different numbers of K clusters, so that the 'best' can be found. The silhouette width can be thought of as the dissimilarity of the cluster, to which an observation belongs, in comparison to its 'neighbouring' cluster (b(i)) minus the average dissimilarity between that observation and all other points within its cluster (a(i)) all divided by the maximum of the two:

$$s(i) = \frac{(b(i) - a(i))}{\max(a(i), b(i))}$$

The 'best' number of clusters is the one which gives the largest average silhouette width, over all the clusters. The plot below shows four possible K-means clusters, two, three, four and five. In each case the silhouette width is shown for each of the individuals, denoted the width of the bar, the cluster and the average silhouette width. From the results (and for testing for much higher levels of K) it appears that it is two clusters which are the 'best' solution. Indeed, notice that the two cluster solution for K-means is identical for the hierarchical clustering results.

While both *dendogram* and *silhouette* plots are useful at providing a visual aid for determining clusters issues emerge due to the nature of their construction. Both methods tend to try to divide the sample into equally sized spherical clusters. What this means is that they are unlikely to find clusters which show dependence between two variables. As this is certainly the case between our variables of interest, then another method can be used to generalise these methods. Rather than relying upon spherical clusters, *finite mixture models* allow these clusters to take any elliptical form. What this method does is to propose a 'mixture' of individual density functions which accurately fit the data, and it is this method which is used in the main analysis.

[9] Quoted from Nema Dean's workshop.

Figure 1.29: Silhouette Plot; 2, 3, 4 and 5 Clusters

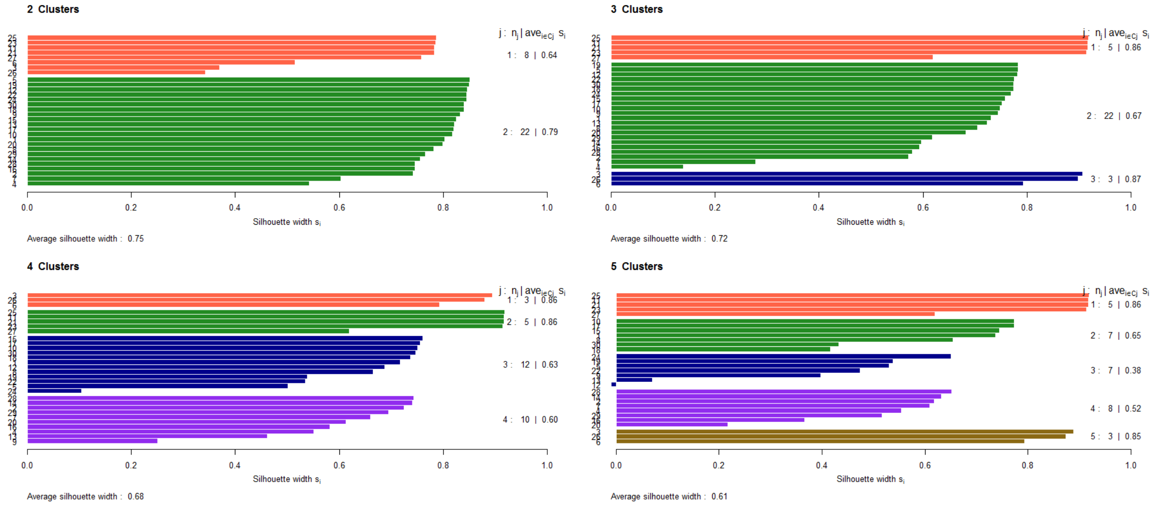

A.8 Relaxing the assumption of $\alpha = 1/3$

A further assumption to be made was that when considering the Veil of Ignorance and Impartial Spectator treatments is that individuals took an α value of 1/3. This assumption is consistent with the models assumptions for others in society, when the individual knows they are Player 1, as the weighting parameter for Player 2 and 3 is equal, and is simply $(1-\alpha)/2$. To further generalise this to the three person case, where the weight is again equal leads to $\alpha = 1/3$. Yet individuals may not exhibit these preferences, there could indeed be biases aligning their preference towards a certain player within the society. By using the Continuous Allocation method, this hypothesis is indeed testable. Rather than setting $\alpha = 1/3$, the estimation procedure can estimate α alongside the other parameters. Then two hypothesis can be tested for both IS and VOI; are the estimated α's significantly different from 1/3 and is the estimated r significantly different from those previously estimated values, when $\alpha \neq 1/3$? The latter of which is highlighted in Figure 1.30, as it shows the distribution of estimated inequality aversion parameters, for these differing assumptions regarding α.

Figure 1.30: Distribution of Inequality Aversion for Alternative α Values

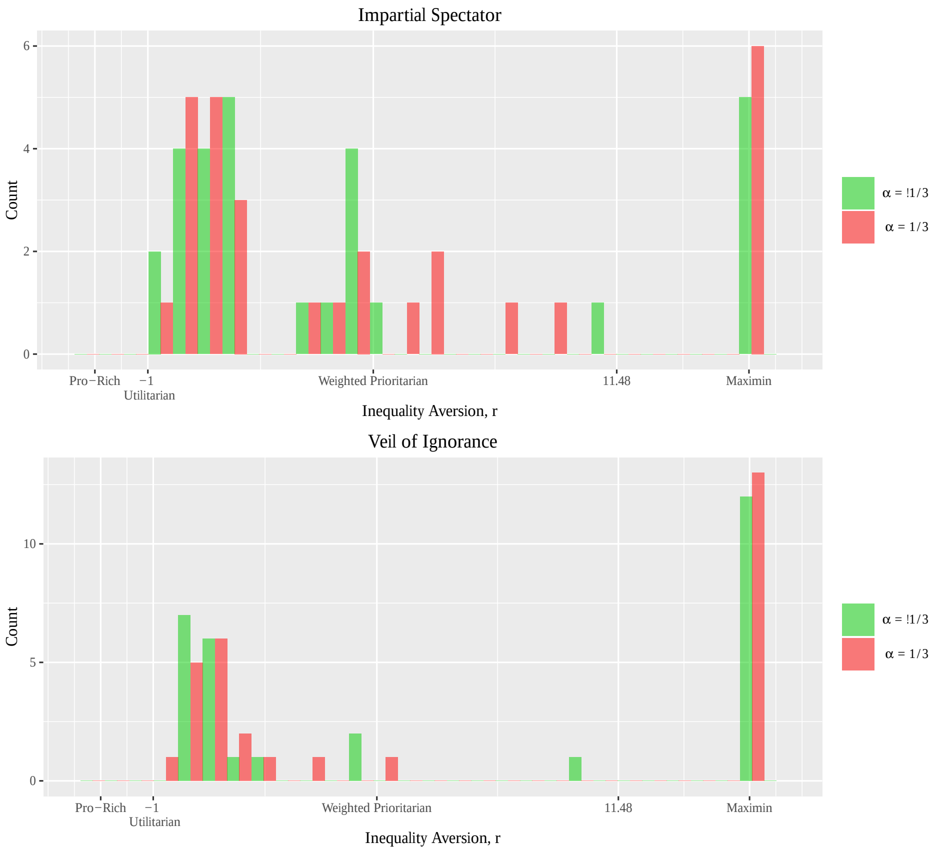

In terms of statistical tests there are two different tests, dependent upon the hypothesis being tested. Firstly, to test if α is significantly different from 1/3 a Paired T-Test was run for both IS and VOI treatments. Secondly, to test if r is significantly different when $\alpha = 1/3$ and $\alpha \neq 1/3$ a Wilcoxon Rank-Sum test was ran, due to issues with r approaching infinity. Table 1.13 shows the results below, in both cases, for both treatments the differences are not significant. The assumption that $\alpha = 1/3$, for VOI and IS, therefore seems like a reasonable assumption to make.

Table 1.13: Paired T-Test and Wilcoxon Rank-Sum Test Results

Paired T-Test	Impartial Spectator	Veil of Ignorance
Self-Interest, α	Mean (Std. Err.)	Mean (Std. Err.)
α=1/3	0.3333	0.3333
	0	0
α!=1/3	0.3729	0.3701
	-0.0358	-0.0446
$Pr(\mid T \mid > \mid t \mid)$	0.2783	0.4128
Obvs.	29	30

Rank-Sum Test	Impartial Spectator	Veil of Ignorance
Inequality Aversion, r	Rank Sum	Rank Sum
α=1/3	884	962
α!=1/3	827	868
$Pr > \mid z \mid$	0.6576	0.4871
Obvs.	29	30

A.9 Varying Parameters for the EDE Regression

One assumption made when comparing treatments concerns the EDE regressions conducted. What was assumed was that the society for which the EDE would be calculated was for an unequal society with a distribution of payoffs: [10, 1, 1], for P1, P2 and P3 respectively; and multiplication factors of [1, 1, 1]. Yet, if this initial distribution were to differ the EDE valuations of the society would change. Table 1.14, Table 1.15 and Table 1.15, show the effects of changing this initial distribution and multiplication factors on the magnitude and significance of the Random Effects model, previously specified. The former shows differences in P1's payoff, while holding that of P2 and P3 constant, while the latter varies the payoffs of P3 while holding the other two constant. The results are perhaps to be expected, when the distribution becomes more unequal (in Table 1.14) both constants and magnitudes increase. When the society is all but equal the average EDE valuations shrink to almost one, and the resulting treatment effects are also reduced. All treatment effects do, however, remain significant throughout. When considering Table 1.15 very little changes, both in regards to significance and magnitude. As long as the distribution is unequal then differences between treatments are clear. Table 1.16 shows a similar pattern, the significance and directions of the coefficients change little, the effects of altering the multiplication factors just multiply the magnitude of the effects. Indeed, it appears that the distribution and multiplication factors assumed make little difference, what is clear is that the magnitudes of the coefficients are only important in the context of the specific distribution assumed.

Table 1.14: Sensitivity Analysis for EDE Regression: Changes in x_1

	(1) [1.1,1,1] Coef./S.E.	(2) [2,1,1] Coef./S.E.	(3) [5,1,1] Coef./S.E.	(4) **[10,1,1]** Coef./S.E.	(5) [20,1,1] Coef./S.E.	(6) [100,1,1] Coef./S.E.
Perspective						
Veil of Ignorance	-0.0486***	-0.4603***	-1.5197***	-3.0730***	-5.9610***	-26.8359***
	(0.003)	(0.030)	(0.130)	(0.313)	(0.683)	(3.671)
Impartial Spectator	-0.0468***	-0.4508***	-1.4908***	-3.0072***	-5.8145***	-26.0023***
	(0.003)	(0.030)	(0.132)	(0.317)	(0.693)	(3.726)
Design						
Discrete Slider	-0.0069**	-0.0634**	-0.3096**	-0.7717**	-1.7402**	-9.8112***
	(0.003)	(0.030)	(0.134)	(0.320)	(0.700)	(3.761)
Continuous Allocation	-0.0064**	-0.0889***	-0.4906***	-1.2203***	-2.7642***	-16.0574***
	(0.003)	(0.030)	(0.134)	(0.321)	(0.702)	(3.771)
Constant	1.0754***	1.6571***	3.1752***	5.4361***	9.7162***	41.8047***
	(0.003)	(0.033)	(0.142)	(0.335)	(0.723)	(3.839)
N	30	30	30	30	30	30
Observations	257	257	257	257	257	257
R-squared	0.5710	0.5096	0.3819	0.3171	0.2773	0.2284
ρ	0.1278	0.1843	0.1636	0.1419	0.1293	0.1154

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.15: Sensitivity Analysis for EDE Regression: Changes in x_3

	(1) **[10,1,1]** Coef./S.E.	(2) [10,1,2] Coef./S.E.	(3) [10,1,4] Coef./S.E.	(4) [10,1,6] Coef./S.E.	(5) [10,1,8] Coef./S.E.	(6) [10,1,10] Coef./S.E.
Perspective						
Veil of Ignorance	-3.0730***	-3.1262***	-3.1104***	-3.0702***	-3.0218***	-2.9677***
	(0.313)	(0.304)	(0.307)	(0.315)	(0.327)	(0.342)
Impartial Spectator	-3.0072***	-3.0571***	-3.0518***	-3.0185***	-2.9718***	-2.9145***
	(0.317)	(0.308)	(0.311)	(0.320)	(0.332)	(0.347)
Design						
Discrete Slider	-0.7717**	-0.7394**	-0.7011**	-0.7012**	-0.7274**	-0.7740**
	(0.320)	(0.311)	(0.315)	(0.324)	(0.336)	(0.351)
Continuous Allocation	-1.2203***	-1.1803***	-1.1530***	-1.1736***	-1.2253***	-1.3002***
	(0.321)	(0.312)	(0.315)	(0.324)	(0.337)	(0.352)
Constant	5.4361***	5.7211***	5.9416***	6.0871***	6.2161***	6.3414***
	(0.335)	(0.330)	(0.340)	(0.354)	(0.370)	(0.388)
N	30	30	30	30	30	30
Observations	257	257	257	257	257	257
R-squared	0.3171	0.3294	0.3155	0.2952	0.2732	0.2508
ρ	0.1419	0.1605	0.1834	0.1969	0.2050	0.2094

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.16: Sensitivity Analysis for EDE Regression; Changes in π

	(1) **[1,1,1]** Coef./S.E.	(2) [1,1,2] Coef./S.E.	(3) [1,1,0.5] Coef./S.E.	(4) [2,1,1] Coef./S.E.	(5) [5,1,1] Coef./S.E.	(6) [10,1,1] Coef./S.E.
Perspective						
Veil of Ignorance	-3.0730***	-3.1262***	-3.0584***	-5.9610***	-13.9563***	-26.8359***
	(0.313)	(0.304)	(0.346)	(0.683)	(1.793)	(3.671)
Impartial Spectator	-3.0072***	-3.0571***	-2.9729***	-5.8145***	-13.5540***	-26.0023***
	(0.317)	(0.308)	(0.348)	(0.693)	(1.820)	(3.726)
Design						
Discrete Slider	-0.7717**	-0.7394**	-0.8448**	-1.7402**	-4.7371***	-9.8112***
	(0.320)	(0.311)	(0.349)	(0.700)	(1.837)	(3.761)
Continuous Allocation	-1.2203***	-1.1803***	-1.3527***	-2.7642***	-7.7235***	-16.0574***
	(0.321)	(0.312)	(0.358)	(0.702)	(1.842)	(3.771)
Constant	5.4361***	5.7211***	5.0490***	9.7162***	21.9271***	41.8047***
	(0.335)	(0.330)	(0.364)	(0.723)	(1.885)	(3.839)
N	30	30	30	30	30	30
Observations	257	257	250	257	257	257
R-squared	0.3171	0.3294	0.2805	0.2773	0.2447	0.2284
ρ	0.1419	0.1605	0.1379	0.1293	0.1212	0.1154

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Chapter 2

Incorporating Oneness: A Ugandan Lab-in-the-Field Experiment

Preferences relating to *inequality aversion*, *self-interest* and *oneness* (the closeness of connection to others) are incorporated in a structural model and estimated in order to explain prosocial behaviour. An incentivised lab-in-the-field experiment was run in Mbale, Uganda (n=156), with both general population and student samples. The experiment was a modified three-person dictator game, run on touch-screen tablets. Decision problems were repeated (54 rounds) to ensure individual-level preferences could be estimated; using the Dirichlet distribution to rationalise *noisy* behaviour. Two within-subject treatments varied if the identity of the 'recipients' was *anonymous* or *known*. Results find extensive heterogeneity in prosocial behaviour, which is accounted for through individual preference parameters. On average, there is a substantial regard for others with a preference for reducing inequality, rather than increasing efficiency. Oneness is found to have large and significant effects on giving; with distinctions between *self-other* and *between-other* trade-offs emerging.

***

1 Introduction

Inequality in society is, in part, determined by the distributional preferences held by individuals. The heterogeneity of such preferences leads to differences in prosocial behaviour. When making distributional decisions, some prefer to allocate resources equally, while others give priority to themselves. Notions of efficiency are important for particular individuals, but of little concern to others. Yet, it is not only distributional preferences which determine prosocial behaviour; the 'recipient's' identity is also an integral consideration. Particular individuals will maintain the same prosocial behaviour regardless of the identity of the 'recipient', while some will exhibit preferential treatment to those they are closely connected to.

Developed here is a utility function which incorporates *inequality aversion*, the trade-off between equality and efficiency, and *self-interest*, the degree to which the 'self' is weighted in relation to the 'other'. Preferences parameters are introduced which reflect behavioural responses to changes in *oneness*, the closeness of connection to others. These parameters distinguish between how elastic *self-other* and *between-other* trade-offs are, to better explain the distributional effects that differential *oneness* can have. To account for 'noise' in decision making, the Dirichlet distribution is proposed as a random behavioural model. Together the proposed utility function and stochastic error model seek to explain individual behaviour, by accounting for heterogeneity in preferences and intuitively modelling 'noise' in decision making.

In order to observe individual behaviour, an incentivised lab-in-the-field experiment was run in Mbale, Uganda. The form of the experiment was a modified three-person dictator game, where participants distributed payoffs amongst themselves and two 'others'. The sample (n=156) was made of two separate groups, general population (n=108) and students (n=48). Two within-subject treatments were administered over 54 rounds. In the *budget* treatment, the 'others' were anonymous and budget (stake-size) was varied. In the *oneness* treatment, anonymity was lifted as the identity of the 'others' was made known. Decisions were made on touch-screen tablets, using a Z-Tree interface (Fischbacher, 2007). Alongside the experiment, an in-depth survey was conducted to establish an extensive list of demographic characteristics.

This work is situated within the *distributional preferences* literature, within which exists a wide body of research.[1] Papers by Andreoni and Miller (2002) and Fisman, Kariv, and Markovits (2007), are most closely related; where modified dictator games are used to observe the prosocial behaviour of subjects. Both estimate preference parameters within CES utility functions; where the former uses two-person dictator games, and the latter extends the analysis to focus on individual preferences and includes a three-person variant of the dictator game.

Several papers have incorporated *social distance* as an explanatory construct within distributional decision making. Leider et al. (2009) use an online field experiment, with real world social network data, to study prosocial behaviour. Using modified two-person dictator games, as in Andreoni and Miller (2002), they find that as social distance increases generosity

[1] Including: E. Fehr and Schmidt (1999); Bolton and Ockenfels (2000); Charness and Rabin (2002); Cox, D. Friedman, and Gjerstad (2007); Cappelen et al. (2007); Dolan and Tsuchiya (2009); Breitmoser (2013) and Jakiela (2013)

decreases. Branas-Garza et al. (2010), Goeree et al. (2010), Ligon and Schechter (2012) and Binzel and D. Fehr (2013) reveal similar trends; where the former two use student samples, in a laboratory setting, and the latter two run lab-in-the-field experiments in rural Paraguay and Cairo, respectively.

While the above predominantly establish social network data through named relationships (i.e. friend, friend-of-a-friend, stranger) an alternative is to characterise the closeness of relationship through 'oneness'; a 'measure of perceived self-other overlap' (Cialdini et al., 1997). Gächter, Starmer, and Tufano (2015) adopt the oneness scale, from the psychology literature; which is calculated as the average of the 'Inclusion of the Other in the Self' (IOS) scale (A. Aron, E. Aron, and Smollan, 1992) and the 'we-scale' (Cialdini et al., 1997). The advantage of such a scale, is that it provides a numerical index of the closeness of connection; without delving into its determinants.

This paper seeks to contribute to the above literature; by proposing a CES utility function which incorporates *oneness levels*, to account for the explanatory power that social distance can have on individual decision making. By estimating preferences relating to *inequality aversion*, *self-interest* and *oneness*, the intricacies of their interactions are explored and individual level behaviour, observed from a lab-in-the-field experiment, is explained.

2 Experiment

2.1 Experimental Session

The general form of the experiment is a modified three-person 'dictator' game. 'Dictators' are given a budget, m, which they must distribute amongst three players. Player 1 denotes the *self*, while Player 2 and 3 are two *other* real participants (the 'recipients'). The *allocations*, x_i, are chosen for each Player i; where $i \in [1, 2, 3]$ and $\sum_i^3 x_i = m$. These *allocations* are then divided by the corresponding *divider*, $1/\pi_i$, to give the *payoff*, $\pi_i x_i$, to each Player i. The *dividers* change the relative prices of giving; meaning that equality-efficiency trade-offs need to be made.

Within the experiment there are multiple rounds, 54 in total. There are two treatments; the *budget* treatment and the *oneness* treatment, each with 27 rounds. In each round the participants are randomly assigned to a group of three. The dividers change every round; ensuring the relative price of giving to each player varies. Table 2.1 shows how the dividers change; only Rounds 1 to 9 are shown, but this pattern is repeated every nine rounds. Note that the dividers are such that each player has the same average divider, over all rounds, and that they each have the same number of 1's, 2's, 3's and 4's that the other has.

Within each experimental session there are six participants. Each of the participants make individual decisions; *as if* they were the 'dictator'. One individual's decisions, from each group of three, is randomly selected (at the end of the experiment) to determine the payoffs of each member of their group. It is from *one randomly selected round* that all participants receive their payoffs, determined by the 'dictators' of that round. In this way, participants are

Table 2.1: Dividers per Round

	Dividers, $1/\pi_i$		
	Player 1	Player 2	Player 3
Round 1	1	1	1
Round 2	1	2	2
Round 3	2	1	1
Round 4	1	2	3
Round 5	2	3	1
Round 6	3	1	2
Round 7	1	2	4
Round 8	2	4	1
Round 9	4	1	2

incentivised; as each distributional decision they make has an equal chance of determining their payoff and the payoffs of two other individuals within the room. Importantly, each decision is entirely anonymous and without feedback; participants neither know the decisions of any other participants nor the identity of the 'dictator' in any round. This removes considerations of reputation and reciprocity, allowing for 'pure' altruism to be identified.

Participants make their decisions using a Z-Tree interface, on touch-screen tablets. They are given extensive instructions, including an interactive on-screen tutorial to enable them to use the tablets. A screenshot of the interface is shown in Figure 2.1. There are three players, Player 1, 2 and 3, amongst whom participants must make allocations, so that the remaining budget reaches zero. Each player has a divider (changing every round), which is used to calculate the payoff to that player. Allocations can be made by using: the slider, arrow keys and written input. The slider (the black bar) can be dragged to make allocations, the arrow keys tapped to make incremental changes (1 or 10), and the written input used to type exact amounts. Calculations of the payoffs are made automatically, and are shown by both the orange numbers and by the height of orange bars. The gap between payoffs, the highest payoff minus the lowest payoff, and the sum of payoffs (aptly named) are shown. All allocations, payoffs and budgets are in Ugandan shillings (*shs*).

2.1.1 Budget Treatment

Within the budget treatment, the budget (or stake-size) is varied; $shs30,000$, $shs60,000$ or $shs90,000$. The order of this variation was random, however, in order to aid the understanding of the participants it was varied only every nine rounds. For the first nine rounds there was a budget of A, B for the next nine and C the last. The order was random for each participant, but each participant had all of the budgets mentioned above. Importantly, within the budget treatment the other participants remained anonymous. Player 2 and 3 were randomly assigned each round, so while they were known to be selected from the other five participants, their identity was *not known*.

Figure 2.1: Z-Tree Interface

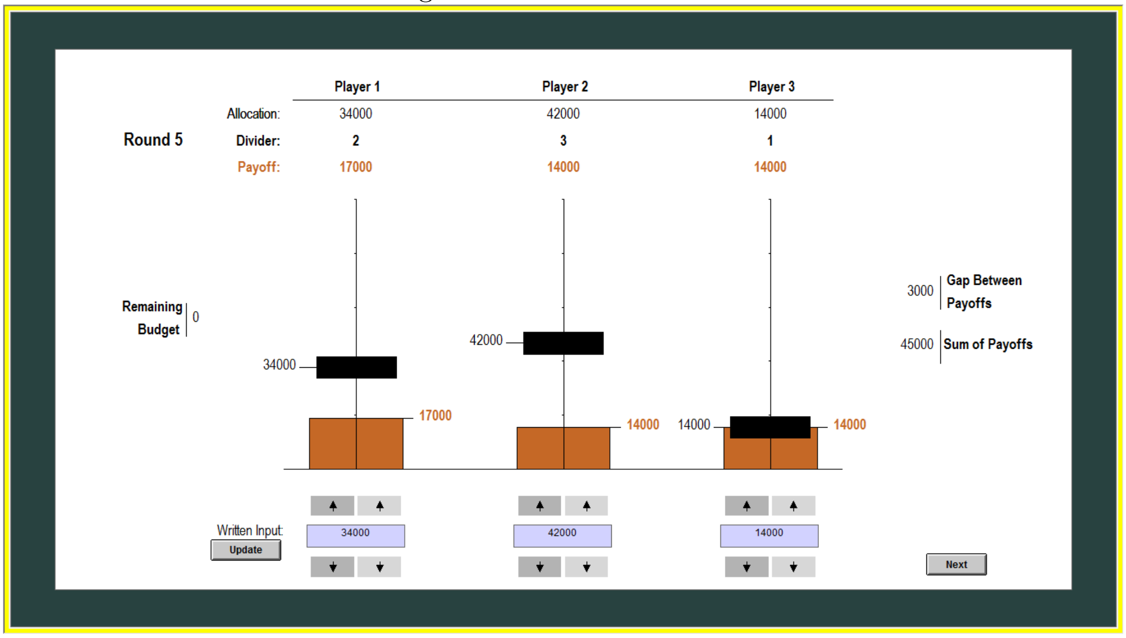

2.1.2 Oneness Treatment

The oneness treatment was preceded by the oneness questionnaire, based on the work by Gächter, Starmer, and Tufano (2015).[2] Oneness denotes the degree of closeness that the respondent (the 'self') feels towards another individual (the 'other'). It is more precisely defined here as the connection the self feels in everyday life with the other. Figure 2.2 shows the computer interface with diagrams which represent the 'oneness level' - an integer scale between 1 ('most distant') and 7 ('closest'). The simple tool effectively captures what a multitude of psychometric questions could do, in a quick, easy and accessible form. Participants were asked the question in Figure 2.2 for each of the other five participants in the room.

Following this they began the oneness treatment. Unlike the budget treatment, here the budget was kept constant (at $shs60,000$). But crucially participants now knew the identity of Player 2 and Player 3, revealed by their desk number. They were encouraged to look around, every round, to see which players were in their group and only then to make their decisions. As before, the participants within each group randomly changed every round and decisions were made anonymously, as before.

2.2 Survey Session

Alongside the experiment an in-depth survey was run. The surveys were implemented with the World Bank's Survey Solutions, on Asus touchscreen tablets, and ran one-on-one; one

[2]The interface is derived from A. Aron, E. Aron, and Smollan (1992), the IOS ('Inclusion of the Other in the Self') Scale. Gächter, Starmer, and Tufano (2015) introduce *oneness* to the economics literature constructing their 'Oneness Index' from both the IOS scale and the We-Scale (Cialdini et al., 1997). Due to language barriers, the We-Scale was dropped from our experimental design, so our 'oneness level' is measured only by the IOS scale. The words 'in everyday life' were used instead of 'before the experiment', as the cultural interpretation of 'before the experiment' would have meant the time immediately before the experiment, which was not desirable.

Figure 2.2: Oneness Questionnaire

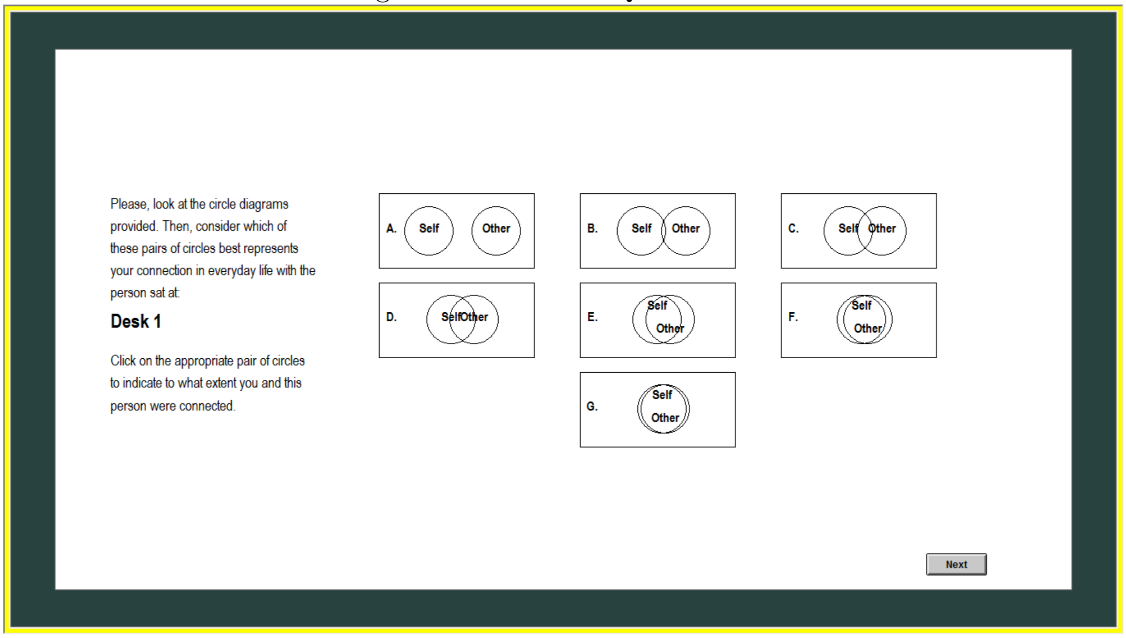

enumerator-one participant. The survey was split into four sections: Individual Characteristics; Household Characteristics; Assets, Wealth and MPI; and Preferences. The survey allowed for the creation of variables which represented individual level characteristics; including the creation of a Wealth Index (WI) and Multidimensional Poverty Index (MPI).

2.3 Sampling and Sample

Two sampling frames were created in order to recruit participants for the experiments. The first, from a general population in the Mbale District, the second, from student records from the Ugandan Christian University (UCU), Mbale. Between experimental days there were two separate samples, the general population (n=102) and the student population (n=48).

Further details of the experimental design (B.1), sampling (B.2), sample characteristics (B.3), wealth and multidimensional poverty indices (B.4) alongside the script (B.5) and tutorial script (B.6), can be found in the Appendix.

3 Theoretical Model

3.1 Utility

The theoretical model proposed takes the form of a Constant Elasticity of Substitution (CES) function, similar to that of Andreoni and Miller (2002), where utility is equivalent to the Equally Distributed Equivalent (EDE).[3] Utility is determined by payoffs, $x_i\pi_i$, distributed

[3]The EDE represents the mean level of payoffs, if equally distributed, which would ensure the individual was indifferent between that and the current distribution of payoffs (Atkinson, 1970). It provides a meaningful ranking of alternative distributions of payoffs (or income) and is, as such, a Social Welfare Function.

amongst the 'self' ($i = 1$) and 'others' ($i > 1$). Individuals striving to maximise utility would, then, make decisions based upon distributing allocations, x_i according to: their preferences, r, α, ϕ and ψ; how *closely connected* they were to the others θ_i; and the multiplication factors: π_i, the reciprocals of the dividers, $1/\pi_i$.

Our utility function is, then:

$$U_1 = \left(\sum_{i=1}^{N} \left(\omega_i (\pi_i x_i)^{-r} \right) \right)^{-\frac{1}{r}} \tag{2.1}$$

Where:

$$\omega_1 = \frac{\alpha}{\alpha + \frac{1-\alpha}{n} \sum\limits_{j=2}^{N} \theta_j^{\phi}}, \qquad \omega_{i \neq 1} = \frac{\theta_i^{\psi}}{\sum\limits_{j=2}^{N} \theta_j^{\psi}} (1 - \omega_1)$$

Self-interest is denoted by α, where $0 \leq \alpha \leq 1$, and can be thought of as the extent to which the individual weights *themselves*, in relation to *others*. Where n is the number of *others* (and $N = n + 1$) the weight to each other is denoted by $(1 - \alpha)/n$. If $\alpha = 1$ (and therefore $1 - \alpha = 0$) individuals are 'egoists', whose utility is only dependant upon their own payoff. When α decreases more weight is put upon others; where $\alpha = 1/N$ reflects equal weighting of themselves and others.

Inequality aversion is represented by r, where $-1 \leq r \leq \infty$ and $r \neq 0$, and can be interpreted as the trade-off individuals are willing to make between *efficiency* and *equality*. When $r = -1$ preferences reflect a concern for *efficiency*, or 'Utilitarianism', where utility is determined by summing (weighted) payoffs. As r increases less concern is given to *efficiency*. 'Cobb-Douglas' preferences are represented when $r \to 0$; which implies that optimal allocations reflect the proportions set by ω_i. As r increases more weight is placed upon the payoff of the worst-off, 'Weighted Prioritarianism' (Parfit, 1997), until $r = \infty$ which represents 'Maximin' preferences, where only increases to the worst-off increase utility (Rawls, 1999).

The closeness of connection to others is characterised by the oneness levels denoted by θ_i, for each respective i.[4] These levels are bounded between 1 ('most distant') and 7 ('closest'). The inclusion of oneness levels, within the functional form, allows for two additional properties: (A) in general, as connectedness to the others increases, *less* weight is given to Player 1; (B) the greater the connectedness to Player 2, relative to Player 3, the *more* weight is given to Player 2, relative to Player 3, and vice versa. The parameters ϕ and ψ are the *oneness elasticities*; determining the responsiveness of the individual weights, ω_i, to changes in oneness levels, where $\phi, \psi \geq 0$. The *self-other oneness elasticity*, ϕ, determines the responsiveness of (A), while the *between-others oneness elasticity*, ψ, determines that of (B). The higher ϕ the greater the willingness to trade-off between the self and others, as oneness varies; the higher ψ the greater the willingness to substitute between others as their relative oneness levels change.

[4]In the Budget treatment, where the identity of each player is unknown, θ_i is calculated as the 'expected oneness level'. An alternative function which could model 'expected utility' rather than 'expected oneness' is formulated in Appendix B.11. But due to reasons of tractability of the model the latter was assumed. Sensitivity analysis in Appendix B.7, compares parameter estimations using data from all 54 rounds and only the 27 rounds of the *oneness* treatment (to circumvent this assumption).

The oneness elasticities, ϕ and ψ, allow distributional decisions to be affected by how closely connected to the others the individual is. Self-interest, α, represents the weight the individual gives to themselves, when they have the least possible connection to the others. The higher ϕ and ψ the greater the weights deviate from α when oneness levels increase. Indeed, when $\theta_2 = \theta_3 = 1$ or $\phi = \psi = 0$, then the weights ω_i are determined solely by α.

Given the above utility function and the budget constraint $m = \sum_{i=1}^{N} x_i$, where m is the budget, the following optimal allocations (which maximise utility) can be obtained, $\forall i$:

$$x_i^* = \frac{m}{1 + \sum_{j \neq i}^{N} \left(\frac{\pi_i}{\pi_j} \left(\frac{\omega_j \pi_j}{\omega_i \pi_i} \right)^{\frac{1}{1+r}} \right)} \tag{2.2}$$

Figure 2.3 illustrates how the optimal allocations change due to different preferences; both are surface plots where the height denotes an optimal allocation for a set of preference parameters. The left panel plots x_1^* for differing α and r values; setting $\theta_2 = \theta_3 = 1$, $\pi = [1, 1, 1]$ and $m = 1$. As α increases, x_1^* increases, until $\alpha = 1$ where individuals take all for themselves. As r increases, x_1^* approaches equal sharing, while as $r \to -1$ efficiency motivates allocations to reflect the weights according to α.

Figure 2.3: Optimal Allocations to P1 and P2: r, α, ϕ and ψ

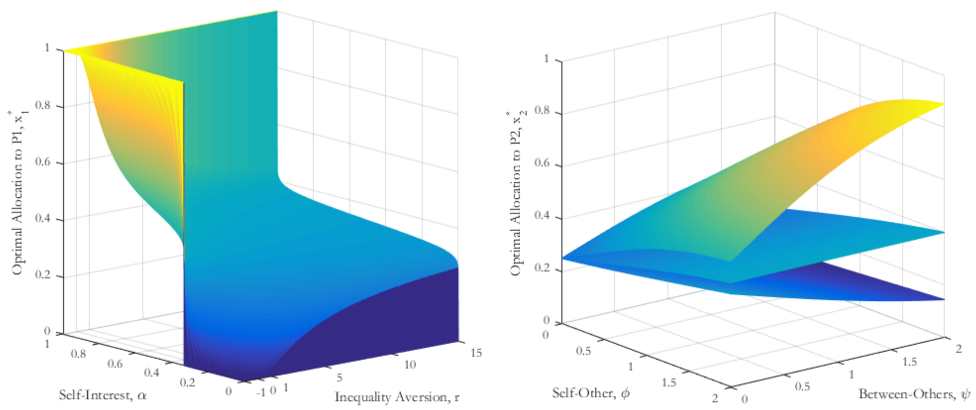

To illustrate the effect of the *oneness elasticities* the right panel shows how x_2^* differs due to ϕ and ψ. Three different oneness levels are represented by the three surfaces; in order from highest to lowest $(\theta_2, \theta_3) = (7, 2), (2, 2)$ and $(1, 2)$. Constant throughout, is: $\alpha = 0.5, r \approx 0$, $\pi = [1, 1, 1]$ and $m = 1$. The surfaces intersect at $x_2^* = 0.25$, when $\phi = \psi = 0$, but as ϕ and ψ increase they diverge due to the different oneness levels. For the highest surface, when $\theta_2 = 7$ and $\theta_3 = 2$, x_2^* increases as ϕ and ψ increase; as the oneness level is higher for P2 the rate of increase is high. For the middle surface, oneness levels are equal and both greater than 1 ($\theta_2 = \theta_3 = 2$); therefore, increases in ψ have no effect on x_2^*, as no between-other substitution occurs, however, as ϕ increases so does x_2^*, as trade-offs are made between the self and others. The lowest surface, when $\theta_2 = 1$ and $\theta_3 = 2$, shows a decrease to x_2^* as ϕ, ψ increase; due to an increase in ω_3 relative to ω_2.

As an alternative to the utility function, a *heuristic model* is formulated in Appendix B.9. Instead of maximising utility, participants are assumed to follow heuristic rules-of-thumb and make decisions sequentially. Individual-level parameters are estimated for the proportion to the self, the proportion between others and rounding precision.

3.2 Dirichlet Error Modelling

While the above utility model provides precise *optimal* allocations, x_i^*, for a particular decision problem and preference set, participants are assumed to make 'error' when calculating, or choosing, these allocations. Instead, we assume they draw their *actual* allocations, x_i, from the Dirichlet distribution (Dirichlet, 1839); where the expected values, $E[X_i]$, equal the *optimal* allocations, x_i^*.

This is an extension of work done by Hey and Panaccione (2011) and Hey and Pace (2014) where the Beta distribution is used for similarly bounded decision problems, for ambiguous and risky decisions. The Dirichlet distribution is a multinomial Beta distribution, allowing for N variables, which here correspond to individual allocations (i.e. $x_1, x_2, ..., x_N$), where $x_i \in (0, 1)$ and $\sum_{i=1}^{N} x_i = 1$. The below formulates the Dirichlet distribution as a random behavioural model.

The following assumptions are made: (1) $E[X_i] = x_i^*$, and (2) $Var(X_i) = \frac{(x_i^*(x_0^* - x_i^*))}{s}$, therefore:

$$E[X_i] = \frac{a_i}{a_0} = x_i^* \tag{2.3}$$

$$Var(X_i) = \frac{(a_i(a_0 - a_i))}{(a_0^2(a_0 + 1))} = \frac{(x_i^*(x_0^* - x_i^*))}{s} \tag{2.4}$$

Where:

$$a_0 = \sum_{i=1}^{N} a_i, \qquad x_0^* = \sum_{i=1}^{N} x_i^*$$

It follows that, $\forall i$:

$$x_i^*(s - 1) = a_i \tag{2.5}$$

The a_i's determine the shape of the Dirichlet probability density function (*pdf*) and represent the weight given to a particular i. Precision, or how *noisy* decisions are, is represented by s. Note that, the higher the value of s, and therefore the higher α_0, the lower the variance will be. To illustrate the above, Figure 2.4 shows the *pdf*'s of alternative Dirichlet distributions, where $N = 3$ and $x_3 = 1 - x_1 - x_2$. The left shows an imprecise individual, $s = 10$, who aims to allocate equally $E[X] = [0.33, 0.33, 0.33]$, with $A = [3, 3, 3]$. Second, with $A = [10, 6, 6]$, an individual allocating slightly more to themselves, $E[X] = [0.45, 0.27, 0.27]$, with a greater deal of precision, $s = 23$. Third, with $A = [4, 1, 1]$ more self-interested preferences, here $E[X] = [0.67, 0.17, 0.17]$, can be represented; with a mode where $x_1 \to 1$ (here precision is low

($s = 7$) but precision can be increased). The flexibility of the Dirichlet distribution is a useful property, and the above derivations allow for easily interpretable parameters to be estimated.

Figure 2.4: Dirichlet Distribution: Probability Density Function

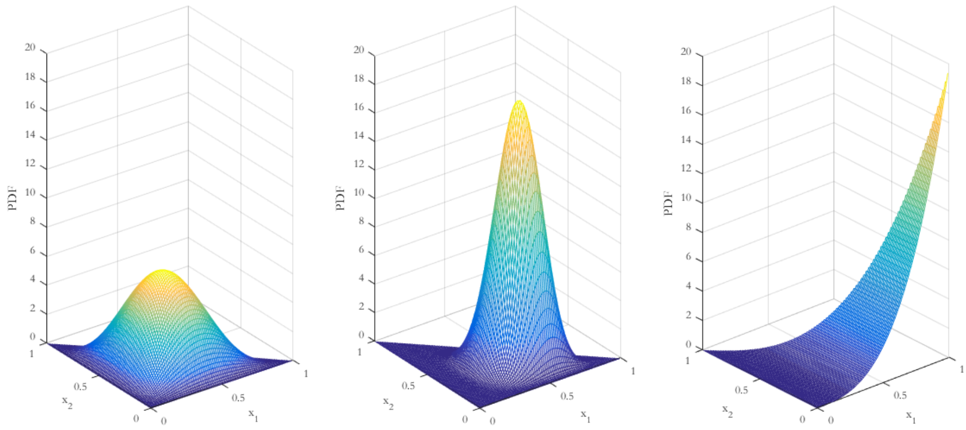

The preference parameters: α, r, ϕ, ψ and s; are estimated, for each individual, through maximising the following log-likelihood function. The preference parameters determine the optimal allocations, x_{it}^*, and consequently the shape parameters, a_{it}, in each round $t \in T$. The multiple integral of the *pdf*, determined by a_{it}, is taken over the n-dimensional 'rounding' interval V_t. V_t is determined by the *observed* decisions, x_{it}; where the 'rounding' interval, around the *observed* decision, is necessary as decisions are not strictly continuous (only to the nearest shs). Estimated parameters are those which maximise the log-likelihood function, hence are the 'most likely' fit for the observed data.

$$\sum_{t=1}^{T} \log \left(\int \cdots \int_{V_t} \left(\frac{1}{B(a_{0t})} \prod_{i=1}^{N_t} \ddot{x}_{it}^{a_{it}-1} \right) d\ddot{x}_{1t} \ldots d\ddot{x}_{nt} \right) \tag{2.6}$$

Where:

$$B(a_{0t}) = \frac{\prod_{i=1}^{N_t} \Gamma(a_{it})}{\Gamma\left(\sum_{i=1}^{N} a_{it}\right)}, \qquad \ddot{x}_{N_t} = 1 - \sum_{i=1}^{n_t} \ddot{x}_{it},$$

$$V_t = \left\{ (\ddot{x}_{1t}, \ldots, \ddot{x}_{nt}) \in \mathbf{R}^{n_t} \ : \ x_{it} - \frac{0.5}{m_t} \leq \ddot{x}_{it} \leq x_{it} + \frac{0.5}{m_t}, \forall i \in [1, n_t] \right\}$$

The multiple integral is reduced to n dimensions (hence $\ddot{x}_{tNt} = 1 - \sum_{i=1}^{n_t} \ddot{x}_{it}$) as $\sum_{i=1}^{N_t} x_{it} = 1$. This ensures the above condition is met and computational demands are lowered. Due to the flexibility of the Dirichlet distribution, if $a_i < 1$, $\forall i$ the PDF is no longer unimodal, leading to singularity at the bounds. A penalty function is used when $a_i < 1$ to ensure parameter estimates exclude this possibility. For sample parameter estimates, the log-likelihood contributions for the decisions of every individual within that sample are summed.

4 Results

4.1 Proportional Payoffs

The Proportional Payoff to P1 (PP to P1) is the share of payoffs given to Player 1 in relation to the total payoffs ($\pi_1 x_1/(\sum_{i=1}^{N} \pi_i x_i)$). Figure 2.5 illustrates the distribution of observed PP to P1; split into per *Round* decisions ($n = 8,046$) and the individual treatment *Average*, the mean over the 27 rounds of each treatment ($n = 298$). The 'grand' mean is 0.5045; on average participants gave 50.45% of the payoffs to themselves. The percentage of the sample who distributed 0.33 to themselves (to the nearest $2d.p.$) was 13.57% and 5.03%, for the round and average respectively. While there were only 3.58% and 0.34%, respectively, who took all payoffs for themselves.

Figure 2.5: Proportional Payoff to Player 1: Per Round and Average

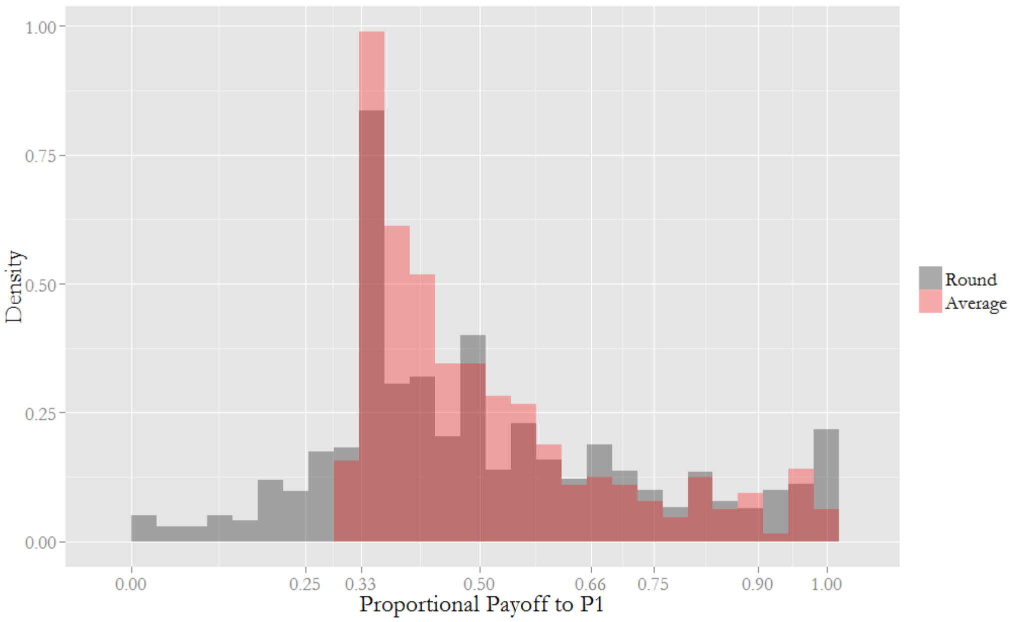

To explain the variation in the PP to P1 (per round), explanatory variables for experimental effects and demographic characteristics are used within a random effects model (with robust standard errors); the results of which are shown in Table 2.2. Observations in the oneness treatment, where others are *known*, are 2.9pp higher than in the budget treatment, where others are *anonymous*. Partaking in experiment in the PM session rather than the AM session has no significant effect, neither does the round number of the decision, nor the presence of a foreign enumerator in the tutorial session. Getting more questions correct in the tutorial has no significant effect and there are no significant differences between the general population and the student samples. The relative divider ($\pi_i/(\sum_{j=1}^{3} \pi_j)$) is included to account for changes in the relative price of giving; a more in depth analysis on the effect of the dividers on allocations is in Section 4.1.4.

Although extensive demographic data has been obtained the PP to P1 is not significantly effected by *any* of the included demographic characteristics, with the exception of the religious activity. Those who had participated in more 'religious activity' in the previous 30 days were more generous, 0.38*pp* less to themselves per activity. From the results, neither the gender, attainment of a higher education degree, the household size, belonging to the Bagisu tribe, being a Christian, the level of wealth nor extended MPI significantly affect the responses. Age also has no significant effect, but is dropped due to missing data. The explanatory variables in Table 2.2, bar the first, are included as controls in all following random effects models; as the effects of some controls become significant in particular models they are included to reduce potentially confounding effects. Further analysis on the relationship between demographic characteristics and preferences is found in Appendix B.8.

Table 2.2: Random Effects Model: Proportional Payoffs to P1

	(1) **PP to P1** Coef.	Std. err.
Experiment Effects		
- Treatment Dummy	0.0290***	(0.0094)
- Round Number	0.0004	(0.0003)
- PM Session Dummy	0.0054	(0.0300)
- Foreign Enumerator Dummy	-0.0573	(0.0378)
- Correct Questions	-0.0242	(0.0390)
- Student Dummy	-0.0221	(0.0367)
- Relative Divider	-0.4403***	(0.0341)
Demographics		
- Gender	0.0051	(0.0259)
- Higher Education Dummy	0.0353	(0.0335)
- Household Size	0.0087	(0.0075)
- Religion Dummy	0.0150	(0.0411)
- Religious Participation	-0.0038**	(0.0016)
- Tribe Dummy	-0.0333	(0.0287)
- Wealth Index	0.1036	(0.1292)
- EMPI Index	0.0658	(0.1872)
Constant	0.5615***	(0.1322)
N	149	
Observations	8046	
R-squared	0.1323	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.1.1 Budget Treatment

Within the budget treatment the effects of a change in the budget upon behaviour were tested. Three budgets - $shs30,000$, $shs60,000$ and $shs90,000$ - were varied, randomly, between rounds for each participant. These budgets were chosen as they represented substantial increases in the amount of wages; approximately two, four and six times the median daily wage of the participants, respectively. The mean PP to P1 was 0.4813, 0.4934 and 0.4953, for the three respective budgets. As the size of the budget is increased, participants are, on average, allocating relatively more to themselves and less to the others. These differences are, however, only significant ($p < 0.1$) between the budgets of $shs30,000$ and $shs90,000$, there is no significant difference between $shs60,000$ and the others.

4.1.2 Oneness Treatment

Within the oneness treatment each of the others were *known*, identified by their desk number. Before the session the oneness questionnaire was conducted to determine the oneness levels between the participant and the other five people in the room. Table 2.3 shows the descriptive statistics for the oneness levels, where 1 is the least connected, and 7 the most. Just under half, 49.3%, of responses show that participants consider the others to be the least connected, with the remaining responses showing some degree of connection. The mean oneness levels are, 2.15, 2.03 and 2.11 for the general, student and total samples, respectively. While these samples are not statistically different (10% level), there is a higher proportion of 'most connected' individuals in the general sample.

Table 2.3: Oneness Levels

	Sample					
	General		**Student**		**Total**	
	No.	%	No.	%	No.	%
Oneness Level						
1 - Least Connected	263	52.4%	102	42.9%	365	49.3%
2	109	21.7%	71	29.8%	180	24.3%
3	45	9.0%	32	13.4%	77	10.4%
4	25	5.0%	15	6.3%	40	5.4%
5	20	4.0%	11	4.6%	31	4.2%
6	13	2.6%	5	2.1%	18	2.4%
7 - Most Connected	27	5.4%	2	0.8%	29	3.9%
Total	606	100.0%	288	100.0%	894	100.0%

Table 2.4: Random Effects Model: Proportional Payoffs to P1, P2 and P3

	(1) **PP to P1**		(2) **PP to P2**		(3) **PP to P3**	
	Coef.	Std. err.	Coef.	Std. err.	Coef.	Std. err.
Oneness Level						
- Player 2, θ_2	-0.0123***	(0.0021)	0.0162***	(0.0021)	-0.0042***	(0.0013)
- Player 3, θ_3	-0.0128***	(0.0022)	0.0004	(0.0014)	0.0123***	(0.0023)
Constant	0.5746***	(0.1409)	0.3802***	(0.0703)	0.3872***	(0.0721)
N	149		149		149	
Observations	3969		3969		3969	
R-squared	0.1394		0.1347		0.1372	
Controls	YES		YES		YES	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Using oneness levels as explanatory variables, the effect of the closeness of connection on distributional decisions can be tested. Table 2.4 shows the effects of the oneness level on the distribution of payoffs amongst P1, P2 and P3. The results show that when participants are more connected to the other players they sacrifice their own payoffs in order to give more to that individual. Each increase of the oneness level decreases the PP to P1 by 1.23*pp* and 1.28*pp*, for P2 and P3 respectively. A change in the oneness level of Player 2, from the least to the most connected, would result in a 7.39*pp* decrease in the PP to P1, and a 7.67*pp* decrease, if it were Player 3. If the closeness to both P2 and P3 changed, as such, there would be a 15.06*pp* decrease.

Similar results (in the opposite direction) can be seen for PP to P2, an increase of 1.62*pp*, and for PP to P3, an increase of 1.23*pp*, for an increase in their respective oneness levels. The results are, however, asymmetric when considering the oneness level of the *alternative other* (i.e. the oneness level of P3 when considering the PP to P2 and vice versa). When distributing to P3, if the oneness level of P2 is higher, there is a significant negative effect. Yet, there is no significant effect on the PP to P2, when P3's oneness level is increased.

Further analysis in Appendix B.10 delves into the origins of oneness; incorporating reciprocated oneness, homophily (interpersonal similarity) and demographic characteristics as explanatory variables. Results show reciprocated oneness, being the same age and belonging to the same religion are positively correlated with oneness, while differences in wealth, participating in more religious activities and working longer hours are negatively correlated. Furthermore, the robustness, and importance of oneness in explaining behaviour is tested, in comparison to homophily. Oneness remains strongly correlated and significant.

4.1.3 Cluster Analysis

In order to identify *differential* effects amongst those in the sample, clusters of individuals can be classified into *types*. A finite mixture model has been run to determine the optimal clusters of responses. The PP to P1, averaged over all rounds and both treatments, has been used to denote decisions. The analysis reveals three *types*, shown in Figure 2.6, denoted as A, B and C. From the clusters we find: mixing proportions of 0.4096, 0.4885 and 0.1019; means of 0.3748, 0.5330 and 0.8893; and variances of 0.0012, 0.0098 and 0.0057 for A, B and C, respectively. In other words, a group of 41% of the sample who share payoffs roughly equally, 49% who weight themselves somewhat higher than others, and the remaining 10% who distribute the majority to themselves.

Table 2.5 models the Proportional Payoff to Player j (PP to Pj), separately for each *type*; where $j \neq 1$. Pooling the data allows for the oneness level for Player j (θ_j) and Player k (θ_k), to be used as explanatory variables, where $j \neq 1$ and $k \neq 1, j$. The results show *differential* effects between types. An increase in θ_j, increases the PP to Pj by (0.44*pp*), (1.81*pp*) and (2.19*pp*), for Type A, B and C, respectively. An increase in θ_k, however, only significantly decreases the PP to Pj for Type A (0.44*pp*) and Type B (0.25*pp*).

This implies that those who share most equally, Type A, alter distributional giving by substituting between others, but are not willing to sacrifice their own payoffs as their connection to others increases. Type C individuals, on the other hand, are not, on average, willing to trade-off between others, but sacrifice their own payoffs in order to increase the payoffs to those they are more closely connected to. Type B lies in between the two, willing to substitute payoffs between themselves and others and between others, but to a lesser magnitude for the latter. To reinforce these observations Table 2.6 shows the effects of increases to θ_2 and θ_3 on the PP to P1. Type A individuals do not, on average, significantly change the payoffs to

Figure 2.6: Clustered Proportional Payoff to Player 1

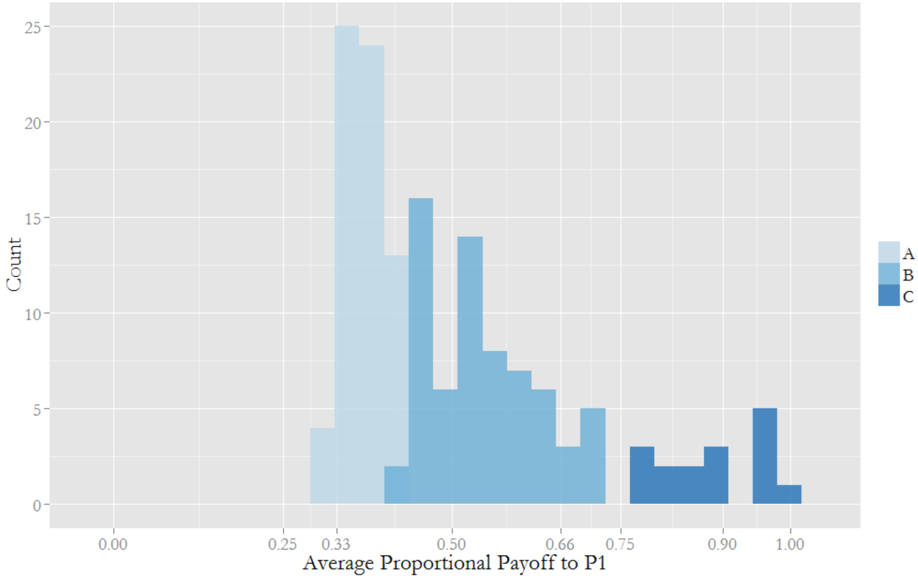

Table 2.5: Random Effects Model: Proportional Payoff to Others, Dependent Upon Cluster

	(1) **Type A** Coef.	Std. err.	(2) **Type B** Coef.	Std. err.	(3) **Type C** Coef.	Std. err.
Oneness Level						
- Player j, θ_j	0.0044***	(0.0013)	0.0181***	(0.0030)	0.0219***	(0.0059)
- Player k, θ_k	-0.0044**	(0.0017)	-0.0025*	(0.0014)	0.0024	(0.0018)
Player 3 Dummy	-0.0017	(0.0039)	-0.0065	(0.0046)	-0.0050	(0.0040)
Constant	0.4239***	(0.0300)	0.3182***	(0.0634)	0.0402**	(0.0181)
N	66		67		16	
Observations	3470		3604		864	
R-squared	0.1390		0.1766		0.3379	
Controls	YES		YES		YES	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

themselves, while Type B do significantly reduce their own proportional payoffs, and Type C does so to an even greater magnitude.

Table 2.6: Random Effects Model: Proportional Payoffs to P1, Dependent Upon Cluster

	(1) **Type A** Coef.	Std. err.	(2) **Type B** Coef.	Std. err.	(3) **Type C** Coef.	Std. err.
Oneness Level						
- Player 2, θ_2	-0.0002	(0.0025)	-0.0162***	(0.0029)	-0.0241***	(0.0045)
- Player 3, θ_3	-0.0027	(0.0022)	-0.0150***	(0.0033)	-0.0285***	(0.0069)
Constant	0.4597***	(0.0544)	0.9464***	(0.1436)	1.0939***	(0.0525)
N	62		58		16	
Observations	1640		1559		432	
R-squared	0.1724		0.2266		0.4425	
Controls	YES		YES		YES	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.1.4 Proportional Allocations

To identify the effects of the dividers on decision making, the proportional allocations ($x_i / \sum_{j=1}^{N} x_j$) are used, rather than the proportional payoffs, as the dividers constitute part of the latter's calculation. Table 2.7 shows how the proportional allocations, change in relation to an increase in the dividers; the inverse of which can be interpreted as an increase in the price of giving. An increase of each respective divider leads to an increase in the proportional allocation of that respective player; 4.59pp, 4.71pp and 4.59pp. The opposite is true for the opposing dividers, for instance an increase of the dividers for P2 and P3 lead to a decrease in the proportional allocations to P1 by 3.31pp and 2.53pp, respectively. These results hold, and are highly significant ($p < 0.01$), for each divider in relation to each proportional allocation. The results show that participants are, on average, more concerned with distributing more equally, than maximising the total surplus. When it is relatively more efficient to give to a particular player, less is given, so those efficiency gains can be redistributed.

Table 2.7: Proportional Allocation to P1, P2 and P3; Random Effects

	(1) **PA to P1**		(2) **PA to P2**		(3) **PA to P3**	
	Coef.	Std. err.	Coef.	Std. err.	Coef.	Std. err.
- Divider for P1	0.0459***	(0.0050)	-0.0205***	(0.0032)	-0.0254***	(0.0022)
- Divider for P2	-0.0331***	(0.0025)	0.0471***	(0.0041)	-0.0140***	(0.0029)
- Divider for P3	-0.0253***	(0.0029)	-0.0206***	(0.0017)	0.0459***	(0.0039)
Constant	0.4318***	(0.1353)	0.2843***	(0.0666)	0.2839***	(0.0699)
N	149		149		149	
Observations	8046		8046		8046	
R-squared	0.1389		0.1735		0.1698	
Controls	YES		YES		YES	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.2 Aggregate Preference Parameters

Preference parameters can be estimated on a sample level; to identify an aggregate preference of the society or group. These parameter values are the 'most likely' given the allocations made by all individuals within the sample. The sample estimates show $r = 0.0191$ and $\alpha = 0.5094$, reflecting a *weakly* 'Weighted Prioritarian', allocating just under half to themselves, and distributing the remainder between the other two. The oneness elasticities, $\phi = 0.1552$ and $\psi = 0.0001$, reflect a greater willingness to trade-off between the self and others, as oneness increases; but less willingness to substitute between others, as the relative oneness between them changes. A low precision parameter, $s = 5.1152$, reflects the 'noise' within decisions, which reflects heterogeneity in decisions.

4.3 Individual Preference Parameters

While the above analysis delves into aggregate trends, it reveals little of the decisions made by individuals. To address this, individual level preferences are estimated; as parameters within the proposed utility function. The results shown below are estimated from the full 54 rounds.

4.3.1 Inequality Aversion and Self-Interest Parameters

Figure 2.7 shows the distribution of each Self-Interest, α, Inequality Aversion, r and Equally Distributed Equivalent (EDE). The top-right and bottom-left histograms show the distribution of α and r, respectively, and map onto the scatter plot in the top-left. The EDE, which conveniently combines both preferences to form an index from 1 to 10, is shown in the bottom-right.

The distribution of r reveals how individuals trade-off between equality and efficiency. We observe 15.44% of the sample are 'Efficiency Prioritarians' ($r < -0.005$), 2.01% are 'Cobb-Douglas' ($-0.005 \leq r < 0.005$), 74.50% are 'Weighted Prioritarians' ($0.005 \leq r < 15$), and 8.05% are 'Maximin' ($r \geq 15$). The median respondent has an r value of 0.69, a 'Weighted Prioritarian'.

Figure 2.7: Distribution of Individual Level Parameter Values: α, r and EDE

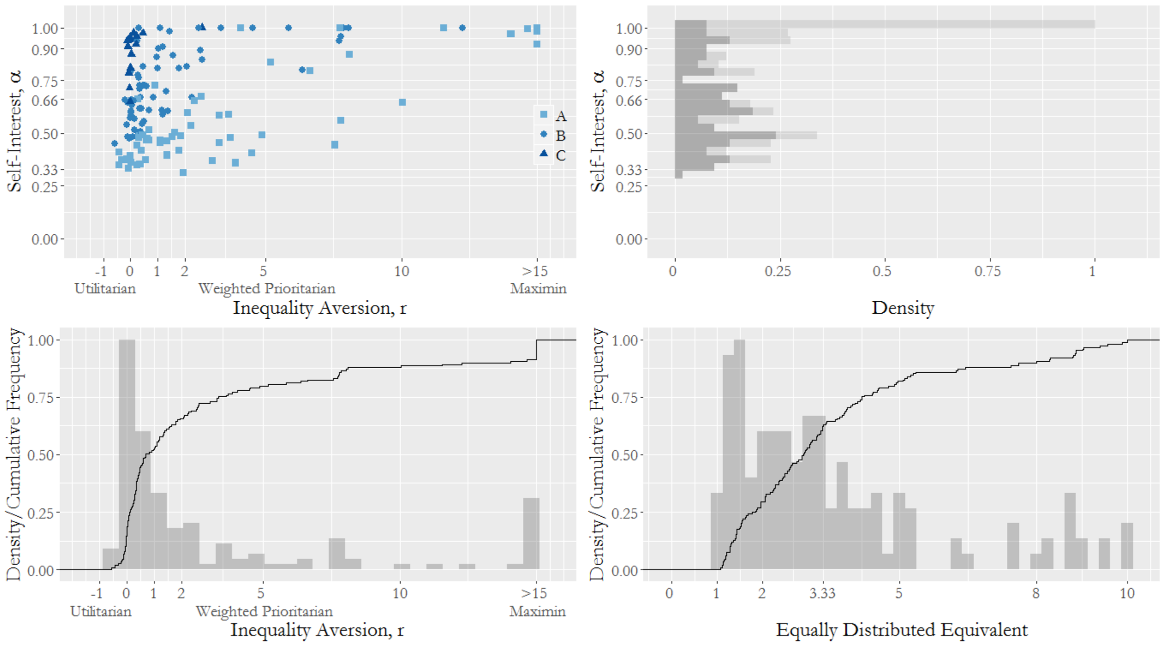

The self-interest parameter is somewhat more difficult to interpret. Within the estimation procedure, as higher values of r are estimated the interpretation of α becomes increasingly irrelevant. This is problematic when interpreting the distribution of α. As a result the histogram for α is shown as a stacked distribution, estimated values of α when $r \geq 3$ are shown in light grey. The dark grey distribution, then more closely resembles what is expected from the average PP to P1 results.

The scatter plot highlights the bivariate distribution of r and α; splitting the observations into the three clusters: A, B and C. The heterogeneity of preferences is further highlighted here, although patterns do emerge from the clustered results. Those who are the least self-interested (Type A), tend to be the most averse to inequality, while the most self-interested (Type C) tend to be the least. The median r, for Type A, B and C respectively, are 2.029 0.496 and -0.024.

The EDE values are an intuitive measure, which combines the two preferences. If each individual were to consider a society where the payoffs to P1, P2 and P3 were: 10, 1 and 1, and they had no connection to either P2 or P3, then according to their preferences they would be indifferent between that distribution and an equal distribution of EDE's. Those whose EDE is 10, reflects pure self-interest, an individual's only concern is their own payoff. If EDE $= 1$ then participants are 'Maximin', caring only for the payoff to the worst-off. Those values in-between reflect 'Weighted Prioritarianism' and/or a regard for others. EDE's which are less than 1.2 make up 6.04% of the sample, 29.53% are less than 2 and 63.09% less than 3.33. Those greater than 9.8 make up 2.01%, while 10.07% are greater than 8. The median EDE is 2.90.

4.3.2 Oneness Elasticities

To illustrate the estimated oneness elasticity parameters Figure 2.8 shows the distribution of ϕ and ψ, in the left panel, and the difference between the two, as the proportion $\phi/(\phi+\psi)$, in the right. Due to lack of variance in elicited oneness levels of a particular individual, 15 individuals are excluded from analysis on oneness elasticity parameters; as no such parameters could be estimated. From the left panel; both ϕ and ψ have similarly right skewed distributions. For ϕ, 33.58% of parameters are less than 0.01, 55.97% are less than 0.5 and 73.88% less than 1. While for ψ, 28.36% are less than 0.01, 64.18% are less than 0.5 and 80.60% less than 1. For the medians: $\phi = 0.36$ and $\psi = 0.23$. As shown by the CDF plot, ψ tends to be lower than ϕ, at the sample level. This observation is confirmed by the right panel, the proportion $\phi/(\phi+\psi)$ shows the difference between the elasticities for each individual. There are 41.04% for whom $\psi > \phi$, with $\psi < \phi$ for 58.96%. Spikes at the extremes imply one elasticity dominates the other for a significant proportion of the population; 22.39% < 0.01 and 17.16% > 0.99.

These results reflect, and explain, those results shown in Table 2.4. On average, the sample is willing to trade-off between the self and others, to a greater extent than between others. This observed trend in the data is because, for the majority of the sample, $\phi > \psi$. Delving further, by identifying the preferences of those individuals belonging to each cluster, we observe the mean $\phi/(\phi + \psi)$ is equal to 0.498, 0.533 and 0.687, for Type A, B and C respectively. This supports the results of Table 2.5 and 2.6, as Type A substitute between-others, but not between themselves and others, while Type C predominantly trades-off between themselves and others, but not between-others, with Type B in-between the two.

Figure 2.8: Distribution of Oneness Elasticities, ϕ and ψ

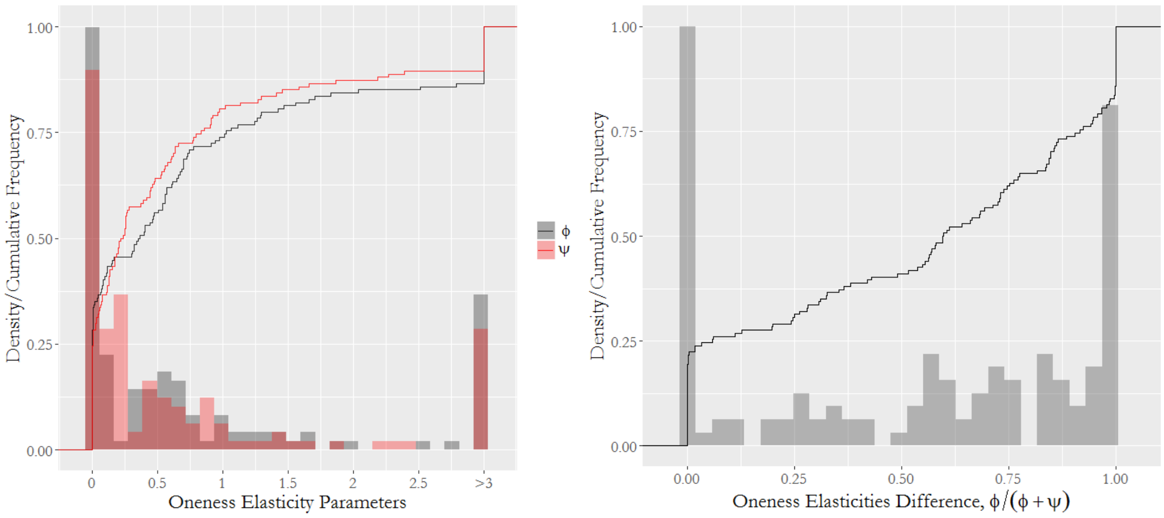

4.3.3 Precision and Fit

In addition to the estimation of preferences parameters, the ability of the utility function to fit individual behaviour also necessitates consideration. Figure 2.9 shows the distribution of s in the left panel; s is the precision, which characterises individual level 'noise'. The higher s the more 'precise' decisions are, as illustrated in Figure 2.4. Those with $s < 7$ make up only 10.07% of the sample, for 40.27% $s < 15$ and for 61.07% $s < 25$. To illustrate how well the model fits the data, the middle and right panel of Figure 2.9, show the distribution of goodness-of-fit measures for each individual. The measures establish the difference between the *observed* allocations x_1 and *optimal* allocations x_1^*, to Player 1.[5] The middle panel provides an intuitive measure for bias; the closer to zero the lower the bias in the optimal allocations. The majority of the sample are clustered around zero, 93.29% are within 0.05, 68.46% are within 0.01 and 46.98% are within 0.005. The right panel highlights efficiency, how close the optimal allocations are to the observed in absolute terms; 89.93% have a value less than 0.2, 64.43% are less than 0.1 and 14.09% are less than 0.05.

4.4 Model Predictions

Once estimated, preference parameters allow for the prediction of behaviour, according to particular environmental factors (i.e. a budget, set of dividers and oneness levels). Rather than relying on the *observed* behaviour, which may be contingent on a particular set of environmental factors, a 'clean' ceteris paribus analysis can be undertaken to identify changes in *predicted* behaviour. It provides a meaningful analysis, based on individual level preferences, to reveal how the distribution responses would change in response to a shift in environmental factors.

[5]The *mean GOF for P1* $= \sum_{t=1}^{54} \frac{(x_{1t} - x_{1t}^*)}{54}$, while the *mean absolute GOF for P1* $= \sum_{t=1}^{54} \frac{|x_{1t} - x_{1t}^*|}{54}$.

Figure 2.9: Distribution of Precision, s, and Goodness of Fit

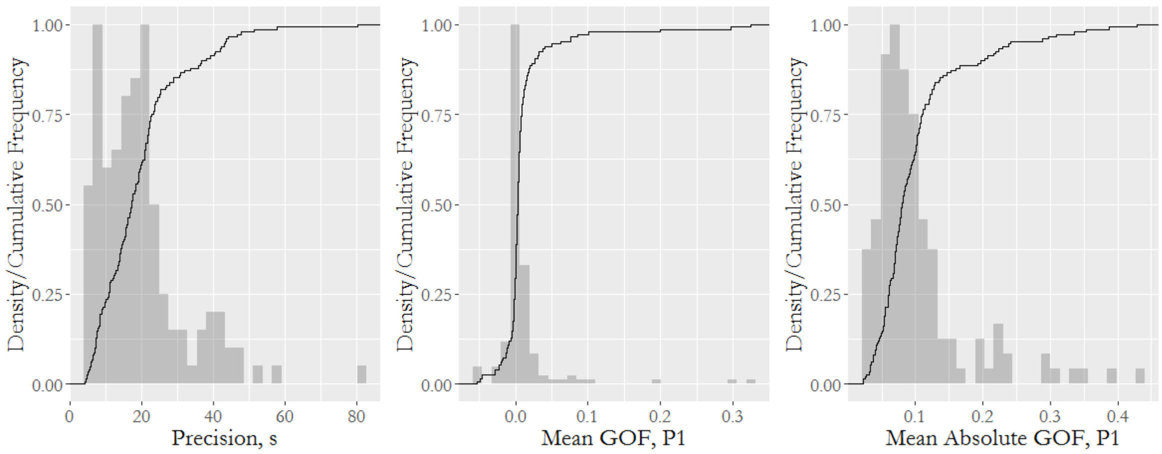

The below analysis will identify how *predicted* behaviour changes, in response to differing oneness levels. Three different sets of oneness levels have been proposed. Where $\theta = [\theta_2, \theta_3]$, in the first set, A, $\theta = [1, 1]$, in the second, B, $\theta = [7, 1]$ and in the third, C, $\theta = [7, 7]$. Given the estimated individual-level preferences, assuming $\pi = [1, 1, 1]$ and $m = 1$, the optimal payoffs can be established for each individual, in each A, B and C. Both *deterministic*, assuming decisions are optimal, and *stochastic*, accounting for the errors each individual would make, forms of analysis will be conducted; the former showing changes in optimal payoffs, the latter illustrating the sample distribution of payoffs.

4.4.1 Deterministic

To identify changes in behaviour, due solely to a change in oneness, the differences between the proportional payoffs, to a particular player, in A and B ($B - A$), and in A and C ($C - A$) can be compared. The former highlights a behavioural change due to an increased connection to one other, while the latter shows the change due to an increased connection to both others.

Figure 2.10: Distribution of Changes in Predicted Optimal Proportional Payoffs; Differing Oneness

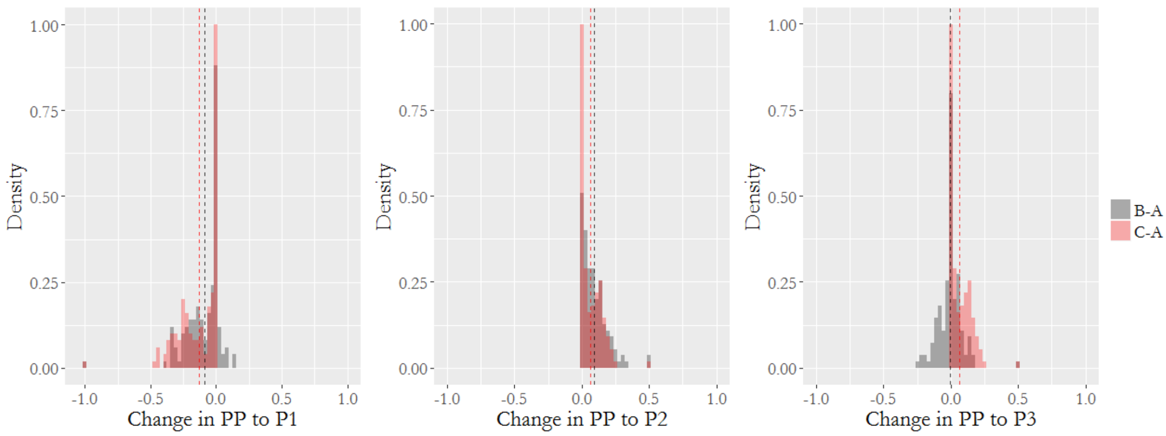

Figure 2.10 shows the distribution of these changes, where the magnitude and sign of the individual change depends on the preferences of that particular individual. A negative change

shows a decrease in the proportional payoff, to that respective player, when moving from A to B (or C), while a positive change shows the opposite. The closer to zero, the less that change has been. Looking at $C - A$ we observe that as the closeness to both players increases, less is given to P1 while more is given to P2 and P3. On average, the PP to P1 decreases by $-13.02pp$, while the PP to P2 and P3 both increase by $6.51pp$. Considering $B - A$, the behavioural changes due to an increase in the oneness to P2, only, can be observed. The change in PP to P2 is as expected, each individual increases the PP to P2 and the magnitude is greater than in $C - A$; $9.40pp$ on average. For majority of the sample, 76.87%, there is a decrease in the PP to P1, however, the remaining minority do actually increase the PP to P1. When considering the PP to P3 a split occurs, 52.24% decrease that to P3, while the remainder, perhaps counter-intuitively, increase the PP to P3; the mean is only slightly negative, $-0.51pp$. This split will be more thoroughly discussed in Section 5.1. These results, similar on average in magnitude and sign to those in Table 2.4, reveal how individual's *optimal* behaviour is shaped by a change in oneness to those to whom they are giving.

4.4.2 Stochastic

While the above shows how the sample *optimally* responds to changes oneness levels; it neglects the *noise* in individual decision making. By combining both utility function and error model a more representative distribution of decision making can be made. Figure 2.11 utilises individual level preferences and precision parameters to predict the distribution of proportional payoffs, according to differing oneness levels. The three sets of oneness levels (A, B and C) are used as above, but the distribution of proportional payoffs, rather than the changes, are plotted. By using a monte-carlo simulation, taking 10,000 random draws each individual level probability density function, the plots reflect an 'asymptotic' distribution of *predicted* decisions.

Figure 2.11: Distribution of Predicted Proportional Payoffs; Differing Oneness

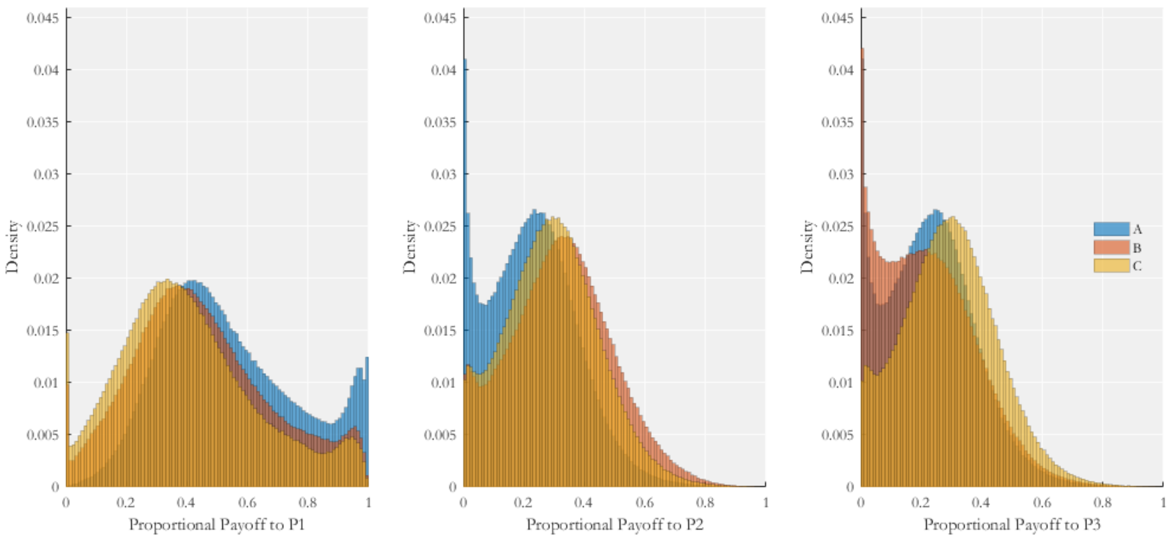

The first observation is that A reflects an 'objective' distribution, where participants have no connection to others. The mean proportional payoffs are [0.539, 0.230, 0.230], for P1, P2 and P3, respectively. As oneness increases the distribution of PP to P1 shifts left, reflected by the decrease in the average to 0.450 and 0.409 for B and C, respectively. For the PP to P2, an increase is observed in both cases, as the plots shift right, however, the shift to B is greater than that of C; with means of 0.324 and 0.295, respectively. For the PP to P3, a decrease is shown, when moving to B, while an increase occurs to C; with means of 0.225 and 0.296, receptively. Of most note are the changes at the extremes. Those observations for which the PP to P1 is greater than 0.98, is 2.28%, 0.44% and 0.33%, for A, B and C respectively. While those observations, for which the PP to P3 is less than or equal to 0.01 is 4.11%, 4.21% and 1.02%, for A, B and C respectively.

5 Discussion

5.1 Oneness Peculiarities

Predictions from the utility function, in Section 4.4.1, have two apparently peculiar tendencies. Both occur when identifying the change in optimal allocations from a situation where $\theta_2 = \theta_3 = 1$ to that where $\theta_2 = 7, \theta_3 = 1$; when, ceteris paribus, the oneness to P2 increases. As expected, x_2^* always increases, and for the majority of individuals, x_1^* and x_3^* decrease. There are, however, two peculiarities, first, with regards to an increase in x_1^* and, second, with an increase to x_3^*.

The former occurs only when $r > 0$ and ϕ/ψ is *sufficiently low.* If $\psi > \phi$, individuals place a higher weight on substituting *between-others*, than on *self-other* trade-offs. If ψ is sufficiently high, then as θ_2 increases relative to θ_3, $\omega_3 \to 0$. When ϕ is low, ω_1 only decreases relatively little. As a result, as $\omega_1, \omega_2 \gg 0$ and $\omega_3 \to 0$; individuals consider the problem as one of distributing between P1 and P2 only. As individuals are inequality averse, $r > 0$, they share between P1 and P2; the increase to P2 is greater, but x_1^* does actually increase. Intuitively, this peculiarity makes sense, when faced with a situation where the individual is disconnected to both P2 and P3, they decide to share between all three (especially as $r > 0$). A change in absolute terms of oneness, as described above, leads to a relative reduction in the connection to P3, as the individual is willing to substitute between others, $\omega_3 \to 0$, but as they are not very willing to trade-off between themselves and others, ω_1 does not. Then, they consider the problem of distributing between themselves and one other, they then give more to P2 but also more to themselves.

The latter occurs when ϕ/ψ is *sufficiently high.* When $\phi > \psi$, individuals place a higher weight on *self-other* trade-offs, than substituting *between-others.* The reason for this is that, in the model, as θ_i increases, $\forall i > 1$, ω_1 decreases; as a result $1 - \omega_1$ increases. If ψ is sufficiently small, little substitution occurs between ω_2 and ω_3, and therefore both ω_2 and ω_3 increase. Intuitively, this could be due to the relative amounts given to P2 and P3. If, due to an increase in θ_2, x_2^* increases, it could lead to an increased generosity to P3. Being close to one other makes me consider that I would want to share with others, rather than giving all to myself; as I now consider sharing more, I will actually give more to P3, even though my closeness to them has not increased.

5.2 Comparability

5.2.1 Giving and Social Distance

Engel (2011) conducts a meta-analysis on dictator games from 129 papers, calculating a 'grand mean' of 0.717, as the proportion kept by the dictator. Our results show a 'grand mean' of 0.505, significantly lower than that calculated by Engel. Two factors may explain this heightened generosity; both of which are stated in Engel's conclusion: that dictators give more 'when they come from a developing country' and 'if there are multiple recipients'.

In support of the latter, Andreoni (2007) finds that as the number of recipients increases dictators give themselves less, however, the rate at which they give is *congested*. For the representative agent a 'gift that results in one person receiving x is equivalent to one in which n people receive $x/n^{0.68}$'. A back-of-the-envelope calculation shows that (all else being equal) the equivalent mean PP to P1, from our experiment, would be 0.603, in a 2-player dictator game. Fisman, Kariv, and Markovits (2007) observe a similar trend, as their mean PP to P1 is 0.75, in the three-player variant, and 0.79 in the two.

In the *social distance* literature dictator games usually have two-players. Considering the most similar treatment in Leider et al. (2009), when the exchange rate is 1:1, the PP to P1 0.761 and 0.875 for the most and least connected respectively. From lab-in-the-field experiments in 'developing economics', when matched with a stranger and a friend, respectively, dictators took 0.636 and 0.5573 on average in Binzel and D. Fehr (2013). In Ligon and Schechter (2012), when the recipient was chosen the average to the self was 0.637, while when they were random it was 0.615. The level of giving is somewhat lower in the later two, giving support to the findings by Engel (2011). While not directly comparable to other papers, combining the unilateral effects above provides an explanation for the difference in the levels of giving.

Aside from the difference in levels, the results appear consistent with those in the *social distance* literature. Giving is shown to increase with *oneness*; the predicted mean PP to P1 is 0.539, when there is the least connection to both others, and 0.408 with the most. Due to the three-player design, this work is able to tease out *between-other*, alongside *self-other*, trade-offs when closeness to others differs. Combining this with the the proposed utility function and error model, we observe a great deal of heterogeneity in preferences, as found in Fisman, Kariv, and Markovits (2007). By accounting for this heterogeneity, and incorporating *oneness* considerations, the average effects of changes in *oneness* can be explained by the distribution of individual behaviour.

5.2.2 Error Models

The random behavioural model proposed here diverges from that which is most common. Usually random behavioural models assume that an individual has some optimal payoff, according to their preferences and the decision problem, say y_i^*. The individual then makes some error calculating, or choosing, y_i^*, meaning the *observed* decisions are $y_i^* + \epsilon_i$, where ϵ_i is normally distributed with mean zero. The issue is that with some probability, $y_i^* + \epsilon_i$ will exceed the bounds of the problem (i.e. 0 and 1, in the case of the share to the self). As a result, censoring must ensue, usually in the form of a tobit model; as in Andreoni and Miller (2002) and Fisman, Kariv, and Markovits (2007). With greater dimensions, a further issue may arise due to the i.i.d nature of the errors; in the three player case (assume $y_3 = (m - y_1^* + \epsilon_1 + y_2^* + \epsilon_2)$), if ϵ_1 and ϵ_2 are sufficiently large, y_3 could be negative.

An alternative is that suggested here, where *observed* decisions are said to be drawn from the Dirichlet distribution, which is bounded between 0 and 1, where the expectation of the distribution is y^* and the sum of *observed* decisions will always equal 1 (or m if scaled). The flexibility of the distribution, intuition of its formulation and generalisability to n dimensions are useful in these bounded decision problems. Further comparison between the Dirichlet distribution and normal error model could identify which of the two best fits individual or aggregate level behaviour. While the Dirichlet distribution could potentially provide a more intuitive fit for individuals, at the aggregate level the normal error model could provide a better fit, the distribution of round decisions in Figure 2.5, for instance, has responses clustered at the upper bound.

5.2.3 Incentives

One point of divergence in this paper, from others in the literature, is the incentive structure. Within both Andreoni and Miller (2002) and Fisman, Kariv, and Markovits (2007), participants receive the payoffs they gave to themselves plus the payoffs other participants gave to them; in effect getting two (or three) sets of payoffs. In this paper, the decision made by only one dictator from each 'group' was carried out, meaning individuals would only get one set of payoffs. In the former incentive mechanism (henceforth (1)), the total experimental payoffs received included their own payoff to themselves plus what the others gave them; and they knew that others had the same payoff structure. Whereas, in the incentive mechanism used here (henceforth (2)) participants knew that, if they were chosen as dictator, then the payoffs they distributed would be the only experimental payoffs that each member of the group would receive.

This subtle difference could lead to a substantial divergence in decision making (if a consequentialist approach is taken). Imagine an individual who is an *egalitarian*; they wish everyone to leave the room with an equal payoff. In (2), individuals would allocate equally, while in (1) more strategic considerations are necessary. If they believe that everyone in the room is an *egalitarian*, then they would share equally, but if they believe everyone else is an *egoist* (who keeps all for themselves) then they would take everything for themselves. Interpreting the behaviour in the latter case would lead to the conclusion that the *egalitarian* was an *egoist*, due to a difference in the beliefs of the behaviour of the others in the room.

A further consideration, relates to the construct of the *impartial spectator* (Smith, 1790). The perspective of the spectator can be considered by putting oneself in the situation of the other; imagining how one's behaviour would be viewed from the standpoint of all members of society. When considering this perspective, behaviour is likely to reflect a greater degree of prosociality.[6] In (1) it is arguably less likely that this perspective would be considered. Participants know that the payoffs they distribute are certain, but that an additional uncertain amount will be added to those payoffs, by others. Participants would consult their beliefs about what others would give when making their decision, but still consider the problem from their individual perspective. In (2) the participant has the same chance of being the recipient, as the others in the room. Moreover, if they are chosen to be the dictator, the payoffs they distribute will be the only experimental payoffs to each player. As participants know this it is easier to imagine the situation of the other (when that other is a recipient of their giving), knowing that the other would receive the exact experimental payoff chosen by the participant. Imagining receiving *only* the payoff given to the other could, then, led to an increase in generosity.[7]

Due to these two observations, a divergence in observed behaviour could occur. For the former, results from (1) could be biased towards more 'apparently' self-interested decisions, while for the latter, results from (2) could be biased toward more equal distributions, than individuals would otherwise have done. This difference could perhaps further explain the differences in average giving (see Section 5.2.1) and is, perhaps, an interesting area for further inquiry.

5.3 Applicability

Two potential applications of this work are first, in the charity sector, and two for government distributional policies. For the former, by estimating oneness elasticities the extent to which charitable donations would fluctuate, due to an increase in oneness could be established. Then, if the cost of increasing an individual's oneness to the charity could be established, a cost-effectiveness analysis could be conducted to estimate the optimal level of investment to increase that individuals connection to the charity, in order to increase donations. For the latter, if individual level preferences were estimated, then a set of alternative government redistributive policies could be ranked. The one which was determined to be optimal could be implemented; rather than striving for one goal of efficiency (such as raising GDP per capita) or equality (increasing the welfare of the worst-off); a policy could be implemented which reflected individual's distributive preferences.

6 Conclusion

To conclude, distributional preferences are an integral component in explaining heterogeneity in prosocial behaviour. Through intertwining distributional preferences with considerations of *oneness*, the impact that *social distance* has on distributional decision making can be

[6]Smith writes of the *social passions*: 'all the social and benevolent affections'. Arguing that to the spectator, the social passions are the most agreeable, as 'we enter into the satisfaction both of the person who feels them, and of the person who is the object of them' (Smith, 1790). Given this, by taking the perspective of the spectator our behaviour should strive to satisfy these passions.

[7]Indeed, it may be easier to imagine the situation of an other to whom you are more closely connected to. If being better able to consider their situation leads to increased giving, this could partially explain why we observe an increase in giving, when oneness increases.

further explored and understood. This paper has formulated a model which incorporates both elements; estimating individual-level preferences to explain behaviour from a lab-in-the-field experiment. Results show that participants tend to favour reducing inequality, as opposed to reducing inefficiencies, and exhibit a substantial regard for others. On average, there is a greater willingness to trade-off between the self and others, when the closeness to others increases, than there is to substitute between-others, when relative oneness levels change.

Heterogeneity in behaviour, is however, the most notable aspect of these findings. There is not a consensus on how to share, nor how to trade-off between equality and efficiency, nor how to redistribute to those who are close. Differences in prosocial behaviour have been shown to depend both upon distributional preferences of individuals and the closeness of connection to others. Accounting for both would, therefore, seem necessary in order to better understand prosocial behaviour.

B Appendices

B.1 Experimental Design

An incentivised lab-in-the-field experiment, with an accompanying survey, was run in Mbale, Uganda. On each experimental day, two parallel sessions were ran with groups of six participants. The first group partook in the *experimental session*, and then the *survey*, while the second group did the survey followed by the experimental session. The experimental session began with the *script*, moved onto the *interactive tutorial* (one-on-one with enumerators with questions to check understanding), then the *budget treatment*. This was followed by the *oneness script*, the *oneness questionnaire*, finishing with the *oneness treatment*. The survey began with a *group discussion*, then the *survey*, one-on-one with an enumerator. The whole experiment lasted between 3 and 5 hours, depending on the speed of the participants. At least six enumerators were present throughout the experiments. Figure 2.12 shows how the experiment was set-up.

Figure 2.12: Photographs of Experiment Set-Up [8]

B.1.1 Scripts

Participants were firstly given instructions in the form of a verbal script, a transcript of which can be found in Appendix B.5. The script was split into two parts, the *general* script and

[8]Photos clockwise from top-left: (1) Experiment set-up, each participant had a Linx Tablet, with Z-Tree software, to make their decisions, tables were numbered from 1 to 6. (2) Chairs set up for General, Part 1 and Part 2 of the script. Semi-circle to ensure equal distance, with A2 posters showing screenshots of the experiment. (3) Survey set-up, one enumerator per participant. Using Asus tablets with Survey Solutions software to conduct the survey. (4) Chairs set up to encourage group discussion for Group 2, before the survey, and a waiting room for those not doing the survey, throughout.

oneness script; the former of which was given when participants entered the experimental session, the later between Part 1 and Part 2. The general script explained all the necessary precursors to the experiment, explained the payoffs, the decisions they would be making and what Part 1 entailed. Throughout the script visual aids were used to enhance the understanding of the participants; a Linx tablet, the round selection bag and a poster. The poster showed a screenshot of the Z-Tree interface, which was used to show the layout of the screen and exactly which decisions they would be required to make. There were no distribution of payoffs, or allocations, shown to avoid any anchoring effects. The oneness script was delivered after the break, before the oneness treatment, where both the oneness decisions and Part 2 were explained. Another poster was used to explain the oneness decisions to be made, and references were made to the previous poster for Part 2. The scripts were delivered by two of the enumerators; Zam the General Script and Isaac the Oneness Script.

B.1.2 Tutorial

To further ensure that participants understood the decisions they were making after the general script they went through an extensive on-screen tutorial. Each participant went through the tutorial one-on-one with an enumerator.[9] The tutorial was split into three parts. The first was a simplified version, with one player (rather than three) which built up each element step-by-step and was led by the enumerator. Starting with only the allocations, then adding the payoffs, alongside differing dividers. The three methods were taught in turn, *slider*, *arrow keys*, then *written input*, before all three could be used. After the first part, the participants had to answer six questions to check their understanding. Enumerators recorded if they had answered correctly, or not, from the first time the participant understood the question. If they got the answer wrong the enumerators explained the correct answer. Once the questions were complete they moved onto a tutorial with the entire interface. Participants then made the decisions, with no input from the enumerators, who were still there to check they understood the decisions they had made. Two practice rounds were given, one where the dividers were $[1, 1, 1]$, the second $[1, 2, 3]$. Once participants were confident with the interface, then the enumerators left the participants to make their decisions for real, in private. A script of the wording of the on-screen tutorial, questions, and instructions to enumerators (in [ ]) can be found in Appendix B.6.

B.1.3 Survey

Alongside the experiment an in-depth survey was ran. The surveys were implemented with Asus touchscreen tablet and ran one-on-one; where one enumerator went through the survey with one participant. The survey was split into four sections: *individual characteristics*; *household characteristics*; *assets, wealth and MPI*; and *preferences*. The *individual characteristics* section comprised predominately of demographic questions, about the participant's: gender, date of birth, religion, tribe, occupation, level of education, hours of work/study,

[9] The exact enumerator was recorded to test for enumerator effects.

income, height and weight. *Household characteristics* asked questions regarding each member of their household: gender, age, relation to the participant, level of education, school attendance and literacy. Alongside this more general household level questions were asked: child mortality, food shortages, bank accounts, health insurance, social insurance, healthcare and unemployment. *Assets, wealth and MPI*, asked questions about their household: household type; housing arrangement; bedrooms; floor, roof and external wall materials; water source and distance; toilet; lighting; cooking fuel and rubbish collection. Further questions on ownership of livestock, vehicles and household assets were made. The final section, *preferences*, asked participants to allocate hypothetical tokens between three pots which denoted the health, education and household assets, of the household. They were asked to allocate the tokens to denote how much importance these three elements had to them.

B.1.4 Group Discussion

In order to test whether having group discussion had an effect on the oneness scores, or the behaviour of participants in the experiment, a group discussion was part of the design. The second group, who did the survey first and experiment second, were firstly sat down to partake in a group discussion. This was informal, and led by the enumerators, but allowed the chance to talk amongst themselves. Further to this, before the experiment, they had lunch together to further any group dynamics, or individual relationships, which could have emerged throughout that time. In contrast, the first group did the experiment as soon as all members of their group arrived.

B.1.5 Pilot Experiments

Prior to the main experiment two pilot sessions were ran. The first, a preliminary paper version was ran with 23 participants in September 2015. The second, was a first run through of the final experiment, ran with 12 participants at UCU, Mbale, in June 2016. The major changes relate to timing, firstly that paper based methods were too time consuming, hence the move to touchscreen tablets, and secondly that only one session could be ran each day, rather than two. Other changes included simplification, ensuring dividers had were equal for each Player, ensuring privacy of decisions was paramount and minor tweaks to the script to enhance understanding.

B.2 Sampling

Two sampling frames were created in order to conduct the experiments. The first, from a general population in the Mbale District, the second, from student records from the Ugandan Christian University (UCU), Mbale. The majority of the sample, 108, were from the general population, the remainder, 48, from the student population.

B.2.1 General Population

In order to create the sampling frame for the general population, mobilizers were employed to gather information for each individual, within each household, within their Cell.[10] There were 9 Cells selected, within two Divisions. From the Industrial Division the cells selected were: Bumasfa (n=169), Pallisa (n=255), Malawa (n=241), Butaleja (n=127), Bugwere (n=248) and Wanyera (n=185). Within the Northern Division the cells were: Mugisu (n=289), Gudoi (n=305) and Nkokonjeru (n=282).[11] A total sampling frame size of n = 2101, where there were 51.5% women and a mean overall age of 26.15.

The characteristics recorded were household ID, gender, age and the relationship to the household head for every member of the cell. For those individuals who were between the age of 18 and 45 further information was taken: if they had a good level of English; if they were in secondary school; if they would be willing to participate in the experiment; their phone number; and the answers to three selected questions. The three questions were: Q1: Do you own a touch-screen phone? Q2: Would you be comfortable using a touch-screen gadget in a study? Q3: Do you regularly use a personal computer?

In order to select a sample from the sampling frame individuals had to meet certain characteristics. They had to: be between 18 and 45; have a good level spoken and written English, not be in secondary school, be willing to participate; and have answered yes to two or more of the three questions. Those who met such requirements numbered 541. Gender characteristic remained similar with 50.01% being female; while the age was censored there was a mean age of 27.96. This criteria was necessary to ensure that participants would be able to participate within the study. With such criteria a bias is indeed created within the sample, while as many steps were made to ensure the sample was random and as unbiased as possible, the claim being made here is not that the sample is representative. A trade-off did have to be made, however, between representativeness and ability to be able to make informed decisions with the tablets provided. Figure 2.13 shows the composition of age and sex for both the "total" sampling frame and the "select" sampling frame. The histogram highlights the age cut-off points and the reduced sample frame size, alongside the population's general age and gender characteristics.

[10] A cell is a geographical area which is smaller than a Division, which is smaller than a District. They are the urban equivalent of a rural village.

[11] There were issues with the sampling frame. The general consensus from the Mobilisers was that they had recorded 70-80% of participants within the cell. Reasons given for the lack of coverage included: Suspicion that the information was being collected for taxation purposes, as the study was going to be in UCU individuals were worried that the aim was to convert them to being born again Christians, some people identify the

Figure 2.13: Composition of Age and Sex: Total and Selected Cells

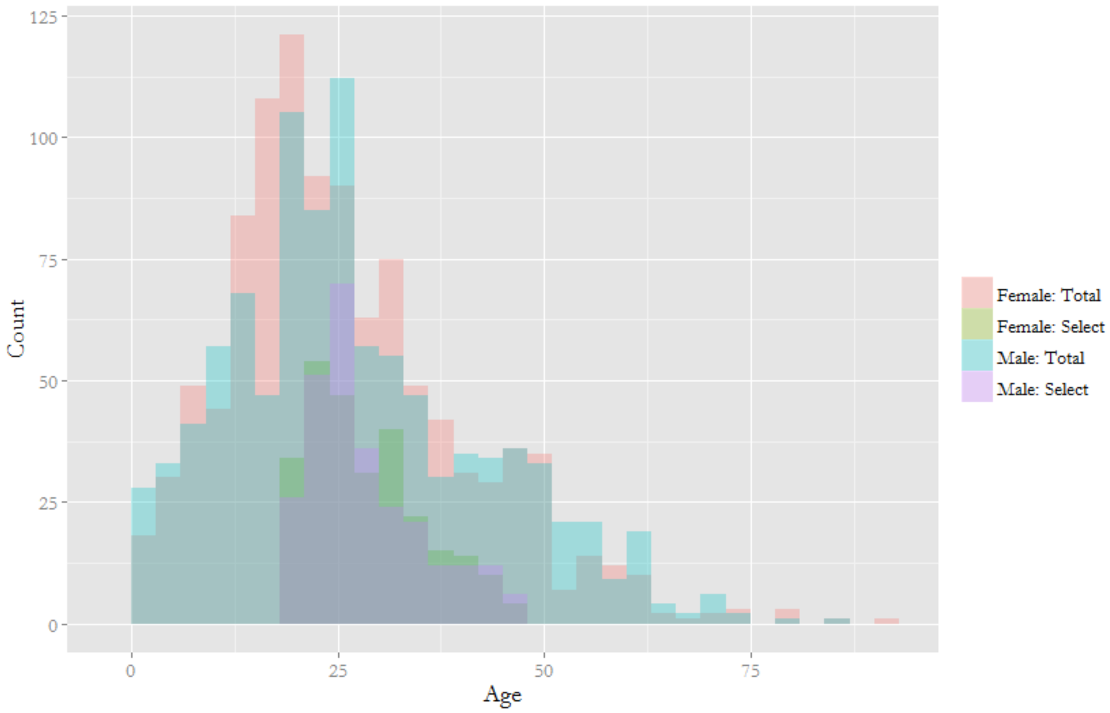

Once the sampling frame was created, and the eligible participants identified, random sampling was conducted to invite participants to the experiment. Two groups of six were selected, from each cell, for each day; alongside a backup list of the next randomly selected participants. Mobilizers were given the information of the first 12 participants there were to invite and bring to the experiment and if difficulties arrived in getting participants then they were replaced by the next participant. Replacements were common, and on average the experiments began an hour late, due to mobilization issues.[12]

B.2.2 Student Population

The student sampling frame was somewhat easier. The Dean of Students, at UCU Mbale, was approached to provide a list of students who were registered. A list of 640 students was acquired and random sampling was carried out to invite the 48 participants to the experiment. A postgraduate student at UCU, was recruited to mobilize the participants.

mobilisers with a particular political party and claim they were tired of government programmes because they only register people but never deliver their promises.

[12]Further issues concerned the gender bias. Although the ratio of men to women, who were randomly selected, was approximately equal, the drops outs were far more often women. To attempt to resolve this issue, when selecting replacements women's names were prioritised.

B.3 Sample Characteristics

Between experimental days there were two separate samples, the general population (n=108) and the student population (n=48). Of these 156, 7 participants were excluded from the analysis: ID 1_11, 6_11 and 8_01 had a clear lack of comprehension (where enumerators had to assist decision making); concerns of contamination between sessions led to the exclusion of 9_03 and 9_08; 1_05 was as an outlier; and data for 3_09 was lost due to a tablet crashing.

Table 3.11 shows the characteristics for the remaining 149, within the two samples. In both samples there are fewer women, 43% on average, than the 51.5% collected from the sampling frame. In both samples the modal age group is $22-25$, but the student population has significantly less over 25 year old participants. Both samples are relatively well educated, 56.4% of the general population are educated to the bachelors level. Every member of the general population have completed at least primary school, but there is a significant proportion who have either not completed secondary, or have only completed secondary. Every participant identified themselves as being religious, not one was agnostic or atheist as is common in western samples. The majority are Christian, with the largest proportion of these being Anglicans in both samples. There is a substantial minority of Muslims within the general population, but not in the student population; as they were sampled from, the Ugandan Christian University. It is important to note that the general population is *not representative*, due to the necessary selection criteria.

Table 2.8: Sample Characteristics

	Sample					
	General		**Student**		**Total**	
	No.	%	No.	%	No.	%
Gender						
Male	59	58.4%	26	54.2%	85	57.0%
Female	42	41.6%	22	45.8%	64	43.0%
Age						
18-21	16	15.8%	10	20.8%	26	17.4%
22-25	37	36.6%	33	68.8%	70	47.0%
26-30	18	17.8%	2	4.2%	20	13.4%
31+	30	29.7%	3	6.3%	33	22.1%
Religion						
Catholic	25	24.8%	12	25.0%	37	24.8%
Anglican	37	36.6%	21	43.8%	58	38.9%
Muslim	16	15.8%	0	0.0%	16	10.7%
Seventh Day Adventists	0	0.0%	1	2.1%	1	0.7%
Born Again	23	22.8%	14	29.2%	37	24.8%
Highest Level of Education						
Secondary: Incomplete	11	10.9%	0	0.0%	11	7.4%
Secondary: Complete	14	13.9%	0	0.0%	14	9.4%
Tertiary: College	19	18.8%	0	0.0%	19	12.8%
Tertiary: Bachelors	56	55.4%	48	100.0%	104	69.8%
Tertiary: Masters	1	1.0%	0	0.0%	1	0.7%
Total	101	100.0%	48	100.0%	149	100.0%

B.4 Wealth and Multidimensional Poverty Indices

B.4.1 Wealth Index

Alongside the experiment an in-depth survey was run. Part of this survey contained variables to assess the wealth of the individual participants. Rather than using income data, which are notoriously unreliable in developing countries, or consumption data, which are very time consuming to collect, a focus upon assets was made. An extensive list of variables has been collected, in order to calculate a relative wealth index; using methods standard to the literature.

Data were collected for 57 variables to be included within the Wealth Index. Dummy variables were established for: access to electricity; material of the floor, roof and exterior walls; water source; toilet type; shared toilet; lighting; cooking fuel and rubbish collection. Numerical values were established for number of bedrooms; distance to water source (mins); livestock; vehicles; and household assets. Each of these variables, split into subcategories, are shown in Table 2.9. The mean value shows either the proportion of the populace who own that asset, for dummy variables, or the mean number owned by the population. The standard error shows the variance of ownership between individuals.

In order to establish the Wealth Index the methods set out in Vyas and Kumaranayake (2006) have been followed. Whereby Principal Components Analysis is run, on the variable list, then the weights from the first component are used to form the base of the wealth index. These weights, shown in Table 2.9, are used to weight each observation for each participant, to provide a single index. This index is then normalised for the sample, providing an index between 0 and 1, where 0 is the least wealthy and 1 is the most wealthy, in relative terms.

What this method reveals is the impact that owning a particular asset has upon the wealth of the individual. The direction and magnitude of the weight implies a relationship between a particular asset and the expected wealth. For instance, having electricity shows that individuals are more wealthy. As do more bedrooms, closer water, rubbish collection, the number of vehicles (apart from bicycles) and the number of household assets (with the exception of handmills). Delving into particular materials it can be observed that the higher quality materials lead to higher wealth. For floor, there is a high positive weight for cement, while dung and earth/sand have negative weight, the former being the higher weight of the two. This can be seen throughout each category: tiled roofs, brick/cement walls, private piped water sources, flush toilets, electric lighting and gas/electric cooking fuel all relate to a higher wealth. On the opposite scale, thatch/straw roofs, mud/pole walls, boreholes as a source of water, pit latrines, paraffin/gas lighting and wood for cooking fuel all lead to lower wealth levels. Interestingly, we observe that owning livestock, apart from hens, leads to lower wealth. This is likely explained by the urban nature of the sample, those who owned large animals are more likely to be rural, and possibly poorer than their urban counterparts, in terms of the other assets mentioned. Hens within a city are more feasible, and owning them is, perhaps, a sign of wealth. The ownership of vehicles increases wealth, with the exception of a bicycle as perhaps the someone with more wealth may dispense with the bicycle in favour

Table 2.9: Wealth Index Construction; PCA

	(1) Weights	Mean	Std. Err.
Electricity	0.245	0.846	(0.363)
Number of Bedrooms	0.042	2.309	(1.456)
Floor Material			
- Earth/Sand	-0.077	0.020	(0.141)
- Dung	-0.240	0.101	(0.302)
- Cement	0.120	0.698	(0.461)
- Ceramic Tiles	0.059	0.067	(0.251)
- Vinyl/Carpet	0.042	0.114	(0.319)
Roof Material			
- Thatch/Straw	-0.133	0.020	(0.141)
- Iron Sheets	0.002	0.872	(0.335)
- Asbestos	0.003	0.034	(0.181)
- Tiles	0.051	0.047	(0.212)
- Cement	0.042	0.027	(0.162)
Exterior Wall			
- Mud/Poles	-0.238	0.101	(0.302)
- Brick/Mud	-0.049	0.040	(0.197)
- Brick/Cement	0.200	0.805	(0.397)
- Cement	0.009	0.054	(0.226)
Water Source			
- Private Pipe/Bottled	0.215	0.617	(0.488)
- Public Taps	-0.033	0.215	(0.412)
- Protected Well	-0.170	0.067	(0.251)
- Borehore	-0.161	0.101	(0.302)
Water Distance (mins)	-0.244	6.544	(14.194)
Toilet			
- Flush	0.180	0.503	(0.502)
- VI Pit Latrine	0.030	0.282	(0.451)
- Pit Latrine	-0.252	0.215	(0.412)
Shared Toilet	-0.015	0.315	(0.466)
Lighting			
- Electric	0.256	0.826	(0.381)
- Paraffin/Gas	-0.228	0.107	(0.311)
- Solar	-0.105	0.067	(0.251)
Observations	149		

	(1) Weights	Mean	Std. Err.
Cooking Fuel			
- Gas/Electric	0.067	0.074	(0.262)
- Charcoal/Coal	0.172	0.725	(0.448)
- Wood	-0.264	0.161	(0.369)
No Food Cooked	0.015	0.040	(0.197)
Rubbish Collected	0.139	0.416	(0.495)
No. of Livestock			
- Heifer/Cow	-0.115	0.624	(2.042)
- Bull/Oxen	-0.089	0.148	(0.608)
- Calves	-0.096	0.114	(0.458)
- Donkey	-0.070	0.007	(0.082)
- Goats	-0.095	0.732	(1.814)
- Sheep	-0.013	0.376	(4.109)
- Pigs	-0.060	0.040	(0.305)
- Hens	0.024	1.295	(8.339)
No. of Vehicles			
- Bicycles	-0.050	0.342	(0.655)
- Motorcycles	0.025	0.208	(0.424)
- Cars	0.108	0.275	(0.624)
- Vans	0.041	0.020	(0.141)
- Trucks	0.056	0.134	(1.239)
- Tractors	0.049	0.013	(0.115)
Household Assets			
- Sofas	0.139	0.893	(0.617)
- Radios	0.064	1.000	(0.678)
- Tables	0.110	1.034	(0.911)
- Fridges	0.167	0.362	(0.548)
- Televisions	0.193	0.926	(0.679)
- Computers	0.134	0.564	(0.738)
- Clocks	0.110	0.913	(0.830)
- Jewellery	0.107	2.027	(2.043)
- Mobile Phones	0.097	4.148	(2.675)
- Handmills	-0.009	0.336	(0.565)
Observations	149		

of a car. While owning all assets, with the exception of a handmill (a traditional pestle and mortar, predominately owned by rural or poorer individuals who do not have more modern electric blenders) also increased wealth.

In order to split the population into socio-economic groups two methods are common in the literature; an arbitrary split, into quintiles for example, or through cluster analysis. Here the later approach has been taken, using K-means with three degrees of freedom to create three cluster: Low, Middle and High. Figure 2.14 shows the overall distribution of wealth, split into these three clusters. Table 2.10 shows the percentage of the population who fall into these three clusters, the mean Wealth Index and the standard deviation. From these results the largest proportion, 51.01%, of the population lies within the high wealth category. The lowest socio-economic group is the smallest cluster, while the middle is in-between.

To provide comparability between this study and potential others a further wealth index was calculated, the International Wealth Index (IWI). The index was proposed by Smits and Steendijk (2015) as a way in which to compare between studies, due to the relative nature of typical wealth indices. The index is made up of a subset of the variables used above: floor material, toilet facility, number of rooms, access to electricity, water source. Alongside

Figure 2.14: Wealth Index Distribution

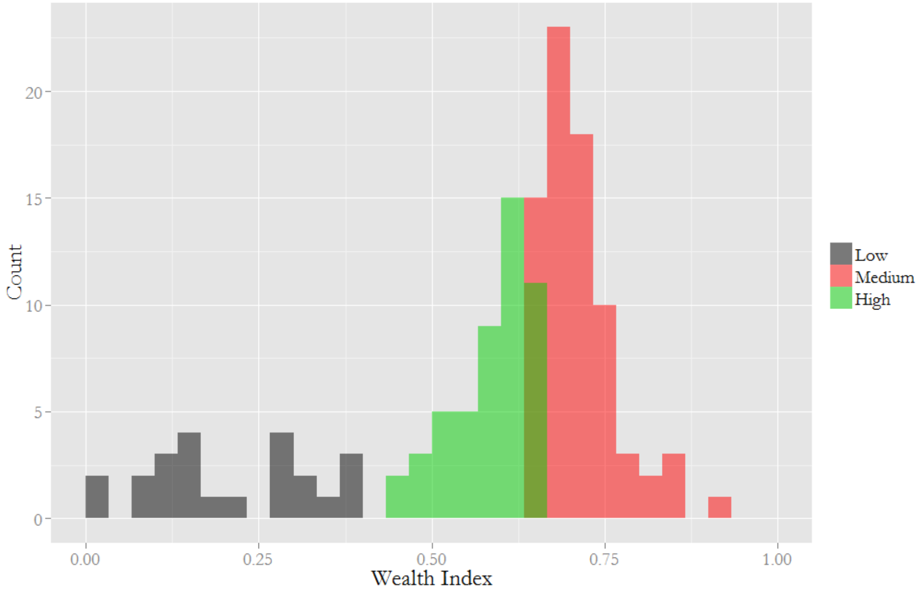

Table 2.10: Wealth Index Categories

	%	Mean WI	Std. Dev.
Low	15.44%	0.209	(0.115)
Middle	33.56%	0.586	(0.0561)
High	51.01%	0.714	(0.061)

dummies for consumer durables: television, refrigerator, phone, car, bicycle, cheap utensils and expensive utensils. The only divergence which was necessary, was that phone had to include only mobile phones and that number of rooms had to be number of bedrooms, due to data limitations.

The mean IWI was 66.308, with a minimum of 12.003, a maximum of 100 and a standard deviation of 20.871. In terms of headcount cut-offs the IWI-30 poverty line is 8.05% and the IWI-50 poverty line is 19.46%; which most closely correlates to the $1.25 and $2.00 poverty cutoff lines, respectively. The IWI and Wealth Index are highly correlated, with a Pearson's Correlation Coefficient of 0.8750.

B.4.2 Multidimensional Poverty Index

The Multidimensional Poverty Index (MPI) is a poverty index, which accounts for severe deprivations in Health, Education and Living Standards. An absolute measure of poverty which considers the necessary capabilities that individuals need in order not to be in poverty. It follows the Alkire and Foster Methodology (Alkire and J. Foster, 2011); where individual households are assessed on nine indicators. The intensity of poverty is calculated by summing the weighted indicators, where each of the dimensions: *health*, *education* and *living standards*; are equally weighted. Education contains two indicators: years of schooling and child school attendance. Health is equally split between child mortality and nutrition. Living standards contains five indicators: electricity, improved sanitation, safe drinking water, adequate flooring, cooking fuel and asset ownership. Data are usually gathered from sources such as DHS, MICS and WHS surveys which are provide large scale, cross-country, datasets. The indicators have been chosen to maximise the sensitivity and usefulness of the index, alongside avoiding issues of missing data.

An alternative, Extended MPI (EMPI) has also been calculated. This makes use of the extensive survey data collected, avoiding issues of missing data that may have limited the number of indicators previously. Adding to the education dimension are the indicators: schooling gap and adult education. Food shortage and health care have been added to the health dimension. While, external wall and roofing have been included within living standards. In addition, a further dimension, *urban*, has been added to account for particular urban characteristics of the sample. Within this dimension the indicators are: overcrowding, unemployment, underemployment, bank account and rubbish collection. These additional dimensions have stemmed from a combination of recent additions to the literature and Mbale specific issues. Other extended indices include the MPI-LA (Santos and Villatoro, 2016), the Mexican MPI (J. E. Foster, 2007), the Colombian MPI (Angulo, Dìaz, and Pardo, 2016), and Bhutan's Gross National Happiness Index (Ura et al., 2012); alongside MPI's in Chile and Ecuador. Specific indicators to Mbale, which could also be widespread issues, include the government clinic access and underemployment issues the town has.

Table 2.11 shows each of the dimensions and indicators. Showing the cutoff thresholds for the cut-off points, alongside the weights for the standard and extended MPI. The weights for the EPMI have been calculated according to similar principles, each dimension has an equal weight, and given a dimension indicator has an equal weight, within. Although the surveys were extensive, they were conducted on an individual level. As a result two indicators, nutrition and underemployment, can only consider the individual rather than the household this may lead to an under-measurement of poverty for both indicators.

A further index, the Preference MPI (PMPI), is calculated; incorporating the *preference* survey responses. In the above MPI calculations the weights between dimensions are chosen to be equal. The choice of weights is, however, open to debate. Is education of equal importance as health and living standards when assessing the intensity of poverty, or do different weights better capture poverty? Alternative weightings could be used, and by using them the resulting indices would differ. An alternative approach, which could be used is to identify the weights that the individuals in question would give to each dimension. The Preference MPI illustrates how this could be done. Using the elicited weightings for health, education and living standards, from the survey data, to calculate the index.

Once the Dimension, Indicators and Weights have been decided upon the statistics relating to the MPI, EMPI and PMPI can be calculated. Figure 2.15 shows the distribution of the uncensored intensity (the proportion of simultaneous deprivations experienced). The average uncensored intensity is 0.126, 0.178 and 0.105, for MPI, EMPI and PMPI, respectively. A percentage of the population classed as MPI, EMPI and PMPI Poor: 6.04%, 10.07% and 5.37%

Table 2.11: MPI Dimensions, Indicators, Deprivation Thresholds and Relative Weights

			Weight	
Dimension	**Indicator**	**Deprived if...**	**Std.**	**Ext.**
Education	Years of Schooling	No household member has completed five years of schooling.	$\frac{1}{6}$	$\frac{1}{16}$
	Child School Attendance	Any school-aged child is not attending school up to class 8.	$\frac{1}{6}$	$\frac{1}{16}$
	Schooling Gap	Any child is over two years delayed with respect to their school grade for age.	.	$\frac{1}{16}$
	Adult Education	Any adult in the household can not read and write, or have not completed secondary school.	.	$\frac{1}{16}$
Health	Child Mortality	Any child has died in the family	$\frac{1}{6}$	$\frac{1}{16}$
	Nutrition	The participant is malnourished, BMI $\leq$ 18.5	$\frac{1}{6}$	$\frac{1}{16}$
	Food Shortage	The household has suffered a food shortage in the last Month	.	$\frac{1}{16}$
	Health Care	The household only has access to a Government Clinic and has no Health Insurance	.	$\frac{1}{16}$
Living Standard	Improved Sanitation	The household's sanitation facility is not "improved", or shared.	$\frac{1}{18}$	$\frac{1}{32}$
	Electricity	The household has no access to electricity.	$\frac{1}{18}$	$\frac{1}{32}$
	Drinking Water	The household does not have access to safe drinking water or it is at least a 30-minute roundtrip away.	$\frac{1}{18}$	$\frac{1}{32}$
	Flooring	The household has a dirt, sand or dung floor.	$\frac{1}{18}$	$\frac{1}{32}$
	External Wall	The household has external walls made from thatch, straw or mud and poles.	.	$\frac{1}{32}$
	Roofing	The household has a roof made from thatch, straw or banana fibres.	.	$\frac{1}{32}$
	Cooking Fuel	The household cooks with dung, wood or charcoal.	$\frac{1}{18}$	$\frac{1}{32}$
	Assets Ownership	The household does not own more than one radio, TV, telephone, bike, motorbike or refrigerator and does not own a car or truck.	$\frac{1}{18}$	$\frac{1}{32}$
Urban	Overcrowding	The household has three or more people per bedroom.	.	$\frac{1}{20}$
	Unemployment	Any household member is long term unemployed and noone receives Social Insurance.	.	$\frac{1}{20}$
	Underemployment	Participant works four hours or less daily.	.	$\frac{1}{20}$
	Bank Account	No household member has access to any form of bank account.	.	$\frac{1}{20}$
	Rubbish Collection	Rubbish is not collected from the household.	.	$\frac{1}{20}$

respectively. These *incidence* rates are rather low in comparison to elsewhere in Uganda, where (in 2011) the average across Uganda was 69.9% and the Urban average is 29.2% (OPHI, 2017).

Two main observations emerge from these results. The first, that the EMPI reveals a higher level of multidimensional poverty than the MPI. Two factors could explain this: (1) the EMPI is a richer index and, therefore, better identifies poverty, implying MPI measures of poverty are biased downward. (2) The thresholds in the additional indicators in the EMPI are 'too low', biasing the EMPI upwards. The second, is that the PMPI reveals a lower level of multidimensional poverty than the MPI. Again two, possible, factors could explain this: (1) As active agents, individual's preferences shape their priorities over dimensions of poverty. As particular dimensions are weighted higher, household resources are directed towards reducing poverty in that dimension, and poverty is therefore reduced. (2) Poverty shapes preferences, households weight lower those dimensions in which they are deprived. Here, these two observations cannot be further investigated, but highlight a potentially interesting area of investigation.

Figure 2.15: Distribution of Uncensored Intensity for MPI, Extended MPI and Preference MPI

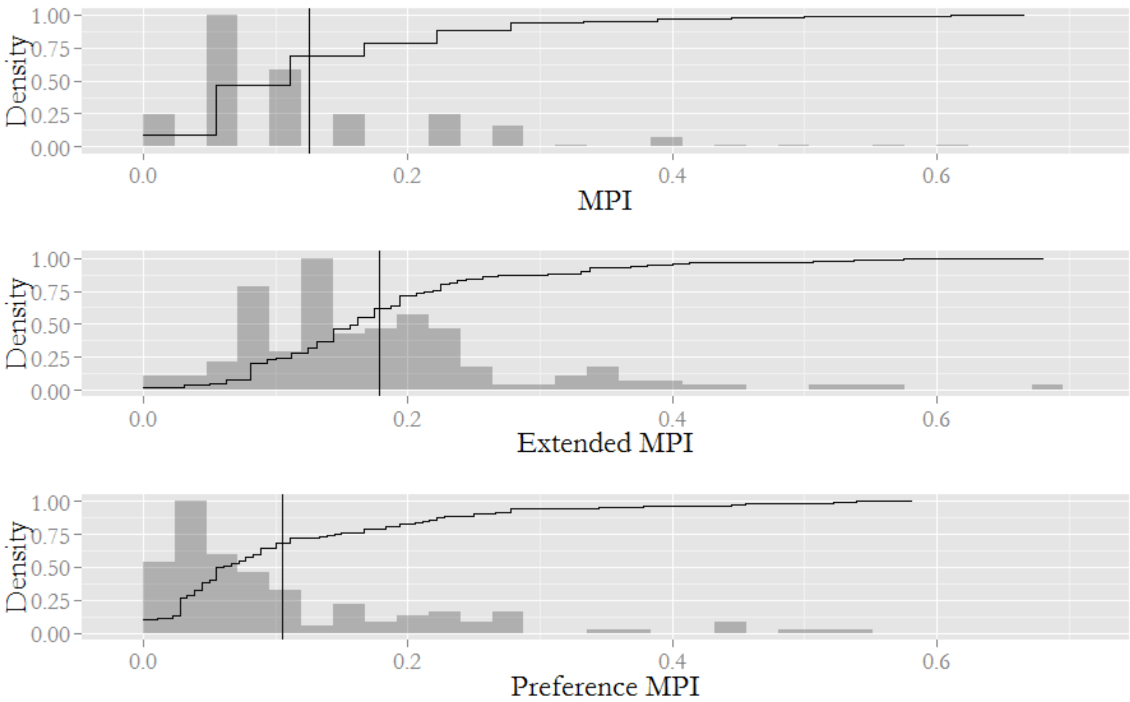

Analysing specific indicators reveal the dimensions in which are deprived in. Censored Headcounts, show the proportion of the sample who are classed as MPI Poor and deprived in that particular indicator. Firstly, considering the MPI, the censored headcount in education is low (1.01%), health the next lowest (3.36%), followed by living standards (3.69%). Years of schooling is particularly low (0%), as every member of the survey had had at least 5 years of schooling themselves, so are not deprived in that measure. Child mortality (5.37%) is somewhat higher than nutrition (1.34%), but this is perhaps affected by the individual level nature of the measure. Households are typically not deprived in assets (0.67%), but indicators such as cooking fuel (6.04%), electricity (4.70%) and flooring (4.70%) score relatively high.

The Extended MPI shows the censored headcounts in each dimension are: education: 2.01%, health: 2.52%, living standards: 3.44% and urban: 2.28%. Adult education is high 5.37%, school attendance is 2.01%, the schooling gap is 0.67% and years of schooling remains at zero. Food shortage (1.34%), nutrition (1.34%) and healthcare (2.01%) are similar to one another, with child mortality (5.37%) remaining high. The external wall (4.03%) appears similar to the other house related indicators, but roofing deprivations (1.34%) occur less often.

Figure 2.16: Percentage of the Population who are MPI Poor and Deprived in each Indicator

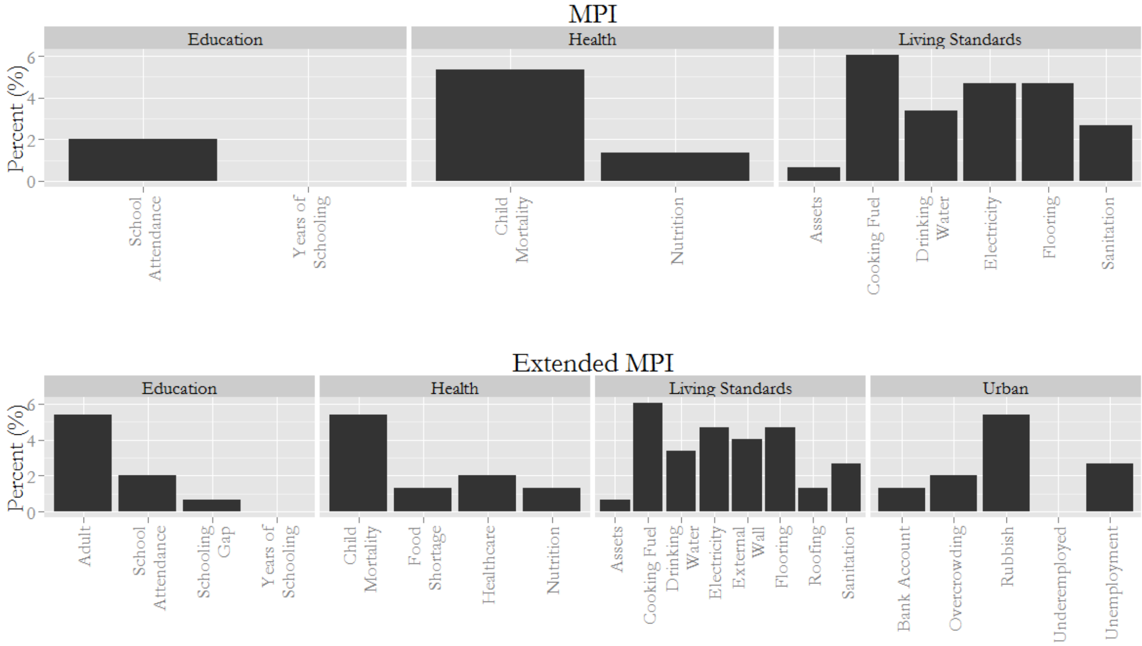

The greatest contributor to the urban dimension is rubbish (5.37%), with unemployment (2.68%) followed by overcrowding (2.01%), then bank account (1.34%). Underemployed has a zero incident rate. To reduce multidimensional poverty, the main considerations are, therefore, the levels of child mortality, cooking fuel, housing materials and drinking water. Alongside these, rubbish collection and unemployment appear the worst within the additional urban dimension.

B.4.3 Comparison of Poverty Indices

The five poverty indices (MPI, EMPI, PMPI, WI and IWI) are different in their construction, but can be compared to identify how correlated they are. Table 2.12 shows a pairwise correlation matrix between each of the indices. Each pairwise relationship is statistically significant at the 0.01% level. The strongest relationships are between the similar indices, MPI, EMPI and PMPI; and IWI and WI, which is to be expected. Strong relationships do also emerge between the wealth and multidimensional poverty indices.

Table 2.12: Pairwise Correlation Matrix of Wealth Indices

	MPI	EMPI	PMPI	IWI	WI
MPI	1.0000				
EMPI	0.7255***	1.0000			
PMPI	0.9275***	0.6523***	1.0000		
IWI	-0.5741***	-0.6628***	-0.4713***	1.0000	
WI	-0.6341***	-0.6812***	-0.5344***	0.8750***	1.0000

B.4.4 Comparison to National and Urban Samples in Uganda

In order to compare our sample with the wider population in Uganda, data from the Uganda National Panel Survey is used (UBOS, 2016). Both a national and urban sample was analysed with a sample size of 3,200 households for the former, and 830 households for the latter. Both the standard MPI and the IWI were calculated, to provide a meaningful comparison with our sample. Figure 2.17 shows the distribution of MPI intensity and IWI levels for the national and urban samples, alongside our experiment sample. The comparison shows that the experimental sample has lower intensity of multidimensional poverty and higher asset wealth levels, than both the national and urban sample. To summarise, the mean MPI intensity is 0.375, 0.257 and 0.136, while the mean IWI levels are 25.66, 41.16 and 73.07, for the national, urban and experiment samples respectively.

Figure 2.17: Wealth and MPI comparisons to National and Urban Sample

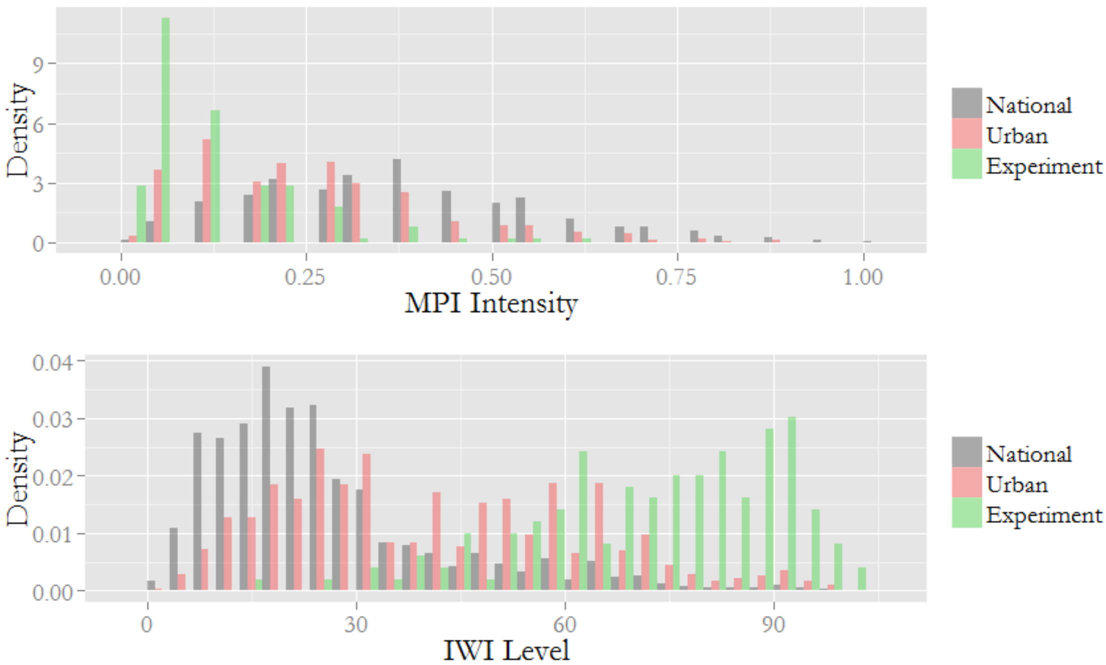

B.5 Script

B.5.1 General

Welcome. Thank you for taking the time to come today. My name is ...... and I am working with researchers from the EXEC laboratory at the University of York. This University is found in the UK. We have invited you here, today, because we want to learn about how people in this area make decisions. We are conducting both this study and a survey, and you will be required to take part in both. This section is the study where you are going to be asked to make decisions that will earn you some money. The money that results from your decisions will be yours to keep.

What you need to do will be explained fully in a few minutes. But first, there are three things I would like to explain to you clearly and you should consider them as very important. First of all, this is not our money. As I told you before, we work for a university and this money has been given to us by that university for this research. Participation is voluntary. You may still choose not to participate in the study. We also have to make clear that this is research about your decisions. Therefore I will not allow you to talk with anyone else during the study. This is very important. I'm afraid that if I find you talking with someone else during the study I will kindly ask you to go back home and what this means is that you will not earn any money. If you have any questions, please ask any of us.

Do you have a mobile phone? [**If yes**] Can I please ask you to switch it off? Make sure that you listen carefully. You could earn a good amount of money today, and it is important that you follow my instructions. There are no right or wrong decisions here, we are interested in exactly what decisions you want to make. But think seriously about your decisions as it will possibly affect how much money you will take home. In this study you will be making decisions about the distribution of money between yourself and other participants in this room. How this is done will be explained to you shortly. You will get an attendance fee of shs5000 in addition to the money you will earn from the study as a result of the decisions that are made.

In order to make your decisions you will be using tablets that we shall provide. This is what they look like [**Show Tablet**]. The decisions you make will be made by interacting with the screen. These tablets are touch screen similar to the smart phones that some of you use. They are easy to use. Do not get worried if you are not familiar with using them because before making your decisions we shall guide on how to use them.

In this study we shall have two parts. Part 1 and Part 2. In Part 1 there will be 27 rounds and in Part 2, there will also be 27 rounds. So in total, there will be 54 rounds. In order to determine your final payment. Out of the 54 rounds, one shall randomly be chosen by asking one of you to pick a number of out this bag [**Show the Round Selection Bag**]. This bag has 54 tokens which represent the 54 rounds. One token will be selected like this [**Pick One Token**] and this will be the round that will determine the final payment that all participants receive. This means that every round has an equal chance of determining your final payment, so consider each decision you make carefully. Everyone will finish at the same time, as you need to wait for every participant to finish before the payments can be determined. The decisions that you make will involve distributing money amongst a group of three participants; yourself and two others in the room. In each round you will be randomly linked to the two other people.

Now listen carefully, I will explain how the distribution of the money works. Do not worry, it is easy to understand, but it is important that you remember it. The money you have to Allocate amongst the group will come from a Budget. You must decide how to Allocate all of the money from the Budget. The Allocations that you make will then be divided by a Divider to give the Payoff that a player will get in a particular round. This will be done by dividing the Allocation you make by the number referred to as the Divider. The Payoff can be thought of as the final amount of money that everyone gets in that round.

The decisions you make in each round will determine the Payoffs for yourself and the other two people in your group would receive. Each of you will make your decisions independently of one another, but only the decision of one person will be selected, randomly, to determine the final Payoffs for all the people in the group. You will not know whose decision has been chosen and will receive your Payoff, individually, at the end. This means that every decision you make is entirely anonymous. Importantly, there are no right or wrong answers, the decisions you make are entirely up to you.

To make your decisions you will be using the tablets, the screen you will see will look like this [**Show Tablet Poster**]. You will always be Player 1, the player on the left [**Point to Player 1**]. Player 2, in the middle [**Point to Player 2**] and Player 3 on the right [**Point to Player 3**] will be the other two participants who you are linked to.

You will be given a budget to allocate among the three of you. This is shown on the left [**Point to Budget Remaining**], and in this example is shs60,000. You must spend the entire budget in each of the rounds. This means the Remaining Budget must be zero. You will be able to make the allocations in three ways. The first is with the sliders [**Point to the Sliders**]; you can drag the sliders to any allocation that you want, for each of the players [**Pretend to Drag the Sliders**]. The second is with the arrow keys [**Point and Pretend to Press the Arrow Keys**]. They allow you to make small increases and decreases in your allocations to each player. The third is the written input [**Point to Written Input**]; you can click in each of the boxes, type your desired allocations to each player with the keyboard and click update [**Point to each Box and then to Update**].

Within each of the rounds there will be different dividers for each player [**Point to the Dividers at the Top**]. The actual payoff that each player will get will be the allocation you give, divided by the divider [**Point to the Payoff at the Top**]. These dividers are important as they change every round. It is the payoff which will be given to each participant so it is important to consider the distribution of payoffs. The payoffs will be always be in shillings and will be shown by the height of the orange bars, the orange numbers beside them and by the numbers at the top of the screen, which you will see later in the tutorial.

You will also see the Sum of Payoffs [**Point to the Sum of Payoffs**] and the Gap Between Payoffs [**Point to the Gap Between Payoffs**]. The Sum of Payoffs is the Payoffs of Player

1, 2 and 3 added together. The Gap Between Payoffs is the highest Payoff minus the lowest Payoff. Notice how these change when making your decisions. You must make a decision in every round, and then click Next or Finish to confirm your decision [**Point to Next/Finish**].

B.5.2 Part 1

In Part 1, the budget you have to allocate will change. The others, Player 2 and Player 3, will always be random and anonymous participants in the room. As earlier mentioned, there will be 27 rounds. After Part 1, you will have a short break, before we explain Part 2. Now we shall guide you on how to use the tablets, and how to participate in the study. One of us will be there with you to guide you. The first tutorial you will see is not the whole picture, it will just explain to you how to use the tablets. The second tutorial will show you the interface which is on the poster here [**Point to Poster**], and explain in more detail. If you have any questions now, or throughout the study, please raise your hand and we will help as best we can. There will be at least two of us in the room at all times.

B.5.3 Part 2

[**Do Not Read Part 2 Until After the Break**]**.** In this part, we shall start by asking you a question on how connected you are to each of the other people in this room in real life. By this we mean how close you are to the other people in everyday life. You will be shown seven diagrams which look like this [**Show Oneness Poster**]. You will find these diagrams on the tablets that you used in Part 1. In each diagram there are two circles, the one on the left denotes yourself [**Point to One Circle on the Left**] and the one on the right show another person in the room [**Point to One Circle on the Right**]. Each diagram represents the degree of closeness between you. If we start at diagram A [**Point to A**] where the circles are very far apart, we see a very distant connection. As we move to G [**Move from A through G**] we see the circles moving closer, representing closer connections, with G as the closest connection.

You will be asked to select the diagram which best represents your connectedness with each person in the room. You will know who the 'other' person is as it will be denoted by their desk number. You will have a sheet of paper on your desk which shows the layout of the desks by number. There is also a number on your desk and each of the other participant's desks which you can see if you look around the room. You should look around to see exactly who is sat at each desk, when making your decisions. In order to select the diagram click on it, it will be highlighted, then you can click on Next in the bottom right corner to move to the next question, about another person.

When you have finished these questions for all 5 other people in the room you will move onto Part 2. Part 2 is similar to Part 1, but in this case you will know the players in the group by their desk number. The layout and interaction with the tablet will be the same as before [**Point to Tablet Poster**]. However, you will now get to know who Player 2 and Player 3 are [**Point to Player 2 and Player 3**] by their Desk Number. Remember to look at exactly who are sat at the desks you are distributing between. As before the decisions and payoffs will remain anonymous and so the other players will not know the decisions you make and you will not know the decisions of others.

B.6 Tutorial Script

B.6.1 Tutorial 1 – One Player

[Enumerators are here to Guide Them One-on-One, Let Them Read the Instructions Themselves, but Carry Out the Instructions in this Script]

Screen 0 - Welcome to the Tutorial. During this tutorial your decisions do not affect your payoffs. They are only to help you learn the experiment. You can click Next and Back in the bottom right corner to move between Tutorials [**Click on Next, then Click Back**]. Follow the instructions to work through the tutorial. Now click Next.

Screen 1 - This is the Allocation Slider. It allows you to distribute the Budget amongst each player. Click the Slider and Drag it upwards. Notice the higher you drag the slider the higher the allocation is, and the lower the Remaining Budget. [**Drag the Slider to Several Points, Each Time Show How the Allocation Increases and the Remaining Budget Decreases, Then Let Them Try**]. When you have done so, click Next.

Screen 2 - Now you have some more information move the Slider again. Notice the Orange Bar, and orange number, these represent the Payoffs. [**Drag the Slider, Pointing Out the Allocation and the Payoffs, both the Numbers and the Heights, Let them Drag the Slider**]. The Payoff is the Allocation divided by the Divider. The Divider is shown at the top, and is 1 in this round. [**Point to the Divider**]. That means that the Allocation and Payoffs are the same. [**Highlight this with an Example**]. When you are finished click Next.

Screen 3 - In each round the Divider will change for each Player. Now the Divider has been increased to 2. [**Point to the Divider**]. Move the Slider and notice the difference to the Payoff. [**Drag the Slider to Several Points, Show Them the Link Between Allocation and Payoff, Emphasize the Payoff is Important, Let Them Drag the Slider**]. It is now always the Allocation divided by 2. When you are finished click Next.

Screen 4 - Alongside the Slider you can also click the Arrow Keys to change the Allocation by small amounts. The Dark Grey arrows change the allocation by 10. The Light Grey arrows change the allocation by 1. Click each of the arrows to see how they change the Allocations. [**Show Them How to Click Up and Down On Both Arrows, Showing How the Allocation/Payoff is Changing**]. Once you have done this, click Next.

Screen 5 - The final way of changing the Allocations is through the Written Input. Click the light blue box, then type an Allocation using the keyboard. Once you have done this, click Update. [**Show Them How to Use the Written Input, Showing How the Allocation/Payoff is Changing**]. After trying the Written Input, click Next.

Screen 6 - Now try and change the Allocations with all three methods. When you finish Allocate all of the remaining Budget and click Finish to end this part of the Tutorial.

B.6.2 Questions

[Enumerators are here to Ensure they Answer Every Question Correctly the First Time they answer, so long as they understand the question. If They Do Not Answer Correctly, Explain Until They Understand, If They are Struggling Go Back to Tutorial]

Q1. Point to the numbers representing the Allocation, Divider and Payoff, on this screen. Then to which of the Slider and Orange Bar represents the Allocation and Payoff.

Q2. What are the three methods of Allocating the Budget?

Q3. If an Allocation is $20,000 and the Divider is 1, as on the screen, what is the Payoff?

Q4. If an Allocation is $20,000 and the Divider is 2, as on the screen, what is the Payoff?

Q5. If the Remaining Budget is $5000 can you proceed to the next Round? Explain why.

Q6. If, in one round, the Allocation is $45,000, the Divider is 3 and the Payoff is $15000, to yourself, what is the Final Amount of Money you will get from that round?

B.6.3 Tutorial 2 – Three Players

[Enumerators are now here to Answer any Questions One-on-One, Let Them Read the Instructions Themselves, and Make any Allocations Themselves].

Screen 1 - Three Person Tutorial. Throughout the experiment you will be allocating between three players, one of which is yourself. You will always be Player 1. While Player 2 and Player 3 will be random anonymous other people in the room. You must allocate all of the budget between the three players. Use the sliders, arrows and written input, as before. This Tutorial round is still a practice round and is not for real. Click Next to proceed. [**Do Not Do Any of the Allocations for Them, Let Them Do it, But be There to Help if They Need**].

Screen 2 - Three Person Tutorial. Now try when the Dividers are different. Notice to the Sum of Payoffs and Gap Between Payoffs, to the right hand side. The Gap Between Payoffs if the difference between the highest payoff and the lowest payoff. [**Point to Sum of Payoffs and Gap Between Payoffs; Do Not Do Any of the Allocations for Them, Let Them Do it, But be There to Help if They Need**].

Screen 3 - Three Person Tutorial. You have finished the tutorial. Now every decision you make will be for real. So take care each round, as it is equally likely to be the actual payoff for yourself and the others in the group. **[Leave Them to Do the Experiment].**

B.7 Sensitivity of Treatment Selection

Within the main analysis one particular assumption is made concerning how participants view *oneness* when others are anonymous. Anonymity here, means that the 'dictator' knows that Player 2 and 3 are others in the experimental session, but does not know their identity. It is assumed that distributional decisions may be affected by how closely connected the 'dictator' is to the others in the group, as a whole; meaning that 'dictators' may be more generous when they are (randomly placed) in a group with whom they feel close. To account for this, the *oneness levels*, which enter into the utility function in the *budget* treatment, are assumed to be the expected value of the elicited oneness levels for the group.

The main analysis incorporates this assumption, so that *oneness* can be incorporated in decision problems where the 'recipients' are anonymous; as is most common in the literature. It could prove to be an interesting area of investigation, for giving in other dictator games.[13] Yet, objections could be made to this assumption. To circumvent this assumption, preference parameters can be estimated using data only from the *oneness* treatment, where the identity of the other players is *known*. Below sensitivity analysis shows the comparison of estimated preferences from the *both* treatments (54 rounds) and the *oneness* treatment (27 rounds).

Figure 2.18: Sensitivity of Treatment Selection: r, α, ϕ and ψ

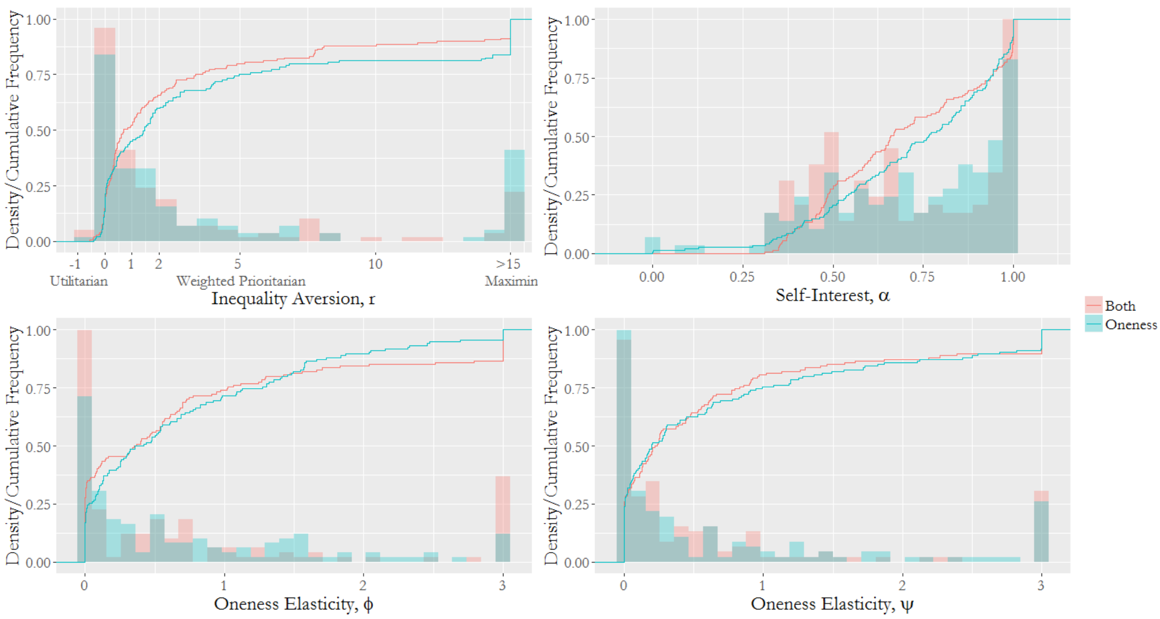

Figure 2.18 shows the distribution r, α, ϕ and ψ; estimated using data from *both* and *oneness* treatments. The comparisons imply that, on average, estimates of inequality aversion and self-interest are higher; when using only the *oneness* treatment data. Oneness magnitude parameters, ϕ and ψ appear very similar, but less extreme values of ψ are estimated using *oneness* treatment data. Kolmogorov-Smirnov corrected test-statistics (non-parametric tests used to identify differences in two arbitrary distributions) show that α distributions are

[13]For example, as the number of participants within a given experimental session increase, and are selected from a more disperse sample, it is likely that average *oneness* will decrease. This predicted decrease could be an interesting source of variation in behaviour, and using this model could address that variation.

significantly different ($p = 0.081$), while r ($p = 0.144$), ϕ ($p = 0.190$) and ψ ($p = 0.960$) distributions are not significantly different. When accounting for the matched nature of the data, the Wilcoxon signed-rank test examines the null hypothesis both distributions are the same. Results show the null cannot be rejected (at 10%) for either α, ϕ or ψ; for r the null can be rejected at 10%, but not at 5%.

Figure 2.19 shows the distribution of the s and the mean Euclidean goodness-of-fit; estimated using data from *both* and *oneness* treatments.[14] Precision, s, estimates are higher and mean Euclidean GOF measures appear lower, when estimated using *oneness* treatment data. However, Kolmogorov-Smirnov corrected test statistics and Wilcoxon signed-rank test statistics reveal that only the distribution of s is significantly different ($p < 0.000$).

Figure 2.19: Sensitivity of Treatment Selection: s and GOF

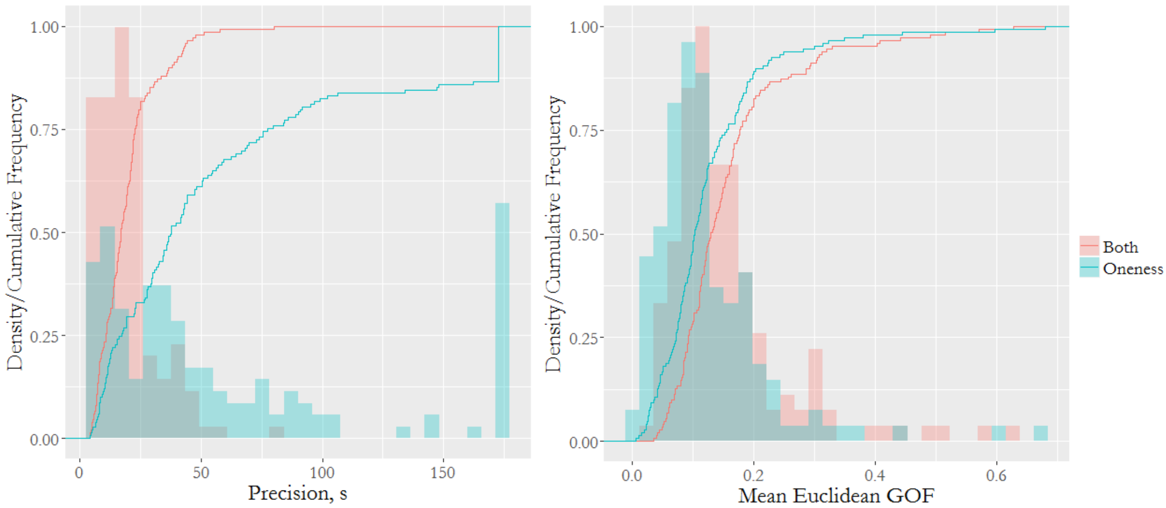

The sensitivity analysis results are reassuring. The preference parameters estimated from the *oneness* treatment alone are similar to those estimated from *both* treatments. When considering that the data are matched pairs, neither α, ϕ nor ψ distributions are significantly different. Only the unmatched distribution of α and matched r, are significantly different, and this only at the 10% level. Higher precision parameters are estimated for the *oneness* treatment, but the mean Euclidean GOF measures are not significantly different. These results imply that the estimation of preference parameters is relatively insensitive to inclusion of treatments where 'recipients' are *known* and *anonymous*. Precision parameters are somewhat higher, but the ability of the model to fit behaviour is similar in both treatments.

[14]The *mean Euclidean GOF* $= \frac{1}{T}\sum_{t=1}^{T}\left(\sum_{i=1}^{N}(x_{it} - x^*_{it})^2\right)^{1/2}$, where: T is the number of rounds in a treatment. It is a measure of goodness-of-fit which incorporates differences in optimal (x^*_{it}) and actual (x_{it}) allocations for Player 1, 2 and 3. The value is less intuitive, than that used in Section 4.3.3, but for the purposes of comparison encompasses more dimensions.

B.8 Parameter Regression Results

To ascertain if any individual characteristics, or experimental variables, are the root of preferences the following ordinary least-squares (OLS) and ordered probit (OP) regressions, in Table 2.13, were ran. The five models keep the independent variables constant, but alter the dependant variable; α, r, EDE, ϕ and ψ, respectively. From the first model those who are: students, female, older and participate in more relgious activities are less self-interested, while those who have a higher level of education and answered more questions in the tutorial correctly are more self-interested. Within Model (2) those who are female are less inequality averse, as are those who are older or have a higher intensity of household education poverty. Those who have larger household sizes, are Christian, belong to the Bagisu tribe or answered more questions correctly are more averse to inequality. Model (3) provides the EDE, which combines the preferences of α and r. Here those with a higher level of education have significantly higher levels of EDE, while those who participate in more religious activities have a lower EDE. Common to bith ϕ and ψ are that those with a larger household size have higher values, while those with a higher intensity of household education poverty have lower values. Students have a significantly lower ϕ parameter, in comparison to the general population. All other variables are insignificant, at the 10% level.

The models used above, while based upon the prior that preferences could be determined by demographic characteristics, are not necessarily optimal models. Using alternate sets of regressors could lead to differing coefficients and levels of significance. There are, however, analytical models which seek the 'optimal' set of regressors to include within the model. Below two alternate methods of achieving this are proposed, and their results shown.

The first method is the LASSO, a concept introduced by Tibshirani (1996), which "minimises the sum of squares subject to the sum of the absolute value of the coefficients being less than a constant" (Tibshirani, 1996). This constant works as a tuning parameter, which 'shrinks' the coefficients towards zero, making some parameters exactly zero. As the tuning parameter is relaxed the coefficients, of the regressors which are most correlated with the dependant, are allowed to increase. When the tuning parameter is sufficiently large the variable coefficients converge to the full OLS estimates. Choosing a tuning parameter is equivalent to choosing a subset of the 'best' regressors. A post-LASSO regression is then ran with that subset of regressors.

By running LASSO, through using adapted STATA code from Christian Hansen (Belloni, Chernozhukov, and Hansen, 2014), for a set of tuning parameters we can observe at which level each regressor will become zero. In doing this, a ranking of the most correlated regressors can be established. Table 2.14 shows the ranking of each of the 13 regressors, used in Table 2.13, for each of the four dependant variables. The four 'best' regressors are: religious participation, age, gender and correct questions, for α; age, correct questions, tribe and gender, for r; religious participation, education level, tribe and being a student, for EDE; being a student, EMPI education, EMPI health and household size, for ϕ; and EMPI education, EMPI health, gender and household size, for ψ. The resulting regressions for these 'best' four regressors are shown in Table 2.15.

Table 2.13: Parameter Regressions

	(1)	(2)	(3)	(4)	(5)
	α	**r**	**EDE**	ϕ	ψ
	Coef.	Coef.	Coef.	Coef.	Coef.
	(Std. err.)	(Std. err.)	(Std. err.)	(Std. err.)	(Std. err.)
Student Dummy	-0.112*	-0.237	-1.060	-0.561*	-0.183
	(0.060)	(0.255)	(0.658)	(0.295)	(0.244)
Gender	-0.064*	-0.362**	-0.005	-0.134	-0.227
	(0.037)	(0.183)	(0.400)	(0.186)	(0.189)
Age	-0.009***	-0.046***	-0.032	-0.010	0.005
	(0.004)	(0.014)	(0.038)	(0.017)	(0.017)
Education Level	0.036*	-0.006	0.466**	-0.013	-0.046
	(0.022)	(0.115)	(0.227)	(0.116)	(0.119)
Household Size	0.014	0.090*	0.082	0.088*	0.090*
	(0.011)	(0.047)	(0.113)	(0.046)	(0.048)
Religion Dummy	0.067	0.559*	0.154	-0.055	-0.029
	(0.054)	(0.303)	(0.577)	(0.328)	(0.377)
Religious Particpation	-0.004*	0.015	-0.057***	-0.006	-0.007
	(0.002)	(0.011)	(0.020)	(0.011)	(0.012)
Tribe Dummy	0.011	0.441**	-0.705	0.037	0.080
	(0.041)	(0.177)	(0.429)	(0.198)	(0.193)
Correct Questions	0.079*	0.445**	-0.038	-0.040	0.112
	(0.042)	(0.203)	(0.478)	(0.201)	(0.218)
Wealth Index	-0.002	-0.720	-0.482	-0.493	-0.244
	(0.133)	(0.653)	(1.774)	(0.587)	(0.610)
EMPI Education	-0.029	-1.219*	0.649	-1.081*	-1.318**
	(0.126)	(0.659)	(1.306)	(0.563)	(0.668)
EMPI Health	0.088	-0.015	-0.954	-0.624	-0.995
	(0.107)	(0.607)	(1.277)	(0.543)	(0.698)
EMPI Urban	-0.055	-0.677	0.014	0.174	0.148
	(0.122)	(0.637)	(1.227)	(0.559)	(0.613)
Constant	0.801***		3.651*		
	(0.161)		(1.904)		
N	147	147	147	147	147
R-squared	0.1378		0.0936		
Pseudo R-Squared		0.0406		0.0189	0.0216
Model	OLS	OP	OLS	OP	OP

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.14: LASSO Rankings of Regressors

Regressors	α	r	EDE	ϕ	ψ
Student Dummy	8	12	**4**	**1**	5
Gender	**3**	**4**	13	=5	**=3**
Age	**2**	**=1**	8	9	7
Education Level	7	11	**2**	13	12
Household Size	9	8	6	**=3**	**=3**
Christian Dummy	5	5	9	12	13
Religious Participation	**1**	6	**1**	=7	8
Tribe Dummy	12	**3**	**3**	11	6
Correct Questions	**4**	**=1**	12	=7	10
Wealth Index	10	9	=10	10	11
EMPI Education	13	7	5	**2**	**1**
EMPI Health	6	13	7	**=3**	**2**
EMPI Urban	11	10	=10	5=	9

Table 2.15: LASSO Parameter Regressions

	(1) α Coef. (Std. err.)	(2) **r** Coef. (Std. err.)	(3) **EDE** Coef. (Std. err.)	(4) ϕ Coef. (Std. err.)	(5) ψ Coef. (Std. err.)
Student Dummy			-0.6924	-0.4166**	
			(0.431)	(0.207)	
Gender	-0.0597*	-0.3181*			-0.2261
	(0.036)	(0.173)			(0.184)
Age	-0.0054**	-0.0355***			
	(0.003)	(0.011)			
Education Level			0.3338		
			(0.206)		
Household Size				0.0791*	0.0739*
				(0.041)	(0.040)
Religious Particpation	-0.0042**		-0.0559***		
	(0.002)		(0.017)		
Tribe Dummy		0.4320**	-0.6108		
		(0.169)	(0.408)		
Correct Questions	0.0660	0.4047**			
	(0.041)	(0.191)			
EMPI Education				-0.9183*	-1.2403**
				(0.524)	(0.629)
EMPI Health				-0.6336	-0.9521
				(0.535)	(0.656)
Constant	0.9085***		3.3187***		
	(0.098)		(0.691)		
N	147	147	149	149	149
R-squared	0.0839		0.0704		
Pseudo R-Squared		0.0230		0.0157	0.0169
Model	OLS	OP	OLS	OP	OP

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The second method is based on the Akaike Information Criterion (AIC) (Akaike, 1998). AIC is a criterion rewards models for goodness-of-fit, but penalises them for a greater number of regressors. By running regressions every possible combination of regressors (2^N), and finding that combination of regressors which minimises the AIC, the 'best quality' model (in relation to the others) can be found.[15] The 'best' models, for each of the dependant variables are shown below.

By comparing the three methods (Full, LASSO and AIC) we can check for robustness and allow a trade-off between fit, and therein bias, and predictability, or variance. The Full model, is that with the highest R^2 statistics, for each model, this greater fit does, however, come at a price. Standard error is, here, greater than in the other two models; prediction is less accurate. This trade-off between fit and predictability is a sliding scale, each criterion is different and none are the absolute 'best'.

As a result we rely upon the robustness of the results. If we focus upon those results which are significant (to the 10% level) we find that, for α: gender, age and religious participation are chosen throughout the three models. Each model shows that those who are female, older and participate in more religious activities are less self-interested. When considering r gender,

[15]Thanks goes to André Casalis, with whom this was discussed over a pint and a curry.

Table 2.16: AIC Parameter Regressions

	(1)	(2)	(3)	(4)	(5)
	α	**r**	**EDE**	ϕ	ψ
	Coef.	Coef.	Coef.	Coef.	Coef.
	(Std. err.)	(Std. err.)	(Std. err.)	(Std. err.)	(Std. err.)
Student Dummy	-0.0971*		-0.7971	-0.5052**	
	(0.050)		(0.497)	(0.208)	
Gender	-0.0644*	-0.3747**			
	(0.035)	(0.177)			
Age	-0.0095***	-0.0451***	-0.0186	-0.0056	0.0101
	(0.003)	(0.013)	(0.032)	(0.016)	(0.015)
Education Level	0.0367*		0.4260*		
	(0.021)		(0.219)		
Household Size	0.0134	0.0609		0.0764*	0.0655
	(0.009)	(0.045)		(0.041)	(0.040)
Religion Dummy		0.5381*			
		(0.298)			
Religious Particpation	-0.0047**	0.0152	-0.0551***		
	(0.002)	(0.010)	(0.017)		
Tribe Dummy		0.4832***	-0.6526		
		(0.170)	(0.409)		
Correct Questions	0.0772*	0.3905*			
	(0.041)	(0.201)			
EMPI Education		-1.0261		-0.9811*	-1.1891*
		(0.635)		(0.521)	(0.625)
EMPI Health					-1.0043
					(0.719)
Constant	0.8606***		3.5066***		
	(0.111)		(0.966)		
N	147	147	147	147	147
R-squared	0.1238		0.0789		
Pseudo R-Squared		0.0373		0.0145	0.0163
Model	OLS	OP	OLS	OP	OP

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

age, tribe and correct questions appear in each model. Those who are female and those who are older have less aversion to inequality, while those who are in the Bagisu tribe and answer more questions correctly are more averse. For EDE only religious participation is significant throughout, showing a higher level of religious participation leads to lower EDEs, which are less self-interested and more inequality averse. Indeed, this result explains that shown in Table 2.2, where it is only religious participation which significantly effects giving behaviour. For ϕ, those who are students or have a higher intensity of EMPI education have lower self-other oneness elasticities, while for those who have larger household sizes it is higher. It is only a higher intensity of EMPI education which leads to a significantly lower ψ throughout all three models.

From the results shown here a mixed picture emerges, if we were to try to explain the preferences of this sample we would perhaps not be doing such a good job. It appears that within these regressors: gender and age explain some variation in α and r, but as the direction is the same, do not explain EDE. It is only religious participation which significantly effects EDE throughout all models. For oneness elasticities, ϕ and ψ, similar regressors seems to explain variation in either, but it is only EMPI education which is significant throughout. Of interest is to note that a lot of characteristics, particularly being wealthy in assets, do not explain differences in preferences. If not these long lists of demographic characteristics, then what (if any) are the factors which formulate our preferences, and value judgements, relating to distributional concerns?

B.9 Heuristic Model

B.9.1 Theoretical Model

Alongside the utility function proposed, an alternative Heuristic function will be proposed. The structure of the model was formulated after the experiment, due to behaviour observed. It aims not only to estimate parameters which explain the decisions, but also to explain the thought process behind the decisions. From observations in the one-to-one tutorial it became clear that some participants considered their allocations to each player in turn; P1 then P2 and then P3 (indeed, due to this the difference between screen position was tested in Chapter 3). Furthermore, dependent on their mathematical ability participants seemed to heavily round their allocations. For example, for a budget of $shs45,000$, $shs20,000$ was given to P1, then out of the remaining $shs25,000$, $shs15,000$ was given to P2, then the remaining $shs10,000$ to P3. Due to observations such as these, the following *heuristic model* will be proposed try to accommodate these decisions.

We assume that participants have some ideal proportion of payoffs amongst the three players, represented by:

$$\tilde{y_1} = \frac{m^*}{p}, \qquad \tilde{y_2} = \frac{m^* - \tilde{y_1}}{q}, \qquad \tilde{y_3} = m^* - \tilde{y_1} - \tilde{y_2} \tag{2.7}$$

Or alternatively, as $\tilde{y_1} = \tilde{x_i}\pi_i, \forall i$:

$$\tilde{x_1} = \frac{m^*}{p\pi_1}, \qquad \tilde{x_2} = \frac{m^* - \tilde{x_1}\pi_1}{q\pi_2}, \qquad \tilde{x_3} = \frac{m^* - \tilde{x_1}\pi_1 - \tilde{x_2}\pi_2}{\pi_3} \tag{2.8}$$

Where $\tilde{y_i}, \tilde{x_i}$ and π_i show the respective payoffs, allocations and multipliers for Player i. Individual preferences are denoted by $p, p \geq 1$, and $q, q \geq 1$. The proportion of the total payoffs, that participants want to distribute to themselves, is represented by $\frac{1}{p}$. While $\frac{1}{q}$ represents the proportion of the remaining payoffs, after $\tilde{y_1}$ has been chosen, to give to Player 2, rather than Player 3. The budget is m, with the constraint of: $m = \tilde{x_1} + \tilde{x_2} + \tilde{x_3}$. It can be shown, due to the above optimal $\tilde{y_i}'s$ and the budget constraint, that the sum of payoffs, m^*, is :

$$m^* = \frac{m}{\frac{1}{p\pi_1} + \frac{1}{q\pi_2}\left(1 - \frac{1}{p}\right) + \frac{1}{\pi_3}\left(1 - \frac{1}{p} - \frac{1}{q} + \frac{1}{pq}\right)} \tag{2.9}$$

Now, we assume that when participants calculate their optimal allocations they make some rounding error, which is denoted by $\delta, 1 \leq \delta \leq m$. Moreover, the decisions they make are not simultaneous but consecutive. The three decision steps are shown below:

Step (1): Participants decide upon a proportion, p, of the budget, m, which they wish to allocate to themselves, Player 1. Then depending on their mathematical ability, round that allocation/payoff to a particular degree, δ.[16] This is then $\hat{x}_1$.

$$\hat{x}_1 = round\left(\tilde{x_1}, \delta\right) \tag{2.10}$$

Step (2): Once $\hat{x}_1$ has been chosen, consider the remaining budget, $m - \hat{x}_1$, and decide upon another proportion, q, to which determines their next allocation, $\hat{x}_2$. Again rounding to a certain level of precision, δ.

$$\hat{x}_2 = round\left(\tilde{x_2}, \delta\right) \tag{2.11}$$

Step (3): The remainder is given to Player 3, $\hat{x}_3$.

$$\hat{x}_3 = m - \hat{x}_1 - \hat{x}_2 \tag{2.12}$$

In order to estimate the parameters p, q and δ, a goodness-of-fit measure was calculated and the minimum value was established using a constrained multivariable solver, *patternsearch*. The goodness-of-fit measure is as follows:

$$gof = \sum_{t=1}^{k} \sum_{i=1}^{3} \left(x_{it} - \hat{x}_{it}\right)^2 \tag{2.13}$$

Where, we take the sum of the squared difference between $\hat{x}_{it}$, our estimated allocation to each player, and x_{it}, the actual allocation the participant made, for each player i. This value gof_t is calculated for each round t, then all rounds, $1 : k$, are summed to give our goodness-of-fit statistic. Intuitively, it represents the total squared difference between our estimated allocation and the actual allocation, over the three players and k rounds.

B.9.2 Heuristic Parameter Results

Within the Heuristic model, three parameters of interest are estimated: the Self Proportion, $1/p$; the Other Proportion, $1/q$; and the Precision parameter, δ. Figure 2.20 shows the distribution of each of these parameters. As $1/p \to 1$, individuals are more self-interested, distributing a greater proportion of payoffs to themselves. If $1/q = 0.5$ then individuals weight the payoffs given Player 2 and 3 equally, however, if $1/q > 0.5$ then P2 is preferred over P3, while if $1/q < 0.5$ the opposite is true. Similar to α the level of $1/p$ varies with a mode around 1/3, a median of 0.4497 and a mean of 0.5127. There are 9.4% of the sample for whom $1/p > 0.9$. For $1/q$, the distribution is concentrated around 0.5, with 63.76% of the sample between 0.45 and 0.55. The median parameter is of 0.5140, which shows a slight preference for

[16]More specifically, δ is used to create a vector $0 : \delta : m^*$, where δ denotes the increments between 0 and m^*, the lower the δ the smaller the increments and the larger the vector. $\tilde{x_1}$ is then rounded to the nearest element of this vector.

P2, over P3. Those more extreme values for $1/q$ are, however, predominantly characterised by those who are more self-interested, as seen in the scatter. Respondents here, are perhaps choosing payoffs for themselves, then quickly allocating amongst the others (indeed, results from the time-taken imply those self-interested individuals take far less time).

The Precision parameter, δ, reveals how finely participants rounded their allocations (in shs). For some individuals, 14.77%, this rounding error was very high, only to the closest $shs10,000$. The median of $shs5281$ shows that half the population were rounding to just over $shs5000$ (roughly equivalent to £1), with only 12.75% rounding to less than $shs2500$. This low level of Precision, and differing levels of $1/q$ may go someway to explaining the differences between allocation to P2 and P3.

Figure 2.20: Distribution of Individual Level Heuristic Parameter Values: p, q and δ

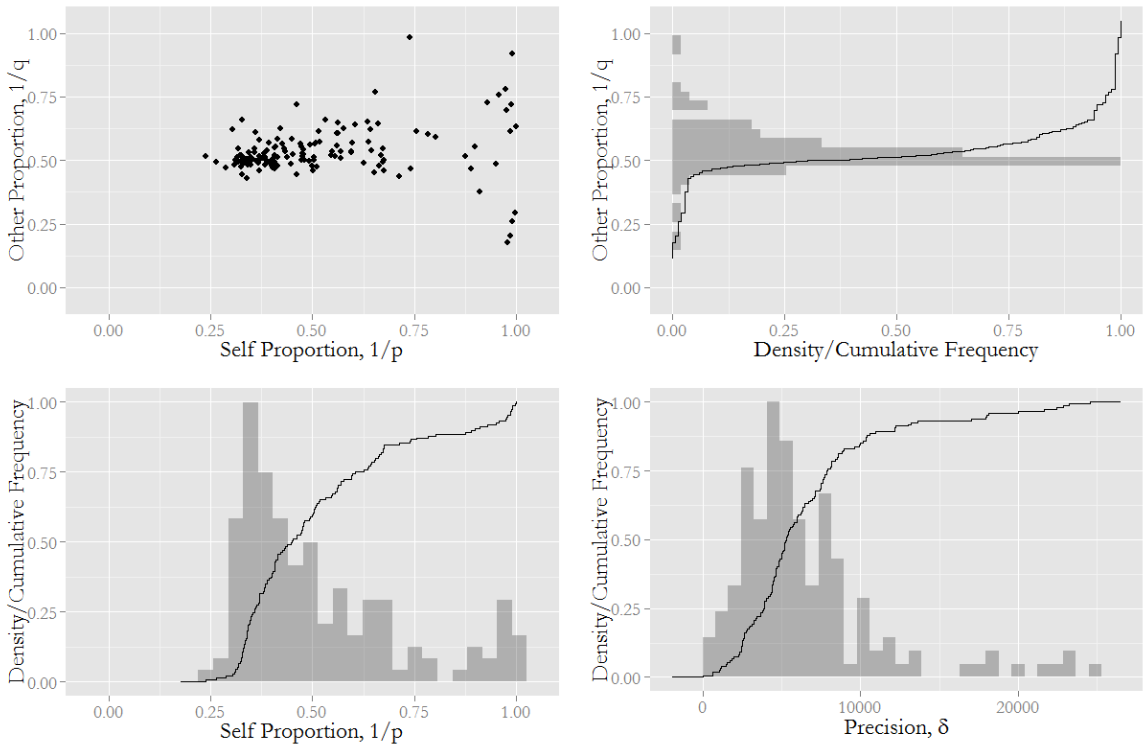

B.10 Oneness Origins and Homophily

Further to the distribution of oneness levels, in Section 2.3, a model can be constructed to identify the *origins* of the oneness levels. Table 2.17 seeks to explain the elicited oneness levels through regressors categorised as *homophily* (similarity between characteristics) and *demographics* are proposed as explanatory variables, alongside the reciprocated oneness level. Three random-effects models are ran: for General Population, Students and Total samples.

Within the Total sample, results show that oneness levels *to* the other participant are significantly and positively correlated with the oneness levels *from* the other participant (reciprocated oneness). We observe that individuals within the same age category or religion are (significantly) more connected, while those with a higher absolute difference in wealth are less connected. Belonging to the same gender, education level category, tribe or having a similar EMPI does not have a significant effect on oneness. For demographic characteristics, significantly lower oneness levels are observed for those who spend more time in religious activities, are more wealthy or work more hours. With the exception of those in the same age category and with a higher level of wealth, the significance levels and coefficient directions hold (or are magnified) for the general population. For the student population, however, there are no explanatory variables significant at the 5% level, with the exception of religious participation. While this is perhaps partially due to sample size, it highlights the importance of using a general population to draw out meaningful connections.

Table 2.17: Origins of Oneness

	(1) **General Pop.**		(2) **Students**		(3) **Total**	
	Coef.	Std. err.	Coef.	Std. err.	Coef.	Std. err.
Reciprocated Oneness	0.4812***	(0.055)	0.0554	(0.072)	0.3883***	(0.049)
Homophily						
- Same Gender	-0.0324	(0.135)	0.1009	(0.132)	-0.0492	(0.103)
- Same Age Cat.	0.2275	(0.154)	0.2295	(0.196)	0.2512**	(0.116)
- Same Education Cat.	-0.0519	(0.175)			-0.0663	(0.175)
- Same Religion	0.2995*	(0.155)	-0.1927	(0.134)	0.1937*	(0.114)
- Same Tribe	-0.0227	(0.151)	0.1652	(0.196)	0.0447	(0.121)
- Absolute Wealth Diff.	-1.2659*	(0.759)	-0.6124	(0.452)	-0.7421*	(0.384)
- Absolute EMPI Diff.	0.7634	(1.126)	0.2842	(0.529)	0.2513	(0.545)
Demographics						
- Gender	0.1470	(0.179)	-0.1227	(0.296)	0.0354	(0.155)
- Age	-0.0055	(0.015)	-0.0153	(0.026)	-0.0037	(0.013)
- Education Category	0.0451	(0.093)			0.0463	(0.088)
- Religious Participation	-0.0180**	(0.007)	-0.0180**	(0.009)	-0.0193***	(0.006)
- Wealth Index	-1.0453	(1.097)	-0.7674	(0.649)	-0.9423*	(0.568)
- Daily Work Hours	-0.0865***	(0.031)	-0.0513	(0.075)	-0.0784***	(0.028)
- Student					-0.2760	(0.258)
Constant	2.3634**	(0.937)	3.3333***	(1.092)	2.5411***	(0.649)
N	97		48		145	
Observations	448		236		684	
R-squared	0.2891		0.0711		0.1950	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

While previous results have shown that demographic characteristics do not explain distributional concerns, by incorporating variables reflecting *homophily* further explanatory power can be gained. Table 2.18 presents results from three random effects models, which regress the Proportional Payoff to Player j, on a set of regressors which include *oneness effects* and *homophily.* This data, from the oneness treatment, considers the individual allocations to both P2 and P3 for each of the 27 rounds. A dummy marking those payoffs to P3, shows there are no significant differences between the payoffs to P2 and P3. Model **(1)** incorporates the oneness effects, the first to the player to whom the payoff is being attributed, Player j, the second the oneness level to the other player, Player k. As before, the higher the oneness level the higher the payoffs, however the effects for the oneness level to Player 3 are not significant at the 10% level ($p = 0.110$). Model **(2)** incorporates the *homophily* variables, showing that those in the same age category receive higher payoffs, while those of the same gender received less. These results hold for Model **(3)**, however, just less than half of the effect from the same age category is absorbed by the inclusion of oneness levels. When considering the R^2 values between Model (1) and (2) it is Model (1) which explains most variance, even though it is the most parsimonious.

Table 2.18: Oneness and Homophily Random Effects Model: PP to Others

	(1) **PP to Pj**		(2) **PP to Pj**		(3) **PP to Pj**	
	Coef.	Std. err.	Coef.	Std. err.	Coef.	Std. err.
Player 3 Dummy	-0.0030	(0.0027)	-0.0036	(0.0028)	-0.0033	(0.0028)
Oneness Levels						
- Player j, θ_j	0.0142^{***}	(0.0019)			0.0138^{***}	(0.0018)
- Player k, θ_k	-0.0018	(0.0011)			-0.0018	(0.0011)
Homophily						
- Same Gender			-0.0088^{**}	(0.0043)	-0.0087^{**}	(0.0041)
- Same Age Cat.			0.0135^{***}	(0.0046)	0.0077^{*}	(0.0040)
- Same Education Cat.			0.0025	(0.0061)	0.0044	(0.0058)
- Same Religion			0.0006	(0.0047)	-0.0025	(0.0041)
- Same Tribe			0.0059	(0.0046)	0.0028	(0.0040)
- Absolute Wealth Diff.			-0.0087	(0.0174)	0.0105	(0.0146)
- Absolute MPI Diff.			0.0117	(0.0219)	-0.0013	(0.0198)
Constant	0.3849^{***}	(0.0701)	0.4327^{***}	(0.0707)	0.3993^{***}	(0.0711)
N	149		147		147	
Observations	7701		7490		7465	
R-squared	0.1367		0.1331		0.1425	
Controls	YES		YES		YES	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.11 Anonymous Decisions Involving Risk

When considering how individuals think about oneness when the others are anonymous a natural extension is to consider decisions under risk. In the main analysis the assumption was that when making decisions within the anonymous setting the oneness levels for P2 and P3 were both designated as the participants *average oneness level* for the group. In the experiment participants knew the five others in the room, so anonymity can be considered as a *gamble* between all possible realised positions (i.e. where the others in the room would be in the position of P2 and P3). Below is a model which incorporates these concerns.

$$U = \left(\frac{1}{K} \sum_{k=1}^{K} (U_k)^{-\gamma} \right)^{-\frac{1}{\gamma}} \tag{2.14}$$

Where:

$$U_k = \left(\sum_{i=1}^{N} \left(\omega_{ik} (\pi_i x_i)^{-r} \right) \right)^{-\frac{1}{r}} \tag{2.15}$$

$$\omega_{1k} = \frac{\alpha}{\alpha + \frac{1-\alpha}{n} \sum_{j=2}^{N} \theta_{jk}^{\phi}}, \qquad \omega_{j \neq 1,k} = \frac{\theta_{jk}^{\psi}}{\sum_{j=2}^{N} \theta_{jk}^{\psi}} (1 - \omega_{1k})$$

Intuitively the above function represents a gamble over all K possible realised Player positions of each participant in the room. With five others there are 20 possible realised player positions (i.e. $(1,2),(2,1),(1,3),...,(5,4)$ represents two desk numbers in the realised position of P2 and P3, respectively). In order to estimate preference parameters the *oneness* treatment data could be used to estimate r, α, ϕ and ψ, as the desk numbers are known. Then, using these estimated parameters, γ is estimated using the data from the *budget* treatment.

Chapter 3

Giving to Varying Numbers of Others

Within a modified N-person dictator game, we test the extent to which giving behaviour changes as the number of recipients varies. Using a *within-subject* design, in an incentivised laboratory experiment, *individual-level* preference parameters are estimated within *five* alternative utility functions. Both *goodness-of-fit* and *predictive accuracy* of each model are analysed, with the 'best' model identified for each individual. The Dirichlet distribution is proposed as a random behavioural model to rationalise *noise*; with parameters accounting for differential error arising from the *complexity* of decision problems. Results show that, on average, participants are willing to give more *total* payoffs to others as the number of players increase, but not maintain *average* payoffs to others. Extensive heterogeneity is found in individual preferences, with no model 'best' fitting all individuals.

***

1 Introduction

Often our behaviour has consequences for the happiness or misery of others in society. In those instances our behaviour is shaped not only by our preferences, but by the number of others whom we affect. This is particularly the case in the context of prosocial behaviour. In giving to a particular individual we forego potential payoffs to ourselves and others. As the number of others increase, so to does the complexity of the decisions we have to make. Not only must we consider the trade-offs we are willing to make between our self and others, but also those between-others. Are we willing to let the average amount to others decrease, in order to maintain the same amount for ourselves, or is there a minimum acceptable level we must give to all others?

This chapter seeks to model prosocial behaviour as the number of recipients of giving increases. Individual-level preference parameters are estimated within five alternate CES utility functions. Preferences accounting for *inequality aversion*, the trade-off between equality and efficiency, and *self-interest*, the weight on the self in relation to others, are central to each functional form. However, additional preference parameters are incorporated within *extended* models, to account for alternative behavioural responses to changes in N. The first models the distinction between *self-other* and *between-other* inequality aversion. The second incorporations *congestion*, the trade-off between *average* and *total* payoffs to others. The third accounts for *minimum threshold* levels of giving, denoting absolute levels of payoffs which are deemed necessary to distribute to each player.

The relative goodness-of-fit and predictive power of each model is tested, allowing for the identification of 'types' of individuals. This approach allows the flexibility to explain heterogeneity in individual behaviour, not only through preference parameters within a particular model, but between different behavioural assumptions made in alternative models. To account for *noise* in decision making, the Dirichlet Distribution is proposed as a random behavioural model. Building upon work in Chapter 2, additional error parameters are incorporated, which allow for differential error as the *complexity* of decision making increases.

To observe prosocial behaviour an incentivised laboratory experiment is run, in the form of a modified *N-person* dictator game. The *within-subject* design of the experiment varies the *number of players*, over 45 rounds of decision problems. Within each session two treatments are run; the *multiple slider* and *single slider* treatments. The former allows for complex between-other distributional decisions to be made, in addition to self-other decisions, for 2, 3 and 4 player groups. The latter simplifies the decision problem, but allows for an increased variation in the number of players: 2, 3, 4, 6 and 12. The experimental design specifically allows for the testing of both goodness-of-fit and predictive accuracy of the alternative utility functions to be compared.

Within the literature surrounding *dictator games*, the number of players within the experiment is often held constant.[1] Some papers have, however, varied the number of players

[1] Two player dictator games are frequently used to identify prosocial behaviour, for example: Forsythe et al. (1994), Hoffman et al. (1994) and Andreoni and Miller (2002). Extensions of such dictator games to incorporate multiple players (for a review see Engelmann and Strobel (2007)) include: Engelmann and

within the experiment. Charness and Rabin (2002) run a host of "simple" experimental games, within which are two and three person dictator games. Fisman, Kariv, and Markovits (2007) investigate modified two and three person dictator games, using budget sets, while Macro and Weesie (2016) use batteries of pairwise questions for two-player and four-player dictator games. The increase in the number of recipients allows for the identification of prosocial behaviour relating to *between other* trade-offs, alongside the usual *self-other* trade-offs. Panchanathan, Frankenhuis, and Silk (2013) increase the number of dictators, and find that dictators transfer less when there are more dictators, while Cason and Mui (1997) run both team and individual dictator games. Schumacher et al. (2017) motivate well an experiment where a 'decider' chooses the provision of a good between themselves and a 'receiver', where such provision comes at a cost to a group of 'payers'. They utilize a general form of the Andreoni and Miller (2002) utility function, and identify a substantial fraction of subjects which are "insensitive to group size", through a between-subject treatment design which varies the number of 'payers'.

The effect of changes in *group size* on behaviour has also been investigated in parallel literatures. Papers by Isaac and Walker (1988) and Isaac, Walker, and A. W. Williams (1994), amongst others, identify group size effects in the context of public goods games. N-person prisoner's dilemmas are studied by many, including Marwell and Schmitt (1972) and Bonacich et al. (1976). The experimental oligopolies literature identifies the effects of group size on cooperation, such as Fouraker and Siegel (1963) and Dolbear et al. (1968). The size of the group, clearly plays an important role in the decision making process. It is, therefore, an integral component within models which strive to explain behaviour.

Andreoni (2007) addresses this observation, proposing a CES utility function to explain prosocial behaviour as the number of recipients of giving increases. A modified N-person dictator game was run, where participants chose to *hold* a number of tokens (from a set budget), *passing* the remainder to a group of other players. The budget sets, prices of giving and number of other participants varied through the 24 rounds. Preferences were estimated within the utility function proposed, which incorporated a *congestion* parameter, b, signifying the extent to which the *total* or *average* payoffs to others were considered.

While the model proposed in Andreoni (2007) allows for extensive heterogeneity amongst individuals, alternative models could better explain the behaviour of particular individuals. Preference parameters can be estimated within alternative functions and the relative goodness-of-fit of each model tested at an individual level. Hey and Orme (1995) compare the goodness-of-fit of five alternative models in the context of risk. Hey and Pace (2014) conduct similar work, focusing upon ambiguity, but highlight the importance of considering both the goodness-of-fit and predictive power of each model. By comparing alternative models the 'best' model can be identified for each individual; enabling the observation of 'types' of individual based on the differing behavioural assumptions made.

Strobel (2004), E. Fehr, Naef, and Schmidt (2006) and Karni, Salmon, and Sopher (2008). Erkal, Gangadharan, and Nikiforakis (2011) and Barr et al. (2015) also utilise four-person dictator games, where entitlements are earned.

This research seeks to contribute to the above literature, by intertwining important considerations of the papers above. We propose a within-subject design which allows for the complexities of *self-other* and *between-other* trade-offs to be observed as the number of recipients increase. Individual-level preference parameters are estimated within five alternative models, allowing for both the goodness-of-fit and predictive power of the respective models to be analysed. In addition, the Dirichlet Distribution is formulated to account for *noise* in decision making, incorporating parameters which model stochastic responses to the *complexity* of decision problems.

2 Experiment

The general form of the experiment is a modified N-person 'dictator' game; where individuals are required to make distributional decisions amongst participants within a group. 'Dictators' are given a budget, m, which they must distribute amongst N players; themselves and n others, the 'recipients'. Dictators choose *allocations*, x_i, for each player in the group; where $i \in [1, ..., N]$ and $\sum_i^N x_i = m$. These *allocations* are then divided by the corresponding *divider*, $1/\pi_i$, to give the *payoff*, $\pi_i x_i$, to each Player i.[2] It is the *multipliers* (π_i) which make the dictator game 'modified', as through them the *relative prices of giving* to each player can vary; meaning that equality-efficiency trade-offs need to be made by participants.

There are two within-subject treatments, across 45 rounds of the experiment; the *multiple slider* treatment (30 rounds) and *single slider* treatment (15 rounds). In each round the participants are randomly assigned to groups, made up of N participants. Between rounds the dividers, $1/\pi_i$, change, ensuring the relative price of giving to each player varies. The budget, m, also changes, varying the average (per player) and total amounts available to distribute.

Participants made their decisions on a computerized Z-Tree interface. They were given extensive paper instructions (found in Appendix C.1), followed by an interactive on-screen tutorial to enable them to use the interface. A screenshot of the interface, from the *multiple slider* treatment, is shown in Figure 3.1. In this example, there are three players, Player 1, 2 and 3, amongst whom the 'dictator' must make allocations, so that the remaining budget reaches zero. Each player has a divider (changing every round), which is used to calculate the payoff to that player. Allocations can be made by using: the slider, arrow keys and written input. The slider (the black bar) can be dragged to make allocations, the arrow keys clicked to make incremental changes (0.01 or 0.1), and the written input used to type exact amounts. The *single slider* treatment differs in that there is only one slider, written input and set of arrow keys; that determines the allocations (and hence payoffs) to the *self*. The remaining budget is then split equally between the *recipients*. Calculations of the payoffs are made automatically, and are shown by both the orange numbers and by the height of orange bars. The payoff gap, the highest payoff minus the lowest payoff, and the total payoffs, the sum of all payoffs, are shown. All allocations, payoffs and budgets are shown in pounds and pence.

[2]The reason that *dividers* ($1/\pi_i$) are used, rather than *multipliers* (π_i) are due to visual constraints on the Z-Tree interface. Multipliers are, however, used throughout the theory, for notational ease.

Figure 3.1: Z-Tree Interface

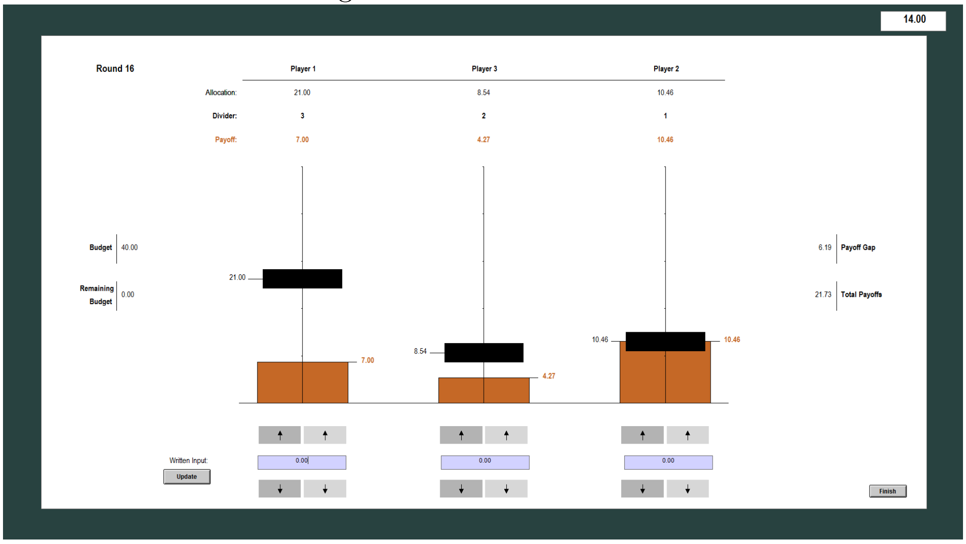

Within each of the seven experimental sessions there were twelve participants. Each participant made individual decisions; *as if* they were the 'dictator'. One individual's decisions, from each group, was randomly selected (at the end of the experiment) to determine the payoffs of each member of their group. Then, *one round* was randomly selected to determine the 'dictators', and their distributional decisions determined the payoffs of the participants in their group. In this way, each distributional decision participants made had an equal chance of determining their payoff and the recipients payoffs, and hence they were fully *incentivised*. Players within each group were randomly matched each round. Each decision made was entirely anonymous and without feedback; participants neither knew the decisions of any other participants nor the identity of the 'dictator' in any round. Removing considerations of reputation and reciprocity, allowing for the identification of 'pure' altruism.

The experiments were run in the EXEC laboratory at the University of York. Randomised invites were sent out, using hroot (Hamburg Registration and Organisation Online Tool), amongst a pool of 2,692 users. Seven experimental sessions were run between the 28th of March and the 6th of April 2017, with twelve participants in each session, to reach a sample size of 84 participants.[3] Each session required twelve participants in order to run. Due to a lack of participants one session had to be cancelled. A further 30 users were invited, as reserve participants. Nine users who showed up could not take part in the experiment, so they (and the six in the cancelled experiment) received show-up fees. The average payoff per participant was £15.45. Details of the demographic characteristics of the sample can be found in Appendix C.2.

2.1 Multiple Slider Treatment

Table 3.1 shows how the design parameters change throughout the 30 rounds of the *multiple slider* treatment. The number of players, N, changes every ten rounds. The change in the

[3]One participant had to be dropped from the analysis due to concerns of contamination.

budget, m, is shown alongside the change in the budget per player, m/N. The dividers, $1/\pi_i$, for each player (1 to 4) are shown, alongside the *relative cost* (of giving). The *relative cost*, p, shows the cost in payoffs to the 'self' of increasing the payoffs to each of the 'others', where $p = \pi_1(\sum_{j=2}^{N} 1/\pi_j)$. The *average relative cost*, p/n, shows this cost per 'recipient'.

Particular variations in the design parameters are ensured, to enable better identification of preferences parameters, in the models shown in Section 3. Between differing N, both overlap and variation is ensured in the budget variables, dividers and relative costs. This ensures that equality-efficiency, self-other and between-other trade-offs need to be made, alongside considerations of average vs total payoffs to others and minimum levels of payoffs to be given to recipients. In particular the design allows for some design parameters to remain identical between problems with differing N, while varying others. The sets (of rounds) [1,20,30], [3, 16, 23], [3, 16, 23] and [2, 22] remain constant in m and p/n. Sets [10, 13], [11, 25] and [2, 15] are constant in m and p. Both m/N and p/n remain constant in [17, 26] and [1, 11, 21], while m/N and p is constant in [6, 27].

Rounds in each set of N maintained the order shown in Table 3.1, to ensure that the grouping procedure was feasible and transparent to participants. The order of N was, however, randomised between experimental sessions, to enable the testing of order effects on decision making. Further randomisation was applied to the screen order of the players, allowing for the effect of screen position (i.e. left, middle, right) and player name (i.e. 2, 3, 4) to be tested.

2.2 Single Slider Treatment

The *single slider* treatment is a simplified version of the above, which increases the variation in the number of players, N. There are 15 rounds within the treatment, where $N \in [2, 3, 4, 6, 12]$ and each N has a set of three rounds. Rather than the more complex decision problem, where the participant must make distributive decisions separately for each player, only one slider (written output and set of arrow keys) is used to make decisions. This slider denotes the share between the self and others, where each of the others will get an equal share of the remainder of the budget, not allocated to the self. As before the budgets and dividers change each round, however, there are only the divider to the self $1/\pi_1$, and divider for each other, $1/\pi_o$, as the dividers are the same for each other player.

Both budget and divider set were randomly generated in each round, for each participant. Each set of three rounds, for each N, consisted of three sets of dividers $[1/\pi_1, 1/\pi_o]$: in the first [1,1], the second [1, A] and the third [B, 1]. The dividers A and B are uniformly and independently drawn from the set [2,3,4]. The budget, m, is similarly uniformly drawn. The set, $\ddot{M}$, from which m is drawn differs between rounds, each $\ddot{M} = \{\ddot{m} - 8N + 2N.i\,,\, i \in \{0, 1, ..., 8\}\}$. The $\ddot{m}$, within the calculation of the set differs between rounds, for each N. For the first rounds of each N, where $[1/\pi_1, 1/\pi_o] =$ [1,1], $\ddot{m} = [24, 26, 48, 72, 144]$, for $N = [2, 3, 4, 6, 12]$, respectively. In both the second and third rounds of each N, $\ddot{m} = [40, 60, 80, 120, 240]$, for $N = [2, 3, 4, 6, 12]$, respectively. This random selection of design parameters is used as the number of rounds is limited by time constraints, but through it the variation allows for aggregate analysis to be undertaken.

Table 3.1: Experiment Design Parameters

Round	Players, N	Budget		Dividers, $1/\pi_i$				Relative Cost	
		m	m/N	P1	P2	P3	P4	p	p/n
1	2	30	15	1	1	.	.	1	1
2	2	40	20	1	2	.	.	2	2
3	2	40	20	2	1	.	.	0.5	0.5
4	2	22	11	1	3	.	.	3	3
5	2	22	11	3	1	.	.	0.33	0.33
6	2	70	35	1	4	.	.	4	4
7	2	70	35	4	1	.	.	0.25	0.25
8	2	12	6	1	3	.	.	3	3
9	2	12	6	3	1	.	.	0.33	0.33
10	2	35	17.5	1	1	.	.	1	1
11	3	45	15	1	1	1	.	2	1
12	3	35	11.67	1	2	2	.	4	2
13	3	35	11.67	2	1	1	.	1	0.5
14	3	40	13.33	1	2	3	.	5	2.5
15	3	40	13.33	2	3	1	.	2	1
16	3	40	13.33	3	1	2	.	1	0.5
17	3	105	35	1	2	4	.	6	3
18	3	105	35	2	4	1	.	2.5	1.25
19	3	105	35	4	1	2	.	0.75	0.38
20	3	30	10	1	1	1		2	1
21	4	60	15	1	1	1	1	3	1
22	4	40	10	1	2	2	2	6	2
23	4	40	10	2	1	1	1	1.5	0.5
24	4	45	11.25	1	2	1	2	5	1.67
25	4	45	11.25	2	1	2	1	2	0.67
26	4	140	35	1	2	3	4	9	3
27	4	140	35	2	3	4	1	4.0	1.33
28	4	140	35	3	4	1	2	2.3	0.78
29	4	140	35	4	1	2	3	1.5	0.5
30	4	30	7.5	1	1	1	1	3	1

The two treatments are run as *within-subject* treatments to allow both goodness-of-fit and predictive accuracy to be analysed. The data from the *multiple slider* treatment is used to estimate preference parameters, allowing for *goodness-of-fit* measures to be constructed. The estimated preference parameters are then used to *predict* behaviour in the *single slider* treatment, where variation in N is increased. The combination of fit and prediction is then analysed for each individual to test between proposed utility models, which are formulated in the following section.

3 Utility Functions

Utility functions are proposed below to model behaviour in the experiment. All are within the family of constant elasticity of substitution (CES) models, incorporate preference parameters associated with prosociality and account for variation in N. Five alternative models are proposed. The first is the *standard* function, derived from Andreoni and Miller (2002), which incorporates preference parameters for inequality aversion, r, and self-interest, α. The three subsequent *extended* utility functions build upon the *standard* function by incorporating addition preference parameters, which account for alternative behaviours. The first of these is derived from Fisman, Kariv, and Markovits (2007), it distinguishes *self-other* inequality aversion, r_1, from *between-other* inequality aversion, r_0. The second accounts for *congestion*, b, and is a generalised form of the model in Andreoni (2007). The third takes the form of a Stone-Geary utility function, which originates from Geary (1950) and Stone (1954), accounting for a *minimum threshold* level, τ. In addition to these models, the *amalgamated* function incorporates each of the above preference parameters into a general model. The models are formally presented below, while graphical analysis in Appendix C.3 illustrates the intuition behind the models.

3.1 Standard

The models assume that decisions are based upon distributing allocations: x_i according to a set of preferences parameters and the multiplication factors: π_i, the reciprocals of the dividers, $1/\pi_i$. It is, then, the payoffs, $x_i\pi_i$, amongst the 'self' ($i = 1$) and 'others' ($i \neq 1$) which determine individual utility. Note that the total number of players, N, is distinct from the number of recipients, n. The standard utility function is as follows:

$$U_1 = \left(\sum_{i=1}^{N} \left(\alpha_i (\pi_i x_i)^{-r} \right) \right)^{-\frac{1}{r}} \tag{3.1}$$

Inequality Aversion is represented by r, where $-1 \leq r \leq \infty$ and $r \neq 0$. When $r = -1$ preferences reflect 'Utilitarianism', where utility is determined by summing payoffs. As r increases more weight is placed upon the payoff of the worst-off, indicating 'Weighted Prioritarianism' (Parfit, 1997), until $r = \infty$ which represents 'Maximin' preferences, where only increases to the worst-off increase utility (Rawls, 1999). Self-interest is represented by α_1, and $\alpha_j = (1 - \alpha_1)/n$, $\forall j \geq 1$ denotes the weight for each 'other', where $\forall i\ 0 \leq \alpha_i \leq 1$ and $\sum_{i=1}^{N} \alpha_i = 1$. As $\alpha_1 \to 1$, preferences reflect *egoism*, where utility is purely a function of the payoffs to the self. As α_1 decreases α_j increases, reflecting an increased regard for others. Note the α and α_1 will be used interchangeably throughout, as $\alpha = \alpha_1$. 'Cobb-Douglas' preferences are represented when $r \to 0$; which implies that optimal distributions reflect the proportions set by α_1. Intuitively, r can be thought of as the trade-off individuals are willing to make between *efficiency* and *equality*, across the entire distribution, while α_1 can be thought of as the extent to which the individual weights *themselves*, in relation to *others*.

Given the above utility function and the budget constraint $m = \sum_{i=1}^{n} x_i$, where m is the budget, the following optimal allocations (which maximise utility) can be obtained, $\forall i$:

$$x_i^* = \frac{m}{1 + \sum_{j \neq i}^{N} \left(\frac{\pi_i}{\pi_j} \left(\frac{\alpha_j \pi_j}{\alpha_i \pi_i} \right)^{\frac{1}{1+r}} \right)} \tag{3.2}$$

3.2 Extended

Building upon the *standard* model are the three *extended* models. Each model adds an additional behavioural assumption, which effects how optimal allocations change as N increases.

3.2.1 Fisman

The first model, derived from Fisman, Kariv, and Markovits (2007) henceforth *fisman*, adds assumptions regarding inequality aversion. Two parameters distinguish between *self-other* inequality aversion, r_1, and *between-other* inequality aversion, r_0. This allows for flexibility in decision making, as different equality-efficiency trade-offs can be made, depending on who the trade-off concerns. For example, an individual may prioritise efficiency between the self and others, but want to ensure equality between others. The model is below:

$$U_{F1} = \left(\alpha_1(\pi_1 x_1)^{-r_1} + \alpha_0 \sum_{i=2}^{N} \left(\alpha_i'(\pi_i x_i)^{-r_0} \right)^{r_1/r_0} \right)^{-\frac{1}{r_1}} \tag{3.3}$$

Similar to the *standard* model $-1 \leq r_1, r_0 \leq \infty$ and $r_1, r_0 \neq 0$. As $r_1, r_0 \to -1$ efficiency is prioritised, while when $r_1, r_0 \to \infty$ equality becomes paramount. Self-interest, α_1, is as before, but now α_0 shows the aggregate regard for others $(1 - \alpha_1)$. Individual *between-other* weights are given by α_i', where $\alpha_i' = \upsilon_i / \sum_{j \neq 1}^{N} \upsilon_j, \forall i > 1$; denoting the relative weight given to each other player, the expected case, which is used throughout the analysis, is that $\upsilon_i = 1, \forall i > 1$, meaning $\alpha_i' = 1/n$. When $r_1 = r_0$ and $\alpha_i' = 1/n$, or $N = 2$ the *fisman* and *standard* models are equivalent.

Optimal allocations are as follows:

$$x_1^* = \frac{m}{1 + \sum_{j \neq 1}^{N} \left(\frac{\pi_1}{\pi_j} \left(\frac{\alpha_0 \alpha_j' \pi_j}{\alpha_1 \pi_1} \right)^{\frac{1}{1+r_1}} \left(\left(\sum_{k=2}^{N} \left(\alpha_k' \left(\frac{\alpha_j' \pi_j}{\alpha_k' \pi_k} \right)^{\frac{r_0}{1+r_0}} \right) \right)^{\frac{r_1 - r_0}{r_0(1+r_1)}} \right) \right)} \tag{3.4}$$

$$x_{j \neq 1}^* = \frac{m}{1 + \frac{\pi_j}{\pi_1} \left(\frac{\alpha_1 \pi_1}{\alpha_0 \alpha_j' \pi_j} \right)^{\frac{1}{1+r_1}} \left(\sum_{k=2}^{N} \left(\alpha_k' \left(\frac{\alpha_j' \pi_j}{\alpha_k' \pi_k} \right)^{\frac{r_0}{1+r_0}} \right) \right)^{\frac{r_0 - r_1}{r_0(1+r_1)}} + \sum_{l \neq 1, j}^{N} \left(\frac{\pi_j}{\pi_l} \left(\frac{\alpha_l' \pi_l}{\alpha_j' \pi_j} \right)^{\frac{1}{1+r_0}} \right)} \tag{3.5}$$

3.2.2 Andreoni

Second is the *andreoni* model, a generalised form the model in Andreoni (2007). Here, a ‘congestion’ parameter, b, is incorporated in the model, where $b \in [0, 1]$. The ‘congestion’ parameter allows for a distinction between considering the *average* or *total* payoffs to others. As an example, if a ‘dictator’ distributes £5, out of £10, to themselves in subsequent decision problems with $N = [2, 3, 4]$ then the total payoffs to others are [£5, £5, £5] while the average payoffs are [£5, £2.5, £1.67], respectively. If a ‘dictator’ wanted to maintain the same average payoffs to others, say £2, they would have to alter the payoff to the self to be [£8, £6, £4], meaning the total payoffs to others would be increasing, as [£2, £4, £6]. The inclusion of b allows preference for trade-offs between the self and average ($b = 0$) or total ($b = 1$) payoffs to others to be incorporated, allowing for the above differential behaviour as N increases. The utility function is as follows:

$$U_{A1} = \left(\alpha_1 (\pi_1 x_1)^{-r} + \sum_{i=2}^{N} \left(\alpha_i (n^b \pi_i x_i)^{-r} \right) \right)^{-\frac{1}{r}} \tag{3.6}$$

The difference between this and the *standard* model is the inclusion of n^b, which is a multiplier of the payoffs to others. If $b = 0$, the models are equivalent, but as $b \to 1$ the two diverge as N increases. The model is identical to that of Andreoni (2007) when $\pi_i x_i = \pi_j x_j, \forall i, j > 1 \,\&\, j \neq i$, and indeed is as such for the *single slider* treatment. The optimal allocations are as follows:

$$x_1^* = \frac{m}{1 + \sum_{j \neq 1}^{N} \left(\frac{\pi_1}{\pi_j} n^{-b} \left(\frac{\alpha_j \pi_j}{\alpha_1 \pi_1} n^b \right)^{\frac{1}{1+r}} \right)} \tag{3.7}$$

$$x_{j \neq 1}^* = \frac{m}{1 + \left(\frac{\pi_j}{\pi_1} n^b \left(\frac{\alpha_1 \pi_1}{\alpha_j \pi_j} n^{-b} \right)^{\frac{1}{1+r}} \right) + \sum_{k \neq 1, j}^{N} \left(\frac{\pi_j}{\pi_k} \left(\frac{\alpha_k \pi_k}{\alpha_j \pi_j} \right)^{\frac{1}{1+r}} \right)} \tag{3.8}$$

3.2.3 Stone-Geary

Third is the *stone-geary* model, with a more general CES form to that derived in Geary (1950). The function incorporates a *minimum threshold* level, τ. This is a level below which negative (or undefined) utility would be obtained; therefore, ensuring τ is distributed to each participant is paramount. Above τ indifference curves take the form of the *standard* function, indeed if $\tau = 0$ the two are equivalent. Below is the model:

$$U_{SG1} = \left(\sum_{i=1}^{N} \left(\alpha_i (\pi_i x_i - \tau_i)^{-r} \right) \right)^{-\frac{1}{r}} \tag{3.9}$$

The model is specified with individual τ_i, but in the analysis we assume $\tau_i = \tau, \forall i$, to reduce the number of estimated parameters. The parameter τ thus signifies a minimum threshold of payoffs for all N. The inclusion of τ allows for behaviour which differs from that in the standard model, as the budget, m, and N change. The higher τ is relative to m the more equally the payoffs will be distributed. Those who have a higher level of self-interest (α_1) will take more for themselves, but only after the minimum threshold has been distributed to all players. The optimal allocations, $\forall i$, are as follows:

$$x_i^* = \frac{m + \sum_{j \neq i}^{N} \left(\left(\frac{\tau_i}{\pi_j} \left(\frac{\alpha_j \pi_j}{\alpha_i \pi_i} \right)^{\frac{1}{1+r}} \right) - \frac{\tau_j}{\pi_j} \right)}{1 + \sum_{j \neq i}^{N} \left(\frac{\pi_i}{\pi_j} \left(\frac{\alpha_j \pi_j}{\alpha_i \pi_i} \right)^{\frac{1}{1+r}} \right)} \tag{3.10}$$

Due to the form of the model, one particular issue emerges. Given that we assume individuals have a 'true' minimum threshold, say τ^*, it is foreseeable that due to budget restrictions, in a particular decision problem, there is not a sufficient budget in order to meet τ^*. In this case utility is undefined (if $|r| < 1$) as $\tau^* > x_i \pi_i, \forall i$. A natural assumption to then make is that if τ^* is greater than the minimum feasible payoff, say $x_i' \pi_i'$, then $\tau = x_i' \pi_i'$. The solution: $x_i' \pi_i' = m / \sum_i^N (1/\pi_i)$, ensures that $x_i \pi_i = x_j \pi_j, \forall i, j$.[4] In each decision problem then $\tau = min(\tau^*, m/\sum_i^N (1/\pi_i))$. It is then the 'true' minimum threshold, τ^*, which is estimated. An alternative solution, using non-negativity constraints, is in Appendix C.4.

3.3 Amalgamated

While the above separately extend the *standard* model to include r_0, b and τ, it is feasible that participant's behaviour can be explained by a combination of those factors. Here the above utility functions are *amalgamated* into a general functional form, which is as follows:

$$U_{1*} = \left(\alpha_1 (\pi_1 x_1 - \tau_1)^{-r_1} + \alpha_0 \sum_{j=2}^{N} \left(\alpha_j' \left(n^b (\pi_j x_j - \tau_j) \right)^{-r_0} \right)^{r_1/r_0} \right)^{-\frac{1}{r_1}} \tag{3.11}$$

The parameters are as explained above. With particular preference parameters the *amalgamated* model reduce to the previous functional forms. If $r_1 = r_0$, $b = 0$ and $\tau = 0$, the model is equivalent to the *standard* model. Differing combinations of these simplifications can draw out which of these considerations are important. For notation we use i to denote the individuals within the set N, j within the set n and k for those in set n excluding j. Given the above utility function and the budget constraint $m = \sum_{i=1}^{N} x_i$ the following optimal allocations can be obtained:

$$x_1^* = \frac{m + \sum_{j=2}^{N} \left(\Phi_j \frac{\tau_1}{\pi_j} - \frac{\tau_j}{\pi_j} \right)}{1 + \sum_{j=2}^{N} \left(\Phi_j \frac{\pi_1}{\pi_j} \right)} \tag{3.12}$$

[4] To relax the assumption of $\tau_i = \tau, \forall i$, giving individual τ_i, if there is only a subset of τ_i where $\tau_{j \neq i} = 0$, then the set of i is reduced to not include j, in the above solution.

$$x^*_{j\neq 1} = \frac{m + \Phi_j^{-1}\frac{\tau_j}{\pi_1} + \sum_{k\neq 1,j}^{N}\left(\frac{\tau_j}{\pi_k}\left(\frac{\alpha'_k\pi_k}{\alpha'_j\pi_j}\right)^{\frac{1}{1+r_0}}\right) - \sum_{k\neq j}^{N}\left(\frac{\tau_k}{\pi_k}\right)}{1 + \Phi_j^{-1}\frac{\pi_j}{\pi_1} + \sum_{k\neq 1,j}^{N}\left(\frac{\pi_j}{\pi_k}\left(\frac{\alpha'_k\pi_k}{\alpha'_j\pi_j}\right)^{\frac{1}{1+r_0}}\right)} \tag{3.13}$$

Where:

$$\Phi_j = \frac{1}{n^b}\left(\frac{\alpha_0 n^b \pi_j \alpha'_j}{\alpha_1\pi_1}\right)^{\frac{1}{1+r_1}}\left(\sum_{k=2}^{N}\left(\alpha'_k\left(\frac{\alpha'_j\pi_j}{\alpha'_k\pi_k}\right)^{\frac{r_0}{1+r_0}}\right)\right)^{\frac{r_1-r_0}{r_0(1+r_1)}}$$

4 Dirichlet Error Modelling

While the above utility models provide precise *optimal* allocations, x_i^*, for a particular decision problem and preference set, participants are assumed to make 'error' when calculating, or choosing, these allocations. Instead, we assume they draw their *actual* allocations, x_i, from the Dirichlet distribution (Dirichlet, 1839); where the expected values, $E[X_i]$, equal the *optimal* allocations, x_i^*.

The Dirichlet distribution is a multinomial Beta distribution, allowing for N variables, which here correspond to individual allocations (i.e. $x_1, x_2, ..., x_N$), where $x_i \in (0,1)$ and $\sum_{i=1}^{N} x_i = 1$. The below formulates the Dirichlet distribution as a random behavioural model, the work follows from Chapter 2, here altering the variance assumption to allow for varying degrees of complexity, κ.[5] The following assumptions are made: (1) $E[X_i] = x_i^*$, and (2) $Var(X_i) = \frac{(x_i^*(x_0^*-x_i^*))}{s\kappa^\gamma}$, therefore:

$$E[X_i] = \frac{a_i}{a_0} = x_i^* \tag{3.14}$$

$$Var(X_i) = \frac{(a_i(a_0 - a_i))}{(a_0^2(a_0+1))} = \frac{(x_i^*(x_0^* - x_i^*))}{s\kappa^\gamma} \tag{3.15}$$

Where:

$$a_0 = \sum_{i=1}^{N} a_i, \qquad x_0^* = \sum_{i=1}^{N} x_i^*$$

It follows that, $\forall i$:

$$x_i^*(s\kappa^\gamma - 1) = a_i \tag{3.16}$$

The a_i's determine the shape of the Dirichlet probability density function (*pdf*) and represent the weight given to a particular i. Precision is represented by s, and is multiplied by κ^γ. The higher the value of $s\kappa^\gamma$, and therefore the higher a_0, the lower the variance will be. The parameter γ allows for flexibility in the estimation procedure, to identify if *variance*

[5]This assumption is relaxed, with two alternative assumptions regarding the variance tested, in Appendix C.5.

increases or decreases as the degree of complexity, κ, increases, independently of the optimal allocations, where $\gamma \in [-1, 1]$. The degree of complexity, κ, denotes how difficult the decision problem is, by accounting for how many allocation decisions are needed to be made (minus that which is the remainder); here for the 2, 3 and 4 player multiple slider rounds $\kappa = 1, 2, 3$, respectively, while $\kappa = 1$, within the single slider treatment. When $\gamma = 0$, n has no effect on variance, independently of x_i^*, while $\gamma < 0$ implies $s\kappa^\gamma$ decreases with κ and if $\gamma > 0$, $s\kappa^\gamma$ increases.

To illustrate the above, Figure 3.2 shows the *pdf*'s of alternative Dirichlet distributions, where $N = 3$, $\gamma = 0$ and $x_3 = 1 - x_1 - x_2$. The left shows an imprecise individual, $s = 10$, who aims to allocate equally $E[X] = [0.33, 0.33, 0.33]$, with $A = [3, 3, 3]$. Second, with $A = [10, 6, 6]$, an individual allocating slightly more to themselves, $E[X] = [0.45, 0.27, 0.27]$, with a greater deal of precision, $s = 23$. Third, with $A = [4, 1, 1]$ more self-interested preferences, here $E[X] = [0.67, 0.17, 0.17]$, can be represented; with a mode where $x_1 \to 1$ (here precision is low ($s = 7$) but precision can be increased). The flexibility of the Dirichlet distribution is a useful property, and the above derivations allow for easily interpretable parameters to be estimated.

Figure 3.2: Dirichlet Distribution: Probability Density Function

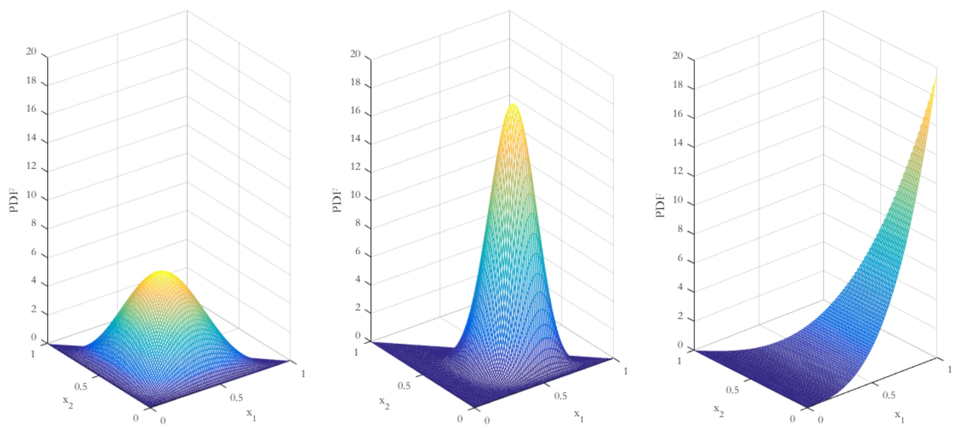

The preference parameters: α, r, s and γ (alongside r_0, b and τ in their respective models) are estimated, for each individual, through maximising the following log-likelihood function. The preference parameters determine the optimal allocations, x_{it}^*, and consequently the shape parameters, a_{it}, in each round $t \in T$. The multiple integral of the *pdf*, determined by a_{it}, is taken over the n-dimensional 'rounding' interval V_t. V_t is determined by the *observed* decisions, x_{it}; where the 'rounding' interval, around the *observed* decision, is necessary as decisions are not strictly continuous (only to the nearest pence). Estimated parameters are those which maximise the log-likelihood function, hence are the 'most likely' fit for the observed data.

$$\sum_{t=1}^{T} \log \left(\int \underset{V_t}{\cdots} \int \left(\frac{1}{B(a_{0t})} \prod_{i=1}^{N_t} \ddot{x}_{it}^{a_{it}-1} \right) d\ddot{x}_{1t} \dots d\ddot{x}_{nt} \right) \tag{3.17}$$

Where:

$$B(a_{0t}) = \frac{\prod_{i=1}^{N_t} \Gamma(a_{it})}{\Gamma\left(\sum_{i=1}^{N_t} a_{it}\right)}, \qquad \ddot{x}_{N_t} = 1 - \sum_{i=1}^{n_t} \ddot{x}_{it},$$

$$V_t = \left\{ (\ddot{x}_{1t}, \ldots, \ddot{x}_{nt}) \in \mathbf{R}^{n_t} \ : \ x_{it} - \frac{0.5}{m_t} \leq \ddot{x}_{it} \leq x_{it} + \frac{0.5}{m_t}, \forall i \in [1, n_t] \right\}$$

The multiple integral is reduced to n dimensions (hence $\ddot{x}_{N_t} = 1 - \sum_{i=1}^{n_t} \ddot{x}_{it}$) as $\sum_{i=1}^{N_t} x_{it} = 1$. This ensures the above condition is met and computational demands are lowered. A penalty function is also applied if $a_i < 0.5, \forall i$, due to the increase in computational demands when calculating triple integrals, at the bounds, when $a_i < 0.5$. In the single slider treatment the number of dimensions of the decision problem is two (hence $\kappa = 1$) for all N, and so allocations to the self (x_1) and total allocations to others (x_o) are modelled, rather than the allocation to each other (x_j). For sample parameter estimates, the log-likelihood contributions for the decisions of every individual within that sample are summed.

5 Results

5.1 Proportional Payoffs

The Proportional Payoff to Player i (PP to Pi), represents the share of payoffs given to a particular player (i). Table 3.2 shows the mean proportional payoffs to each player, given a particular group size and the type of slider used. In general, we observe a decrease in the PP to P1 (the *self*) as the group size increases. On average, participants are willing to sacrifice their own payoffs to increase the total given to the *others*. This increase to others does not, however, maintain the same average level of giving to each *other*. Payoffs between multiple and single sliders are not significantly different (10% level) for any player or group size.

Table 3.2: Average Proportional Payoffs; Players and Sliders

	Multiple Slider				Single Slider	
N Players	P1	P2	P3	P4	PS	PO
2	0.709	0.291			0.712	0.288
	(0.008)	(0.008)			(0.015)	(0.015)
3	0.624	0.188	0.188		0.626	0.187
	(0.010)	(0.006)	(0.006)		(0.018)	(0.009)
4	0.586	0.139	0.138	0.137	0.604	0.132
	(0.011)	(0.005)	(0.005)	(0.004)	(0.021)	(0.007)
6					0.562	0.088
					(0.024)	(0.005)
12					0.489	0.046
					(0.026)	(0.002)

Figure 3.3 shows the distribution of PP to P1, given the group size and slider type, at the per *round* level (n = 3,611). The top panel shows results from the *multiple* slider, while the bottom panel shows those from the *single* slider. The underlying reasons for the differences in the averages, shown above, emerge. In both panels, as N increases the PP to P1 decreases.

The shift in the average is, however, predominantly due to those who are sharing equally between themselves and others. The modal spike at 1 (an average of 23.5% and 26.4% of the sample, for multiple and single sliders, respectively) shifts little as N increases. However, the second model spike, at equal sharing ($\approx 1/N$) shifts proportionately as N increases. Indeed, the same shift is found in both slider treatments.

Figure 3.3: Distribution of PP to P1: Differing Sliders and N

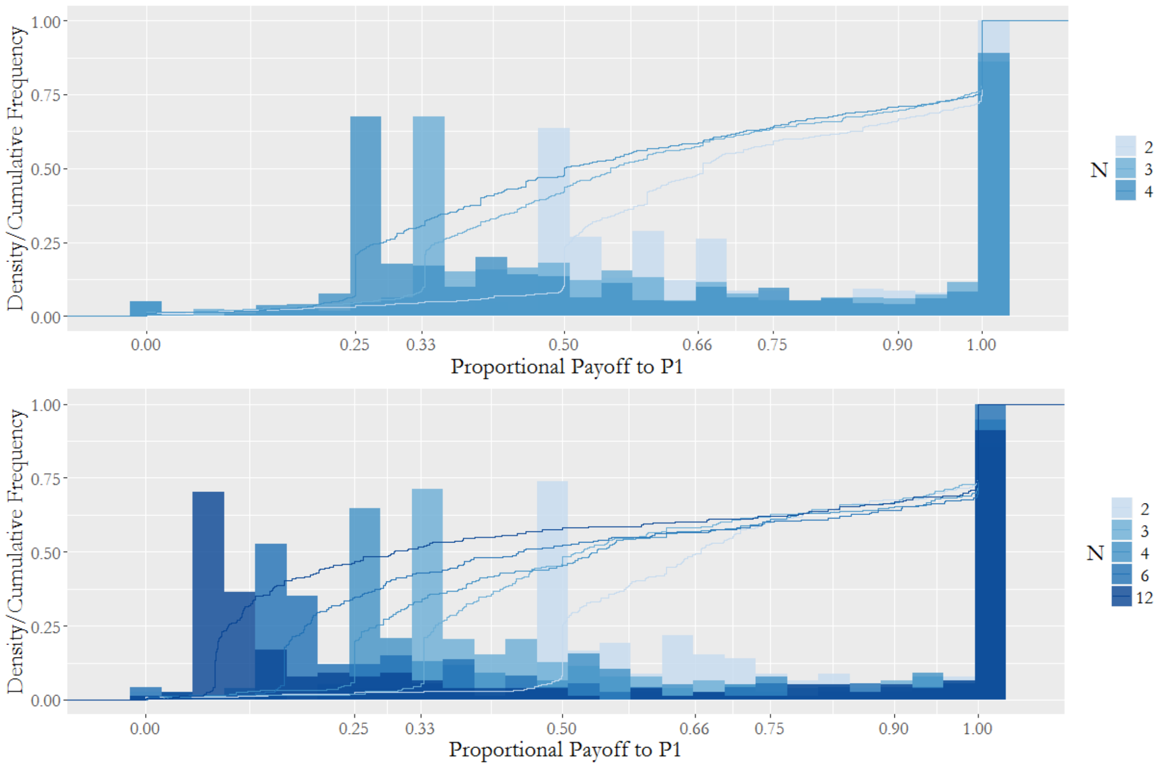

Table 3.3 shows results from two random effects models, with the Proportional Payoff to Player 1 (the *self*) and Player j (each *other*) as dependent variables. Results show that as N increases there are large and significant effects on the PP to P1 and Pj. With $N = 2$ as the reference category, we observe that an increase in the number of others leads to a reduction of the PP to P1 and a decrease in the PP to Pj. This reveals that, on average, participants are willing to significantly reduce their own payoffs, and therefore increase the total payoffs to others, but not to the extent that the average payoffs to others remain constant. We observe that switching from the multiple to single slider has no significant effect on behaviour, and neither does the round number. The relative multiplier for P1 ($\pi_1/(\sum_{j=2}^{N} \pi_j/n)$) is included as a control, and is positively correlated. An increase in the budget, standardised within each N, is shown to have a negative effect on PP to P1, and a positive effect on PP to Pj.

Neither the sign nor significance of any coefficient changes when those who on average keep more than 0.99 of the proportional payoffs to themselves are excluded from the analysis. The same is true when an extensive list of demographic characteristics (excluding parental income), 'oneness' levels and opinion questions are included as controls; with the exception of the significance levels of the budget levels, which decrease. Further analysis of the design parameters, including the player name (i.e. 2, 3, 4), screen position (i.e. left, middle, right)

Table 3.3: Random Effects Model: Proportional Payoff to P1

	(1) **PP to P1**		(2) **PP to Pj**	
	Coef.	Std. err.	Coef.	Std. err.
N Players				
- 3	-0.0717***	(0.0103)	-0.1100***	(0.0092)
- 4	-0.1037***	(0.0134)	-0.1625***	(0.0116)
- 6	-0.1493***	(0.0203)	-0.2023***	(0.0146)
- 12	-0.2137***	(0.0245)	-0.2476***	(0.0172)
Single Slider Dummy	0.0092	(0.0163)	0.0050	(0.0095)
Relative Multiplier P1	0.0560***	(0.0077)	-0.0316***	(0.0046)
Standardised N Budget	-0.0270**	(0.0119)	0.0124**	(0.0057)
Round Number	-0.0003	(0.0006)	-0.0002	(0.0004)
Constant	0.6406***	(0.0244)	0.3354***	(0.0213)
N	83		83	
Observations	3611		5985	
R^2 Within	0.2019		0.2659	
R^2 Between	0.0023		0.1044	
R^2 Overall	0.0839		0.1667	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

and the randomised order of N (all of which are found to have insignificant effects) are found in Appendix C.2, alongside analysis of the effects of demographic characteristics, 'oneness' levels and opinion questions on giving behaviour.

6 Preference Parameters

While the above analysis describes the observed aggregate-level behaviour and treatment effects, it lacks an explanation of why such behaviour is observed. Behaviour depends on preferences, and aggregate-level behaviour ultimately depends on the nature and distribution of individual preferences. Through estimating preference parameters, by assuming participants are behaving *as if* they are (*noisily*) maximising a utility function, the preferences held by participants can be characterised. This characterisation enables intuitive insights into the reasons why we observe such behaviour. Below, aggregate-level preference parameters are estimated to characterise the preferences of the representative agent. Then, individual-level preference parameters are estimated; the distribution of which accounts for the aggregate trends observed.

6.1 Aggregate Preference Parameters

At the aggregated level preference parameters can be estimated for a representative agent, within each of the utility functions proposed. The following results characterise how the sample behaves on average, but also identify how additional preference parameters affect the estimates of those in the simpler models. Table 3.4 shows the estimated preference parameters and parameters within the error model.

Results from the *standard* model shows $\alpha = 0.328$ and $r = 0.143$, showing a high regard for others and weakly weighted prioritarianism. Self-interest parameters are similar for both *fisman* and *andreoni* models, but are lower than those estimated in the *stone-geary* and *amalgamated* models. This difference is perhaps explained by the inclusion of τ, as behaviourally individuals will allocate equally until each player have more payoffs than the minimum threshold, and then distribute according to their α (and other parameters).

We observe that the estimates for r are lower than the estimates of r_1 value of 6.652, within the *fisman* model. This difference perhaps accounts for the split of r into r_1 and r_0, showing that individuals are more averse to inequality between themselves and others, but will slightly prioritise efficiency between individuals. Estimates of b, show that participants will increase the total payoffs to others as N increases, but will not maintain the same average payoff. The inconsistency within the estimated parameters lie in the estimated r_1 in the *fisman* and *amalgamated* models, where the r_1 in the latter implies efficiency prioritisation between the self and others.

Estimates of the error parameters, are relatively consistent. We observe low values of s, which are expected due to the pooled nature of the data. The elasticity of precision, γ, is positive and relatively high in each estimation. This implies that as the degree of *complexity* increases the variance of X_i will decrease.

Table 3.4: Sample Level Estimates of Parameter Values

	Preference Parameters					**Error Parameters**	
	α	r	r_0	b	τ	s	γ
Standard	0.328	0.143	.	.	.	3.547	0.600
Fisman	0.331	6.652	-0.093	.	.	3.880	0.516
Andreoni	0.346	0.685	.	0.821	.	4.802	0.547
Stone-Geary	0.492	0.793	.	.	2.451	2.927	0.958
Amalgamated	0.490	-0.594	-0.087	0.515	5.839	4.319	0.581

6.2 Individual Level

Preference parameters are estimated at the individual level, for each of the five models. For the analysis that follows there are seven individuals who are excluded, as they made purely self-interested decisions in every round. They are classed as 'egoists' who have $\alpha = 1$. The remaining 76 participants have individual-level preference parameters estimated.

Figure 3.4 shows the distribution of inequality aversion, r, and self-interest, α, estimated with the *standard* model. The top-left and bottom-right panels show histograms (and cdf plot) of r and α, respectively, while the bottom-left shows a scatter plot of the two variables. According to r the individuals classified into five different categories. There are 22.37% classified as 'Efficiency Prioritarians' ($r < -0.01$), 4.95% exhibit preferences close to 'Cobb-Douglas' ($-0.01 \geq r \leq 0.01$), 55.26% who are 'Weighted Prioritarians' ($0.01 > r < 15$) and 18.42% who are 'Maximin' ($r \geq 15$). The median value of $r = 1.08$. As r increases α becomes increasingly difficult to interpret, as a result the histogram of α shows a stacked histogram, where the lighter grey plot shows the distribution of α where $r > 10$ and the darker plot where

$r <= 10$. Of those, where $r <= 10$, 15.5% have $\alpha < 0.5$, 44.8% have $\alpha < 0.75$ and 70.7% have $\alpha < 0.9$. Note that, in addition to this, seven 'egoists' are omitted, who have $\alpha = 1$. These results show that, the majority of the sample have a substantial regard for others and are willing to sacrifice total payoffs in order to increase the payoffs of the worst-off. Yet there are significant minorities within the sample who are predominantly self-interested alongside others who prioritise efficiency.

Figure 3.4: Distribution Standard Preference Parameters

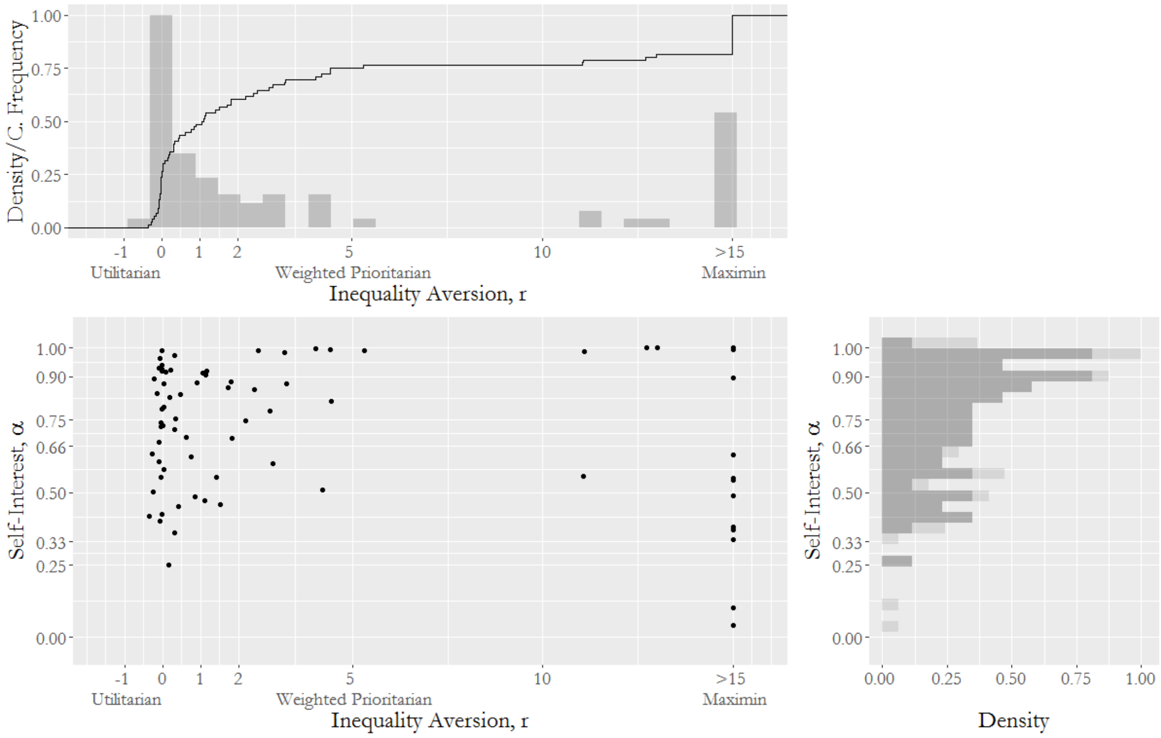

Alongside preference parameters, the error parameters s and γ are estimated for each individual. Figure 3.5 shows the distribution of the two parameters. The left panel shows s, the precision parameter. The higher is s the lower the variance X_i. There are 23.68% of the sample with $s \leq 5$, 42.11% with $s \leq 10$ and 71.05% with $s \leq 50$. Elasticity of precision, γ, identifies how the variance of X_i changes as N increases; if $\gamma < 0$, ceteris paribus, variance decreases as N increases, while if $\gamma = 0$ there is no change, and if $\gamma > 0$ there is an increase. The right panel show the distribution of γ. Only 7.89% of the sample have $\gamma < 0$, $\gamma > 0$ for 92.11%, and $\gamma > 0.99$ for 40.79%. These results imply that, for the majority, as N increases, the variance of X_i decreases. In other words, individuals draw their actual allocations, x_i, closer to the optimal allocations, x_i^*, more frequently as the number of recipients increases.

6.2.1 Between Model Parameter Comparisons

Within each of the five utility functions preference parameters for inequality aversion, r, and self-interest, α, are estimated, alongside other parameters of interest. Of interest, is the difference between the estimated parameters, as the incorporation of alternative (potentially omitted) parameters may effect the estimates. Table 3.5 shows the p-values of a one-sided Sign-

Figure 3.5: Distribution Standard Precision Parameters

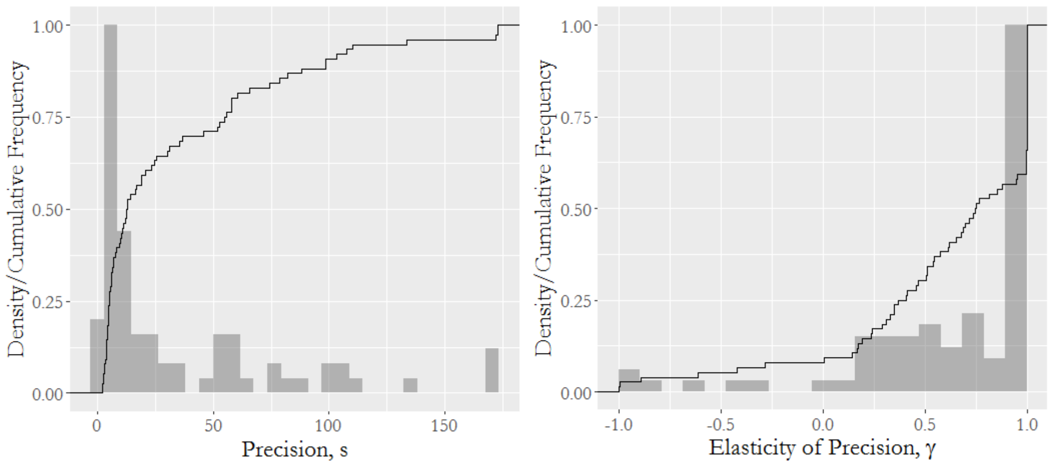

Test of Matched Pairs, between pairs of estimates from alternate utility functions. The test is used as it accounts for the matched nature of the data, and makes no assumptions about the distribution of the parameters. The null hypothesis is that the median of differences, between the parameters, is zero. The alternative hypothesis is that the median of the difference in parameters is less than zero. A low p-value, therefore, rejects the null, showing that the parameters estimated from the first model (denoted in the column) are lower than the second model (denoted in the row).

Table 3.5: Between Model Comparison of r and α: Sign-Test of Matched Pairs (p-values)

	Self-Interest, α				
	Std	Fis	And	SG	Amal
Std	1.000	0.849	0.789	0.849	0.634
Fis	0.211	1.000	0.634	0.546	0.546
And	0.283	0.454	1.000	0.546	0.546
SG	0.211	0.546	0.546	1.000	0.454
Amal	0.454	0.546	0.546	0.634	1.000

	Inequality Aversion, r				
	Std	Fis	And	SG	Amal
Std	1.000	0.634	0.717	1.000	0.986
Fis	0.454	1.000	0.634	1.000	0.998
And	0.366	0.454	1.000	1.000	1.000
SG	0.000	0.000	0.000	1.000	0.151
Amal	0.025	0.004	0.001	0.897	1.000

	Precision, s				
	Std	Fis	And	SG	Amal
Std	1.000	0.008	0.151	0.849	0.004
Fis	0.996	1.000	0.932	0.986	0.008
And	0.897	0.103	1.000	0.789	0.000
SG	0.211	0.025	0.283	1.000	0.000
Amal	0.998	0.996	1.000	1.000	1.000

	Elasticity of Precision, γ				
	Std	Fis	And	SG	Amal
Std	1.000	0.634	0.789	0.454	0.897
Fis	0.454	1.000	0.789	0.283	0.932
And	0.283	0.283	1.000	0.366	0.634
SG	0.634	0.789	0.717	1.000	0.932
Amal	0.151	0.103	0.454	0.103	1.000

Results show that there are no significant differences between estimates of either α or γ, for any pair of models. There are also no significant differences in r between the *standard*, *fisman* and *andreoni* models. We do, however, observe that the estimates of r for *stone-geary* are significantly lower than those in the *standard*, *fisman* and *andreoni* models. Likewise the estimates for the *amalgamated* model are shown to be significantly lower than those in the *standard*, *fisman* and *andreoni* models. This difference is apparent when considering the proportion of the sample classed as ‘maximin’ ($r \geq 15$), by each model; with only 9.3% and 7.9% for *stone-geary* and *amalgamated*, respectively, compared to 18.4%, 14.5% and 17.1%, for *standard*, *fisman* and *andreoni* respectively. The *stone-geary* and *amalgamated* estimates of r are not significantly different. These results imply that the inclusion of τ partially accounts for the equal-sharing behaviour, which the other models explain through a higher r.

Results from comparisons of s show that the parameters estimated from the *standard* model are higher than those in the *fisman* model, and those estimated in the *stone-geary* model are lower than those in the *fisman* model. The parameters within the *amalgamated* model are significantly higher than those in each of the other models, indicating that the additional parameters within the model allow for the flexibility for a more precise fit.

6.3 Extended

In addition to the preference parameters described above, the extended preference parameters r_0, r_1, b and τ are estimated; the distributions of which are shown in Figure 5. The left panel shows the distribution of r_0 and r_1, estimated from the *fisman* model, the middle panel shows b, from the *andreoni* model, and the right, τ, from the *stone-geary* model.

The distributions of r_0 and r_1 show potential differences in *self-other* and *between-other* equality-efficiency trade-offs. The distributions are similar, but r_0 tends to take more extreme values. 23.68% of the sample have $r_0 < -0.01$, compared to 15.79% for whom $r_1 < -0.01$. Similarly, $r_0 \geq 15$ for 17.11%, while 14.47% have $r_1 \geq 15$. The two preferences are strongly correlated, with a spearman's rank correlation coefficient of 0.786. Those who have r_0 and r_1 with the same sign make up the majority of the sample; for 53 individuals $r_0, r_1 \geq 0$, while $r_0, r_1 < 0$ for 9 individuals. There are, however, 4 for whom $r_0 \geq 0$ and $r_1 < 0$ and 10 for whom $r_0 < 0$ and $r_1 \geq 0$.

Figure 3.6: Distribution Additional Preference Parameters

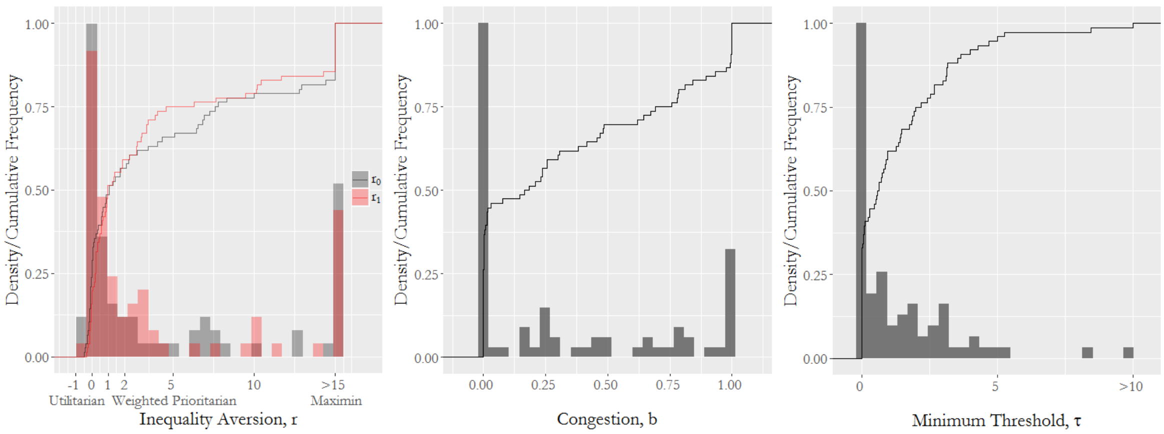

The distribution of b shows that the 'average' and 'total' payoffs to others matter for different individuals. For a large proportion of the sample (39.47%) it is 'average' payoffs which matter ($b < 0.01$), however, a significant amount (13.16%) consider the 'total' payoffs ($b > 0.99$) and do not reduce the payoffs to the self as n increases. Those who have a parameter between 0.01 and 0.99 make up the remaining 47.37%. The mean value of b is 0.336. Minimum thresholds, τ, also vary between individuals. For 73.68% $\tau > 0.01$, $\tau > 1$ for 38.16%, $\tau > 3$ for 17.11% and $\tau > 5$ for 5.26%. Showing that for the majority of the sample there is a minimum threshold which they will allocate before considering other self-other and equality-efficiency trade-offs. The median value of τ is $62p$.

The differences in r_1 and r_0, alongside parameter values of b and τ show that the behaviour of participants diverges from that predicted in the *standard* model. Indeed, by performing likelihood-ratio tests (to identify if the additional preference parameters increase the goodness-of-fit significantly, at the 10% level) individuals can be separated into either *standard* or (one of the) *extended* types. The preference parameters (estimated from the *extended* model) can then be compared. We observe a median τ of 1.881 for those the *stone-geary* model fits best, with a lower 0.127 for those for whom the *standard* model fits best. The median congestion, b, is 0.672 and 0.000 for those in the *andreoni* and *standard* models respectively. For comparing r_1 and r_0, we are interested in the difference between the parameters, so calculate the weighted euclidean distance from the estimated (r_0, r_1) to the closest point on the line where $r_0 = r_1$ as: $d = (((r_0 - r_1)/2)^2 + ((r_1 - r_0)/2)^2)^{1/2}/((|r_1| + |r_0|)/2 + 1)$. We observe a median distance, d, of 0.616 and 0.220, for those in the *fisman* and *standard* models, respectively. The preference parameters which ensure the *extended* model diverges most from the *standard*, are observed to a greater degree when the *extended* models fit individual behaviour better than the *standard* model.

6.4 Amalgamated

Of interest is not only the distribution of individual level preference parameters, but their relation with one another. Figure 3.7 shows the histograms (combined with cumulative frequency plots) of each preference parameter on the diagonal. In the bottom-left triangle, are the scatterplots of each corresponding pair of preference parameters. The top-right triangle shows Spearman's rank correlation coefficients. Loess (local regression) fitted curves are shown (with 95% ci) if coefficients (from the mirrored panel) are significant at the 5% level. The distribution of preference parameters, is somewhat similar to those from the above individual utility functions. As shown in the between model comparisons, there are no significant differences in α, but $r(_1)$ is significantly lower.

Correlations between preference parameters can also be established. First, inequality aversion parameters, r_0 and r_1, are positively correlated (0.72^{***}); participants between-other preferences appear to be closely related to their self-other preferences. Congestion, b is negatively correlated with both r_1 and r_0 (-0.429^{***} and -0.322^{***}, respectively), this implies that as the number of others increases, those who are more efficiency seeking would sacrifice their own payoffs to maintain the total to others. The minimum threshold, τ, is negatively correlated to α, -0.44^{**}; this result is partially driven by the few individuals with very high τ, as they are almost precisely equally distributing payoffs equally each round.

To summarise the estimated preference parameters, and to relate the estimations to the increased complexity that the amalgamated model incorporates, Table 3.6 tabulates individuals for whom the *extended* parameters are 'negligible' or not. The 'negligible' *extended* preference parameters are those which would collapse the amalgamated function to a more simple functional form. Those classed as 'negligible' are when $b \leq 0.01$, $\tau \leq 0.1$ and $d \leq 0.25$,

Figure 3.7: Distribution and Correlation of Amalgamated Preference Parameters

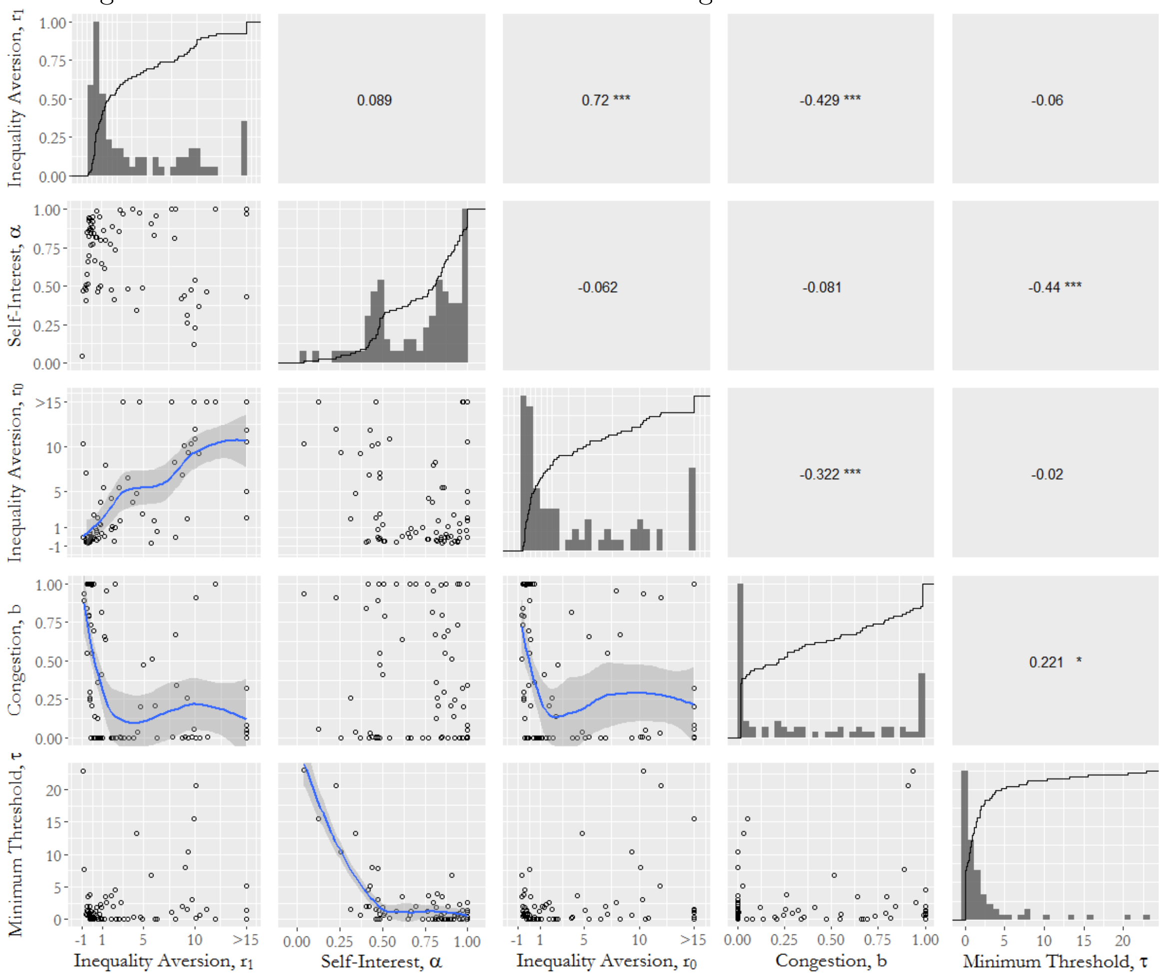

the eight possible combinations of parameters being 'negligible' or not are shown in Table 3.6. The top-left results shows that all parameters are 'negligible' for 3 individuals, while the bottom-right shows 14 individuals for whom all parameters are 'non-negligible'. Results between the two extremes show the combinations of which *extended* preferences are important. There are 36 for whom differences in r_0 and r_1 are large enough, 47 for whom $b > 0.01$ and 52 for whom $\tau > 0.1$. For those with only $d > 0.25$ and $b > 0.01$ there are 8, only $d > 0.25$ and $\tau > 0.1$ there are 6, while there are 20 with only $b > 0.01$ and $\tau > 0.1$. This heterogeneity points to models which could distinguish between either having any one, a combination of two or all three *extended* preference parameters accounted for.

Table 3.6: Summary of Amalgamated Preference Parameters

		Minimum Threshold, τ				
		$\tau \leq 0.1$ **Congestion, b**		$\tau > 0.1$ **Congestion, b**		Total
		$b \leq 0.01$	$b > 0.01$	$b \leq 0.01$	$b > 0.01$	
Inequality Aversion, d	$d \leq 0.25$	3	5	12	20	40
	$d > 0.25$	8	8	6	14	36
Total		11	13	18	34	76

Further analysis in Appendix C.6 uses a finite mixture model to identify 'clusters' of individuals. This allows for an intuitive summary of the high dimensional preference parameters, characterising groups of participants with similar preferences.

7 Goodness-of-Fit and Predictive Accuracy

Analysis can be conducted on both goodness-of-fit and predictive accuracy to determine how well the utility functions proposed explain individual behaviour. The 'best' utility model can be identified for each individual, splitting the sample into different 'types'. The alternative utility functions can be ranked, by comparing the maximised log-likelihood (MLL) values. The MLL is a measure which accounts for the stochastic nature of individual behaviour, as the measures are constructed of the likelihood of observing the actual behaviour, given the preferences estimated and error model assumed.

Due to the experimental design both *goodness-of-fit* and *predictive accuracy* can be analysed. MLL values can be calculated for *multiple slider* treatment on which the preference parameters are estimated, determining goodness-of-fit, and for the *single slider* treatment, using those estimated parameter values, to determine predictive accuracy. The ability of a model to both fit and predict behaviour is important, therefore, analysis of the two separately and as a combined measure 'Both' (a weighted average of the two) is conducted to identify if particular models are 'best' in either criteria.

An issue with comparing the 'raw' MLL is that alternative models may have a differing number of parameters. Models with a larger number of parameters are more flexible so should fit behaviour better; yet, if the difference is small the additional complexity of the model is perhaps not warranted. Several measures of *information criterion* seek to address this trade-off between fit and model complexity. Three commonly used alternatives are the Akaike information criterion (AIC) (Akaike, 1998), Bayesian information criterion (BIC) (Schwarz et al., 1978) and Hannan–Quinn information criterion (HQI) (Hannan and Quinn, 1979).[6] The three criterion may give slightly alternative rankings, due to differences in their calculation and different implicit trade-offs being made between the fit and model complexity. To sidestep such differences, the three criterion are calculated for each of the five models, for each individual, and a composite criterion, the *information criterion* (IC), is constructed whereby a model is 'best' if two or more of the criteria rank that model highest.

Table 3.7 tabulates the above. Results from the MLL are shown to the left of the IC, each split into three columns: goodness-of-fit, predictive accuracy and both. The results show the importance of comparing goodness-of-fit and predictive accuracy, as well as accounting for the trade-off between fit and complexity, as mismatches in the rankings occur. The *amalgamated* model shows this most starkly. In the MLL GOF it is the modal 'type', with 29 individuals for whom it fits 'best'. This number drops to only six and seven in predict and both, respectively. The higher number of parameters allows the flexibility to fit data well, but this comes at a

[6] The information criteria statistics are as follows: AIC $= 2k - 2(MLL)$, BIC $= \ln(n)k - 2(MLL)$ and HQI $= 2k.\ln(\ln(n)) - 2(MLL)$, where k = number of estimated parameters and n = number of observations.

cost of predictive power. Furthermore, when penalising the function for the higher number of parameters the information criteria finds there are no individuals for whom the *amalgamated* function is 'best', in either 'Predict' or 'Both'. Results are opposite for the *standard* model, there are less for whom the model is 'best' in GOF compared to 'Predict', and in MLL compared to IC. The three models with five parameters, tend to lie somewhere in between these extremes.

The results of most interest are in the final column. These rankings are those which will be used to determine the 'type' of each individual. The modal 'type' is the *standard* model and no individuals are classed within the *amalgamated* model. A substantial proportion of the sample are classed as *extended* types, with 13 *fisman* types, 21 *andreoni* type and 9 *stone-geary* types.[7]

Table 3.7: Utility Types: Ranked by Log-Likelihood and Information Criterion

	Log-Likelihood			**Information Criterion**		
	GOF	Predict	Both	GOF	Predict	Both
Standard	9	16	19	27	41	33
Fisman	14	17	17	11	11	13
Andreoni	21	24	23	26	19	21
Stone-Geary	3	13	10	9	5	9
Amalgamated	29	6	7	3	0	0

7.1 Likelihood Proportions and R-Squared

While the above analysis shows how the models do relative to one another, it reveals little about how well the model performs in absolute terms. The standard metric to analyse performance of a model is R^2, which determines how much of the sample variation in the variable of interest, is explained by the model. The *likelihood proportion*, ι, is an alternative metric, which focuses on likelihood contributions. In each decision problem, t, the *likelihood contribution*, l_t, is calculated as the area under the probability density function (given by estimated preference and error parameters) within the 'rounding' interval, around the *observed* decision (see Section 4). Intuitively, l_t, denotes the likelihood of observing the decision made, given the error model. The uniform likelihood contribution, lU_t, can likewise be derived from assuming that the probability density function takes the form of a uniform distribution. This denotes the likelihood of observing the decision made, given uniformly random draws are made. The *likelihood proportion* in each decision problem is defined as, $\iota_t = l_t/(l_t + lU_t)$. The measure shows how much 'more likely' the observed behaviour is in the specified model, in relation to the uniform distribution. If $\iota_t > 0.5$ the proposed model does 'better' at explaining behaviour than uniformly random draws, if $\iota_t = 0.5$ then the two are equal, while if $\iota_t < 0.5$ the uniform distribution 'better' explains behaviour. The summary measure $\iota = \sum_1^T(\iota_t)/T$ shows how well the proposed model explains behaviour for each individual, on average.

[7]Appendix C.7 discusses and analyses mismatches between rankings; firstly by using RSS and secondly with preferences estimated using alternative error modelling.

Figure 3.8 shows the distribution of ι, in the left panel, and R^2, in the right, across individuals in the sample. The model assumed for each individual is that based on their 'type', established in Table 3.7. The measures are calculated for each the *goodness-of-fit*, *predictive accuracy* and *both*. For both ι and R^2 the measures within the goodness-of-fit measures tend to be higher than the predicted accuracy, with both lying between. The mean values for $\iota = 0.806, 0.696$ and 0.751, for GOF, pred and both, respectively, with mean $R^2 = 0.806, 0.770$ and 0.796, respectively. There are 2, 11 and 5 individuals for whom $\iota < 0.5$ and 1, 3 and 2 for whom $R^2 < 0$, for GOF, Pred and Both, respectively. The central panel shows a scatter plot of ι and R^2, highlighting the strong correlation between the two measures (with Spearman's rank correlation coefficients of 0.90, 0.88 and 0.87 for GOF, Pred and Both, respectively).

Figure 3.8: Distribution of 'Best' Likelihood Proportions and R^2

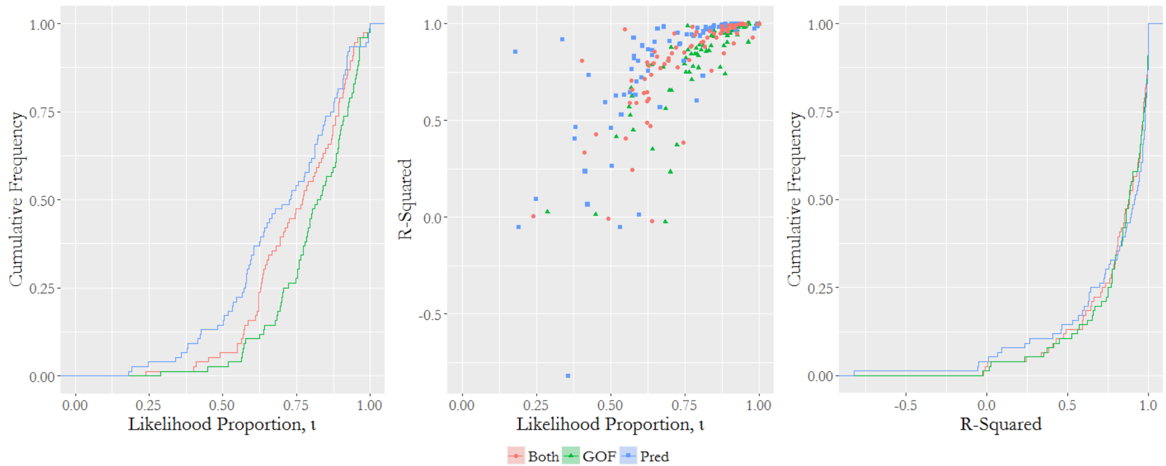

The two measures similarly aim to measure the strength of the models proposed. R^2 focuses on how *close* observed decisions are to the optimal decisions proposed by the utility function. The *likelihood proportion*, ι, however, incorporates the stochastic assumptions made, identifying how *often* the model proposed would predict the observed behaviour. Both measures show that for the majority of the sample, the models proposed and preferences estimated explain well the observed behaviour.

8 Discussion

8.1 Comparing Giving

Results from the *single slider* treatment are particularly comparable with those from Andreoni (2007). Table 3.8 shows the mean PP to P1 and PP to PO, where PO represent the average payoff to others, for differing N.[8] Results show that in the two-player game participants in our experiment are less generous, than those in Andreoni (2007). As N increases, however, while participants give a lower proportion to themselves in our results, the PP to

[8]Results from Andreoni (2007) are calculated from individual level data from: http://econweb.ucsd.edu/ jandreon/WorkingPapers/GARPN%20cesEstimates%20APX%20table.htm

P1 do not decrease, and indeed appear to have an upward trend, in Andreoni. The PP to PO follow a similar, but opposite, trend, with the average PP to others being approximately equal in our twelve-person treatment as the six-person treatment in Andreoni (2007).

While these differences are interesting, they should be approached with caution. The distributional decisions of participants is heavily dependent upon the experimental design parameters, the particular choice of dividers/multipliers, budgets and incentives will have differential effects on raw giving, depending on the preferences of participants. Indeed, this is one reason why estimating preference parameters is important; if preferences are estimated then behaviour in differing experimental designs can be predicted to identify differences not purely based on experimental design. One difference between the designs is the difference in average budgets as N increases; within our design the average budget remains the same, while in Andreoni's design it decreases. Similarly, the incentive structure leads to different behaviour. In Andreoni's set-up the participant knows they will receive the payoff they give to themselves, plus the 'Pass' payoffs from each of the n other participants in their group.

Table 3.8: Comparing Average Proportional Payoffs

	Andreoni (2007)		**Robson**	
N Players	P1	PO	P1	PO
2	0.622	0.378	0.712	0.288
3	0.710	0.145	0.626	0.187
4	0.688	0.104	0.604	0.132
5	0.695	0.076	.	.
6	0.756	0.049	0.562	0.088
10	0.727	0.030	.	.
12	.	.	0.489	0.046

Results from Fisman, Kariv, and Markovits (2007), do however, appear to be more in line with our results. Comparing results to the *multiple slider* treatment, the equivalent mean PP to P1 is 0.79 and 0.75, in the two and three person treatments, of their experiment. Comparing this to our 0.71 and 0.62, we observe that 'dictators' take less for themselves as N increases; however, both the absolute level of generosity and the change in giving are higher in our experiment.

8.2 Comparing Preference Parameters

8.2.1 Self-Other and Between-Other Inequality Aversion

While Fisman, Kariv, and Markovits (2007) (FKM) run both two and three-person dictator games, preference parameters are estimated separately for each treatment. We compare classifications of r_1 and r_0 with those estimated in their three-person treatment, and the r estimated in their two-person treatment. To make estimates comparable, we use their classifications, and exclude those 'selfish' individuals with an average PP to P1 greater than 0.95 or who are not 'consistent'.[9] Our total sample of participants with "consistent nonselfish

[9] In their paper they calculate Afriat's Critical Cost Efficiency Index (CCEI) and exclude those individuals with $CCEI < 0.8$, as they behave in a manner 'inconsistent' with utility maximisation. We do not calculate

preferences" is 63, with 33 from the three-person and 45 from the two-person treatments in FKM; the percentages shown below in Table 3.9 refer to these totals.

Results in Table 3.9 show the categorisation of inequality aversion parameters in FKM and this study. We observe that for FKM the majority of the sample are either 'utilitarian' or 'efficiency prioritarians' for both r_1 (66.7%) and r_0 (66.7%), although there is a lower proportion within this categorisation in the two-player experiment for r (53.3%). The opposite is true from our results, with the majority of the sample being either 'weighted prioritarians' or 'maximin', for r_1 (85.7%) and r_0 (74.6%). This reversal shows a much higher weight on *efficiency* concerns for the FKM sample, in contrast to a higher concern for *equality* in our sample.

Table 3.9: Comparison of Inequality Aversion, r_1 and r_0

	FKM			**Robson**	
	r (2P)	r_1 (3P)	r_0 (3P)	r_1	r_0
Utilitarian ($r < -0.9$)	4.4%	12.1%	15.2%	0.0%	0.0%
Efficiency Prioritarian ($-0.9 < r < -0.1$)	48.9%	54.6%	51.5%	6.4%	9.5%
Cobb-Douglas ($-0.1 < r < 0.1$)	11.1%	3.0%	9.2%	7.9%	15.9%
Weighted Prioritarian ($0.1 < r < 0.9$)	31.1%	21.2%	6.1%	22.2%	14.3%
Maximin ($r > 0.9$)	4.4%	9.0%	18.2%	63.5%	60.3%

There could be several reasons for these differences. The first, is the sample. Participants in the UK are perhaps more averse to inequality than their US counterparts. The second, the differences in experimental design alter individual behaviour. In our design participants had to individually allocate to each individual, with a slider, while in their design a single point on a budget line was clicked. The latter allows for quicker and easier decisions to be made, while the former requires more effort. In itself, this could lead to different responses; on the one hand the former method could lead to more 'considered' distributions, accounting for each of the other participants, on the other the ease of clicking a single point could allow for more time to consider the efficiency implications of the choices made. This, however, should then appear in the distributional decisions between the *single* and *multiple* slider treatment, which it does not.

A further difference in design, is the incentive structure. In our design one 'dictators' choice is picked at random to determine the payoffs of all in the group, while in FKM each participant receives the payoffs they gave to themselves, plus the payoffs others gave them. This may have an impact on average giving (as discussed in Chapter 2), but also on trade-offs between equality and efficiency. On average, participants know that if everyone distributes efficiently then payoffs will be greater, but in FKM this carries a much lower risk of particular individuals receiving a low payoff. Other difference include: the explicit statement of the 'Payoff Gap' and 'Total Payoffs' (representing the trade-off between equality and efficiency) in our design; the explicit statement of the 'Dividers' opposed to the difference in graphical

CCEI values, but instead use the likelihood-proportion value, ι to exclude those with $\iota < 0.5$; which (while it is a test dependent upon the utility function chosen) excludes individuals for whom random behaviour better explains their behaviour.

representation; and the difference in language between 'allocations' to each player (implying the budget is a common good) compared to 'hold' and 'pass' (implying the budget belongs to the 'dictator', which they can choose to share).

Consistent between our findings is that there are strong within-subject correlation between r_1 and r_0. With FKM there was 63.6% of the sample with $r_0, r_1 \leq 0$, while $r_0, r_1 < 0$ for 24.2%. With only 6.1% with $r_0 \geq 0$ and $r_1 < 0$ and 6.1% with $r_0 < 0$ and $r_1 \geq 0$. In our (similarly reduced) sample there are 81.0% of the sample with $r_0, r_1 \geq 0$, while $r_0, r_1 < 0$ for 4.8%. With 4.8% with $r_0 \geq 0$ and $r_1 < 0$ and 9.5% with $r_0 < 0$ and $r_1 \geq 0$. This means that there are 87.9% and 85.7% of the sample, for FKM and our study respectively, with both *self-other* and *between-other* inequality aversion in the same direction.

8.2.2 Congestion

Andreoni (2007) estimates the congestion parameter, b, at both the sample and individual level. At the sample level, the representative b estimated was 0.68, which is slightly lower than our estimated value of 0.82, but not extensively so. At the individual level Andreoni (2007) estimates preferences for 109 participants, with 11 participants identified as 'perfectly selfish'. Of those 109 participants $b = 0$ was estimated for 25%, while $b = 1$ for 17% and $0 < b < 1$ for the remaining 58%. From our estimates, there are 39.5% of the sample for whom $b < 0.01$, 13.2% with $b > 0.99$ and the remaining 47.4% with $0.01 \leq b \leq 0.99$. The results are somewhat similar, spikes at either extreme, where the modal group has $b \to 0$; but the majority exhibits some degree of congestion.

8.2.3 Minimum Threshold Levels

Comparison with the τ preference parameter within the Stone-Geary function is limited. Its use is more common in other literatures, such as the time and risk preferences. Andreoni and Sprenger (2012) estimate Stone-Geary "consumption minima" (ω_1), within a CRRA utility function with quasi-hyperbolic discounting. While contextually different, the experimental set-up is somewhat similar, with convex time budgets. Their aggregate estimate of $\omega_1 = \$1.35$, when $\omega_1 = \omega_2$ is assumed (the hypothesis of which is not rejected), which lies somewhere between our median individual estimate of 62p and aggregate estimate of £2.45. Of interest, however, is that they find the estimates of other preference parameters (especially curvature) depend on the assumed ω_i, a result we also find (with significantly lower estimates of r in the *stone-geary* in comparison to the *standard* model, in Table 3.5). Andersen et al. (2008) also use a similar functional form, but do not estimate a minimum threshold, instead utilising the average value of daily consumption in Denmark as the threshold.

8.3 Charity Fundraising

While the main focus of this paper is somewhat technical and abstract, the methods used can readily be applied to the domain of charity fundraising. This section provides an illustration of how the estimation of preferences could increase charitable giving, if projects rather than people are assumed to be the *others*.

Imagine a charity. Within the charity there are four *projects*: Water, Education, Shelter and Medication. The aim of the charity is to raise money to enable the projects to be funded. In order to do so, there are alternative fundraising *campaigns* which can be undertaken, which encourage people to donate. Each campaign advertises alternate bundles of the projects. There are sixteen possible campaigns which can be delivered to potential donors:

1. Water
2. Education
3. Shelter
4. Medication
5. Water, Educ
6. Water, Shelter
7. Water, Medication
8. Education, Shelter
9. Education, Medication
10. Shelter, Medication
11. Water, Education, Shelter
12. Water, Education, Medication
13. Water, Shelter, Medication
14. Education, Shelter, Medication
15. All Projects
16. Do Nothing

In order to advertise a campaign there is a cost of £3 per person, with the exception of (16) in which no campaign is run. Within each campaign information will be provided about the respective projects. Each project has differential fixed costs which determine the 'cost-effectiveness' of that project; stated as "for every pound given the amount of money going directly to that project is X", where X is 50p for Water, 33.3p for Education, 25p for Shelter and 25p for Medication. The charity's task is then to deliver the campaigns which raise the most amount of money. The following analysis addressed this problem.

Donors are assumed to have a budget, m, which they can distribute between consumption (which is entirely cost-effective) and donations to particular charity projects. Using estimated preferences parameters for each of the 83 participants of the experiment (according to each individual's 'type') predictions of how they would optimally allocate between themselves (the 'donors') and each project, within a given campaign, can be made. The advertising costs of the campaigns can be deducted and the average profit per person calculated.

The left panel of Figure 3.9 shows the calculation of average profit, for varying levels of m, for three alternative methods of choosing the fundraising campaign. The *random* method denotes the profits that would be made if the charity had no information about the preferences of the donors. Here, as there is no information, charities would randomly choose a campaign to send. The *sample* method uses the representative agent preferences (from the *standard model* in Table 3.4) to identify the optimal campaign to advertise, for each m. The *individual* method uses individual-level preference parameters to establish the optimal campaign to advertise to each individual. The results show that the profits from the *individual* method dominates the *sample* method, which in turn dominates to the *random* method. For low values of m the *random*, and even *sample*, method give negative profits, as the low budget means that the advertising costs are not exceeded by the donations. As m increases the *random* method diverges from the *sample* and *individual*, making relatively lower profits.

The right panel shows the proportion of each campaign advertised, under the *individual* method, for differing m. The four campaigns selected are: (1) Water; (10) Shelter, Medication;

Figure 3.9: Fundraising Profit and Distribution of Campaigns

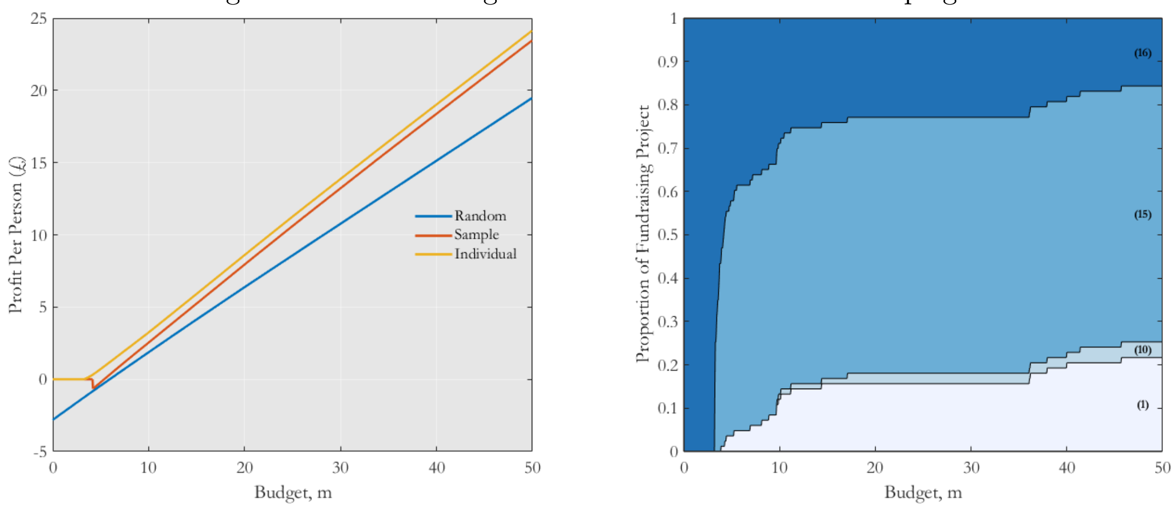

(15) All Projects; and (16) Do Nothing. At very low m the optimal campaign is to do nothing, as the donations do not exceed advertising costs. As m increases campaigns (1), (10) and (15) are sent to particular individuals, when $m = 50$, the campaign with the largest proportion is (15), with 59.04%, next is (1) with 21.69%, followed by (16) with 15.66% with 3.61% being selected for (10). The reasons for the differences lie in individuals preferences. Those for whom (16) is optimal tend to be self-interested, the mean self-interest parameter of the group is 0.976. Individuals who donate most in (1) are all efficiency prioritarians, with a mean inequality aversion parameter of -0.112. The three individuals within (10) are all classed as *andreoni* types, who are (slightly) weighted prioritarians with high values of b (0.965 on average), meaning they consider *total* rather than *average* payoffs to others. Those within (15) tend to be weighted prioritarians, with substantial regard for others; 26 of whom are *standard* types with positive r, 8 are *fisman* types with positive r_1, 7 are *andreoni* types with low b, and 7 who are *stone-geary* types with a positive τ. Being able to account for individual preferences allows for selection, which in turn allows for an increase in profit per person.

While this section is primarily illustrative, there are a number of extensions which could be conducted to make it more applicable and realistic. The first relates to error. In the analysis above the assumption is that donors act optimally and according to the preferences estimated, however, there could be error in those predictions. By incorporating the error model proposed, a monte-carlo simulation could be run, to establish the optimal campaigns to run, given the error made. The second extension relates to the fundraising aim of the charity. Two types of funding are commonly found in charitable giving, *restricted* and *unrestricted* funding. The above assumes that donations are unrestricted, meaning that the charity can allocate resources to any project they need. However, (especially with a move towards the tracking and accountability of individual donations) donors may give restricted funding, meaning that only those projects they give directly to can be allocated that funding. These considerations can be incorporated into the analysis, selecting the optimal set of campaigns to increase the funding of the 'worst-off' charity, rather than maximising the total profits (equality vs efficiency criteria). Finally, the incorporation of the value of acquiring information is important. While it is clear that the *individual* method performs the best it may be more costly to acquire information on individual level preferences. Collecting information at an aggregate level (perhaps one decision problem, rather than 30) could prove to be less costly, but if this information cost exceeds the gains made above the *random* method then it is counter productive to gather such information. By accounting for the value of information the choice of method optimised at different budget levels.

While experiments run in the laboratory may appear abstract, external parallels do emerge. By utilising the methods proposed and accounting for individual preferences real world charitable giving could perhaps be increased.

9 Conclusion

To conclude, through running a modified N-person dictator game both *between-other* and *self-other* distributional trade-offs have been investigated as the number of players increases. Results have found that, on average, the proportional payoffs given to the self decrease, as the number of others increases, but not to the extent that the proportion of payoffs to each others remains constant. The majority of the sample are shown to have other-regarding preferences (91.6%), where the majority are classed as 'Weighted Prioritarians' (55.3%), with significant proportion classed as either 'Efficiency Prioritarians' (22.4%) or 'Maximin' (18.42%).

The importance of estimating preferences within alternative utility functions has been shown, with intuitive *extended* preference parameters of: *self-other* and *between-other* inequality aversion, *congestion* and *minimum thresholds*; better explaining the behaviour of particular individuals. The importance of incorporating both goodness-of-fit and predictive accuracy has been shown, alongside considerations of 'information criteria'. The *amalgamated* model (the most complex) provided the 'best' fit for the modal group of participants; however, when accounting for predictive accuracy and 'information criteria' it performed 'best' for no individuals in the sample. Splitting the sample into 'types', of the 83 participants, we observe 33 individual's behaviour is best explained by the *standard* model, 13 by the *fisman*, 21 by *andreoni* and 9 by *stone-geary*; with 11 individuals being classed as *egoists*. Values from the *likelihood proportion* reveal the 'best' utility functions, combined with the Dirichlet error model, well fit and predict individual-level behaviour, with only 5 participants with $\iota < 0.5$.

Prosocial behaviour and distributional preferences have been shown to be extremely heterogeneous. Not only do particular preferences within utility functions best explain their certain individual's behaviour, but alternative models best suit different individuals. Varying the number of players to whom participants can give to may complicate modelling decision making, but it is something we regularly do as humans and is, therefore, something worthy of striving to explain.

C Appendices

C.1 Instructions

Instructions

Welcome. Thank you for coming today.

Please **Read These Carefully**.

Everyone Will Receive the **Same** Instructions.

General Instructions

In this experiment you will be making **decisions** about the **distribution of payoffs** between **yourself** and **other participants** in this room. These **payoffs** are in addition to your **turn-up fee of £3**.

There will be **two** different **stages**, each made up of multiple **rounds**. Your **actual payoff** will be determined from **one randomly selected round**. It is from this **one round** that **all participants** will receive their **payoff**. This means that **every round** has an **equal chance** of determining your **final payoff**, so consider each choice you make **carefully**. **Everyone** will finish at the **same time**, as you will need to wait for **every participant** to finish **each round** before you can move onto the next.

The **individual choices** you make will involve payoffs for **multiple players**. You will make choices which concern the **distribution of payoffs** between those players. In **each round**, you will be **randomly** grouped with some **other** players. Each of you will make your choices **independently** of one another, but only **one** of the player's **individual choices** will be selected, randomly, to provide the payoffs for **all players within that group**. You **will not** know whose decision has been chosen and will receive your payoff, in private, at the end. This means that **every decision** you (and the others) make is **entirely anonymous**. There are **no right or wrong answers**, the decisions you make are entirely up to you and will determine the potential payoffs for you and the others in that group.

The money you have to **allocate** amongst the group will come from a **Budget**. You must decide how to allocate **all** of the money from the **Budget**. For each player in your group, the **Allocation** that you make to them will then be divided by a **Divider** to give their **Payoff** in that round. The **Payoff** can be thought of as the final amount of money that each player gets in that round.

To make your decisions you will be using a **computer interface**, a screenshot of one of the rounds is shown on the next page. Importantly, **you will always be Player 1**, and the other players are **real other participants** in the room.

The **on-screen order** of each player **will vary**. In the example shown here, Player 1 is in the middle.

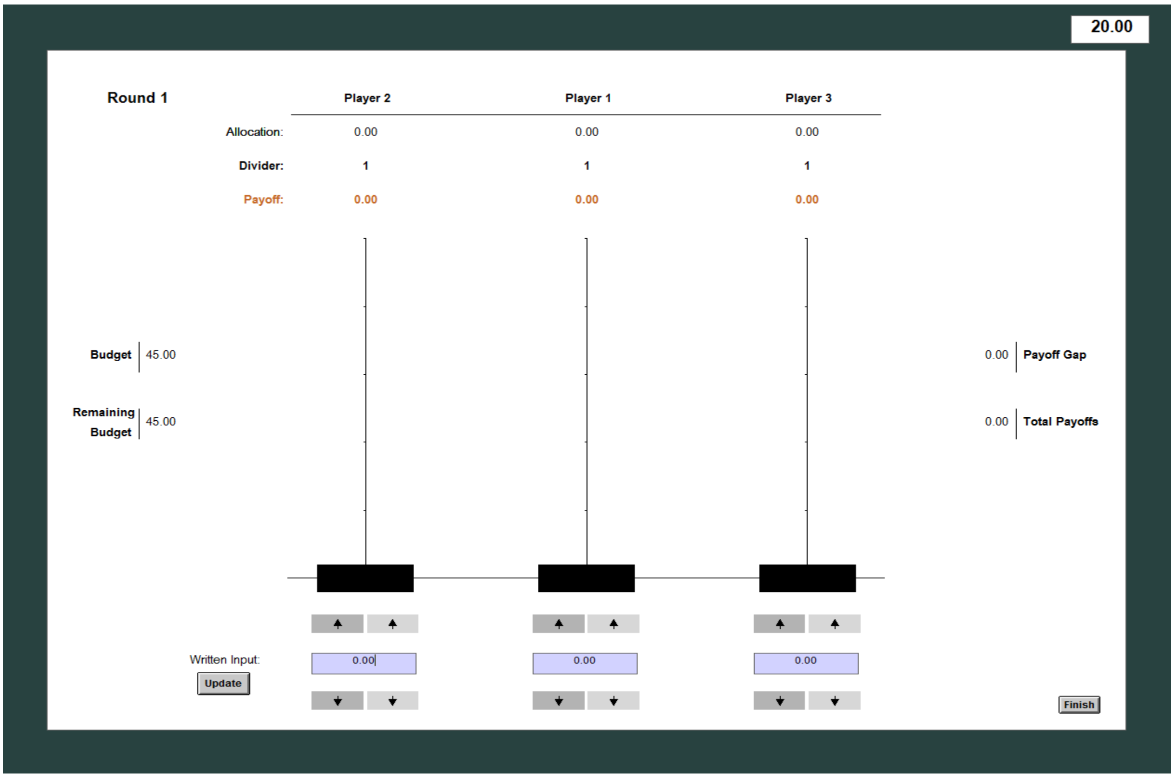

You will be given a **Budget** to **Allocate** amongst the group. This is shown on the left of the screen, and in this example is £50. You **must** spend the **entire budget** in each round. This means the **Remaining Budget must be zero**. You will be able to make the allocations in three ways. The first is with the **sliders**; you can **drag the sliders** to any allocation that you want. The second is with the **arrow keys**. They allow you to make increases and decreases of 10p and 1p, respectively. The third is the **written input**; you can click in each of the blue boxes, **type** your desired allocations and **click update**.

Within each of the rounds there will be **different Dividers** for each player. The actual **Payoff** that each player will get will be the **Allocation** you give, divided by the **Divider**. For example, if you give an Allocation of £10, and the Divider is 2, the Payoff will be £5. These **Dividers** are important as they **change every round**, but are **predetermined** and **not dependent** upon your choices.

The **Payoffs** are the final **amount of money** which **will** be given to **each participant**; they will be always be in **pounds** and will be shown by the **height of the orange bars**, the **orange numbers** beside them and the numbers at the **top** of the screen.

Throughout the **rounds**, two elements will change. The first is the **Budget**, so be sure to consider exactly how much the Budget is **before** beginning each decision, as it will vary by a **considerable amount**. The second is the **number of players** in your group. This will change as you go through the experiment. You will see **2, 3, 4, 6** and **12 players** in the groups, throughout various rounds. So remember that **each** of these players is a **real participant** in the room, who will be **anonymous** and **randomly** chosen for each round.

There are **two stages** in this experiment. The first is where you will have **multiple sliders**, one for each Player. The second is where you will have a **single slider** which determines the share of the Budget you choose to give to Player 1, where the other players allocations are equalised.

Remember you are always Player 1. Take note especially in the first stage, as the **order of the players on-screen changes between rounds**.

You will also see the **Total Payoffs** and the **Payoff Gap**. The **Total Payoffs** is the Payoffs of all players added together. The Payoff Gap is the highest Payoff minus the lowest Payoff. Notice how these change when making your decisions. You **must** make a decision in **every** round, and then click **Next** or **Finish** to confirm your decision.

A **minimum time** will be displayed in the **top right corner** in every round, in black. This time **must have elapsed** before you progress to the next round. There will also be a **maximum time**, in red, which will be **double the minimum time**, you **must make** a decision in this time and **click Finish**. **If not**, will receive a **Payoff of zero** for that round and one of the other participants in your group will be the individual whose decisions will count for that round.

After the experiment you will be asked to fill out a **questionnaire**. Your responses from the **questionnaire**, and from the **entire experiment**, will be treated **anonymously**.

After reading these instructions you will go through an **on-screen tutorial**, which will explain how to use the **computer interface** and the exact nature of the experiment. You will then be allowed several **practice rounds** (which **will not** affect your payoff) before making your decisions for **real**.

If you require help at any time, **please raise your hand**.

Please proceed to the On-Screen Tutorial.

C.2 Design and Demographic Differences in Proportional Payoffs

C.2.1 Design

By running separate random effects models for each N, and focusing upon the *multiple* slider treatment, further analysis can be conducted on more specific design effects. Table 3.10 models PP to Pj, for $N = 2, 3$ and 4, incorporating player specific multipliers, time taken, screen position, the player name and the order of N, alongside the standardised budget. The index j denotes a particular 'other' player, where $j \neq 1, \in N$. Considering the multipliers (π_i), k and l are the 'alternative others', where $k, l \in N$, k is the lowest number that satisfies $k \neq 1, j$, and $l \neq 1, j, k$.

Table 3.10: Random Effects Model: Proportional Payoff to Pj, Design Effects

	(1) 2 Players		(2) 3 Players		(3) 4 Players	
	Coef.	Std. err.	Coef.	Std. err.	Coef.	Std. err.
Multiplier						
- Player 1	-0.1569***	(0.0354)	-0.0908***	(0.0180)	-0.0557***	(0.0123)
- Player j	0.0631***	(0.0198)	0.1140***	(0.0201)	0.1004***	(0.0193)
- Player k			-0.0526***	(0.0184)	-0.0368***	(0.0124)
- Player l					-0.0342***	(0.0106)
Time Finished						
- Average	0.3757***	(0.1299)	0.3273***	(0.0847)	0.1927***	(0.0692)
- Mean Diff: Positive	0.1393***	(0.0454)	0.0238	(0.0226)	0.0315	(0.0251)
- Mean Diff: Negative	-0.0729	(0.0814)	-0.0332	(0.0324)	-0.0144	(0.0226)
Screen Position						
- 2	-0.0000	(0.0124)	0.0083	(0.0075)	0.0024	(0.0063)
- 3			-0.0038	(0.0059)	0.0027	(0.0042)
- 4					-0.0014	(0.0056)
Player Name						
- Player 3			0.0015	(0.0037)	-0.0003	(0.0027)
- Player 4					-0.0023	(0.0023)
N Order						
- Second	0.0598	(0.0419)	-0.0404	(0.0306)	-0.0214	(0.0253)
- Third	-0.0224	(0.0478)	-0.0156	(0.0314)	0.0035	(0.0221)
Standardised N Budget	0.0041	(0.0126)	0.0016	(0.0069)	-0.0022	(0.0073)
Constant	0.2288***	(0.0605)	0.1312***	(0.0383)	0.1056***	(0.0301)
N	83		83		83	
Observations	818		1576		2379	
R-squared	0.1452		0.1496		0.1228	
Between-Subject Variance	0.1662		0.1057		0.0840	
Within-Subject Variance	0.1588		0.1117		0.0880	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Results show that as π_j increases the PP to Pj increases, and conversely as $\pi_{i \neq j}$ increases the PP to Pj decreases. The average time participants took to finalise their decision is shown to be strongly positively correlated with the PP to Pj. However, this correlation is perhaps one of reverse causality. The time it takes to allocate all to the self, is much less than ensuring payoffs are distributed equally. The mean difference variables, are constructed to identify within-subject timing differences. For the decisions which participants took longer than their individual average time (within each N) participants give more to others, while when they take less than their average the give less. This relationship is, however, only significant for positive differences in the two player treatment. The effect of the screen position, on giving, is also tested. The base case, is where the screen position is on the left (1), while increasing numbers denote a position further to the right. Results show that there are no significant differences for screen position within either N, for any position. Similar results hold for the name of the 'other' player. The order in which participants made decisions for each N was also varied, the dummy variables show if the order of the N (of that model) was either second or third, with the base case being first. Results show there are no significant order effects. Within N, the budget (which is standardised within each N) has no effect of giving.

C.2.2 Demographics

Alongside the experiment a questionnaire was conducted to establish the demographic characteristics of the participants, alongside questions regarding 'oneness', political persuasion, altruism and beliefs of others giving. The demographic composition of the sample is shown in Table 3.11. Further to this results from simple random effects models are shown in Table 3.12. The models include design control variables (number of players, single slider dummy, relative multiplier for P1, standardised budget within each N and round number) and a 'demographic' variable of interest, in order to determine if such variables explain the PP to P1 (over all 45 rounds). A separate model is ran for each, and the resulting coefficient for the variable of interest, alongside standard errors, number of participants and R^2. These are run as separate models for two reasons. The first, missing data. For particular questions a significant proportion of the sample did not answer (in particular parental income and degree subject). The second, is multicollinearity between particular variables, in particular the final four variables concerning altruism and beliefs of the payoffs others gave. As a result a simple modelling approach has been taken, allowing for the comparison of coefficients, significance and model fit; while being wary that these results are prone to omitted variable bias.

Results show that neither age, being an undergraduate, studying science, being of British or Asian nationality, having parents with higher incomes or education, being more right wing have a significant effect on giving to the self. Surprisingly, neither does hypothetical donations, nor willingness to donate to good causes. Females are somewhat more generous, as are art/humanities students, those who are religious, who come from a larger family or are more liberal. While having more friends in the session does not increase giving, a greater 'oneness' (the closeness of connection to others) to the group does. The hypothetical slider questions on "how do you believe the others in this session distributed payoffs" and "what do you believe is a fair distribution of payoffs between yourself and one other" are highly correlated with giving. Those regressors which give the most explanatory power are the fair payoffs, beliefs of payoffs given by others and the 'oneness' to others in the group.

Table 3.11: Sample Characteristics

	No.	%		No.	%
Gender			**Subject**		
Male	38	47.5%	Arts and Humanities	18	27.3%
Female	42	52.5%	Science	22	33.3%
Total	80	100.0%	Social Science	26	39.4%
Age			**Total**	66	100.0%
18-21	24	29.6%	**Degree Level**		
22-25	30	37.0%	Postgraduate	39	49.4%
26-29	13	16.0%	Undergraduate	40	50.6%
30+	14	17.3%	**Total**	79	100.0%
Total	81	100.0%	**Religion**		
Nationality			Agnostic/Atheist	10	14.3%
Asian	37	46.3%	Christian	10	14.3%
British	32	40.0%	Muslim	12	17.1%
European	8	10.0%	None	32	45.7%
Other International	3	3.8%	Other	6	8.6%
Total	80	100.0%	**Total**	70	100.0%

Table 3.12: Random Effects Model: Proportional Payoff to P1, Demographic Effects

		(1)		
		PP to P1		
	Coef.	Std. err.	N	R2
Age	0.0022	(0.0045)	81	0.0757
Gender	-0.1043**	(0.0520)	80	0.1034
Undergraduate Dummy	0.0447	(0.0519)	79	0.0891
Arts/Humanities	-0.1110*	(0.0607)	66	0.0897
Science	-0.0324	(0.0660)	66	0.0650
Social Science	0.1223**	(0.0583)	66	0.1008
British	0.0786	(0.0540)	80	0.0920
Asian	-0.0722	(0.0520)	80	0.0890
Religious	-0.1149**	(0.0562)	70	0.1034
Parental Income	-0.0001	(0.0232)	49	0.0718
Parental Education	0.0533	(0.0782)	70	0.0874
Family Size	-0.0236**	(0.0116)	80	0.0923
Oneness - Group	-0.0449***	(0.0137)	83	0.1275
Friends in Session	-0.0082	(0.0184)	80	0.0779
Authoritarian - Liberal	-0.0369***	(0.0130)	80	0.1128
Left - Right	-0.0098	(0.0144)	80	0.0764
Donate	-0.0001	(0.0001)	82	0.0855
Good Cause	-0.0152	(0.0095)	83	0.0986
Fair Payoffs	0.8202***	(0.0844)	82	0.3660
Belief Others Payoffs	0.5937***	(0.1347)	82	0.1455
Controls	YES			

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C.3 Graphical Intuition of Utility Functions

In addition to the formal notation in Section 3 the following surface plots illustrate the graphical intuition behind the preference parameters. Each figure plots the optimal proportional payoffs for each Player i ($\pi_i x_i^*$) for particular preference sets and numbers of recipients.

Figure 3.10 plots $\pi_1 x_1^*$ in the left panel and $\pi_2 x_2^*$ for different parameter values of self-interest, α_1, and inequality aversion, r. For simplicity the design parameters are set as: $N = 2$, $\pi_1 = 1$ and $\pi_2 = 0.5$. In general as α_1 increases the payoffs to Player 1 increase, while those to Player 2 decrease. Indeed at the extremes, when $\alpha_1 = 1$ then $\pi_1 x_1^* = 1$ and when $\alpha_1 = 0$ then $\pi_1 x_1^* = 0$. The extent to which α_1 changes behaviour depends on r. When $r \to -1$ efficiency concerns are important and so optimal payoffs reflect the highest weighted payoffs that can be obtained. As $r \to 0$ preferences approach Cobb-Douglas, where allocations (x_1^*) are directly proportional to α_1. As $r \to \infty$ equality is the primary concern, and so payoffs become more equal.

Figure 3.10: Inequality Aversion and Self-Interest: Standard Model

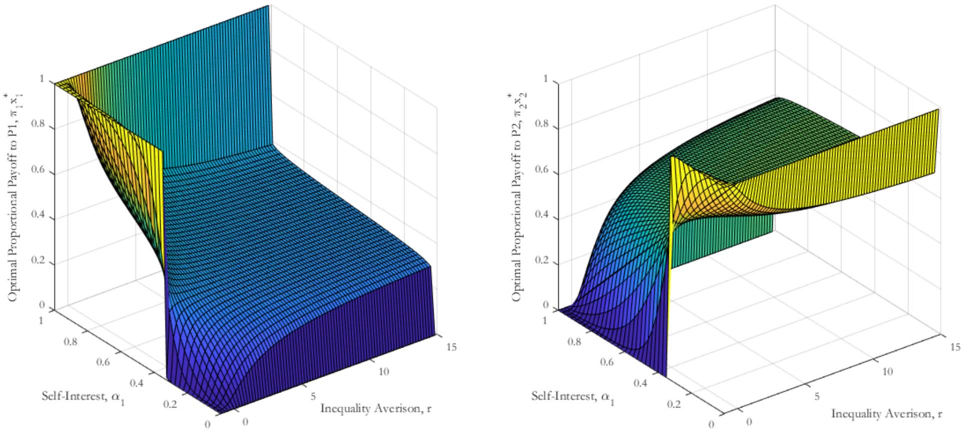

The *fisman* model allows for the distinction between *self-other* and *between-other* inequality aversion. As the experimental design allows $N > 2$ (and enables participants to distribute between others) this allows for differential behaviour to be observed, and explained by the model. To illustrate this Figure 3.11 shows the optimal PP to P1, P2 and P3 for varying values of r_1 and r_0; where $\alpha_1 = 0.5$ and the design parameters are: $N = 3$, $\pi_1 = 0.5$, $\pi_2 = 1$ and $\pi_3 = 0.25$. The variation in π_1 allows for differing behaviour to be predicted. As $r_1 \to \infty$ the payoffs are equally distributed between the self and others, and as $r_0 \to \infty$ payoffs are equally distributed between others. When $r_1, r_0 \to 0$ allocations are proportionate to α, so $\pi_1 x_1^* = 0.25$, $\pi_2 x_2^* = 0.25$ and $\pi_3 x_3^* = 0.0625$. If $r_0 \to -1$ the most efficient allocation between-others is preferred, so the share between Player 2 and Player 3 goes entirely to Player 2. If $r_0 > -1$ and $r_1 \to -1$ then all payoffs are allocated to Player 1. At the extreme when $r_1 = r_0 = -1$ individuals are technically indifferent between payoffs to P1 and P2, as $\pi_2 > \pi_3$ and $\alpha_1 = \pi_1/\pi_2 = 0.5$. When $r_1 = r_0$ then behaviour follows that in the *standard* model.

To illustrate how the *andreoni* function models behaviour as n increases Figure 3.12 plots $\pi_1 x_1^*$ for differing levels of congestion, b, and number of participants n. To simplify, we assume that $\alpha_1 = 0.5$, $\pi_1 = 1$ and $\pi_j = 0.5$ ($\forall j > 1$) but vary r across the three panels, where $r = -0.5$

Figure 3.11: Self-Other and Between-Other Inequality Aversion: Fisman Model

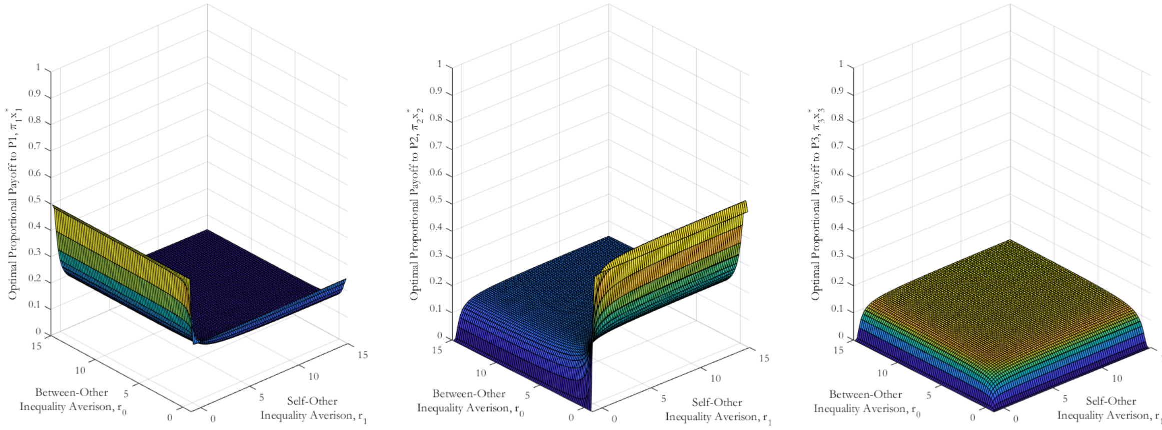

in the left, $r \to 0$ in the middle and $r = 2$ in the right. If $b = 1$, participants consider the *total* given to others, and so as n increases $\pi_1 x_1^*$ remains the same, and therefore $\pi_j x_j^*$ decreases. Indeed, the payoffs to Player 1 are the same for any n if $b = 1$ as when $n = 1$, for any value of b. The different absolute levels between the three panels emerge, as $\pi_1 > \pi_2$; as those prioritising efficiency give more to Player 1, Cobb-Douglas preferences mean that allocations are proportionate to $\alpha_1 = 0.5$ and total payoffs are sacrificed in order to reduce inequality when $r > 0$. As b increases changes in behaviour depend on n and r. If $r < 0$ payoffs to the self increase, if $r \to 0$ then $\pi_1 x_1^*$ remains the same, while if $r > 0$ payoffs to the self decrease. Total payoffs to others actually decrease as n increases when $r < 0$ and $b > 0$ because the weight to each other $((1-\alpha_1)/n)$ decreases, and so the more 'efficient' choice is to give more to the self. Conversely, when $r > 0$ and $b > 0$, the total payoffs to others increases, as individuals consider the weighted payoff to each other and prefer to weight higher the worst-off.

Figure 3.12: Congestion: Andreoni Model

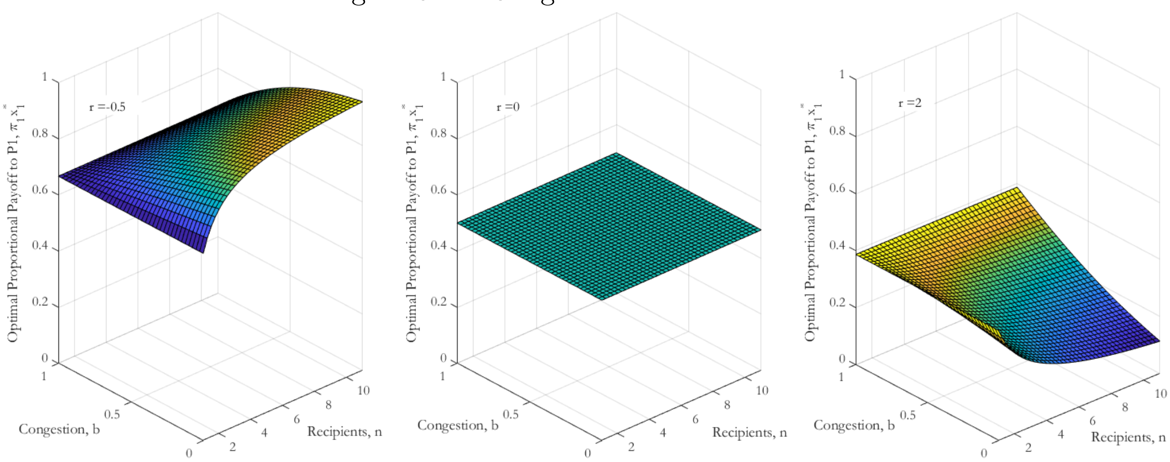

The above CES models concern only relative payoffs, and so assume distributions are proportionate to the budget. By incorporating an absolute minimum threshold, τ^*, the effects of a change in the total and average budgets, on behaviour, can be modelled. Figure 3.13 illustrates how $\pi_1 x_1^*$ changes in relation to τ^* and n (note that the n axis is reversed). An increase of n entails a reduction in the average budget available, and so the effects of differing

τ^* can be observed for each n. As above we assume $\alpha_1 = 0.5$, $\pi_1 = 1$ and $\pi_j = 0.5$ ($\forall j > 1$), with r varying across the three panels ($r \approx -0.5, 0, 2$). In each panel for high levels of τ^* the distribution of payoffs are equal between each participant, which ensures $\pi_1 x_1^*$ decreases as n increases as $m = 1$ throughout. When $\tau^* = 0$ the predictions converge to the predictions of the *standard* function. For values of τ^* between the two points, r affects decisions made. The minimum threshold level, in effect, allows for participants to distribute equally to a point, and then distribute according to their other preferences. The scope for this latter distribution depends on the available total and average budget.

Figure 3.13: Minimum Threshold Levels: Stone-Geary Model

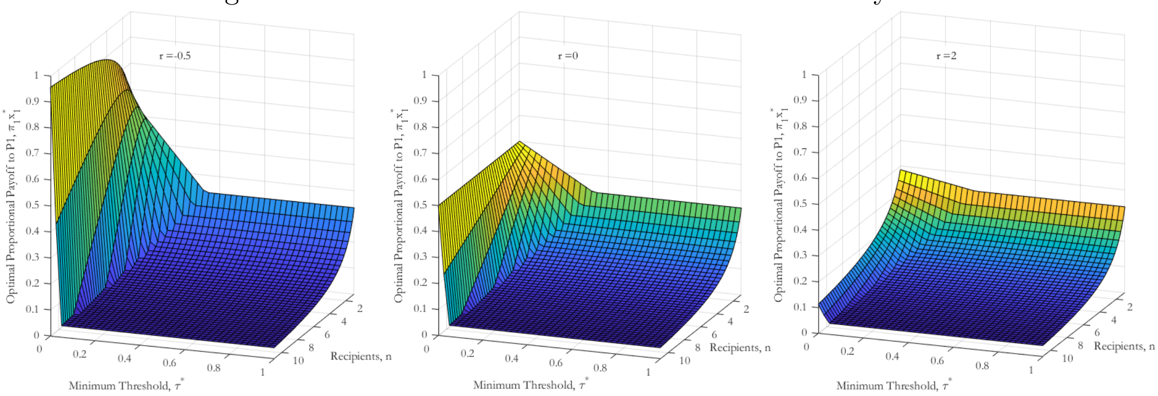

The above illustrates the intuition behind the derivations of the optimal distributions of payoffs, for each of the four models. The *amalgamated* model allows for these *standard* and *extended* behavioural concerns to be combined. This flexibility in modelling enables the extensive heterogeneity in behaviour to be accounted for; when the budget, number of recipients and prices of giving change.

C.4 Stone-Geary Non-Negativity Conditions

Due to the inclusion of τ some optimal allocations may lead to allocations where $\pi_i x_i - \tau_i < 0, \forall i$, which is not feasible. To solve this issue in the main analysis the assumption that $\tau = min(\tau^*, m/\sum_i^N \frac{1}{\pi_i})$ is used. An alternative solution is to make no such assumptions and instead incorporate non-negativity conditions. This does not restrict $\tau_i \leq x_i\pi_i$ explicitly, but provides a set of optimality conditions, within which a subset will ensure $\tau_i \leq x_i\pi_i$. Through this approach, the optimal allocations are as follows:

$$x^*_{i\neq k} = \frac{m + \sum_{j\neq i,k}^{N}\left(\left(\frac{\tau_i}{\pi_j}\left(\frac{\alpha_j\pi_j}{\alpha_i\pi_i}\right)^{\frac{1}{1+r}}\right) - \frac{\tau_j}{\pi_j}\right) - \sum_{k\neq i,j}^{N}\left(\frac{\tau_k}{\pi_k}\right)}{1 + \sum_{j\neq i,k}^{N}\left(\frac{\pi_i}{\pi_j}\left(\frac{\alpha_j\pi_j}{\alpha_i\pi_i}\right)^{\frac{1}{1+r}}\right)}, \qquad x^*_k = \frac{\tau_k}{\pi_k} \tag{3.18}$$

The set of optimality conditions can be concisely written by incorporating k. To solve, we use a solution similar to Lagrange's theorem with non-negative variables (see Dixit (1990) p. 28), which provides equations $\partial L/\partial x_i \leq 0, x_i \geq \tau_i/\pi_i, \forall i$, with complementary slackness, and $\partial L/\partial\lambda = 0$, in order to solve for optimal allocations. In other words, $\forall i$ either $\partial L/\partial x_i = 0$ or $x_i = \tau_i/\pi_i$ (or both). There are $2^N - 1$ combinations of equations, which provide optimal allocations (i.e. for $N = 2$: $[\partial L/\partial x_1 = 0, \partial L/\partial x_2 = 0], [\partial L/\partial x_1 = 0, x_2 = \tau_2/\pi_2], [x_1 = \tau_1/\pi_1, \partial L/\partial x_2 = 0]$). Vector k, then, indexes all those instances where $x_k = \tau_k/\pi_k$. In order to find the optimal, a series of if conditions are formulated (from each $2^N - 1$ combinations of Equation (3.18)) to enable the optimal to be found, while not violating the above conditions. The ordering of the if statements is, however, crucial as often several optimality condition ensure $\tau_i \leq x_i\pi_i, \forall i$. While this approach is feasible, the additional complexity, computing time needed and issues of ordering meant that it was not used.

C.5 Alternative Error Modelling

Within the specification of the error modelling two main assumptions are made. (1) $E[X_i] = x_i^*$, and (2) $Var(X_i) = \frac{(x_i^*(x_0^* - x_i^*))}{\hat{s}}$. While (1) is not particularly controversial, objections could be made to (2) when considering the nature of the problem, with differing degrees of complexity, κ. The variance of X_i could indeed depend on how many allocation decisions need to be made, independent of x_i^*. As a result three alternative error models, which define s differently, have been used to estimate preference parameters:(A) $\hat{s} = s$, (B) $\hat{s} = s.\kappa$ and (C) $\hat{s} = s.\kappa^{\gamma}$. The three specifications allow for differing variance assumptions to enter into the error modelling. (A) estimates s, with no consideration of the differing κ, while (B) assumes a positive linear relationship between the precision parameter $\hat{s}$ and κ. (C) parameterises this relationship, including an additional parameter for estimation, γ, where $\gamma \in [-1, 1]$, allowing for flexibility which can be captured by γ. The choice of the error model has consequences for the estimation of preference parameters, through its effect on the shape of the probability density functions from which the log-likelihood is calculated.

Figure 3.14 shows individual-level goodness-of-fit results for the three error models, from estimates within the *standard* model. The left panel shows the distribution of the likelihood proportion, ι, for each error model. The higher ι the better the model is explaining individual behaviour, relative to the uniform distribution, where $\iota < 0.5$ implies that drawing randomly (from a uniform distribution) better explains an individual's behaviour. The distribution shows that (A) performs worse that (B) and (C), while (B) and (C) are closely matched, with the exception of the worst explained, for which (B) performs somewhat better. There are 10.53% of the sample for whom $\iota \leq 0.5$, in (A), 1.32% in (B) and 5.26% in (C).

Figure 3.14: Goodness-of-Fit for Alternative Error Models; Standard Model

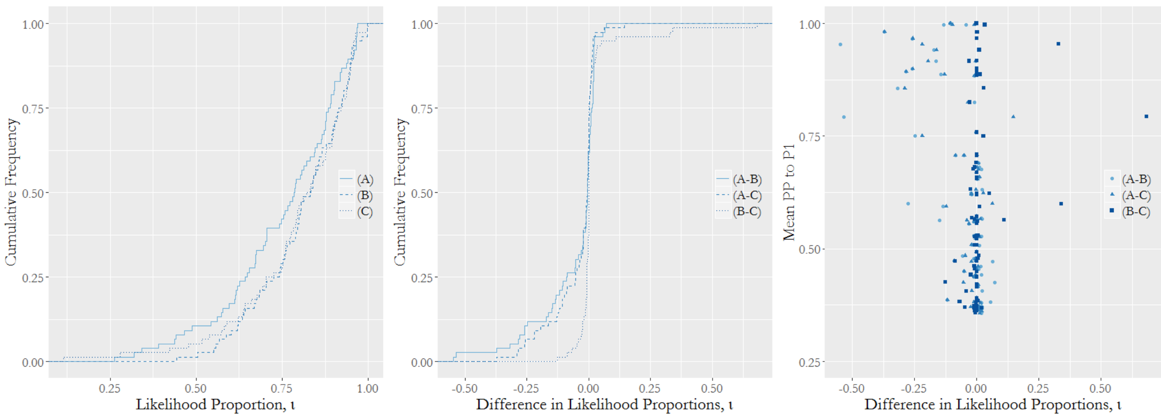

While these distributions show only the aggregate distributions the middle panel shows the distribution of the difference in ι between the models, for each individual. Take the solid line for example, (A-B) shows the difference between ι_A and ι_B, the lower (and negative) the value the better B performs, the higher (and positive) the better A performs. Both (A-B) and (A-C) show that A tends to perform worse, with 61.84% and 63.16% higher ι values for B and C, respectively. The third cdf (B-C) shows a similarity of performance between the two, with C outperforming B for 52.63%. To identify for whom the error models perform better the right panel shows the aforementioned differences plotted against the mean PP to P1. Here we observe that it is the more self-interested individuals for whom A tends to perform worse. B outperforms A to a greater extent than C, for those individuals, however, it also tends to perform worse than A for those who share more equally. C does not tend to have this latter issue.

The reason for the differences in the right panel, can be explained by considering the formula $a_i = x_i^*(\hat{s} - 1)$. Model (A) has an issue when participants have a high degree of self-interest. It models well behaviour for a particular N, as $a_1 > 1$ and $a_{j>1} < 1$, meaning the pdf asymptotes at $x_1 = 1$. However, as $x_{j>1}^*$ may be higher in rounds with lower N, $a_{j>1} > 1$, meaning the pdf becomes uni-modal at an interior allocation, meaning at the bound (where $x_1 = 1$) the likelihood value is very low. Model (B) solves this issue, allowing a_i to vary with κ, however, it does so at the expense of those who are allocating more equally. For them variance is perhaps not decreasing as κ increases, as indeed their decision problem becomes more difficult to distribute equally. Model (C) then allows for the flexibility of estimation, which ensures that the behaviour of the more self-interested is not modelled badly, but that compensating for that does not lead to worse estimates for those who share more equally. An additional parameter does need to be estimated in (C), but due to the above issues and the additional information that γ carries (C) has been chosen for the main analysis.

C.6 Clustering Amalgamated Preferences

Due to the high dimensionality of the preferences in the *amalgamated* utility function visualising and describing the estimated parameters can be difficult. An alternative method of understanding the distribution of preferences is through cluster analysis. A 'mixture' of individual density functions which accurately fit the data can be estimated within a finite mixture model (more details are found in Section 5.4 of Chapter 1). The multidimensional ellipses, with specific mean and variance in each dimension, capture patterns in the distribution of preferences. Below results are shown from finite mixture model results for r_1, α, r_0, b and τ.[10] Table 3.13 shows results where three clusters are optimal (allowing *mclust* to search between 0 and 5 clusters), while Table 3.14 shows results where eight clusters are optimal (allowing *mclust* to search between 0 and 15 clusters).

In Table 3.13 we observe that the sample is split into three clusters, the largest is Cluster C (49%), followed by Cluster A (30%), with Cluster B as the smallest (21%). Cluster A consists of those who are slightly averse to inequality, with a high level of congestion and some concern for a minimum threshold. This cluster are the most tightly packed, with the lowest variance for all parameters, bar b, in particular with regard to the inequality aversion parameters. Cluster B consists of weakly 'weighted prioritarians', who have zero congestion and, again, and some concern for a minimum threshold. For b variance is very low, but for each other parameter it lies between A and C. Cluster C captures those who are the most averse to inequality and have the largest τ. There is a degree of congestion, more than B, but less than A, while the variance is the largest for all preferences. The mean self-interest remains similar across the clusters, but is the lowest in C. The mean levels and variance of r_1 and r_0 are similar within each cluster.

Table 3.13: Finite Mixture Model for Amalgamated Preference Parameters: Three Clusters

	Cluster		
	A Mean (Var.)	**B** Mean (Var.)	**C** Mean (Var.)
Ineq. Aversion, r_1	-0.17	1.77	7.03
	(0.07)	(2.76)	(26.31)
Self-Interest, α	0.77	0.77	0.65
	(0.03)	(0.04)	(0.08)
Ineq. Aversion, r_0	-0.26	1.59	7.57
	(0.11)	(3.24)	(28.50)
Congestion, b	0.69	0.00	0.33
	(0.12)	(0.00)	(0.13)
Min. Threshold, τ	1.04	0.84	3.69
	(1.02)	(1.28)	(33.24)
Proportions	0.30	0.21	0.49

In Table 3.14 an extended number of clusters accounts for a greater extent of the heterogeneity in preferences. As above the clusters are ordered in relation to r_1, with Cluster A capturing a efficiency prioritarian standpoint, while Cluster H encompassed an extreme

[10] To reduce issues of outliers the r_1 and r_0 estimates greater than 15, where capped at a value of 15.

aversion to inequality. As before r_0 and r_1 appear similar within most clusters, with the exception of Cluster E which captures those who have a higher self-other inequality aversion and Cluster H within which between-other inequality aversion is higher. Self-interest varies to a greater extent between clusters, with the lowest α in Cluster F, this is however, mainly due to the large τ which we observe. Congestion is very low in Clusters D and G and high in Cluster A, while the minimum threshold levels are highest in C and F. Interesting differences between similar clusters can be observed. Clusters D and G have similarly low b and τ, but G has higher inequality aversion and lower self-interest. Cluster C and D have similar levels of inequality aversion, but C has much higher congestion, higher τ and lower α. Clusters G and H have similar levels of r_1, but H has much higher α and a higher r_0, while congestion is higher and τ is lower in H.

Table 3.14: Finite Mixture Model for Amalgamated Preference Parameters: Eight Clusters

	Cluster							
	A	**B**	**C**	**D**	**E**	**F**	**G**	**H**
	Mean (Var.)	Mean (Var.)	Mean (Var.)	Mean (Var.)	Mean (Var.)	Mean (Var.)	Mean (Var.)	Mean (Var.)
Ineq. Aversion, r_1	-0.29	0.04	1.29	2.10	6.18	8.25	10.01	10.35
	(0.04)	(0.11)	(3.70)	(2.77)	(25.64)	(19.16)	(0.65)	(22.83)
Self-Interest, α	0.68	0.88	0.50	0.72	0.89	0.27	0.45	0.99
	(0.03)	(0.00)	(0.00)	(0.04)	(0.01)	(0.02)	(0.00)	(0.00)
Ineq. Aversion, r_0	-0.25	0.05	1.76	1.81	3.82	9.43	10.47	14.37
	(0.11)	(0.66)	(6.69)	(3.57)	(8.51)	(15.13)	(6.98)	(2.38)
Congestion, b	0.97	0.30	0.66	0.00	0.41	0.28	0.00	0.24
	(0.00)	(0.08)	(0.05)	(0.00)	(0.11)	(0.14)	(0.00)	(0.11)
Min. Threshold, τ	1.03	0.96	3.05	0.81	1.10	12.14	1.49	0.38
	(0.32)	(1.57)	(6.72)	(1.33)	(4.10)	(48.01)	(1.02)	(0.29)
Proportions	0.16	0.19	0.08	0.17	0.14	0.11	0.07	0.09

Through using finite mixture models the complexities of heterogeneous multidimensional preferences can be more easily summarised, and through it interesting differences within the sample observed.

C.7 Mismatches Between Rankings

While the main analysis identifies 'types' of individuals by considering the information criteria for a combined measure of goodness-of-fit and prediction, there are concerns related to the mismatch in the 'ranking' of alternative utility functions. Here a deeper look into the mismatches between rankings based on the differences between Residual Sum of Squares and Log-Likeihood and amongst alternative error models.

C.7.1 Residual Sum of Squares or Log-Likelihood

The Residual Sum of Squares (RSS), identifies the difference between the optimal and observed allocations. Likelihood, identifies the probability that the allocation is observed, given the utility and error model assumed. The former asks how close, the second how probable. Through this alternative criteria differences in which model is considered best will inevitably emerge. In the main analysis the log-likelihood was the metric used, as indeed the estimation procedure was based on maximising the log-likelihood. Here, the differences between the two can be analysed.

Table 3.15 shows the rankings of utility functions, if the utility functions had been compared using RSS, rather than the log-likelihood values (in Table 3.7). Comparing Table 3.15 with Table 3.7, we observe similar trends. The *amalgamated* model does better within *raw* RSS fit, than prediction, and in RSS compared to IC, while the *standard* model does the opposite. Final IC results for both are somewhat similar, the modal type is *standard*, with the lowest being *amalgamated*. *Standard*, *stone-geary* and *amalgamated* are 'best' for more individuals, while *fisman* and *andreoni* are best for fewer individuals (compared to the log-likelihood IC). When comparing matching within-individuals we observe 51.31%, 59.21% and 59.21% of the sample have matched rankings for GOF, Pred and Both, respectively.

Table 3.15: Utility Types: Ranked by Residual Sum of Squares and Information Criterion

	Residual Sum of Squares			**Information Criterion**		
	GOF	Predict	Both	GOF	Predict	Both
Standard	6	20	9	39	46	40
Fisman	21	18	19	5	10	10
Andreoni	12	14	13	16	12	13
Stone-Geary	12	12	14	11	5	11
Amalgamated	25	11	20	5	3	2

C.7.2 Alternative Error Modelling

In Appendix C.5 three alternatives error models, (A) $\hat{s} = s$, (B) $\hat{s} = s.\kappa$ and (C) $\hat{s} = s.\kappa^{\gamma}$ are discussed. Using preference estimates from each of (A), (B) and (C), similar rankings to those in Table 3.7 can be conducted. The result of assuming an alternative error model may, lead to different compositions of 'types' in the sample. Table 3.16 shows the composition of 'types', by using the information criteria, for each alternative error model.

Table 3.16: Mismatches in Error Models: Ranked by Information Criterion

	(A)	(B)	(C)
Standard	49	42	33
Fisman	9	14	13
Andreoni	11	12	21
Stone-Geary	6	3	9
Amalgamated	1	5	0

Results show that, in comparison to (C) which is used in the main analysis, there is one more *amalgamated* type in (A), five more in (B). There are lower numbers of *andreoni* types in both alternatives, while the number *fisman* types is the lower in (A), and higher in (B). The number of *standard* types is higher in both (A) and (B), while *stone-geary* types are both lower. While these results are sample aggregates, of most interest is how many subjects are classed as the same 'type' in the alternative models. Those of the same 'types' in (A) and (B) are 65.8%, with 55.3% between (A) and (C) and 59.2% between (B) and (C). There are 69.7% for whom two or more models designate the same 'type', with 44.7% who have the same type in all three.

While it is clear that mismatches between rankings do occur, whether that be the metric used to identify the 'goodness' of the model or from differing estimates according to alternate error models, one main result remains. There is still heterogeneity in which models are 'best'. In none of the specifications does one particular utility function dominate and best explain all individual's behaviour.

Conclusion

Preferences relating to inequality aversion and self-interest have been explored throughout this thesis. Through running incentivised laboratory and lab-in-the-field experiments prosocial behaviour has been observed and preferences at the individual, cluster and sample level have been estimated to account for such behaviour. The perspectives from which decisions are made, alongside the particularities of the design of the experiment, are shown to have large effects on the behaviour of individuals and the preference parameters subsequently estimated. The effect of *oneness*, the closeness of connection to others, was investigated, modelled and found to have a large and significant impact on giving behaviour; where distinctions between *self-other* and *between-others* trade-offs emerged. Finally, the effects of increasing the number of recipients on giving behaviour has been investigated, with both the goodness-of-fit and predictive accuracy of alternative utility functions compared. This conclusion will draw out the contributions of this thesis, explore some limitations and possible extensions, before discussing proposed future research.

***

In general, this thesis contributes to our understanding of prosocial behaviour and other-regarding preferences. The incentivised experiments are run to provide novel datasets allowing for the observation of prosocial behaviour for both student and general population samples. Extensive individual-level observations are gathered within each experiment, across different treatments and decision problems. In each instance the experimental data is complemented by in-depth survey data. Throughout, structural models are formulated which build upon existing models; either by providing generalisations, incorporating additional behavioural assumptions or amalgamating functional forms. Within these models preference parameters are estimated, to strive to explain the observed behaviour and providing insights into the value judgements that participants have. The stochastic nature of behaviour is further modelled, with the Beta and Dirichlet distributions proposed as random behavioural models. These error models allow for a greater flexibility in the stochastic nature of decision making, are generalisable to N dimensions and can be easily extended to other constrained allocation problems.

In particular, the main contributions of the three chapters are as follows. The first chapter contributes to existing research by identifying within-subject treatment effects between *multiple designs* and *alternative perspectives*. The three-person design of the experiment enables the estimation (and elicitation) of preferences parameters when participants are both *partial* and *impartial* agents. The second chapter contributes to the literature by proposing a CES utility function which incorporates *oneness levels*, to account for the explanatory power that social distance can have on individual decision making. By estimating preferences relating to *inequality aversion*, *self-interest* and *oneness*, the intricacies of their interactions are explored and individual-level behaviour, observed in a lab-in-the-field experiment, is explained. The third chapter contributes by adopting a within-subject design which allows for the complexities of both *self-other* and *between-other* trade-offs to be observed as the number of recipients increase. Individual-level preference parameters are estimated within five alternative models, allowing for both the goodness-of-fit and predictive power of the respective models to be analysed.

Methodologically this thesis has contributed in several ways. In terms of experimental design, novel computer interfaces have been programmed for each of the experiments run; with multiple designs being tested and compared within the first chapter. The precision and complexities of the laboratory were extended to a lab-in-the-field setting, with the use of touch-screen tablets. This removed barriers of computer literacy (due to the intuitive nature of touch-screen technology) and social desirability bias (caused by a lack of anonymity, from experimenters recording responses face-to-face) in the lab-in-the-field context. The programming code produced throughout the thesis also provides a contribution. The code for the experiments allows others to replicate and adapt the experiments ran, while the analytical code enables others to perform similar descriptive, econometric and structural estimation (particularly in relation to the error modelling). Finally, the range and differing intensity of the analytical methods used is important. Behaviour and preferences have been analysed at each the sample, cluster and individual level, using descriptive, econometric and structural methods. Dependent on the context, one approach may be more appropriate for a particular strand of research than another; it is, therefore, hoped that of the methods used here, some can be appropriate for future researchers needs.

***

As in all research there are limitations to this thesis. The limitations discussed below relate to the representativeness and size of the sample, the domain of the incentives and computational burden. The first set of limitations are frequently observed in experimental research. Within the laboratory experiments there is a lack of representativeness in the sample, as the majority of participants are students. This means that generalisations of behaviour and preferences cannot be made. To solve this issue the general population could be recruited to the experiment, however, this is often impractical and costly (both in terms of organisation and monetary payoffs). Within the second chapter, the general population were recruited within the lab-in-the-field experiment, increasing the representativeness of the sample. This was feasible as the issues of costs of monetary payoffs are minimised, due to the lower average earnings in the developing context. The sample is, however, still far from being representative, with the urban populace, more highly educated and wealthy being over represented. In the case of the latter the sample could be weighted to ensure a more accurate representation of the general population. However, this was not done. The sample size is still not particularly large, and so such methods are not very reliable, especially with such a high degree of heterogeneity in behaviour. As a result the thesis does not claim that the preference parameters estimated are population preference parameters. The aim is rather to test economic theory and establish the models, and preference parameters within those models, which best fit the observed behaviour.

The second limitation relates to the domain in which participants are incentivised. Throughout each experiment monetary incentives are used. Therefore, the preference parameters estimated relate to money in particular. Participants could have alternative preferences over different domains. For instance, the degree of self-interest and level of inequality aversion could differ in the domains of health, education and income. Indeed, comparisons in Chapter 1 show substantial differences between inequality aversion over income and health. It is, however, difficult to incentivise other domains in an experimental setting, without breaking ethical practices.[11] If these domains cannot be incentivised, an alternative could be to use surveys which ask participants for their stated preference. This alternative does, however, raise concerns from experimentalists, as the potential for particular biases creep in when such decisions are not incentivised; meaning that observed differences could be due to such biases rather than an underlying difference. The point highlights another limitation of the thesis, as the survey data, which accompanies the experimental data, is not incentivised. As a result biases could lead to a misreporting of the demographics and opinions elicited in the survey.

Finally, a limitation which would perhaps not be apparent is that of computational burden. Throughout the analysis of the experimental data, one particular issue was that of the computation time taken to estimate structural preference parameters. As the thesis progressed the analysis became more sophisticated and due to work in the third chapter, the computational time needed to re-estimate structural parameters in the first and second chapter was significantly reduced. In the third chapter, however, the high dimensionality of the optimisation procedure led to a high computational burden. If higher dimensions were required, then this could lead to further issues.

***

In addition to the limitations of the thesis, extensions of the thesis are also proposed. Throughout the thesis, the theoretical framework considers one family of functional forms, namely constant elasticity of substitution (CES). An extension could be to compare these models to others in the literature; for example, E. Fehr and Schmidt (1999), Bolton (1991) or Cappelen et al. (2007). Within other models alternative behavioural assumptions are used, and therefore different preference parameters can be estimated. The goodness-of-fit and predictive accuracy could be compared, identifying which of the functional forms best models the data. Similarly, alternative error models could be compared. Within the random behavioural framework, a standard (truncated) normal error model could be compared to the Beta and Dirichlet distributions, or indeed the random behavioural models could be compared to either random utility or random preference models. A broader approach, before estimating

[11] As voiced by John Hey, one cannot go around amputating limbs in experiments.

structural preference parameters, would be to conduct tests of *rationality.* Measuring the extent of GARP violations, using the Critical Cost Efficiency Index (CCEI), is one example of such a test. The scope for additional research adopting alternative models is extensive, but with the experimental data gathered is entirely plausible.

The incorporation of preferences relating to risk into models of other regarding preferences could be a further extension. Some of the distributional decision problems presented within the experiments could have been interpreted within the context of risky decision making. In the first chapter, decisions made under the veil of ignorance could have been considered as a risky decision between the possible realised positions of either player. In the budget treatment of the second chapter, participants knew the five others in the room, but did not know to whom they were allocating. Therefore, behaviour could be modelled considering the expected utility of a gamble between the possible realisation of the five other participants being either Player 2 or Player 3, rather than assuming that participants weight the payoffs to the others with the expected oneness levels. These extensions have been tentatively explored in the respective chapters, but further more rigorous analysis could be conducted.

A methodological extension relates to cluster analysis. The cluster analysis within the thesis relies either on splitting the sample into clusters based on average levels of giving and then estimating representative preference parameters, or on estimating individual level preference parameters and then clustering to characterise the estimates. An extension to this would be to estimate the clusters and preferences simultaneously. This would provide a more complete analysis than the former analysis used and require less extensive individual data, than the latter. Mixtures of Dirichlet distributions could be formulated, rather than mixtures of Gaussian densities, enabling greater flexibility and explanatory power.

While the work in the thesis is somewhat abstract, one possible extension would be into real world charitable donations. The various elements which constitute giving behaviour explored in this thesis are important considerations of real world donations. The 'abstract' preferences explored underpin giving behaviour in general. So, through running representative, or donor specific, experiments which have concrete examples of charities programmes and incorporate these 'abstract' elements, insights into real world charitable giving could be gained. Fundraising packages could then be tailored to individuals, or clusters of individuals, to increase charitable donations, for example. A speculative analysis of this extension is carried out in the third chapter, but this could be extended further to incorporate oneness, perspective and concrete real world examples.

***

Building upon the research in the thesis I aim to extend the consideration of 'distributional' aspects of inequality to account for those aspects which are 'dynamic'. The two are inextricably interwoven and through incorporating both elements interesting behavioural insights can be gained. Central to this research will be two papers. The first, "Giving to Others and the Future Self", focuses on the relation between other-regarding preferences and time preferences. By incorporating closeness measures to others and the future-self, the extent to which behaviour changes due to the strength of connection to either entity can be explored. The trade-offs that individuals make between the self, future-self and others will be modelled and these models tested using data from incentivised laboratory experiments. To initialise investigations into this research I conducted a pilot study at the Choice Lab, NHH. The results are promising, and this provides a platform from which to begin postdoctoral research. The second paper, "Intergenerational Allocations in an Uncertain World", delves into dynamic decision making. This work will consider how individuals make distributional decisions, both between and within 'generations'. By nesting a modified dictator game within a multi-period risky intergenerational problem, individual behaviour will be observed and preferences, synonymous with ethical value judgements, will be estimated in incentivised laboratory and lab-in-the-field experiments.

***

To conclude, humans are complex, varied and fallible. By utilising structural models which incorporate behavioural constructs, accounting for heterogeneity in preferences and rationalising noise in decision making, economists can strive to better explain human behaviour. It is clear that other-regarding preferences are an integral component of human nature, and influence how we act in certain circumstances. This work aims to further this understanding, allowing economists to move away from supposing that self-love is the only motivating factor of humanity.

Economists are born free; and everywhere they are in chains (cf. Rousseau, 1651). Chained to suppose that self-love alone defines humanity. But fools, "Wisedome is acquired, not by reading of Books, but of Men" (Hobbes, 1651). Through relaxing the constricting assumption of self-love economists have nothing to lose but their chains (cf. Marx and Engels, 1846). They have the world to explain.

Bibliography

Abasolo, Ignacio and Aki Tsuchiya (2004). “Exploring social welfare functions and violation of monotonicity: an example from inequalities in health”. In: *Journal of Health Economics* 23(2), pp. 313–329.

Abasolo, Ignacio and Aki Tsuchiya (2008). “Understanding preference for egalitarian policies in health: are age and sex determinants?” In: *Applied Economics* 40(19), pp. 2451–2461.

Abásolo, Ignacio and Aki Tsuchiya (2013). “Is more health always better for society? Exploring public preferences that violate monotonicity”. In: *Theory and Decision* 74(4), pp. 539–563.

Akaike, Hirotogu (1998). “Information theory and an extension of the maximum likelihood principle”. In: *Selected Papers of Hirotugu Akaike*. Springer, pp. 199–213.

Alkire, Sabina and James Foster (2011). “Counting and multidimensional poverty measurement”. In: *Journal of Public Economics* 95(7), pp. 476–487.

Amiel, Yoram and Frank Cowell (1999). *Thinking about inequality: Personal judgment and income distributions*. Cambridge University Press.

Amiel, Yoram, Frank A Cowell, and Wulf Gaertner (2009). “To be or not to be involved: a questionnaire-experimental view on Harsanyi’s utilitarian ethics”. In: *Social Choice and Welfare* 32(2), pp. 299–316.

Andersen, Steffen et al. (2008). “Eliciting risk and time preferences”. In: *Econometrica* 76(3), pp. 583–618.

Andreoni, James (2007). “Giving gifts to groups: How altruism depends on the number of recipients”. In: *Journal of Public Economics* 91(9), pp. 1731–1749.

Andreoni, James and John Miller (2002). “Giving according to GARP: An experimental test of the consistency of preferences for altruism”. In: *Econometrica* 70(2), pp. 737–753.

Andreoni, James and Charles Sprenger (2012). “Estimating time preferences from convex budgets”. In: *The American Economic Review* 102(7), pp. 3333–3356.

Angulo, Roberto, Yadira Dìaz, and Renata Pardo (2016). “The Colombian multidimensional poverty index: Measuring poverty in a public policy context”. In: *Social Indicators Research* 127(1), pp. 1–38.

Aron, Arthur, Elaine Aron, and Danny Smollan (1992). “Inclusion of Other in the Self Scale and the structure of interpersonal closeness.” In: *Journal of Personality and Social Psychology* 63(4), pp. 596–612.

Atkinson, Anthony B (1970). “On the Measurement of Inequality”. In: *Journal of Economic Theory* 2(3), pp. 244–263.

Bardsley, Nicholas et al. (2010). *Experimental economics: Rethinking the rules*. Princeton University Press.

Barr, Abigail et al. (2015). "Economic status and acknowledgement of earned entitlement". In: *Journal of Economic Behavior & Organization* 118, pp. 55–68.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen (2014). "High-dimensional methods and inference on structural and treatment effects". In: *The Journal of Economic Perspectives* 28(2), pp. 29–50.

Binmore, Ken (n.d.). "Interpersonal Comparison of Utility". In: *The Oxford Handbook of Philosophy of Economics*.

Binzel, Christine and Dietmar Fehr (2013). "Giving and sorting among friends: evidence from a lab-in-the-field experiment". In: *Economics Letters* 121(2), pp. 214–217.

Bolton, Gary E (1991). "A comparative model of bargaining: Theory and evidence". In: *The American Economic Review*, pp. 1096–1136.

Bolton, Gary E and Axel Ockenfels (2000). "ERC: A theory of equity, reciprocity, and competition". In: *American Economic Review*, pp. 166–193.

Bonacich, Phillip et al. (1976). "Cooperation and group size in the n-person prisoners' dilemma". In: *Journal of Conflict Resolution* 20(4), pp. 687–706.

Bosmans, Kristof and Erik Schokkaert (2004). "Social welfare, the veil of ignorance and purely individual risk: An empirical examination". In: *Research on Economic Inequality* 11, pp. 85–114.

Branas-Garza, Pablo et al. (2010). "Altruism and social integration". In: *Games and Economic Behavior* 69(2), pp. 249–257.

Breitmoser, Yves (2013). "Estimation of social preferences in generalized dictator games". In: *Economics Letters* 121(2), pp. 192–197.

Cappelen, Alexander W et al. (2007). "The pluralism of fairness ideals: An experimental approach". In: *The American Economic Review* 97(3), pp. 818–827.

Cason, Timothy N and Vai-Lam Mui (1997). "A laboratory study of group polarisation in the team dictator game". In: *The Economic Journal* 107(444), pp. 1465–1483.

Charness, Gary and Matthew Rabin (2002). "Understanding Social Preferences with Simple Tests". In: *Quarterly Journal of Economics* 117(3), pp. 817–869.

Chatfield, Mark and Adrian Mander (2009). "The Skillings–Mack test (Friedman test when there are missing data)". In: *The Stata Journal* 9(2), p. 299.

Cialdini, Robert B et al. (1997). "Reinterpreting the empathy–altruism relationship: When one into one equals oneness." In: *Journal of Personality and Social Psychology* 73(3), pp. 481–494.

Conte, Anna and Peter G Moffatt (2014). "The econometric modelling of social preferences". In: *Theory and decision* 76(1), pp. 119–145.

Cox, James C, Daniel Friedman, and Steven Gjerstad (2007). "A tractable model of reciprocity and fairness". In: *Games and Economic Behavior* 59(1), pp. 17–45.

Crosetto, Paolo and Antonio Filippin (2015). "A theoretical and experimental appraisal of four risk elicitation methods". In: *Experimental Economics* 19(3), pp. 1–29.

Dirichlet, Peter Gustav Lejeune (1839). "Sur une nouvelle méthode pour la détermination des intégrales multiples". In: *Journal de Mathématiques Pures et Appliquées* 4, pp. 164–168.

Dixit, Avinash K (1990). *Optimization in economic theory.* Oxford University Press on Demand.

Dolan, Paul and Aki Tsuchiya (2009). "The social welfare function and individual responsibility: some theoretical issues and empirical evidence". In: *Journal of Health Economics* 28(1), pp. 210–220.

Dolan, Paul and Aki Tsuchiya (2011). "Determining the parameters in a social welfare function using stated preference data: an application to health". In: *Applied Economics* 43(18), pp. 2241–2250.

Dolbear, F Trenery et al. (1968). "Collusion in oligopoly: an experiment on the effect of numbers and information". In: *The Quarterly Journal of Economics* 82(2), pp. 240–259.

Engel, Christoph (2011). "Dictator games: A meta study". In: *Experimental Economics* 14(4), pp. 583–610.

Engelmann, Dirk and Martin Strobel (2004). "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Reply". In: *The American Economic Review* 94(4), pp. 857–869.

Engelmann, Dirk and Martin Strobel (2007). "Preferences over income distributions: Experimental evidence". In: *Public Finance Review* 35(2), pp. 285–310.

Erkal, Nisvan, Lata Gangadharan, and Nikos Nikiforakis (2011). "Relative earnings and giving in a real-effort experiment". In: *The American Economic Review* 101(7), pp. 3330–3348.

Erlei, Mathias (2008). "Heterogeneous social preferences". In: *Journal of Economic Behavior & Organization* 65(3), pp. 436–457.

Fechner, Gustav (1965 [1860]). *Elements of psychophysics.* Holt, Rinehart & Winston.

Fehr, Ernst, Michael Naef, and Klaus M Schmidt (2006). "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment". In: *The American Economic Review* 96(5), pp. 1912–1917.

Fehr, Ernst and Klaus M Schmidt (1999). "A theory of fairness, competition, and cooperation". In: *Quarterly Journal of Economics* 114(3), pp. 817–868.

Fehr, Ernst and Klaus M Schmidt (2006). "The economics of fairness, reciprocity and altruism–experimental evidence and new theories". In: *Handbook of the economics of giving, altruism and reciprocity* 1, pp. 615–691.

Fischbacher, Urs (2007). "z-Tree: Zurich toolbox for ready-made economic experiments". In: *Experimental Economics* 10(2), pp. 171–178.

Fisman, Raymond, Shachar Kariv, and Daniel Markovits (2007). "Individual preferences for giving". In: *The American Economic Review* 97(5), pp. 1858–1876.

Florek, K et al. (1951). "Sur la liaison et la division des points d'un ensemble fini". In: *Colloquium Mathematicae.* Vol. 2. 3-4, pp. 282–285.

Forsythe, Robert et al. (1994). "Fairness in simple bargaining experiments". In: *Games and Economic Behavior* 6(3), pp. 347–369.

Foster, James E (2007). "A report on Mexican multidimensional poverty measurement". In: *Oxford Poverty and Human Development Initiative Working Paper* 40.

Fouraker, Lawrence and Sidney Siegel (1963). "Bargaining and group decision making". In: *New York/Toronto/London 196o.*

Friedman, Milton (1937). "The use of ranks to avoid the assumption of normality implicit in the analysis of variance". In: *Journal of the american statistical association* 32(200), pp. 675–701.

Gächter, Simon, Chris Starmer, and Fabio Tufano (2015). "Measuring the closeness of relationships: a comprehensive evaluation of the 'inclusion of the other in the self' scale". In: *PloS one* 10(6), e0129478.

Gaertner, Wulf and Erik Schokkaert (2012). *Empirical social choice: questionnaire-experimental studies on distributive justice*. Cambridge University Press.

Geary, Roy C (1950). "A note on "A constant-utility index of the cost of living"". In: *The Review of Economic Studies* 18(1), pp. 65–66.

Goeree, Jacob K et al. (2010). "The 1/d law of giving". In: *American Economic Journal: Microeconomics* 2(1), pp. 183–203.

Hannan, Edward J and Barry G Quinn (1979). "The determination of the order of an autoregression". In: *Journal of the Royal Statistical Society. Series B (Methodological)* 41(2), pp. 190–195.

Harsanyi, John C (1953). "Cardinal utility in welfare economics and in the theory of risk-taking". In: *Journal of Political Economy* 61(5), pp. 434–435.

Harsanyi, John C (1975). "Can the maximin principle serve as a basis for morality? A critique of John Rawls's theory". In: *American Political Science Review* 69(2), pp. 594–606.

Herne, Kaisa and Maria Suojanen (2004). "The role of information in choices over income distributions". In: *Journal of Conflict Resolution* 48(2), pp. 173–193.

Hey, John D and Noemi Pace (2014). "The explanatory and predictive power of non two-stage-probability theories of decision making under ambiguity". In: *Journal of Risk and Uncertainty* 49(1), pp. 1–29.

Hey, John D and Luca Panaccione (2011). "Dynamic decision making: what do people do?" In: *Journal of Risk and Uncertainty* 42(2), pp. 85–123.

Hobbes, Thomas (1651). *Leviathan.* Ed. by John Charles Addison Gaskin. Oxford World's Classics (2009).

Hoffman, Elizabeth et al. (1994). "Preferences, property rights, and anonymity in bargaining games". In: *Games and Economic Behavior* 7(3), pp. 346–380.

Isaac, R Mark and James M Walker (1988). "Group size effects in public goods provision: The voluntary contributions mechanism". In: *The Quarterly Journal of Economics* 103(1), pp. 179–199.

Isaac, R Mark, James M Walker, and Arlington W Williams (1994). "Group size and the voluntary provision of public goods: Experimental evidence utilizing large groups". In: *Journal of public Economics* 54(1), pp. 1–36.

Jakiela, Pamela (2013). "Equity vs. efficiency vs. self-interest: on the use of dictator games to measure distributional preferences". In: *Experimental Economics* 16(2), pp. 208–221.

Karni, Edi, Tim Salmon, and Barry Sopher (2008). "Individual sense of fairness: an experimental study". In: *Experimental Economics* 11(2), pp. 174–189.

Kolm, Serge-Christophe (1976). “Unequal Inequalities. I”. In: *Journal of Economic Theory* 12(3), pp. 416–442.

Leider, Stephen et al. (2009). “Directed altruism and enforced reciprocity in social networks”. In: *The Quarterly Journal of Economics* 124(4), pp. 1815–1851.

Ligon, Ethan and Laura Schechter (2012). “Motives for sharing in social networks”. In: *Journal of Development Economics* 99(1), pp. 13–26.

Macro, David and Jeroen Weesie (2016). “Inequalities between others do matter: evidence from multiplayer dictator games”. In: *Games* 7(2), p. 11.

Marwell, Gerald and David R Schmitt (1972). “Cooperation in a three-person Prisoner’s Dilemma.” In: *Journal of Personality and Social Psychology* 21(3), p. 376.

Marx, Karl and Friedrich Engels (1846). *The Communist Manifesto.* Penguin Classics (2002).

Melville, Herman (1851). *Moby-Dick; or, the Whale.*

OPHI (2017). “Uganda Country Briefing: Multidimensional Poverty Index Data Bank”. In: *Oxford Poverty and Human Development Initiative, University of Oxford* 1, p. 2017.

Panchanathan, Karthik, Willem E Frankenhuis, and Joan B Silk (2013). “The bystander effect in an N-person dictator game”. In: *Organizational Behavior and Human Decision Processes* 120(2), pp. 285–297.

Parfit, Derek (1997). “Equality and priority”. In: *Ratio* 10(3), pp. 202–221.

Rawls, John (1999). *A Theory of Justice; Revised Edition.* Cambridge, MA: Belknap.

Rawls, John (2001). *Justice as Fairness: A Restatement.* Harvard University Press.

Robson, Matthew et al. (2017). “Eliciting the level of health inequality aversion in England”. In: *Health Economics* 26(10), pp. 1328–1334.

Rousseau, Jean-Jacques (1651). *The Social Contract.* Ed. by Christopher Betts. Oxford World’s Classics (2008).

Santos, Maria Emma and Pablo Villatoro (2016). “A multidimensional poverty index for Latin America”. In: *Review of Income and Wealth* 64(1), pp. 52–82.

Schumacher, Heiner et al. (2017). “One, two, many—Insensitivity to group size in games with concentrated benefits and dispersed costs”. In: *The Review of Economic Studies* 84(3), pp. 1346–1377.

Schwarz, Gideon et al. (1978). “Estimating the dimension of a model”. In: *The annals of statistics* 6(2), pp. 461–464.

Skillings, John H and Gregory A Mack (1981). “On the use of a Friedman-type statistic in balanced and unbalanced block designs”. In: *Technometrics* 23(2), pp. 171–177.

Smith, Adam (1790). *The Theory of Moral Sentiments.* Penguin Classics (2009).

Smits, Jeroen and Roel Steendijk (2015). “The international wealth index (IWI)”. In: *Social Indicators Research* 122(1), pp. 65–85.

Sokal, Robert R (1958). “A statistical method for evaluating systematic relationship”. In: *University of Kansas Science Bulletin* 28, pp. 1409–1438.

Sørensen, Thorvald (1948). “A method of establishing groups of equal amplitude in plant sociology based on similarity of species and its application to analyses of the vegetation on Danish commons”. In: *Biol. Skr.* 5, pp. 1–34.

Stone, Richard (1954). "Linear expenditure systems and demand analysis: an application to the pattern of British demand". In: *The Economic Journal* 64(255), pp. 511–527.

Sugiura, Nariaki (1978). "Further analysts of the data by akaike's information criterion and the finite corrections: Further analysts of the data by akaike's". In: *Communications in Statistics-Theory and Methods* 7(1), pp. 13–26.

Tibshirani, Robert (1996). "Regression shrinkage and selection via the lasso". In: *Journal of the Royal Statistical Society. Series B (Methodological)* 58(1), pp. 267–288.

Traub, Stefan et al. (2005). "Friedman, Harsanyi, Rawls, Boulding–or somebody else? An experimental investigation of distributive justice". In: *Social Choice and Welfare* 24(2), pp. 283–309.

Tsuchiya, Aki and Paul Dolan (2007). "Do NHS clinicians and members of the public share the same views about reducing inequalities in health?" In: *Social Science & Medicine* 64(12), pp. 2499–2503.

UBOS, Uganda Bureau of Statistics (2016). *Uganda National Panel Survey (UNPS): 2013-2014, UGA_2013_UNPS_v01_M.* http://microdata.worldbank.org/index.php/catalog/2663/get_microdata. Accessed: 03/11/2016.

Ura, Karma et al. (2012). *A short guide to gross national happiness index.* The Centre for Bhutan Studies.

Vyas, Seema and Lilani Kumaranayake (2006). "Constructing socio-economic status indices: how to use principal components analysis". In: *Health Policy and Planning* 21(6), pp. 459–468.

Ward, Joe H (1963). "Hierarchical grouping to optimize an objective function". In: *Journal of the American Statistical Association* 58(301), pp. 236–244.

Williams, Alan, Aki Tsuchiya, and Paul Dolan (2004). "Eliciting equity-efficiency trade-offs in health". In: *Health Policy and Economics*. Ed. by Peter Smith, Laura Ginnelly, and Mark Schulpher, pp. 64–87.