What is a Fair Innings?

Preferences for Prioritisation by Age and Health

Matthew Robson[†] Owen O'Donnell[‡] Erik Schokkaert[§]
Tom Van Ourti[¶][*]

Work-in-Progress
September 24, 2025

Abstract

The fair innings principle prioritises healthcare for patients who are younger or who would have fewer quality-adjusted life years (QALYs) without treatment. We gauge support for different interpretations and intensities of this prioritisation in a United Kingdom general public sample (n = 230). We use a novel *localised* convex budget-set experiment (observations = 20,700) to estimate participant-specific parameters of social welfare functions (SWFs). There is overwhelming support for the principle, broadly defined. Prioritarian SWFs, with continuously diminishing marginal welfare from QALYs, fit the data of most participants better than linear SWFs with age-weights, or than SWFs that are linear in QALYs up to a fair innings threshold at which there is a downward discontinuity in marginal welfare. We use the estimated parameters to obtain welfare weights by quality-adjusted life expectancy and illustrate their policy consequences. We compare our results with other justifications of age-based weighting in health policy decisions.

Keywords: Healthcare, Resources, Inequality Aversion, Equity, Experiment.

JEL: C90, D30, D63, I14, I38.

[†]Corresponding Author: Erasmus School of Economics, Tinbergen Institute, Erasmus University Rotterdam, Netherlands. Email: robson@ese.eur.nl. Website: www.mrobson92.com.

[‡]Erasmus School of Economics, Erasmus School of Health Policy and Management, Tinbergen Institute, Erasmus University Rotterdam, Netherlands. Email: odonnell@ese.eur.nl.

[§]Department of Economics, KU Leuven, Leuven, Belgium. Email: erik.schokkaert@kuleuven.be

[¶]Erasmus School of Health Policy and Management, Erasmus School of Economics, Tinbergen Institute, Erasmus University Rotterdam, Netherlands. Email: vanourti@ese.eur.nl.

[*]This study was funded by the Erasmus Initiative Smarter Choices for Better Health, The Netherlands and the Wellcome Trust (205427/Z/16/Z). Ethics approval given by the Internal Review Boards of the Erasmus Research Institute of Management and Erasmus School of Economics. We have no conflicts of interest to disclose. All errors remain our own.

1 Introduction

Age-weighting, in which a larger weight is given to the young than to the old, is hotly debated in health policy making in general and more specifically in the actual practice of priority setting in health care (Reckers-Droog et al., 2019; Attema et al., 2022). There are different possible justifications for such age-weighting; one of which is "fair innings" ageism.[1] This is an equity argument: it broadly states that those who, without intervention, would have died young should receive priority over those who would have lived a longer, healthier, life (Harris, 1985; Williams, 1997). While this argument is intuitively appealing, its specifics are more difficult to pin down. The nature and extent of this priority for the young is debatable, and any given interpretation has consequences for policy decisions which affect the distribution of health across society.

There are two broad issues. The first relates to what we understand by "age". The fair innings principle is sometimes interpreted as a prioritisation of the young over the old, i.e. prioritisation on the basis of the age of the patients at the moment of the intervention. The alternative is to argue that prioritisation should, primarily, be a function of the number of Quality Adjusted Life Years (QALYs), since it is the number of QALYs which contributes to lifetime welfare. Age at the moment of the intervention and number of QALYs are related, but distinct concepts, is illustrated by the following example:

> *Imagine three individuals: Peter, Joyce and John. Peter and Joyce are 30 years old, whilst John is 60. Without treatment, Peter and John will die later today, but Joyce will live until she is 70. All live in perfect health throughout their lives.*

[1]Tsuchiya (1999) and Tsuchiya et al. (2003) make a useful distinction between three possible approaches. The first is health maximisation ageism, in which priority is given to survival of the young because young people are further removed from the moment of death and saving them thus leads to a larger (expected) gain in Quality Adjusted Life Years (QALYs). The second is productivity ageism, giving a larger weight to people who are (socially) more productive, i.e. young adults (see Murray and Acharya (1997), and for a critical discussion, Bognar (2008)). The third one is "fair innings" ageism.

If prioritisation depends on age, Peter and Joyce should receive equal priority, and should both be prioritised over John. If, instead, prioritisation depends on QALYs, then Peter should be prioritised over John, who should be prioritised over Joyce. Most proposals about fair innings, including the seminal one in Williams (1997), focus on QALY's, rather than on actual age. Thought-experiments (and empirical experiments) which do not carefully disentangle the two may find intuitions (or evidence) which seem to support the prioritisation of the young, even when age is only an indirect indicator of QALY's.

Second, how should the idea of fair innings be modelled? The philosophical literature distinguishes between *threshold* and *prioritarian* perspectives. The threshold perspective "requires that everyone be given an equal chance to have a fair innings, to reach the appropriate threshold but, having reached it, they have received their entitlement" (Harris, 1985) and those that reached the threshold can, thus, be given lower priority than those who have not reached it.[2] The *prioritarian* perspective, on the other hand, argues that we "should assign different weights to additional life years, but their value should be diminishing as they are added to the lives of older persons" and that these weights should arise from "an increasing, strictly concave function" (Bognar, 2015). This approach was already proposed in Williams (1997) and is formalised by Adler et al. (2021).[3] Social welfare functions (SWFs) can provide representations of these alternative fair innings arguments. *Threshold* functions require a threshold value and a weight for those below this threshold. *Prioritarian* functions include concavity parameters, which determine the rate at which marginal welfare of health gains diminish with quality-adjusted life expectancy.

[2] To be fair, Harris (1985) defends the idea of a threshold only for choices between treatments aiming at prolonging the length of life, and not for treatments improving the quality of life.

[3] Adler et al. (2021) introduce uncertainty and distinguish between ex ante and ex post prioritarianism. In the former, social welfare is a (concave) function of the expected number of QALYs for all individuals, in the latter social welfare is the expected value of the (concave) transformations of the individual QALY's. They show that in both approaches there will be priority for the young. We neglect uncertainty in the main part of the paper, but we will return to this choice in the discussion section 6.

To check whether the fair innings idea is popular among the population, and if so, which formal representation or which parameter values should be used, useful information can be obtained from the preferences of the individuals. Questionnaire studies find mixed results about the acceptance of the idea of fair innings (Tsuchiya et al., 2003; Nord, 2005; Oliver, 2009; Adler et al., 2025). The formulation of the questions and the context in which they are set seem to matter a lot: as an example, letting people choose whether a young or an old patient should get a life saving treatment, generates different answers than letting them choose between different options to improve health-related quality of life. With the exception of Adler et al. (2025), none of these studies distinguishes carefully between the different formal approaches to fair innings.

An extensive experimental literature elicits preference parameters within health-related SWFs to provide evidence of public preferences. Health inequality aversion parameters are estimated within "prioritarian" (Dolan and Tsuchiya, 2011; Robson et al., 2017; Pinho and Botelho, 2018; Hurley et al., 2020; Robson et al., 2024; Robson et al., 2025), or rank-dependent (Hardardottir et al., 2021) SWFs. Other experimental work identifies equity weights relating to age (and severity) (among many others Nord et al., 1996; Johannesson and Johansson, 1997; Rodriguez and Pinto, 2000; Dolan et al., 2005; Reckers-Droog et al., 2019; Reckers-Droog et al., 2021; Attema et al., 2022). Whilst these papers provide extensive empirical evidence on aversion to health inequality, fair innings or age-weighting, they either focus on one functional form or do not have experimental designs which are capable of distinguishing between alternatives.[4]

We contribute to this literature in three ways. First, drawing upon the function proposed by Köbberling and Wakker (2005) in the context of reference dependence and loss aversion, we offer a coherent formalisation of the various philosophical approaches. The loss aversion function of Köbberling and Wakker (2005) covers the threshold and the prioritarian model

[4]Robson et al. (2024) and Robson et al. (2025) are exceptions, but only use "prioritarian" functions.

as special cases. We will focus on the case of absolute inequality aversion (Kolm, 1976). For comparison purposes, we also consider a model of simple linear age weighting, with rank dependent weights (Bleichrodt et al., 2004). In addition to these basic models, we consider some extensions: a *threshold-prioritarian* model which gives, in addition to a concavity parameter, higher weights above a threshold, and an *age-weighted prioritarian* model, which combines a concavity parameter with differential weights by age. We relax the assumption of constant inequality aversion with a *discontinuous prioritarian* model which allows for differential concavity above and below a threshold, and with the even more flexible *expo-power* model (Saha, 1993), allowing for increasing, constant and decreasing absolute inequality aversion.

Second, we propose a novel experimental design - using *localised* convex budget sets - that allows for the estimation of participant-level preferences across these alternative interpretations of the fair innings principle. In our experiment, participants take the role of a social decision maker allocating resources to determine the health of three hypothetical individuals. Across rounds, the current age and resource productivity of these individuals is exogenously varied. The choices participants make reveal their willingness to prioritise individuals by age and the trade-offs they are willing to make between equalisation and maximisation of health. We model participant's choices using a random utility model and estimate preference parameters using maximum likelihood estimation. The design can be extended to other settings (e.g. prospect theory and loss aversion) where discontinuous thresholds are often difficult to identify.

Third, we run the experiment with a UK adult sample ($n = 230$) and provide empirical evidence on the precise form and intensity of the UK public's preferences over fair innings. We identify which models best fit participants' choices, alongside providing estimates and welfare weights for use in health policy evaluation and formulation. We find overwhelming support for the fair innings principle, broadly defined, as opposed to standard health maximisation. Age-weights alone do not best explain the choices of any participant. Prioritarian models fit (much) better than threshold models, and more general forms of prioritarian models - threshold-prioritatian and age-weighted, fit somewhat better for some participants. A substantial minority of the respondents has increasing absolute inequality aversion. We illustrate how our estimated preference parameters can be used in policy decisions, and provide welfare weights which identify the priority given to one individual (or group) compared to another.

In section 2 we introduce the alternative formal models of fair innings that are estimated in our experimental work. We then present our experimental design (section 3), our estimation procedure (section 4) and our results (section 5). Limitations and possible extensions of our approach, and the relationship with previous work, is discussed in section 6. Section 7 concludes.

2 Theory: formal models of fair innings

We assume that, from the perspective of a social decision maker (SDM), welfare generated by the population distribution of health is

$$W = \frac{1}{N}\sum_{i}^{N} \omega_i U(h_i), \tag{1}$$

where $U(h_i)$ is utility the SDM assigns to the health (h_i) of individual i, $\omega_i \geq 0$ is a Pareto weight and N is the population size.[5] Health is number of quality-adjusted life years (QALYs). Current age will be represented as y_i. If we assume that all years lived have been lived in perfect health, then $h_i = y_i$. This is basically the simplification that will be

[5]We will show that this approach in terms of social maximization can capture many alternative interpretations of the fair innings idea. It cannot represent very well the so-called prudential lifespan account of Daniels (1988), however, because Daniels reformulates the problem of justice between different age groups as an individual decision making problem about how to allocate resources over different periods of life.

assumed in the empirical work. We will return to the implications of this choice in sections 5.3 and 6.

We distinguish between interpretations of the fair innings principle through specifications of $U(h_i)$ and ω_i. We consider specifications that are reasonable representations of the most common interpretations of the fair innings that have been proposed in the literature. We specify *age-weighted, threshold* and *prioritarian* models, alongside *threshold-prioritarian* and *age-weighted prioritarian* generalisations of these models. *Discontinuous prioritarian* and *expo-power* models relax the assumption of constant absolute inequality aversion.[6]

2.1 Age-weighted models

The simplest interpretation of the fair innings idea, often used for health policy, is that priority should be given to the young.[7] This can be represented in eq. (1) by assuming that welfare is linear in QALYs, but with differential weights by current age. We formalize this approach with an age rank-dependent model:[8]

$$\omega_i = \frac{[P(Y \geq y_i)]^\delta - [P(Y > y_i)]^\delta}{N \times P(Y = y_i)}, \qquad U(h_i) = h_i, \tag{2}$$

where y_i is current age, $P()$ is a probability and individuals are ordered from youngest to oldest, $y_i \leq y_{i+1} \, \forall_i$. With $\delta > 1$, ω_i decreases monotonically with increasing rank in the age distribution, giving greater weight to younger individuals. As $\delta \to \infty$, the weights approach zero for all but the youngest individual. This weighting scheme is sufficiently flexible to also

[6]In line with the dominant welfare economic approach, we focus in the main text on health outcomes rather than on the health *gains* achieved by an intervention. We discuss the gains model in Appendix D.8.

[7]Health technology assessment in Norway and the Netherlands implement some form of age-weighting (see respectively Ottersen et al. (2016), and Reckers-Droog et al. (2018)).

[8]These rank-dependent weights are derived from the Generalised Gini index (Donaldson and Weymark (1980); Donaldson and Weymark (1983)). In our setting, the natural choice is to rank individuals by current age. Other rank-dependent models in the health economics literature rank individuals by QALYs (Bleichrodt et al., 2004) or income (Wagstaff et al., 1991).

capture preferences of SDMs who wish to prioritise older individuals.[9] With $0 < \delta < 1$, the weights increase monotonically with increased age rank. With $\delta = 1$, the weights are constant and there is no prioritisation by age.

2.2 Threshold and prioritarian models

In the philosophical literature the main competitors to represent the idea of fair innings are the threshold and the prioritarian approaches.

Threshold.—One interpretation of the fair innings principle is that QALY gains below and above some threshold QALY level, τ, should get different weights (Harris, 1985):[10]

$$\omega_i = \begin{cases} \beta & \text{if } h_i < \tau \\ 1 & \text{if } h_i \geq \tau, \end{cases} \qquad U(h_i) = h_i \,. \tag{3}$$

with as a natural possibility $\beta > 1$.

Prioritarian.—The alternative approach is the prioritarian one (Williams, 1997; Bognar, 2015; Adler et al., 2021). Prioritisation of the less healthy can be captured with constant weights if utility is strictly concave in QALYs, such that they generate diminishing marginal welfare:[11]

$$\omega_i = 1\,, \qquad U(h_i) = \begin{cases} \dfrac{1 - e^{-\mu h_i}}{\mu} & \mu \neq 0 \\ h_i & \mu = 0. \end{cases} \tag{4}$$

[9]Empirical work (e.g. Adler et al. (2025)) indeed finds that a substantial minority of respondents give priority to the old.

[10]This model is adapted from the loss aversion literature (Köbberling and Wakker, 2005). A simpler interpretation of the *threshold* model, akin to the poverty-gap model, is presented in Appendix D.9.

[11]We restrict attention to a Kolm (1976) SWF, as it is nested within the more general function suggested by Köbberling and Wakker (2005) which allows for weights and concavity parameters to vary below and above a reference point (or threshold). An alternative function would be a power function, (i.e. the Atkinson (1970) SWF). However, this function has issues for values close to zero (i.e. it has extreme derivatives).

A SDM aiming to maximise aggregate QALYs has $\mu = 0$. As μ increases, there is greater willingness to sacrifice the QALYs aggregate to reduce absolute inequality by increasing the QALYs of those who have less. As $\mu \to \infty$, preferences approach maximin (Rawls, 1971).[12]

2.3 Extensions of the prioritarian model

The different intuitions underlying age-weighting, threshold and prioritarian approaches can be combined by extensions of the prioritarian model.

Age-Weighted Prioritarian.— Combining age-weighting with concavity in health allows for a SWF which exhibits diminishing marginal welfare in QALYs, alongside increasing, constant or decreasing weights with age-rank, as follows:[13]

$$\omega_i = \frac{[P(Y \geq y_i)]^\delta - [P(Y > y_i)]^\delta}{N \times P(Y = y_i)}, \qquad U(h_i) = \begin{cases} \dfrac{1 - e^{-\mu h_i}}{\mu} & \text{if } \mu > 0 \\ h_i & \text{if } \mu = 0. \end{cases} \tag{5}$$

Threshold Prioritarian.—Including weights, β, which give differential weights above a threshold, τ, alongside the concavity parameter μ, allows for a discontinuity in marginal weights at the threshold:

$$\omega_i = \begin{cases} \beta & \text{if } h_i < \tau \\ 1 & \text{if } h_i \geq \tau, \end{cases}, \qquad U(h_i) = \begin{cases} \dfrac{1 - e^{-\mu_k (h_i - \tau)}}{\mu_k} & \mu_k \neq 0 \\ h_i & \mu_k = 0, \end{cases} \tag{6}$$

where $\mu_k = \mu$ if $h_i \geq \tau$ and $\mu_k = -\mu$ if $h_i < \tau$.

[12]We use the terminology "prioritarianism" in a loose way. Following the stricter and more correct interpretation, proposed by Adler (2019) (see also the many applications in Adler and Norheim (2022)), eq. (1) is just a version of generalised utilitarianism, even if $U(.)$ is a concave function. Eq. (4) can only be interpreted as prioritarian if we interpret h as utility and $U(.)$ as a concave transformation of this utility.

[13]Note that this functional form is equivalent to that of Bleichrodt et al. (2005) or Robson et al. (2024), with age ranks rather than QALY or income ranks.

2.4 Non constant absolute inequality aversion

It is also possible to relax the assumption of constant absolute inequality aversion.

Discontinuous prioritarian.— One approach is to allow for a discontinuity of μ at the threshold, τ:

$$\omega_i = 1\,, \qquad U\left(h_i\right) = \begin{cases} \dfrac{1 - e^{-\mu_k (h_i - \tau)}}{\mu_k} & \mu_k \neq 0 \quad \begin{cases} \mu_k = -\mu_b & \textit{if } h_i < \tau \\ \mu_k = \mu_a & \textit{if } h_i \geq \tau \end{cases} \\ h_i - \tau & \text{otherwise.} \end{cases} \tag{7}$$

This allows absolute inequality aversion (AIA), $^{-U''}/_{U'}$, to remain constant, decrease or increase on crossing the threshold. In the first case, $AIA = -\mu_b = \mu_a$, and utility is equally concave (inequality averse) on each side of the threshold if $\mu_b < 0$ and $\mu_a > 0$. With these signs of the two parameters, $-\mu_b > \mu_a > 0$ implies a discontinuous drop in AIA on crossing the threshold from below, while $-\mu_b < \mu_a > 0$ corresponds to a discontinuous rise in AIA at the threshold. If $\mu_b > 0$ and $\mu_a > 0$, then utility is convex - inequality seeking - below the threshold and concave above it.

Expo-power.—A more general approach for testing whether participants exhibit increasing, constant or decreasing absolute inequality aversion is to use the expo-power functional form of Saha (1993):

$$\omega_i = 1\,, \qquad U(h_i) = \left(\frac{1 - e^{-\mu\left(h_i^{1-\gamma}\right)}}{\mu}\right). \tag{8}$$

The parameter γ determines whether we have decreasing ($\gamma > 0$), constant ($\gamma = 0$) or increasing ($\gamma < 0$) absolute health inequality aversion. We assume this to be bolted onto the standard prioritarian function, with no threshold, equal weights and one concavity parameter, alongside the γ parameter.

2.5 Interpretation

Figure 1 illustrates the intuition behind the three basic models of social welfare, with the top panel showing (normalised) welfare from an individual's health and the bottom panel showing social indifference curves between the health of two individuals. The intuition for the other models is attained by the appropriate combination of the models.

Figure 1: Social Utility and Social Welfare Curves

(a) Welfare from individual health

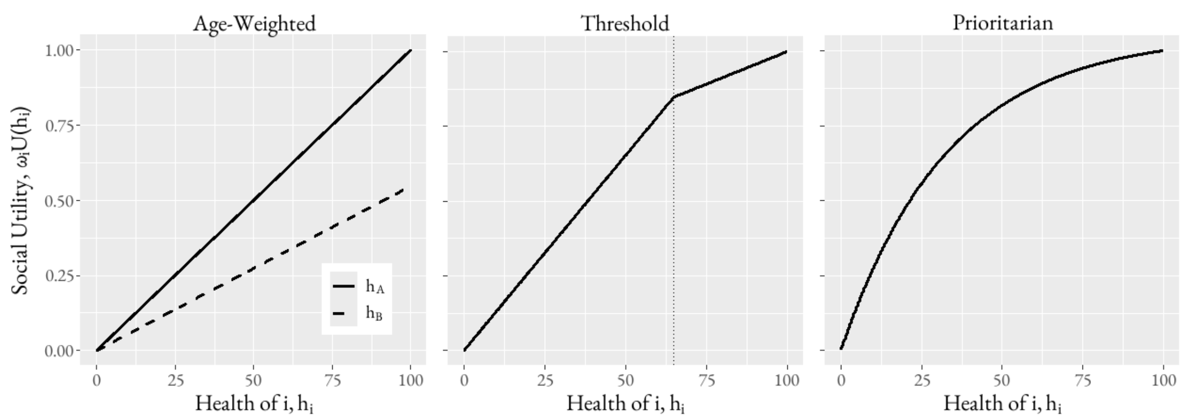

(b) Social indifference curves

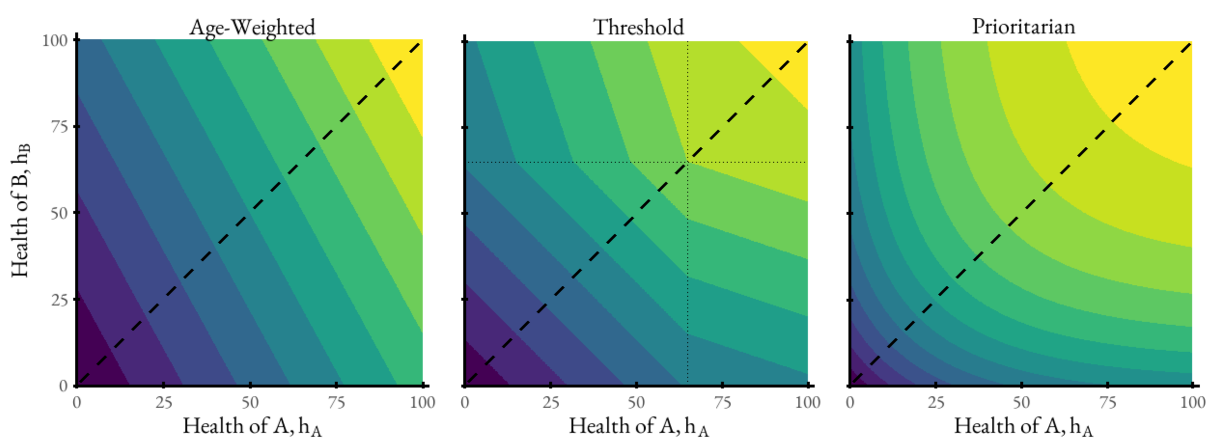

The age-weighted function shows the case in which greater weight is given to the health of a younger individual (A). The threshold model is drawn with a threshold value at 65 QALYs. Threshold and prioritarian preferences exhibit diminishing marginal welfare from health, although this is not continuous in the threshold model. The slopes of indifference curves correspond to social marginal welfare weights (SMWWs) – the rate of trade-off between the health of two individuals holding welfare constant – that reflect both the Pareto weights and concavity of utility. The age-weighted model allows weights to be non-symmetric, providing higher (or lower) weight for individuals with a lower current age. The threshold SWF only gives higher priority to an individual if they are below the threshold while the other is above. The prioritarian SWF gives a constantly-increasing higher weight to the individual with the lowest health, as the absolute health gap between two individuals increases. These models provide alternative social welfare valuations across the health distribution, and thus provide alternative prioritisation of individuals with differing health levels and ages.

3 Experiment

3.1 Design[14]

In an online experiment, participants allocate constrained resources between three hypothetical individuals to determine their lifetime QALYs. Individuals can differ by QALYs gained from a given amount of resources (productivity) and their current age. Current age is defined as the number of years they have lived for (in full health). Allocations reveal a participant's willingness to sacrifice maximisation of aggregate QALYs for less inequality and their willingness to prioritise individuals by age; together this provides evidence on the extent and nature of their support for the fair innings principle.

In each round, a participant has a resource *budget* (m) to allocate between the three individuals. Each *individual* (i) is characterised by their *current age* (y_i) and by a productivity factor, aka the *multiplier* (p_i). By allocating *resources* (x_i) to an individual, participants

[14]A more detailed description of the experimental design, including the complete text of the instructions and the tutorial, can be found in Appendix A.

extend the number of quality-adjusted life years that individual will live for (by $x_i p_i$). If no resources are allocated to an individual, they die having lived for y_i years in full health. The *health* ($h_i = y_i + p_i x_i$) of each individual is thus determined by an individual's current age and the product of their multiplier and a participant's allocations to them.

Figure 2 shows the experimental interface. Participants use sliders to allocate resources to the three individuals. Current age levels are shown by grey bars. Interactive graphics display resources allocated to each individual and the resulting health (grey plus black bars). Measures of resource and QALY gaps (max - min) between individuals, the total QALYs and the remaining budget, which must be exhausted, are also shown.

Figure 2: Experiment Interface

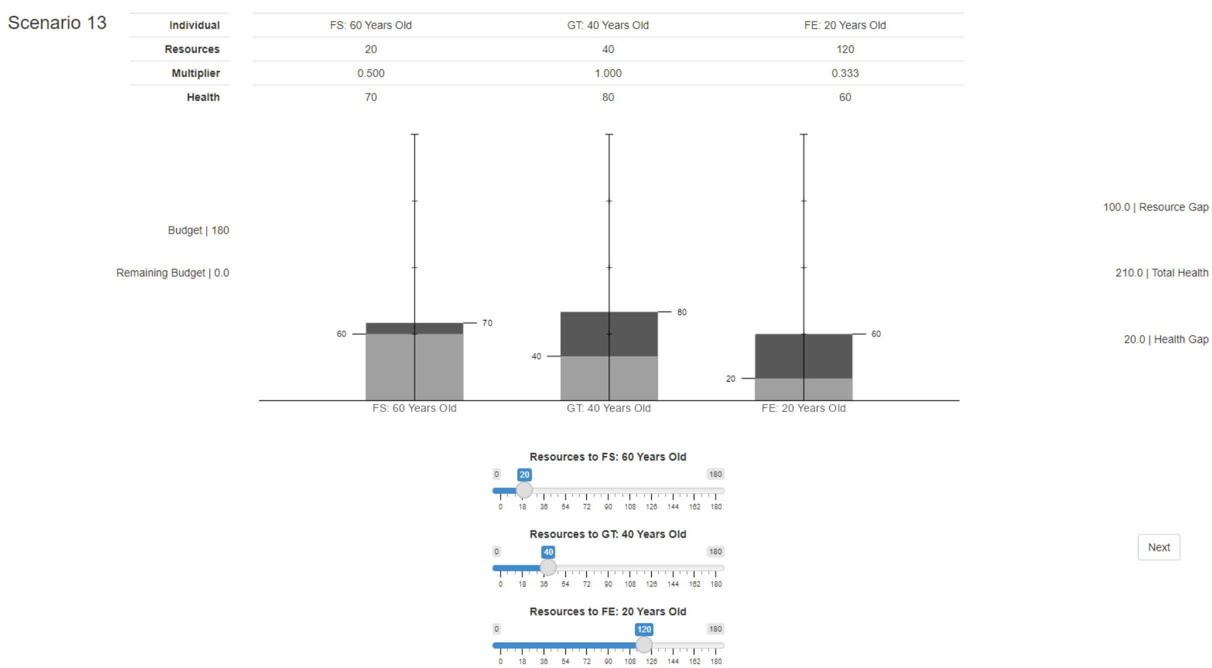

3.2 Treatments

Participants make decisions across three within-subject treatments. Each treatment has 10 rounds. Budgets, multipliers and current age levels are exogenously varied across treatments and rounds. All treatments use the same sets of multipliers in the respective 10 rounds (Table A1). In one round, the multipliers are equal: $p_i = 1 \forall i$. In three rounds, one of the three individuals has a lower multiplier: $\mathbf{p} = (1, 1, 1/3)$. In six rounds, each individual has a different multiplier: $\mathbf{p} = (1, 1/2, 1/3)$. For each screen position, the distribution of the multiplier over the 10 rounds is the same, which ensures that screen position does not bias comparisons. Within each treatment, the order of the 10 rounds is randomised over participants. In each round for each participant, a set of initials is drawn randomly to label the individuals.

Variation in the budget and current age levels depends on the treatment: *budget*, *age-equal* or *age-varying*. In each round of the *budget* treatment the budget was randomly drawn from $m \in \{180, 240, 360\}$, whilst the current age y_i was zero for all individuals across all rounds. In each round of the *age-equal* treatment, the same current age was randomly drawn for all three individuals from $y_i \in \{0, 5, 20, 40, 60, 80\}$, such that $y_i = y_j \; \forall i, j$, whilst the budget, $m = 180$, was held constant across all rounds. In each round of the *age-varying* treatment, a current age level was randomly drawn, without replacement, for each individual from $y_i \in \{5, 20, 40, 60, 80\} \, \forall i$, such that $y_i \neq y_j \, \forall i \neq j$, whilst the budget, $m = 180$, was held constant across all rounds. Further details are found in Table A1.

Figure 3 shows simplifications of the convex budget set problems the participants face in each treatment. Each line traces the budget set of one decision problem. It shows the feasible distributions of QALYs between individuals i and j if half of the available budget were allocated to individual k.[15] The relative multipliers, $-p_j/p_i$, determine the slope of each line. Its length increases with the size of the budget (distinguished by colour on the left). Current ages determine the end points, e.g. $min(h_i) = y_i$ (distinguished by colour in the middle and right). The importance of this design becomes clearer, when considering Figure 3 in conjunction with Figure 1, which shows how, depending on the functional form

[15] A participant's feasible options are actually on a plane $(\mathbf{h_i}, \mathbf{h_j}, \mathbf{h_k})$ rather than a line.

of the social welfare function, the slopes of indifference curves differ across localised regions. Optimal allocations are such that the social indifference curves are tangential to the budget set.

Figure 3: Convex budget sets by treatment

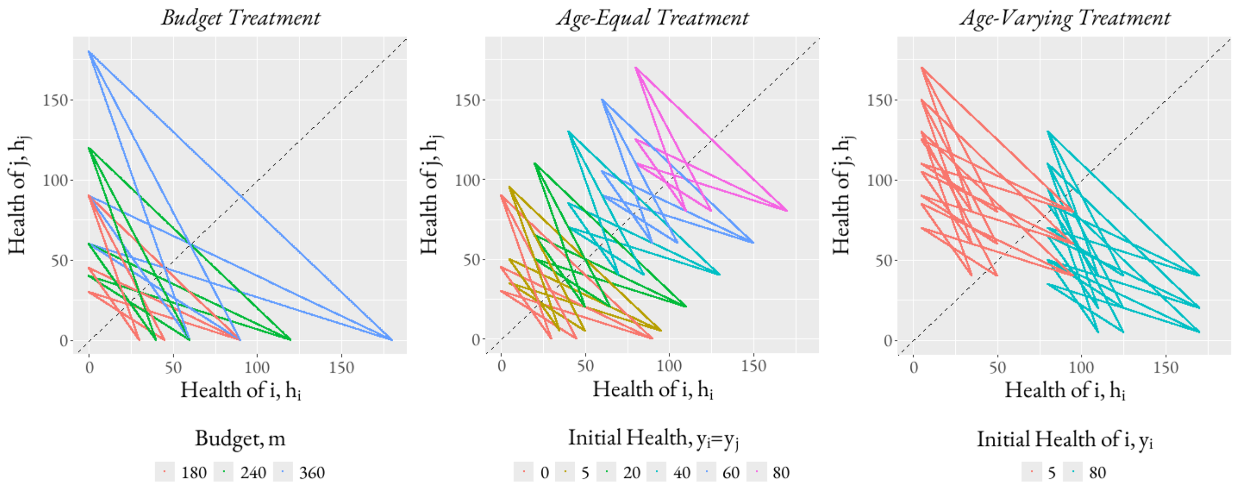

Varying the multipliers identifies the trade-off a participant is willing to make between maximising aggregate QALYs and reducing inequality. When $p_i \neq p_j$, the QALY aggregate is maximised at an extreme. To reduce inequality, the participant must allocate to move along the budget line towards the 45 degree line. This holds irrespective of the budget and current age levels. In the budget treatment, increasing the budget changes the maximum feasible QALY levels, which allows identification of whether inequality aversion changes with the maximum. The age-equal treatment is similar, in that increasing current age levels increases the maximum feasible QALYs. However, in this case, the minimum feasible QALYs also increase and the budget remains constant. The age-varying treatment shifts the feasible regions above (red) and below (blue) the line of equality. Together these *localised* budget sets allow identification of marginal rates of substitution between individuals whose feasible QALYs are both low, both high, or one high and one low.[16]

[16]Note that the relative multipliers and budget levels are orthogonal in the budget treatment, whilst the relative multipliers and current age levels are orthogonal in the age-equal and age-varying treatments.

3.3 Sample

Participants were recruited through Prolific. They were intended to be representative of the UK adult population by sex, age and ethnicity. The experiment was conducted in two sessions in the period April 7-30, 2022. Of the 267 participants who completed the first session, 23 (8.6%) did not complete the second session. We drop an additional 14 participants who gave incorrect answers to at least 3 of 5 comprehension questions after a tutorial in the first session. This leaves an analytical sample of 230 participants with complete data from both sessions. This sample is representative with respect to sex but is younger and has a lower proportion of white individuals (see Table B2).

The 230 participants made choices in 30 rounds across three treatments, with each round involving allocations to three individuals. This gives a total of 20,700 (= 230 × 30 × 3) observations. Table 1 summarises the data used.

Table 1: Summary of experimental data

Variable	**Notation**	**Definition**	**Mean**	**Range**	**Obs.**
Experimental					
Budget	m	m	206.157	[180-360]	20,700
Resources	x_i	x_i	68.719	[0-321.2]	20,700
Resource Share	$\tilde{x}_i$	x_i/m	0.333	[0-1]	20,700
Multiplier	p_i	p_i	0.7	[0.333-1]	20,700
Relative Multiplier	$\tilde{p}_i$	$p_i/\sum p_i$	0.333	[0.143-0.545]	20,700
Age	y_i	y_i	25.036	[0-80]	20,700
Relative Age	$\tilde{y}_i$	$y_i/\sum y_i$	0.333	[0.034-0.762]	20,700
Health	h_i	$y_i + x_i p_i$	43.522	[0-274.5]	20,700
Health Share	$\tilde{h}_i$	$h_i/\sum h_i$	0.333	[0-1]	20,700

Note: Number of participants is 230, across 30 rounds. $N = 3$ is the number of individuals in each round.

The median completion time was 26.9 minutes in the first session and 24.5 minutes in the second. Participants were paid £3.50 for the first session and £5 for the second. The average payment was £9.98 per hour over the two sessions.

4 Estimation

We use a random utility model to estimate participant-specific parameters of each social welfare function specification. We assume a participant allocates resources, $\mathbf{x} = (x_1, x_2, x_3)$, with the aim of maximising Eq. (1) subject to the budget constraint, $m = \sum_i^N x_i$, and non-negativity constraints, $0 \leq x_i \, \forall i$. However, the perception of welfare generated by an allocation, $V(\mathbf{x})$, is allowed to diverge randomly from actual welfare: $V(\mathbf{x}) = W(\mathbf{x}) + \epsilon$, where the last term is a random error.

We model the probability of a participant choosing a particular allocation, $\mathbf{x}_q$, from all feasible allocations, G: $P\left(\mathbf{x} = \mathbf{x_q}\right) = P\left(V(\mathbf{x_q}) = max\left\{V(\mathbf{x}_1), V(\mathbf{x}_2), \ldots, V(\mathbf{x}_G)\right\}\right)$. We assume ϵ are i.i.d. mean-zero, Gumbel distributed errors: $\epsilon \sim Gumbel(0, \sigma)$, where σ is a noise parameter (McFadden, 1972). Then, the probability of choosing $\mathbf{x_q}$ is given by a conditional logit model: $P\left(\mathbf{x} = \mathbf{x_q}\right) = \frac{\exp\left[\frac{W(\mathbf{x_q})}{\sigma}\right]}{\sum_{g=1}^{G} \exp\left[\frac{W(\mathbf{x_g})}{\sigma}\right]}$. For each participant, the log-likelihood across all rounds $r = 1, 2, ..., R\,(=30)$ is

$$LL = \sum_{r=1}^{R} \sum_{g=1}^{G_r} \log\left(\mathfrak{I}\{\mathbf{x_r} = \mathbf{x_{rq}}\}\left(\frac{\exp\left[\frac{W(\mathbf{x_{rq}})}{\sigma}\right]}{\sum_{g=1}^{G_r} \exp\left[\frac{W(\mathbf{x_{rg}})}{\sigma}\right]}\right)\right), \tag{9}$$

where $\mathfrak{I}\{\}$ is the indicator function. We maximise this log-likelihood to obtain participant-specific estimates of parameters for each SWF specification: β and τ for the threshold model (3) , μ for the prioritarian model (4), δ for the age-weighted model (2), μ, β and τ for the threshold prioritarian (6), δ and μ for the age-weighted model (5), μ_a, μ_b and τ for the discontinuous prioritarian model (7), μ and γ for the Saha-model (8), as well as the noise parameter (σ) in each case. We do this by numerically solving for parameter values that maximise the log-likelihood. We use the likelihood values to calculate the corrected Akaike Information Criterion ($AICc$) and compare goodness-of-fit across the SWF specifications.[17]

5 Results

5.1 Raw data analyses

We estimate, for each participant i, across all rounds of decision problems t, the following two regression models of shares of resources $\left(\tilde{x}_i = {}^{x_i}/\sum_j^3 x_j\right)$ and health (QALYs) $\left(\tilde{h}_i = {}^{h_i}/\sum_j^3 h_j\right)$, respectively, on the relative multipliers $\tilde{p}_i = \left({}^{p_i}/\sum_j^3 p_j\right)$ and the relative current age levels $\tilde{y}_i = \left({}^{y_i}/\sum_j^3 y_j\right)$ of the recipients:

$$\tilde{x}_{it} = \alpha_{0i} + \alpha_{1i}\tilde{p}_{it} + \alpha_{2i}\tilde{y}_{it} + \varepsilon_i \tag{11}$$

$$\tilde{h}_{it} = \beta_{0i} + \beta_{1i}\tilde{p}_{it} + \beta_{2i}\tilde{y}_{it} + \varepsilon_i \tag{12}$$

The coefficients identify the expected change in the resource and health shares in response to higher relative multipliers and relative age, for each participant. Figure 4 plots these coefficients. There is considerable individual heterogeneity in the responses to the multipliers and current age levels. Still, a vast majority of participants give significantly (5% level) less resources to individuals who are younger (92%) and have higher multipliers (84%). More than four fifths (82%) give less resource both to those are older and to those with higher

[17]Before estimating, we normalise the welfare generated by each individual, $W_i = \omega_i U(h_i)$ to ensure that it lies in the unit interval for the majority of choices:

$$\tilde{W}_i = \frac{\omega_i U(h_i) - \omega_0 U(0)}{\omega_{100} U(100) - \omega_0 U(0)}, \tag{10}$$

where ω_0 and ω_{100} are weights consistent with the minimum ($h_i = 0$) and "maximum" ($h_i = 100$) values, respectively. This affine transformation has no effect on the ranking of alternative allocations. It maintains the same scaling of welfare across alternative SWF specifications, which is useful for numerical estimation and ensures that σ is comparable across models.

multipliers.[18] However, these reductions are not sufficient to remove all inequalities in health between young and old or more or less productive, as older individuals and individuals with higher relative multipliers, still have significantly higher health shares.[19]

Figure 4: Participant-level Coefficients

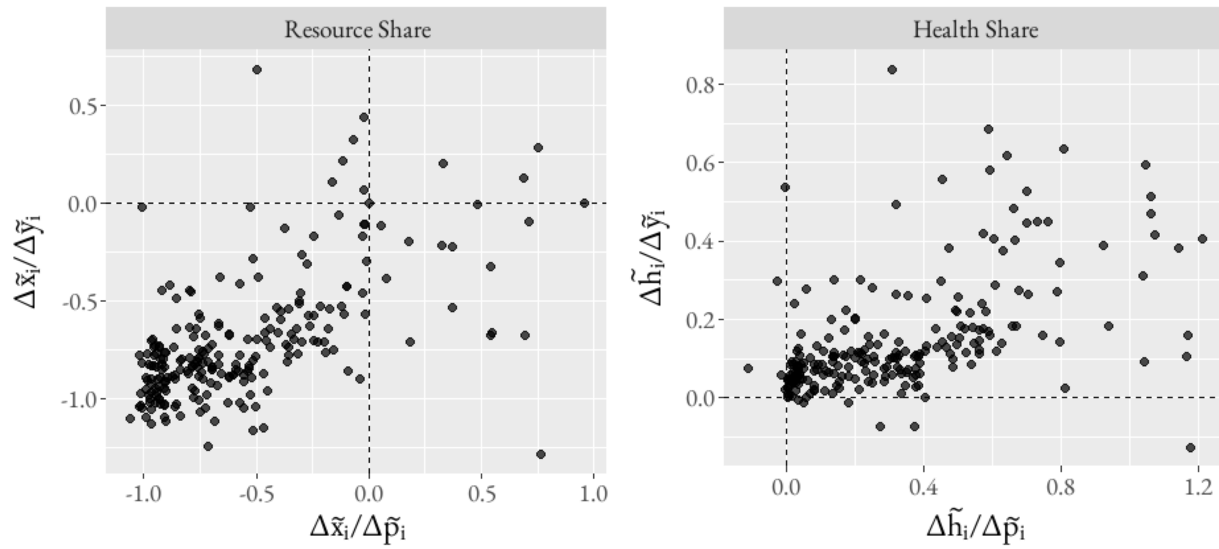

This general picture is confirmed by Figure 5. Using data from the age-varying treatment, this figure plots kernel densities of resource shares and health shares conditional on age and relative multipliers. A share of a third always indicates receipt of the mean resources or health within a round. As already revealed by the regressions, resource shares decrease with age, while health shares increase with age. The distributions do show heterogeneity in each of these relationships. Across all possible multipliers, about 86% of allocations give individuals who have lived for only 5 years (in full health) more than a mean share of resources ($\tilde{x}_i > 1/3$), but they get less than the mean share in around 12% of allocations. These individuals still

[18]Appendix C.1 presents regressions of individual coefficients to provide a sense check and identify an associations with socio-demographic characteristics. Reassuringly, results show positive associations with self-reported views about the importance of reducing health inequality versus maximising aggregate health and the stated aim to prioritise younger or older individuals. We find no significant associations socio-demographic characteristics, with the exception of political affiliation. Interestingly, age of the participant is not significantly associated with prioritisation by age.

[19]These results are consistent with those reported previously with different data (Robson et al., 2024).

end up with a less than average share of health ($\tilde{h}_i < 1/3$) in about 65% of allocations, but in about one third they get a larger than mean share. At the other extreme, individuals who have already lived for 80 years in full health get less than an equal share of resources in about 91% of allocations. In more than one fifth (23%) of the allocations, they get no resources at all. In about three quarters of all allocations, these individuals still end up with more than a mean share of health, but they actually get less than a mean share in 23% of allocations.

Figure 5: Resource and health shares by age and relative multiplier

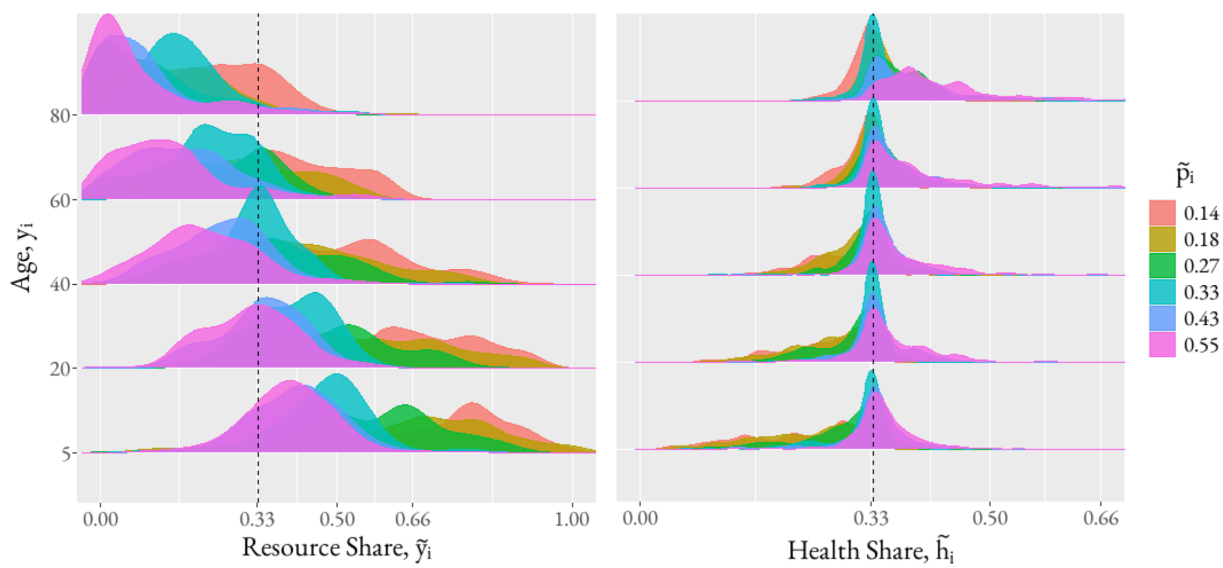

Further non-structural results (about the evidence of discontinuities and the health differences in the chosen allocations) are found in Appendix C.

5.2 Model estimates

We estimate participant-level preference parameters within each alternative social welfare functions and assess the goodness-of-fit of each model. Table 2 and Figure 6 summarise these results, and our main findings are discussed in detail below. Table 2 provides an overview of the estimated participant-level preference parameters of the different models. Median preference parameters are shown, alongside the 5th and 95th percentiles. Full distributions of estimated participant-level preferences parameters and noise parameters are shown in Appendix D and Appendix E, respectively. Figure 6 shows the goodness-of-fit of each model. In the left panel, Mean Proportional Likelihoods (MPL) show the how likely each models fits choices when compared to a uniform random draw.[20] In the right panel, the percentage of participants for whom each model fits "best" is shown.[21]

Table 2: Estimated Participant-Level Preference Parameters: Median and Percentiles

	Age-Weighted	*Threshold*		*Prioritarian*
	δ	β	τ	μ
Median $[5th, 95th]$	30.898 $[1.24, 35909]$	86.246 $[4.58, 2353034]$	60.45 $[32.54, 79.83]$	0.052 $[0.01, 0.11]$

	Age-Weighted Prioritarian		*Threshold Prioritarian*		
	δ	μ	β	τ	μ
Median $[5th, 95th]$	1.151 $[0.44, 2.31]$	0.051 $[0.02, 0.11]$	2.534 $[0.20, 45.89]$	62.9 $[54.16, 90.26]$	0.056 $[0.01, 0.12]$

	Discontinuous Prioritarian			*Expo-Power*	
	$-\mu_b$	τ	μ_a	μ	γ
Median $[5th, 95th]$	0.050 $[0.00, 0.11]$	64.814 $[58.52, 92.39]$	4.401 $[0.02, 72.66]$	0.056 $[0.00, 0.30]$	-0.026 $[-0.87, 0.41]$

Result 1 *We find overwhelming support for the fair innings principle, broadly defined.*

The estimated parameters indicate that, regardless of the assumed model, the majority of participants' preferences align with the fair innings principle. The median parameters of $\delta = 30.898$, $\beta = 86.246$ and $\mu = 0.052$, indicate substantial weight for those who are younger, below the threshold or have lower QALYs, respectively. For the *age-weighted* model there

[20] Where the Mean Proportional Likelihood (MPL), is the average Proportional Likelihood (see eq.9) across all participants, $r \in R$, and rounds, $t \in T$: $MPL = \frac{1}{TR}\sum_r^R \sum_t^T \frac{L_{rt}}{L_{rt}+L_{rt}^{UNI}}$.

[21] The "best" model is that which has the highest AICc, where AICc $= -2\log L + 2k + \frac{2k(k+1)}{n-k-1}$, where L is the likelihood of a given model, pooled across all participants, k is the number of estimated parameters, and n the number of observations.

are 95.2% of participants who give priority to the young ($\delta > 1$), and for the *threshold* and *prioritarian* models the vast majority give priority to those below threshold (99.6% with $\beta > 1$) or those who are worse-off (98.3% with $\mu > 0.01$).

Similar results hold with the extended models. Median parameters from the *age-weighted* prioritarian model indicate a (mild) priority to the young ($\delta = 1.151$) combined with priority to those with lower QALYs ($\mu = 0.051$), whilst those from the *threshold prioritarian* model show higher weights ($\beta = 2.534$) for those below the threshold ($\tau = 62.9$) in addition to the priority to the worse off ($\mu = 0.056$). Median concavity parameters in both the *discontinuous prioritarian* and *expo-power* models also show this support.

Figure 6: Goodness of Fit of Alternative Models

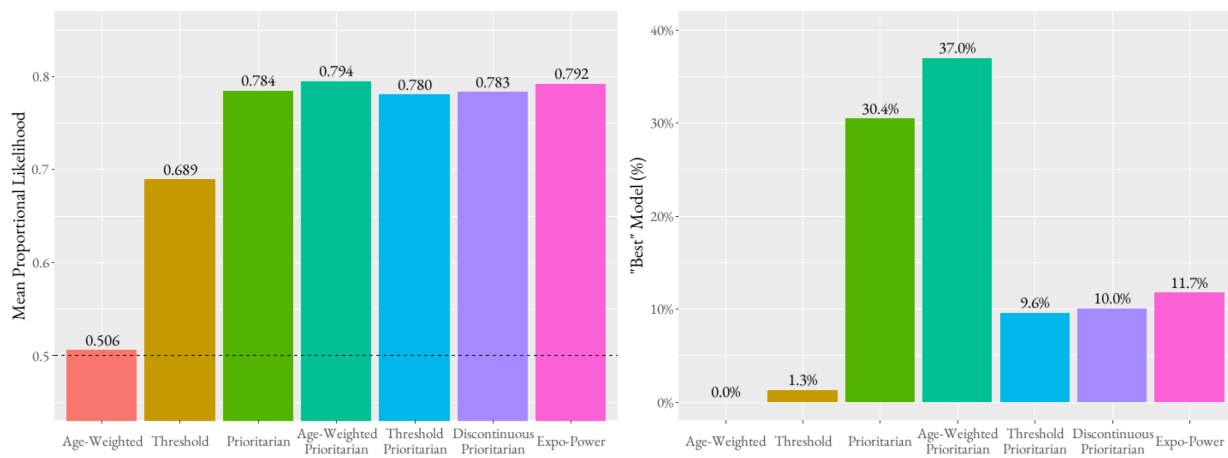

Note: Mean Proportional Likelihoods (MPL) are shown for each model. An MPL of 0.5 shows that the model does as well as a uniform random draw, higher values show a superior fit. The "best" model is that which has the lowest participant-level AICc, amongst the seven alternative models, for each participant.

Result 2 *Age-weighted linear models do not work.*

As shown in Figure 6, the linear age-weighted model provides the worst fit, among all models compared. Indeed, the MPL of 0.506 shows that choices from a uniform random draw would be almost as likely to fit participants choices as the age-weighted model. Moreover, there is not a single participant for whom this model provides the best fit. The age-weighted models, often used in practice for distributional cost-effectiveness analysis, are therefore rejected in our sample.[22]

Result 3 *The prioritarian model does much better than the threshold model.*

As we have seen, the philosophical debate focuses on the choice between models with a specific threshold and so-called prioritarian models. Figure 6 shows that the latter approach is the clear winner in our sample. The MPL is 0.689 for the *threshold* model compared to 0.784 for the *prioritarian* model. Moreover, when comparing which model fits each individual "best" we find that the threshold model only best fits the choices of 1.3% of participants, with the remaining models which include some form of concavity through μ. Indeed, if we were to only compare the simpler *age-weighted*, *threshold* and *prioritarian* models, there are 0.0%, 8.7% and 91.3% of participants for whom these models best fit, respectively. Continuously diminishing marginal weights by QALYs appear to fit choices better than a discontinuous kink at a fair innings threshold.

Result 4 *Extensions of the prioritarian model do not add very much.*

Extending the prioritarian model by either introducing age-weighting (eq. (5)) or a threshold (eq. (6)) adds flexibility and does improve fit. However, these gains are modest with $MPLs$ remaining close to 0.784, at 0.794 and 0.780. There is heterogeneity in which model provides the best fit for each individual, with 30.4%, 37.0% and 9.6% of participants classified into the *prioritarian*, *age-weighted prioritarian* and *threshold prioritarian* models, respectively.

In the age-weighted prioritarian model, the distribution of μ is similar to the *prioritarian* model, with a median of 0.051. We find higher weights for initially *younger* individuals

[22]The even more popular cost-effectiveness approach without age-weighting is even more decisively rejected. Table E6 shows substantially lower log-likelihoods than any other model, with no participant's choices being best explained by this model.

($\delta = 1.151$), but again with extensive heterogeneity. Note that the median value of δ is much smaller than the one in the linear age-weighted model, because its "distributional" role is taken over by the concavity. This also has to be taken into account for the interpretation of the distribution of δ. Across all participants, 26.09% can be characterised as giving priority to the young ($\delta > 1$), whilst 23.04% give priority to the old ($\delta < 1$). Interestingly, but not surprisingly, we find that those who prioritise the young generally have a higher μ (median: 0.079) compared to those who prioritise the old (median: 0.037). It is important to note that no participants have $\mu = 0$. This confirms the lack of support for the *linear age-weighted* models, often used in policy decisions.

For the threshold prioritarian model, the median threshold of 62.9 is very close to the threshold in the simple threshold model (60.5), the median value for the concavity parameter μ (0.056) is close to that in the simple prioritarian model (0.052). The estimated slope parameter β (median value 2.534), however, is much smaller than in the simple threshold model (86.25). The explanation is the same as for the parameter δ in the age-weighted approach. Introducing concavity is the most important feature to capture distributional considerations. The parameters δ and β introduce some additional ("second order") flexibility to refine the fit of the model for the detailed response patterns of the participants. However, introducing additional complexity in the prioritarian model only leads to a relatively minor gain in fit.

Result 5 *A fraction of the respondents has increasing absolute inequality aversion.*

The prioritarian models discussed until now all have the parameter of absolute inequality aversion constant. Relaxing this assumption, the *discontinuous-prioritarian* model (7) and the *expo-power* model (8) provide a better fit than the *prioritarian* model for some individuals. Indeed, the expo-power model even has the best fit, by log-likelihood and AICc, of all the models we have estimated (see Table E6).

For the *discontinuous-prioritarian* model we find a single peaked distribution of τ, with a median of 64.81. This is again close to the threshold estimated in the simpler models with a threshold. The distribution of concavity parameters below this threshold ($-\mu_b$) is similar to the estimated μ in the *prioritarian* model, with a median of 0.050. However, we find that the μ_a parameters are much higher, with a median of 4.401. We find a majority of participants (73.04%) whose choices are best explained by *constant* absolute inequality aversion (i.e. the traditional *prioritarian* model), but a sizeable minority (26.09%) with higher concavity parameters above the threshold.

The results with the *expo-power* model show a similar picture. The median μ of 0.056 is similar to that estimated in the prioritarian model. However, there is a higher concentration of μ at the extremes (closer to 0 or above 0.2). The distribution of γ indicates a lack of consensus in whether participants have increasing-, constant- or decreasing-absolute inequality aversion. By comparing participant AICs with the prioritarian model, we do, however, find that 65.7% of participants have preferences in line with constant ($\gamma \approx 0$), whilst 25.7% exhibit increasing ($\gamma < 0$) and only 8.7% exhibit decreasing ($\gamma > 0$) absolute inequality aversion. As shown in Figure D10, those with lower μ generally exhibit increasing, whilst those with higher μ exhibit decreasing absolute inequality aversion. These results suggest that the notion of increasing absolute inequality aversion - seldom utilised in this field - could be investigated further in future research. To some extent, it captures the intuition behind the threshold approach to fair innings: inequalities become less acceptable at high ages (or high QALY levels).

In sum, all models show that the majority of participants give priority to the worst-off, providing resounding support for the fair innings argument. However, we find that the *prioritarian* model provides a substantially better fit than the *threshold* model: it is not a sharp discontinuity in marginal weights at a threshold, but a continuously decreasing marginal weight that best characterises participants choices. While extensions of the simple prioritarian model indeed improve the fit, the gain is not very substantial. For pragmatic reasons, it seems adequate to focus on the simple prioritarian model. The notion of increasing absolute inequality aversion deserves more attention in future work.

5.3 Hypothetical Choice and Preference Predictions: Age or QALYs?

Let us now return to the issue that was raised in the introduction: does the preferred fair innings interpretation prioritise the young, or those that have already achieved a smaller number of QALYs? The finding that age-weighting does not at all explain the choices of the participants in our experiment strongly suggests that the latter interpretation is closer to their intuitions. We further illustrate this finding with focusing on a specific set of hypothetical choices, similar to the example from the introduction.[23]

Imagine there are two patients - a 30-year-old and a 60-year-old - who need treatment to extend their lives. However, only one treatment is available and no other treatments will become available later. Assume that both individuals live in perfect health throughout their lives, and treatment will extend the number of years they would live in perfect health. The *older* individual would die later today without treatment or live for another 10 years with treatment. For the *younger* individual, we distinguish three cases:

A. Die later today without treatment or live or another 10 years with treatment,

B. Die later today without treatment or live for another 1 year with treatment, or

C. Die in 40 years without treatment or live for another 50 years with treatment.

The question now is whether, in these three cases, you would give treatment to the older or the younger individual. Table 3 presents these three choices, in the top panel. Those

[23] We take inspiration from the cases described by Tsuchiya (1999), Bognar (2015) and Adler et al. (2025), but focus on the distinction between current age and QALYs.

who adhere to the fair innings principle would, generally, provide treatment to the younger individual in **A**, irrespective of the form of prioritisation by age or QALYs. **B** forces a trade-off between equality and efficiency, and the choice to provide treatment to the young depends on the intensity of the prioritisation by age. **C** highlights the distinction between prioritisation by actual age or by QALYs, as the (currently) young would receive treatment with the former interpretation, but not with the latter.

Table 3: Choices and Predictions

	Current		**Younger Treated**		**Older Treated**	
	Age	QALYs	Gain	QALYs	Gain	QALYs
A	30, 60	30, 60	10, 0	40, 60	0, 10	30, 70
B	30, 60	30, 60	1 , 0	31, 60	0, 10	30, 70
C	30, 60	70, 60	10, 0	80, 60	0, 10	70, 70

	Predicted Preference (%)							
	Age-Weighted		*Threshold*		*Prioritarian*		*'Best'*	
	Younger	Older	Younger	Older	Younger	Older	Younger	Older
A	95.2	4.8	89.6	3.5	100.0	0.0	89.1	8.7
B	89.1	10.9	42.6	57.4	40.0	60.0	46.5	51.3
C	95.2	4.8	0.4	50.4	0.0	100.0	17.0	82.2

In the lower panel of Table 3, using participant-level preference parameters, we predict the choices of participants across these three examples. In the first three columns we assume that participants behave according to a particular model (threshold, prioritarian, age-weighted) with their individual estimated parameters.[24] In the last column we predict the choices by using the "best" model for each participant. All models show that the majority of participants would prioritise the younger individual in **A**, and that the number who prioritise the younger individual drop as it become less efficient to do so in **B**. An important distinction between these models is shown in **C**. The *threshold* and *prioritarian* models

[24]The shares of participants giving priority to the young or the old do not always sum to 1 for the threshold model, because, depending on the size of their preferred threshold, participants can be indifferent between the two options.

predict that the majority would give treatment to the older individual *because* they have lower lifetime QALYs. For the *age-weighted* model, the majority prioritises the younger individual; irrespective of their lifetime QALYs. However, this latter prediction diverges from the predictions of the 'best' model, because the age-weighted model does not fit well the choices of any participant. According to our participants, it is not the young who should be prioritised, but those who will live for fewer QALYs.

Whilst these are simple examples, designed to test our intuitions, they illustrate how the estimated preferences could be used to aid policy decisions. Combined with the distribution of age and resulting QALYs from alternative policies, a prediction of the ranking by each participant can be computed and the percentage of the preference for each policy calculated.

5.4 QALEs and Social Marginal Welfare Weights

Our results can also be used in a more macro perspective by identifying potential population groups to prioritise for prospective health policies. The closest approximation of the health indicator h in our theoretical model is the number of quality-adjusted life years. At the level of subgroups, the expected number of QALY's (QALE) can be predicted using population data. Focusing on these expected values means that we neglect the interindividual differences within each group, and therefore also the uncertainty at the individual level. Moreover, we also take an ex ante perspective at the level of the different age groups: when the young and the old have the same QALE, it is still possible that a young individual will never reach the age of the old, while the older individual is sure to have survived until their age.

The top panel of Figure 7 shows the estimated Quality Adjusted Life Expectancy (QALE) by age, sex and deprivation quantile (Index of Multiple Deprivation IMD), for the English population, estimated by McNamara et al. (2023). The bottom panel shows the corresponding median Social Marginal Welfare Weights[25] (henceforth, welfare weights), by age, sex and deprivation, which are derived from our ("best") estimated participant-level preference parameters. This is the first empirical estimation of these graphs, which were proposed by Williams (1997).

Figure 7: QALEs and Social Marginal Welfare Weights

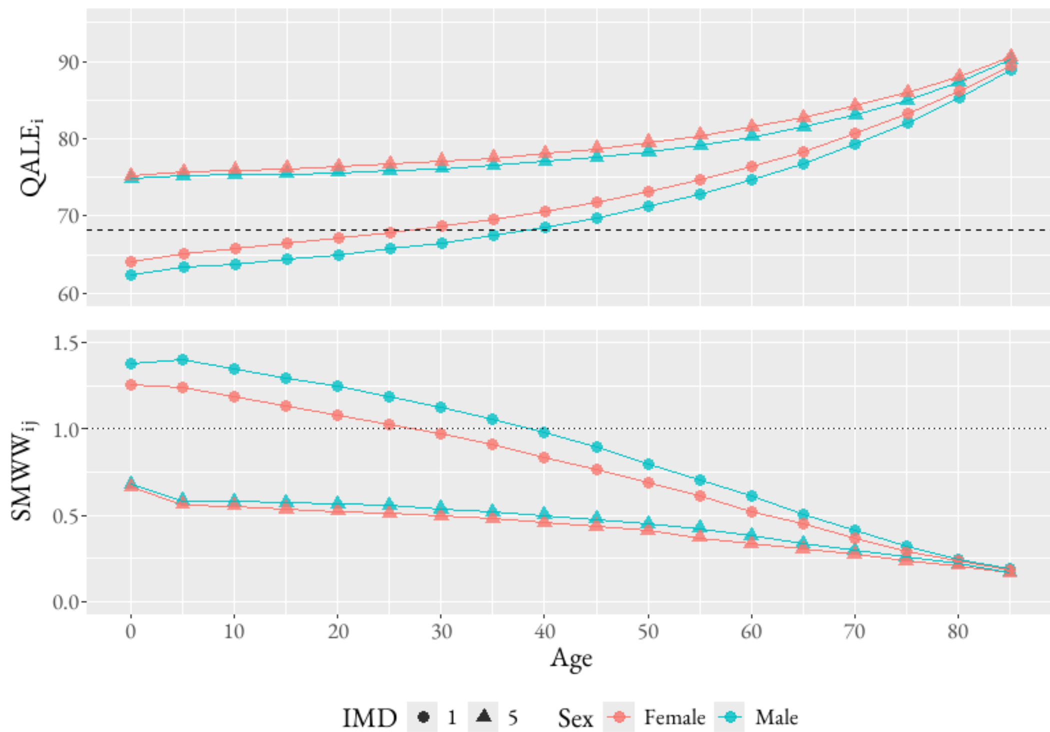

The top panel shows the QALEs are higher for older, richer and female individuals. The largest differences are by age and by IMD, and we see a reduction in inequalities in QALE by IMD as individuals age. The bottom panel shows how median welfare weights change by age. These welfare weights are relative to the QALE at birth across the English population (68.23, dashed line in the top panel). Results show that, compared to an individual with a QALE

[25]Social marginal welfare weights (SMWWs) are the relative weight given to one individual, i, compared to another, j, given their levels of health:

$$SMWW_{ij} = \frac{\partial W}{\partial h_i} \Big/ \frac{\partial W}{\partial h_j} \tag{13}$$

of 68.23, poor males, who have just been born, would receive the highest welfare weights (1.35), whilst a rich female 85 year old would received the lowest welfare weights (0.155). We see that differences in welfare weights by deprivation are highest at birth, and close to zero for those aged 85, and that differences by sex are highest amongst poorer and younger individuals. These results indicate that, if health investments are equally productive across groups, policy makers wishing to respect the public's social preferences should prioritise policies or interventions which improve the QALE of younger and poorer individuals.

Note again the distinction between age and QALEs. Age is indeed a predictor of QALEs. But it is the QALEs themselves, as predicted by age, sex, IMD (or indeed life threatening illness), which contribute to welfare in these models. For policy purposes, knowing that we can predict QALEs and target groups defined by these observable characteristics is useful to improve overall social welfare; but predictors themselves should, arguably, not be the sole reason for prioritisation.

6 Discussion

The flexible and rich design of our experiment has allowed us to compare the performance of different formalisations of the fair innings idea. The participants in our sample strongly support to give priority to those who have achieved a lower level of QALYs. More specifically, threshold approaches are far less popular than functional forms of the prioritarian type. Our design delivers estimates of the relevant parameters at the level of the individual participants. In a less sophisticated design, Adler et al. (2025) estimated for a fraction of their respondents what they call the "age premium", and converting this into a concavity parameter yields results that are not too far removed from the estimates in our paper. Yet we offer an estimate of the complete distribution of all the crucial individual-specific parameters, also for other SWF's than the simple prioritarian one.[26]

Adding additional complexity to the prioritarian SWF improves the statistical performance, but rather marginally. An intriguing finding is that about one quarter of the participants have increasing absolute inequality aversion.[27] One possible interpretation is that this to some extent captures the intuition behind the threshold idea. Making it more important to reduce inequality at higher levels of QALY's, may boil down to pushing persons with a very high amount of QALYs in the direction of a threshold. It would be interesting to explore this idea further, also in other contexts than health.[28]

None of our participants prefers linear age-weighting and, a fortiori, simple maximisation of total health. Moreover, there is only a minority whose answers can best be explained by a gains model: most participants focus on final outcomes. The incremental approach with as main objective health maximization, which is the underlying framework in cost-effectiveness analysis, is therefore decisively rejected. While this is a common finding in the literature, we add that age weighting will not come to the rescue. Moreover, our results also show that "health maximisation" ageing is not adequate to support the prioritisation of the young.

The rejection of age-weighting also supports the conclusion that QALYs, not age, should be considered welfare improving. This is further illustrated by the hypothetical example in section 5.3. It is, however, important to note that age is a strong predictor of QALYs, and so it may indeed be the case that younger individuals are prioritised, not directly because they are young, but indirectly because of their lower QALYs. Moreover, age is an easily observable, albeit imperfect, indicator of QALYs achieved, and can therefore be used as such

[26] Adler et al. (2025) focus on an Atkinson SWF, which makes the comparison somewhat tricky.

[27] Some raw results of Adler et al. (2025) (see their Table 1) go in the same direction, as the "priority for the young" becomes stronger at higher levels of QALYs (with the difference between the young and the old constant).

[28] As an example, the "threshold" idea has also been introduced as "limitarianism" in the context of income and wealth inequality (Robeyns (2024)).

by policy makers. As a caveat, our strong findings in this regard can be driven to some extent by the way we designed the experiment. Whilst we emphasise differences by "current age", we also tell participants to assume that individuals have reached that age in perfect health, such that current age equals initial health. As a result, our age weights can also be interpreted as initial health weights, and it may have been difficult for participants to grasp the difference. A possible extension would be to present scenarios where individuals had differential health-related quality-of-life across their lives.

These last remarks point to a possible limitation of our design. On the one hand, because it is highly structured and rather abstract, it allows to estimate specific functional forms for the SWF at the individual participant level. On the other hand, precisely because it is so highly structured and abstract, it may be less suited to capture additional arguments and intuitions that have played an important role in the fair innings discussion.[29] The empirical literature has shown that the design of the experiment may strongly influence the outcomes, certainly when participants do not have stable and well-defined preferences (Oliver (2009), Adler et al. (2025)). Our work makes no exception on that rule.

Much empirical work (and many examples in the philosophical literature) tries to sharpen intuitions by focusing on a discrete choice problem, in which participants must choose between two treatments, one for a young person and another one for an old person. Our example in section 5.3 was similar, but the literature adds more information to the description of the cases. A relevant consideration is the choice between the treatment of an old severely ill patient versus a young patient with a mild disease. Nord (2005) argues that in this case the actual severity of the disease will be more crucial to determine the priority than

[29]We noted already before that the framework of an individual prudential life time account (Daniels, 1988) is very different from a framework in which a SWF is maximised.

the number of QALYs obtained in the past. Introducing actual severity in the picture would require another design than the one we used in this paper.[30]

This issue of the relevance of severity brings us to another limitation of our design. The philosophical literature, and certainly Harris (1985), makes a sharp distinction between quality of health and length of life. He proposed the threshold idea only for the latter, i.e. in the context of life-prolonging treatments. He did not state that older patients should get less health care to improve the quality of their lives. Nord (2005) argues that the idea that young people should have priority over old people when it comes to functional improvements and symptom relief for non-fatal conditions runs counter to moral intuitions. Throughout our theory and experiment we focus on quality-adjusted life years (QALYs): the product of health-related quality-of-life and length of life. As a result, we cannot disentangle whether prioritisation, or more specifically concavity parameters, would differ between the two. It would be useful to refine our design in that direction.

Our framework is fully consequentialist, with a SWF defined only over outcomes in terms of QALYs. Sen (2002), among others, noted that health equity is a multidimensional notion, that should also capture deontological aspects. In fact, Adler et al. (2025) find that about one third of the participants in their French sample aim at giving the same treatment to the old and the young.[31] This notion of equal treatment can be translated as equal resource shares in our design. Individuals following this rule would have $\alpha_{1i} = \alpha_{2i} = 0$ in eq. (11), and Figure 4 shows that this pattern does hardly occur in our sample. On the other hand, equal health shares (i.e. $\beta_{1i} = \beta_{2i} = 0$) are also rare (see the right panel in Figure 4). This can point to a trade-off between total health and health equity, but it could also reflect some consideration for the deontological aspect. Whatever the interpretation, and while it is possible to make no-discrimination choices in our design, it seems fair to concede that this

[30] Results in the empirical literature are mixed (see Reckers-Droog et al. (2019), Reckers-Droog et al. (2021)), but this does not detract from the importance of the issue.

[31] Similar findings are reported by Nord et al. (1996)).

possibility is not made salient for the participants. The trade-off between consequentialist and deontological arguments is an intesting topic for further research.

Another important topic for further research is the introduction of uncertainty. We present QALY outcomes for each individual with certainty. In the empirical application of section 5.4, we worked with expected quality-adjusted life expectancies (QALEs), rather than with probabilistic distributions of QALYs. This is an implicit focus on ex-ante, rather than ex-post, distributions. This ex ante approach has obvious limitations. While we know for sure that the older person has survived past those younger ages, there is still a positive probability that the young person's life will be cut-short (and they won't get their fair innings). We may therefore be willing to prioritise the younger individual, and actually increase inequalities in QALEs in order to compensate the younger individual for the risk of dying young. Further research could investigate preferences over ex-ante vs ex-post inequality aversion.

A further limitation relates to our sample. Although we recruit a general population sample, the sample is not perfectly representative, is somewhat small, and may exclude those who do not have access to the internet. Further research could gather larger, more representative samples, or indeed samples from other countries. Finally, our experiment is hypothetical. The lack of incentivised choices may leave participant's responses prone to social desirability bias or inattention.

7 Conclusion

We find resounding support for the fair innings argument: the vast majority of participants give, often substantial, priority to those who have a lower quality adjusted life expectancy. However, the form of the argument matters. Prioritarian social welfare functions are shown to provide a substantially better fit than the threshold model or age-weighted model. This indicates that support for the fair innings argument is captured by smoothly diminishing marginal welfare returns to QALEs, rather than a discontinuous drop above a fair innings threshold. We find that more general models - threshold-prioritarian and age-weighted prioritarian - do provide better fit for some participants, but the gains in fit are not very large. An intriguing finding is that for about a quarter of our participants, their answers can be rationalised better with a model with increasing absolute inequality aversion.

The preference parameters we estimate can be used to rank alternative health policies, in accordance to the views of the public, thus aiding difficult evaluation decisions where equity and efficiency concerns may conflict. Welfare weights can additionally be used to provide intuition for which groups could be prioritised in order to maximise social welfare (primarily, the young and poor). Underlying these decisions, is a trade-off between 'first best', but complicated, models, and 'second best', but simpler, models. We find that the simplest models (threshold, efficiency maximisation, and age-weighting) have very little support. Yet the prioritarian function provides a reasonably good fit, as compared to the more general models, and maintains a degree of simplicity which should aid its more widespread use in policy making.

References

Adler, M. (2019). *Measuring social welfare: An introduction*. Oxford University Press.

Adler, M., Ferranna, M., Hammitt, J., de Laubier, E., & Treich, N. (2025). Fair innings: An empirical test. *Health Economics*, *34*, 1350–1364.

Adler, M., Ferranna, M., Hammitt, J., & Treich, N. (2021). Fair innings? the utilitarian and prioritarian value of risk reduction over a whole lifetime. *Journal of Health Economics*, *75*(102412).

Adler, M., & Norheim, O. (2022). *Prioritarianism in practice*. Cambridge University Press.

Atkinson, A. B. (1970). On the measurement of inequality. *Journal of Economic Theory*, *2*(3), 244–263. https://doi.org/10.1016/0022-0531(70)90039-6

Attema, A. E., Brouwer, W. B., & Pinto-Prades, J. L. (2022). Reference-dependent age weighting of quality-adjusted life years. *Health Economics*, *31*(12), 2515–2536.

Bleichrodt, H., Diecidue, E., & Quiggin, J. (2004). Equity weights in the allocation of health care: The rank-dependent qaly model. *Journal of health economics*, *23*(1), 157–171.

Bleichrodt, H., Doctor, J., & Stolk, E. (2005). A nonparametric elicitation of the equity-efficiency trade-off in cost-utility analysis. *Journal of Health Economics*, *24*(4), 655–678. https://doi.org/10.1016/j.jhealeco.2004.10.001

Bognar, G. (2008). Age-weighting. *Economics and Philosophy*, *24*, 167–189.

Bognar, G. (2015). Fair innings. *Bioethics*, *29*(4), 251–261.

Clarke, P., & Erreygers, G. (2020). Defining and measuring health poverty. *Social science & medicine*, *244*, 112633.

Daniels, N. (1988). *Am i my parent's keeper? an essay on justice between the young and the old*. Oxford University Press.

Dolan, P., Shaw, R., Tsuchiya, A., & Williams, A. (2005). Qaly maximisation and people's preferences: A methodological review of the literature. *Health Economics*, *14*(2), 197–208. https://doi.org/10.1002/hec.924

Dolan, P., & Tsuchiya, A. (2011). Determining the parameters in a social welfare function using stated preference data: An application to health. *Applied Economics*, *43*(18), 2241–2250. https://doi.org/10.1080/00036840903166244

Donaldson, D., & Weymark, J. A. (1980). A single-parameter generalization of the gini indices of inequality. *Journal of Economic Theory*, *22*(1), 67–86. https://doi.org/10.1016/0022-0531(80)90065-4

Donaldson, D., & Weymark, J. A. (1983). Ethically flexible gini indices for income distributions in the continuum. *Journal of Economic Theory*, *29*(2), 353–358. https://doi.org/10.1016/0022-0531(83)90053-4

Foster, J., Greer, J., & Thorbecke, E. (1984). A class of decomposable poverty measures. *Econometrica: journal of the econometric society*, 761–766.

Hardardottir, H., Gerdtham, U.-G., & Wengström, E. (2021). Parameterizing standard measures of income and health inequality using choice experiments. *Health Economics*, *30*(10), 2531–2546. https://doi.org/10.1002/hec.4395

Harris, J. (1985). *The value of life: An introduction to medical ethics*. Routledge, London, New York.

Hurley, J., Mentzakis, E., & Walli-Attaei, M. (2020). Inequality aversion in income, health, and income-related health. *Journal of Health Economics*, 102276. https://doi.org/10.1016/j.jhealeco.2019.102276

ISER, I. (2023). Understanding society: Waves 1-12, 2009-2021 and harmonised bhps: Waves 1-18, 1991-2009. *UK Data Service*, 102276. https://doi.org/10.5255/UKDA-SN-6614-18

Johannesson, M., & Johansson, P.-O. (1997). Is the valuation of a qaly gained independent of age? some empirical evidence. *Journal of Health Economics*, *16*, 589–599.

Köbberling, V., & Wakker, P. P. (2005). An index of loss aversion. *Journal of Economic Theory*, *122*(1), 119–131.

Kolm, S.-C. (1976). Unequal inequalities. i. *Journal of Economic Theory*, *12*(3), 416–442. https://doi.org/10.1016/0022-0531(76)90037-5

McFadden, D. (1972). Frontiers in econometrics. In P. Zarembka (Ed.).

McNamara, S., Schneider, P. P., Love-Koh, J., Doran, T., & Gutacker, N. (2023). Quality-adjusted life expectancy norms for the english population. *Value in Health*, *26*(2), 163–169.

Murray, C., & Acharya, A. (1997). Understanding dalys. *Journal of Health Economics*, *16*, 703–730.

Nord, E. (2005). Concerns for the worse off: Fair innings versus severity. *Social Science & Medicine*, *60*(2), 257–263.

Nord, E., Street, A., Richardson, J., Kuhse, H., & Singer, P. (1996). The significance of age and duration of effect in social evaluation of health care. *Health Care Analysis*, *4*, 103–111.

Oliver, A. (2009). A fair test of the fair innings? *Medical Decision Making*, *29*(4), 491–499.

Ottersen, T., Forde, R., Kakad, M., Kjellevold, A., Melberg, H. O., Moen, A., Ringard, A., & Norheim, O. F. (2016). A new proposal for priority setting in norway: Open and fair. *Health Policy*, *120*(3), 246–251.

Pinho, M., & Botelho, A. (2018). Inference procedures to quantify the efficiency–equality trade-off in health from stated preferences: A case study in portugal. *Applied Health Economics and Health Policy*, *16*, 503–513. https://doi.org/10.1007/s40258-018-0394-6

Rawls, J. (1971). *A theory of justice* [Eleventh printing, 1981]. The Belknap press of Harvard University Press. https://doi.org/10.2307/j.ctvjf9z6v

Reckers-Droog, V., van Exel, J., & Brouwer, W. (2018). Looking back and moving forward: On the application of proportional shortfall in healthcare priority setting in the Netherlands. *Health Polcy*, *122*(6), 621–629.

Reckers-Droog, V., van Exel, J., & Brouwer, W. (2019). Equity weights for priority setting in healthcare: Severity, age, or both? *Value in Health*, *22*(12), 1441–1449.

Reckers-Droog, V., van Exel, J., & Brouwer, W. (2021). Willingness to pay for health-related quality of life gains in relation to disease severity and the age of patients. *Value in Health*, *24*(8), 1182–1192.

Robeyns, I. (2024). *Limitarianism. the case against extreme wealth*. Penguin Books.

Robson, M., Asaria, M., Cookson, R., Tsuchiya, A., & Ali, S. (2017). Eliciting the level of health inequality aversion in england. *Health Economics*, *26*(10), 1328–1334. https://doi.org/10.1002/hec.3430

Robson, M., O'Donnell, O., & Van Ourti, T. (2024). Aversion to health inequality — pure, income-related and income-caused. *Journal of Health Economics*, *94*, 102856. https://doi.org/10.1016/j.jhealeco.2024.102856

Robson, M., O'Donnell, O., & Van Ourti, T. (2025). Responsibility-sensitive welfare weights for health. *Journal of Health Economics*, *102*, 103018.

Rodriguez, E., & Pinto, J. (2000). The social value of health programmes: Is age a relevant factor? *Health Economics*, *9*, 611–621.

Saha, A. (1993). Expo-power utility: A 'flexible'form for absolute and relative risk aversion. *American Journal of Agricultural Economics*, *75*(4), 905–913.

Sen, A. (2002). Why health equity? *Health Economics*, *11*(8), 659–666.

Tsuchiya, A. (1999). Age-related preferences and age weighting health benefits. *Social Science & Medicine*, *48*(2), 267–276.

Tsuchiya, A., Dolan, P., & Shaw, R. (2003). Measuring people's preferences regarding ageism in health: Some methodological issues and some fresh evidence. *Social Science & Medicine*, *57*, 687–696.

Wagstaff, A., Paci, P., & Van Doorslaer, E. (1991). On the measurement of inequalities in health. *Social Science & Medicine*, *33*(5), 545–557. https://doi.org/10.1016/0277-9536(91)90212-U

Williams, A. (1997). Intergenerational equity: An exploration of the 'fair innings' argument. *Health economics*, *6*(2), 117–132.

APPENDICES

Appendix A Experimental Design

Figure A1 provides an overview of the experiment. All participants went through two sessions. In the first session, participants received instructions, followed an interactive tutorial, answered follow-up questions of comprehension, did the *budget* treatment (10 rounds), and completed a questionnaire about sociodemographic characteristics and beliefs. In the second session, they followed a shorter tutorial, then the *age-equal* (10 rounds) and *age-varying* (10 rounds) treatments – the order of these two treatments was randomised between participants – and then completed a final questionnaire. All participants completed all sections, consisting of 30 rounds of decision problems and 71 questions in total.

Figure A1: Experiment Overview

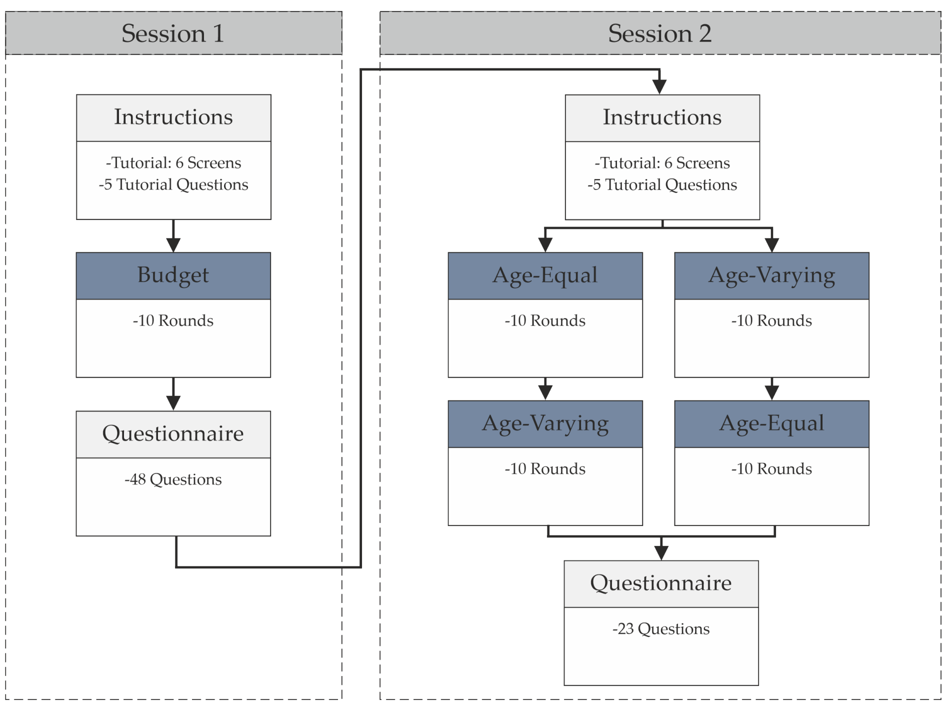

Table A1 provides an overview of the experimental design parameters used across all rounds and treatments of the experiment. Each treatment has 10 rounds of decisions problems. In all treatments, there are the same 10 rounds of multipliers, the order of which is randomised within each treatment and between participants. There is one round with equal multipliers for all individuals: $p_i \in \{1, 1, 1\}$, three rounds with a lower multiplier for one individual: $p_i \in \{1, 1, 1/3\}$, and five rounds with varying multipliers across all three individuals: $p_i \in \{1, 1/2, 1/3\}$. The distribution and expectation of p_i over the 10 rounds is the same irrespective of the screen position (left, middle, right). In each round of the *budget* treatment the budget was randomly drawn from $m \in \{180, 240, 360\}$, whilst the current age was zero for all individuals $y_i = 0\,\forall i$ across all rounds. In each round of the *age-equal* treatment one current age was randomly drawn for all three individuals from $y_i \in \{0, 5, 20, 40, 60, 80\}$, such that $y_i = y_j\,\forall i, j$, whilst the budget, $m = 180$, was held constant across all rounds. In each round of the *age-varying* treatment a current age was randomly drawn, with replacement, for each individual from $y_i \in \{5, 20, 40, 60, 80\}\,\forall i$, such that $y_i \neq y_j\,\forall i \neq j$, whilst the budget, $m = 180$, was held constant across all rounds.

Table A1: Experimental Design Parameters

All Treatments			**Budget**	**Age-Equal**			**Age-Varying**		
Multipliers, p_i			Age, y_i	Age, y_i			Age, y_i		
1	1	1	0	0	0	0	5	20	40
1	1	0.33		5	5	5	5	20	60
1	0.33	1		20	20	20	5	20	80
0.33	1	1		40	40	40	20	40	60
1	0.33	0.5		60	60	60	20	40	80
0.5	1	0.33		80	80	80	40	60	80
0.33	0.5	1	Budget, m	Budget, m			Budget, m		
1	0.5	0.33	180	180			180		
0.33	1	0.5	240						
0.5	0.33	1	360						

A.1 Instructions and Tutorial Script

The text for the instructions, tutorial and tutorial questions below are shown to all participants in Session 1, on screen within the experiment. The instructions give an overview of the experiment to come. The six stages of a tutorial explain how to use the on-screen interface; each of the scripts are followed by an interactive on-screen tutorial. Finally, five tutorial questions are presented to check and reinforce understanding.

A.1.1 Instructions

Welcome. Thank you for taking part today.

Please Read These Instructions Carefully.

You will be asked to make decisions which determine the health of hypothetical individuals in society.

You will be given a "Budget" that you must divide between these individuals. The Budget is the total amount of "Resources" available to spend.

Resources determine "Health". Health is the number of years a person lives, adjusted for illness or disability. For example, consider someone who reached the age of 70 without any illness or disability, who then lived for a further 10 years with an illness which reduced their quality of life to half of what it was before. That person might be said to have lived for the equivalent of 75 years in full health. For shorthand, we refer to this as "Health".

Giving more Resources to an individual increases their Health. The impact of Resources on Health is determined by a number referred to as the "Multiplier". The higher the Multiplier, the higher the level of Health achieved from a given number of Resources.

On the screen, you will distribute Resources between three Individuals. You will do this a number of times. Each screen will show a different scenario. The choices you make on one screen will not affect the scenarios that follow.

There are no right or wrong answers. We are interested in the choices you make, whatever they are.

You will now go through a tutorial, which will explain how to use the computer interface and the exact nature of the experiment.

Please click Next to continue.

Tutorial 1

This tutorial will show you how to use the on-screen interface.

You will first get practice in giving Resources to only one individual, who is identified by initials (e.g. CS). Drag the horizontal slider at the bottom of the next screen to the right to give more Resources to the individual.

The amount of Resources you give is shown by the height of the blue bar in the chart above the slider. Resources are also shown by the number to the left of the blue bar and by the number in the table at the top of the page.

Once you have dragged the slider, you can use the left and right arrow keys to make precise changes to the amount. Press the arrow key for a change of 0.1 and hold the arrow key for changes of 1.

The Resources you give to an individual are taken from the Budget, which is shown on the left of the screen. As you increase the Resources, the Remaining Budget will decrease.

You must always use all of the Budget, so that the Remaining Budget is zero. When there is only one individual, this means dragging the slider all the way to the right. Later you will have to distribute the Budget between individuals.

Press Next to try out the slider. When you are done, allocate all of the Budget (100) and press Next.

Tutorial 2

The Resources you give to an individual determines their "Health". Health is the number of years a person lives, adjusted for illness or disability. Health is equal to the Resources multiplied by a number we call the "Multiplier".

The Multiplier is shown in the table at the top of the screen. When you give Resources by moving the slider, the resulting Health is shown by the height of the grey bar. The number to the right of this bar is the amount of Health, and this is also shown in the table at the top of the screen.

For this first individual the Multiplier is 1. So, if you give all of the Budget of 100 to the individual, their Health will be 100. They will live 100 years in full health.

Press next and see how Health changes as you adjust the Resources given to the individual. When you are done, allocate all of the Budget and press Next.

Tutorial 3

The Multipliers can vary from individual to individual. In the previous scenario, the Multiplier was 1. In the next scenario, it is 0.5.

Press Next and see how Health changes as you give more Resources to this individual. Notice that there is now a gap between the Resources given (blue bar) and the Health achieved (grey bar). If you give all the Budget of 100 to this individual, their Health will be 50. They will live 50 years in full health.

When you are done, allocate all of the Budget and press Next.

Tutorial 4

In each round of the experiment, there are three individuals. Individuals are identified by their initials (e.g. CS, SJ and TD) and change between rounds.

On the next screen, there are three sliders at the bottom of the screen that you can use to give Resources to each individual and so determine their Health.

The Resources and Health of each individual are shown by the blue and grey bars. The table at the top also shows the Resources, Multiplier and Health for each individual.

Now you must allocate the Budget across the three individuals. In doing so, you determine the Health of each one. In the example on the next screen, the Multipliers are the same for all three Individuals.

You must use all of the Budget, so that the Remaining Budget (on the left) equals zero.

Press Next and then give Resources to the three individuals. Remember: you can use the left and right arrow keys to make small changes to the amount of Resources. When you have used all of the Budget on the next screen, press Next.

Tutorial 5

The three individuals change from round to round.

On the previous screen, all three individuals had a Multiplier of 1. But the Multipliers can differ between individuals, as on the next screen.

Move the sliders to give Resources to the three individuals and notice how the Health achieved depends on the Multiplier of each individual. If you give the three individuals the same Resources, their Health will differ.

Take note of the size of the Budget, which can change from screen to screen.

If you are having difficulty seeing both the table and the graph on your screen, zoom out on your web browser by holding "Ctrl" and pressing "-". Hold "Ctrl" and press "+" to zoom in.

Press Next and then give Resources to the three individuals. When the Remaining Budget is zero, press Next.

Tutorial 6

The right of the screen shows further information.

"Resource Gap" is the gap between the largest and smallest amounts of Resources you give to the individuals.

"Total Health" is the total amount Health of the three individuals (e.g. Health to CS + Health to SJ + Health to TD).

"Health Gap" is the gap between the largest and smallest amounts of Health achieved by the individuals.

If you have used the whole Budget, you will not be able to move any slider to the right. If you want to give more Resources to one individual, you will need to give less to another individual first.

Press Next and then distribute Resources across the three individuals. When you have allocated all of the Budget, press Next.

Tutorial Questions

Following the tutorial, participants answered five questions to reinforce and check understanding. The questions are shown below, with the correct response in bold. After submitting answers participants were given feedback about the correct response for each question.

1. On each screen, you will give Resources to how many individuals? - Options: 2; **3**; 4; Not Sure.
2. You can make and adjust the Resources you give by (tick all that apply): **Clicking and Dragging the Allocation Sliders**; **Using the Arrow Keys**; Moving the Vertical Bar; Not Sure.
3. If you give 100 Resources to an individual with a Multiplier of 1, then the Health of that individual will be? - Options: 25; 50; **100**; Not Sure.
4. If you give 100 Resources to an Individual with a Multiplier of 0.5, then the Health of that individual will be? - Options: 25; **50**; 100; Not Sure.

5. Once you have finished giving the Resources, you proceed to the next screen by: - Options: Clicking Next; **Ensuring the Remaining Budget = 0, then Clicking Next**; Waiting; Not Sure.

In Session 2, a modified version of the above instructions and tutorial are shown. First, to remind participants of the experiment, and second, to highlight the additional information on the initial age. They are told that additional information is provided on "Current Age (in years)" and that "They have reached this age without illness or disability (i.e. in full health)". It is pointed out this is shown, for each individual, by the "height of the light grey bar, and the number to the left of the bar".

Appendix B Descriptives

Table B2 shows descriptive statistics of our analytical sample of 230 participants.The right-hand column gives (weighted) means of some of these characteristics estimated from wave 10 (2018/19) of the UK Household Longitudinal Study (ISER, 2023), which is representative of the UK adult population. The variables included have comparable definitions in the two studies. Our sample is representative of the UK adult population with respect to sex, but is younger and has a lower proportion of white individuals. Our sample also has fewer individuals who are married or born in the UK. The average household size and country of residence are representative, but participants generally have higher incomes, are more likely to be employed and less likely to be retired. Although not perfectly representative, our sample is, at least, heterogeneous across a host of sociodemographic characteristics, health outcomes and political views.

Table B2: Characteristics of Analytical Sample and UK Representative Sample

	Sample				UKHLS
	Obs	Min	Max	Mean	Mean
Demographics					
Female	229	0	1	0.48	0.52
Age	229	18	78	44.79	51.08
Married	228	0	1	0.55	0.76
Born in UK	227	0	1	0.86	0.90
White	230	0	1	0.82	0.93
Household Size	230	0	7	2.71	2.74
Country					
England	230	0	1	0.82	0.84
Wales	230	0	1	0.05	0.05
Scotland	230	0	1	0.11	0.08
N. Ireland	230	0	1	0.02	0.03
Income					
Income (£k)	209	2	175	27.15	23.17
Labour Market Status					
- Employed	229	0	1	0.61	0.56
- Unemployed	229	0	1	0.11	0.04
- Retired	229	0	1	0.16	0.28
- Student	229	0	1	0.08	0.04
- Other	229	0	1	0.04	0.09
Occupation					
- Manager/Professional	222	0	1	0.44	.
- Intermediate	222	0	1	0.30	.
- Drivers/Labourers	222	0	1	0.03	.
- Never Worked	222	0	1	0.09	.
- Other	222	0	1	0.18	.
Highest Education					
- Postgraduate	228	0	1	0.22	.
- Undergraduate	228	0	1	0.39	.
- A-Level	228	0	1	0.29	.
- Secondary/Primary	228	0	1	0.11	.
Health					
Self-Assessed Health	230	1	4	2.22	.
Health: Likert 0-100	228	7	100	68.20	.
Cigarettes Smoked Per Day	227	0	30	1.31	.
Views					
Left-Right	224	0	10	4.04	.

Appendix C Additional Non-Structural Results

C.1 Regressions on Individual Coefficients

This appendix presents two regression analyses on the individual coefficients presented in Figure 4. The first provides a sense check, using questionnaire responses, the second estimates associations with socio-demographic characteristics. The sense check is passed, and we find almost no association with socio-demographic characteristics.

Table C3 shows results from two regressions. The first shows how the expected change in resource shares in response to higher relative multipliers ($\Delta\tilde{x}_i/\Delta\tilde{p}_i$) is associated with responses to the question "On a scale of 0 to 10, where would you place your views? where "0" is "Reducing inequalities in health is more important than improving total population health" and '10' is "Improving total population health is more important than reducing inequalities in health."' Reassuringly, we find a positive and significant coefficient. Those who reported improving total health was most important made choices which were more efficiency seeking (higher $\Delta\tilde{x}_i/\Delta\tilde{p}_i$). The second shows how the expected change in resource shares in response to higher relative age ($\Delta\tilde{x}_i/\Delta\tilde{y}_i$) is associated with responses to the question "In general, did you try to distribute Resources to give more (total) Health to Younger or Older individuals?". We find $\Delta\tilde{x}_i/\Delta\tilde{y}_i$ is positively associated with those who reported they allocated more to older individuals and a negative (though insignificant) association with those reporting higher allocations to younger individuals. Again, reassuring, as higher $\Delta\tilde{x}_i/\Delta\tilde{y}_i$ indicates a higher propensity to prioritise older individuals.

Table C3: Regression of Individual Coefficients on Sense Checks

	Dependent variable:	
	$\Delta\tilde{x}_i/\Delta\tilde{p}_i$	$\Delta\tilde{x}_i/\Delta\tilde{y}_i$
	(1)	(2)
Reduce Ineq. vs Max. Total	0.056***	
	(0.011)	
More to Younger		−0.053
		(0.042)
More to Older		0.702***
		(0.160)
Constant	−0.850***	−0.700***
	(0.054)	(0.037)
Observations	221	225
R^2	0.107	0.204
Adjusted R^2	0.103	0.197
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table C4 presents regression results of socio-demographic characteristics on participant-level estimates of $\Delta\tilde{x}_i/\Delta\tilde{p}_i$ and $\Delta\tilde{x}_i/\Delta\tilde{y}_i$. Results show that, with the exception of left-right political affiliation, there are no significant associations with any of the included socio-demographic characteristics. In particular, there is no effect of age on $\Delta\tilde{x}_i/\Delta\tilde{y}_i$: neither the old nor young are prioritising their own age groups. Those who report political alignment with the left are less efficiency seeking and more likely to prioritise the young.

Table C4: Regression of Individual Coefficients on Demographics

	Dependent variable:	
	$\Delta\tilde{x}_i/\Delta\tilde{p}_i$	$\Delta\tilde{x}_i/\Delta\tilde{y}_i$
	(1)	(2)
Female	−0.002	0.062
	(0.059)	(0.044)
Age	−0.001	−0.001
	(0.002)	(0.002)
Married	−0.019	0.066
	(0.058)	(0.040)
Household Size	0.005	0.012
	(0.025)	(0.017)
England	0.042	−0.00003
	(0.066)	(0.055)
Undergrad	−0.083	−0.061
	(0.068)	(0.053)
Employed	−0.003	0.026
	(0.061)	(0.048)
Manager/Professional	−0.021	0.014
	(0.060)	(0.043)
Self-Assessed Health	−0.038	0.003
	(0.038)	(0.028)
Left-Right	0.026*	0.037***
	(0.015)	(0.012)
Constant	−0.523***	−0.910***
	(0.161)	(0.116)
Observations	216	216
R^2	0.041	0.091
Adjusted R^2	-0.005	0.047
Note:	*p<0.1; **p<0.05; ***p<0.01	

C.2 Evidence of Discontinuities?

If people were allocating according to a threshold social welfare function, then we would expect to see evidence of discontinuities in the levels of health, at those thresholds. Figure (C2) show the distribution of health levels, by treatment. We do observe some evidence of discontinuities, but not much. The most suggestive evidence of discontinuity arises at max health = 80, but this is, in part, due to the zero allocations to individuals with a current age of 80, and a prioritarian function would have predicted a similar pattern.

Figure C2: Discontinuities in Health, by Treatment

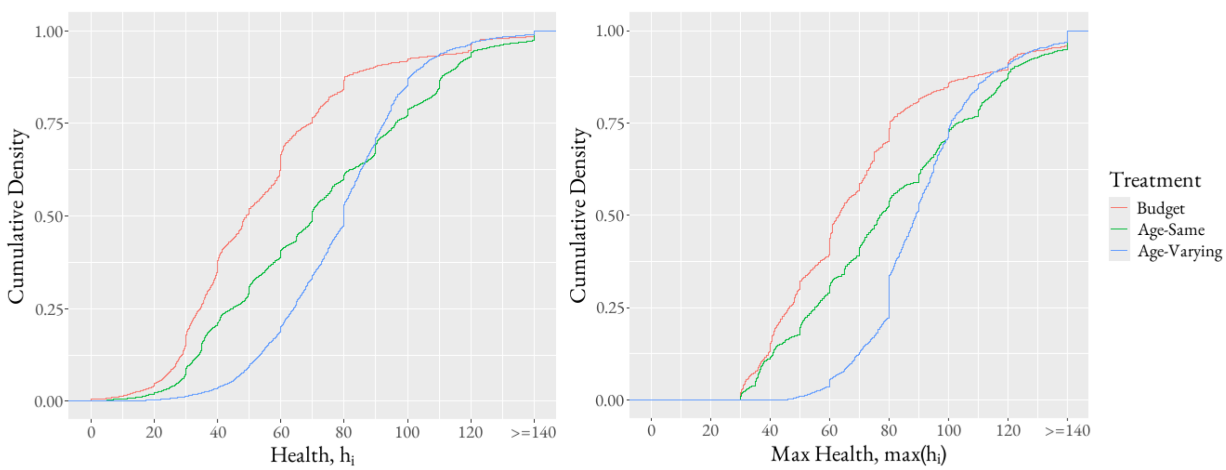

C.3 Health Differences

Figure (C3) shows the distributions of health differences, h_{ij}^{Δ}: the mean difference in health between individual i and the other two individuals, $j \neq i$, $h_{ij}^{\Delta} = \sum (h_i - h_j)/2)$. If $h_{ij}^{\Delta} > 0$, then individual i receives more health than the other two individuals, $j \neq i$, in a given round, whilst $h_{ij}^{\Delta} < 0$ implies lower health, and $h_{ij}^{\Delta} = 0$ means the health level of individual i equals the mean. Distributions of h_{ij}^{Δ} across participants reveal how much priority is given to i, and how much inequality participants allow, for a given multiplier, budget and current age level.

Figure C3: Distribution of Health Differences, by Multipliers, Budgets and current age

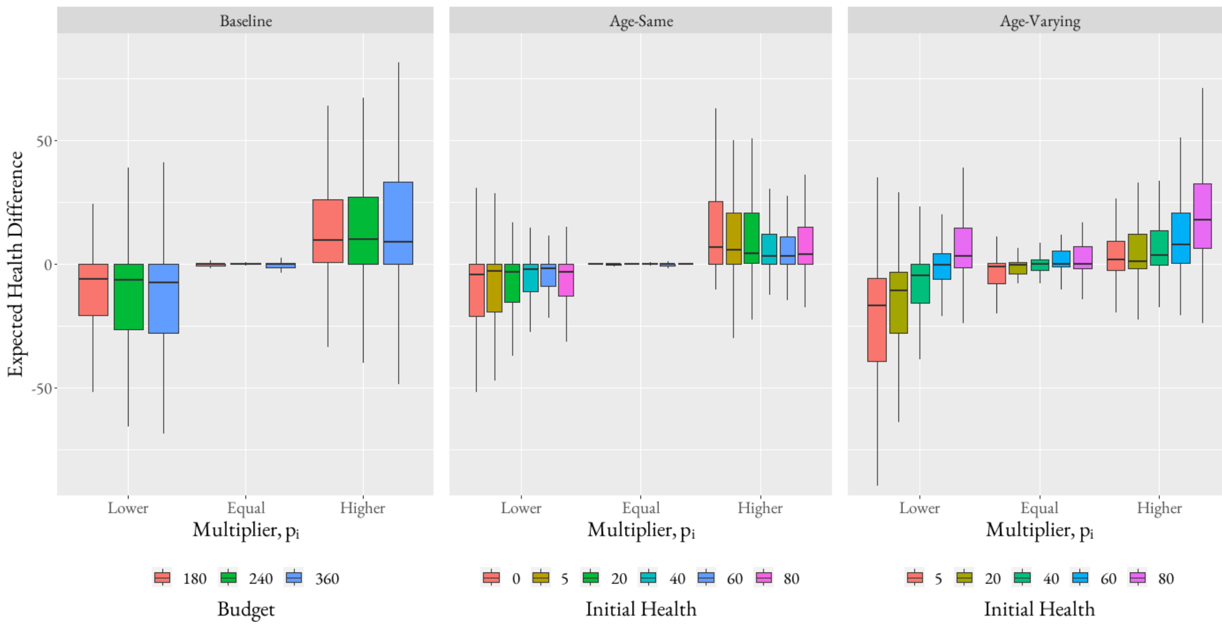

Results show that when individuals have the same multipliers or current age levels (equal multipliers in budget or age-equal treatments) participants unanimously allocate health equally amongst individuals, regardless of the level of the budget or current age. Generally, individuals with lower multipliers attain lower health levels. In the budget treatment, when budgets are higher there are higher (absolute) inequalities in health. In the age-equal treatment, when current age is higher there are (generally) lower (absolute) inequalities in health. In the age-varying treatment, whilst health is higher for those with a higher current age, when individuals with a lower current age have a higher multiplier, then they are given more health than average.

C.4 Sensitivity analysis: graphical interface

One difference between the budget treatment and the age treatments, is the inclusion of a blue bar in the budget treatment and exclusion of this in the age treatment. The table below includes only rounds which are comparable between the two treatments (budget = 180, current age = 0,0,0), and regresses allocation and health shares against relative multipliers, conditional on the treatment. Results show the coefficients are almost identical, showing that the blue bar makes little difference.

Table C5: Sensitivity to Graphical Interface

	Dependent variable:	
	Resource Share	Health Share
	(1)	(2)
Budget X p_i	-0.424^{***}	0.540^{***}
	(0.016)	(0.014)
Age-Equal X p_i	-0.424^{***}	0.540^{***}
	(0.017)	(0.014)
Constant	0.475^{***}	0.153^{***}
	(0.006)	(0.004)
Observations	3,495	3,495
R^2	0.227	0.414
Note:	*p<0.1; **p<0.05; ***p<0.01	

Appendix D Additional Structural Results

D.1 Age-Weighted

Figure D4 shows the distribution of δ parameters for the *age-weighted* model. Results show that the vast majority (95.2%) have $\delta > 0$, which indicates higher weights to the young. These weights are, in general, extremely high; with median parameter of 30.89. Very few (only 4.78%) give priority to older individuals.

Figure D4: Parameter distribution for age-weighted linear model

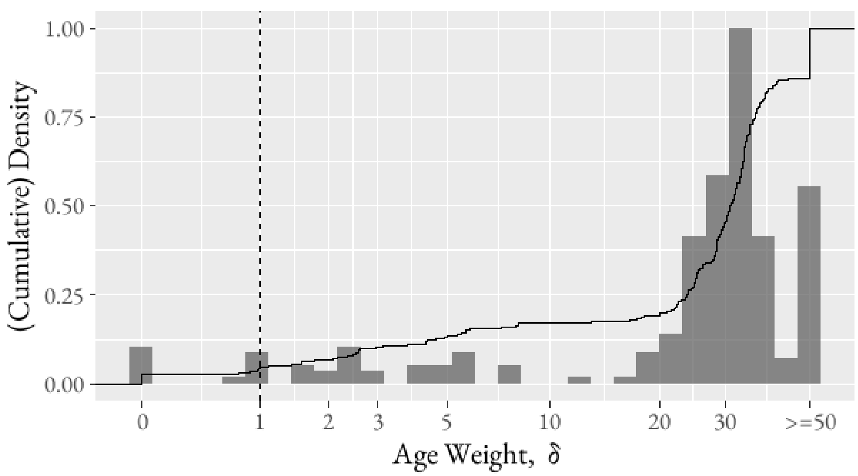

D.2 Threshold

Figure D5 shows the distribution of β and τ parameters for the *threshold* model. Results show substantially higher weights for those below the fair innings threshold, with a median β of 86.24 and τ of 60.45. Indeed, the vast majority of participants (99.56%) give higher weights to those below the threshold (with $\beta > 1$). Weights of some participants are extreme, implying an almost absolute priority in ensuring individuals reach the fair innings threshold. Whilst the distribution of τ has a modal spike around 60-70, there are a substantial number of participants with thresholds lower than 50. In later, better fitting models –*threshold-prioritarian* and *discontinuous prioritarian* – we find these lower thresholds disappear.

Figure D5: Parameter distribution for threshold model

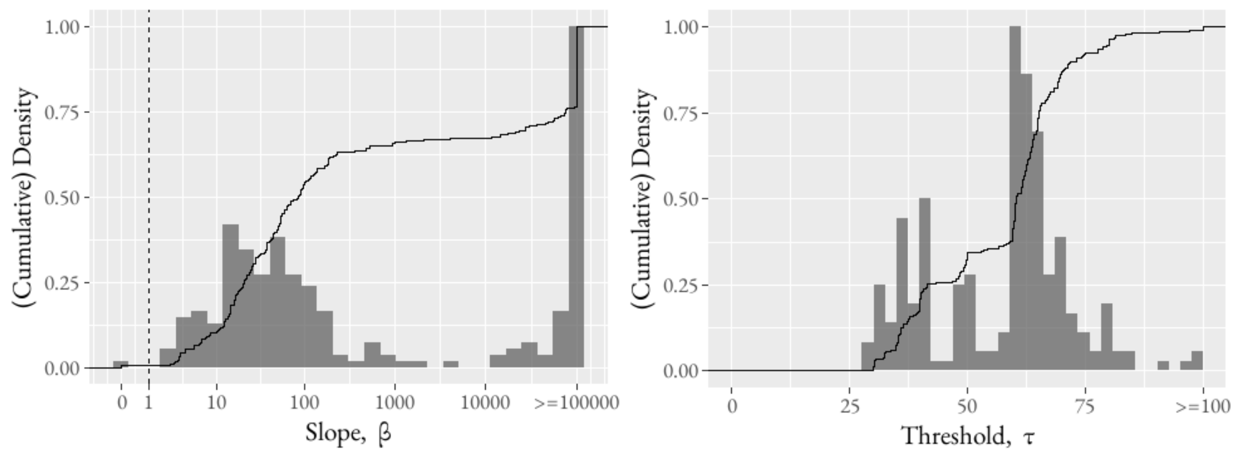

D.3 Prioritarian

Figure D6 shows the distribution of μ parameters for the *prioritarian* model. We find a median μ of 0.051, implying substantial priority to the worse-off. There is heterogeneity across different levels of μ, however, there are only 1.73% with $\mu < 0.01$, implying that almost all participants give some priority to the worse off.

Figure D6: Parameter distribution for prioritarian model

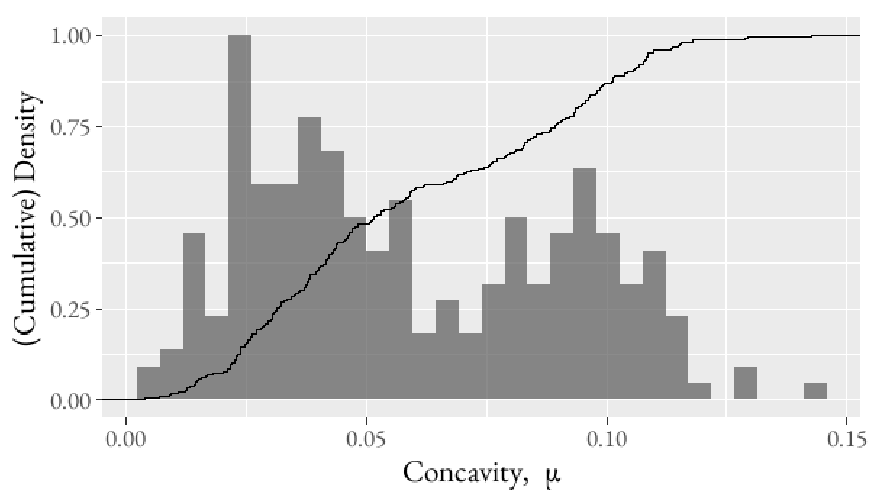

D.4 Age-Weighted Prioritarian

Figure D7 shows the distribution of concavity μ and age-weights δ for the *age-weighted prioritarian* model. The distribution of μ parameter is similar to the *prioritarian* model, with a median of 0.051. However, the distribution of δ indicates extensive heterogeneity in age-weighting parameters.

Figure D7: Parameter distribution for age-weighted prioritarian model

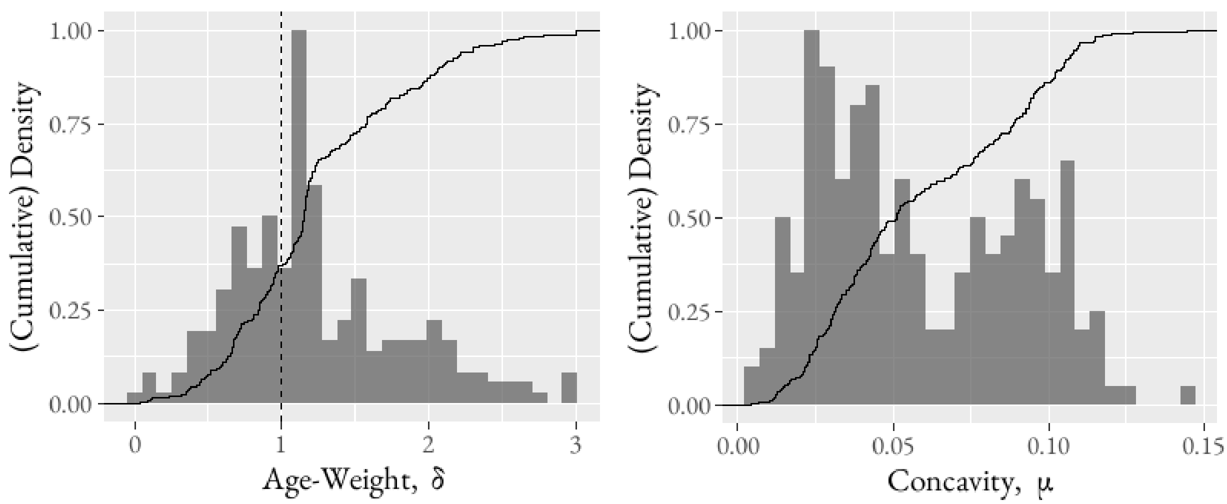

D.5 Threshold Prioritarian

Figure D8 shows the distribution of β, τ, and μ for the *threshold prioritarian* model. The median threshold of 62.9, whilst the median value for the concavity parameter μ is 0.056. The median estimated slope parameter β is 2.534.

D.6 Discontinuous Prioritarian

Figure D8 show the estimated $-\mu_b$, μ_a, and τ parameters. We find a single peaked distribution of τ, with a median of 64.81. The distribution of concavity parameters below this threshold ($-\mu_b$) is similar to the estimated μ in the *prioritarian* model, with a median of

Figure D8: Parameter distribution for threshold prioritarian model

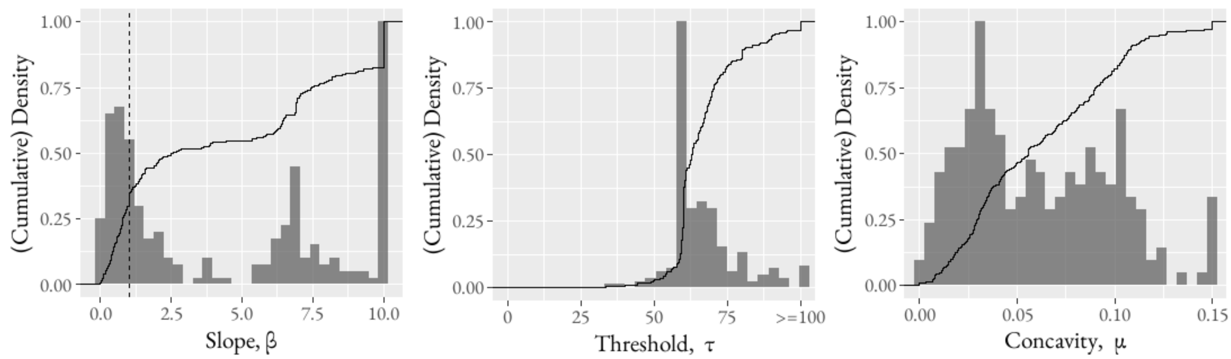

0.050. However, we find that the μ_a parameters are much higher, with a median of 4.401. To identify whether the flexibility provided by μ_b and μ_a increases fit, we can compare participant-level AICs between the *prioritarian* and *discontinuous-prioritarian* models. We find a majority of participants (73.04%) whose choices are best explained by *constant* absolute inequality aversion (i.e. the *prioritarian* model), but a sizeable minority 26.08% with higher concavity parameters above the threshold, and few with the opposite 0.08%. This finding is consistent with the expo-power model in the next appendix.

Figure D9: Parameter distribution for discontinuous prioritarian model

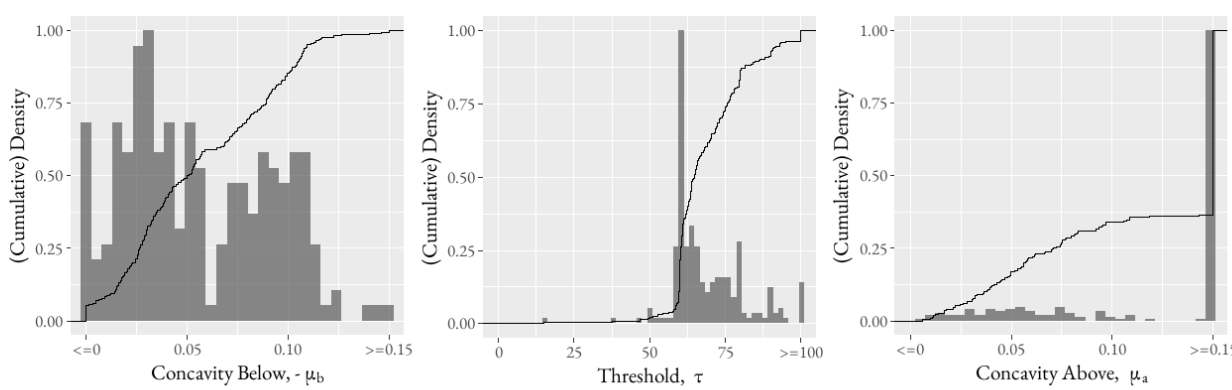

D.7 Expo-Power

Figure D10 shows the distribution of participant-level preferences, estimated within the expo-power model. The median μ of 0.056 is similar to that estimated in the *prioritarian* model, however, there is a higher concentration of μ at the extremes (closer to 0 or above 0.2).

Figure D10: Parameter distribution for expo-power model

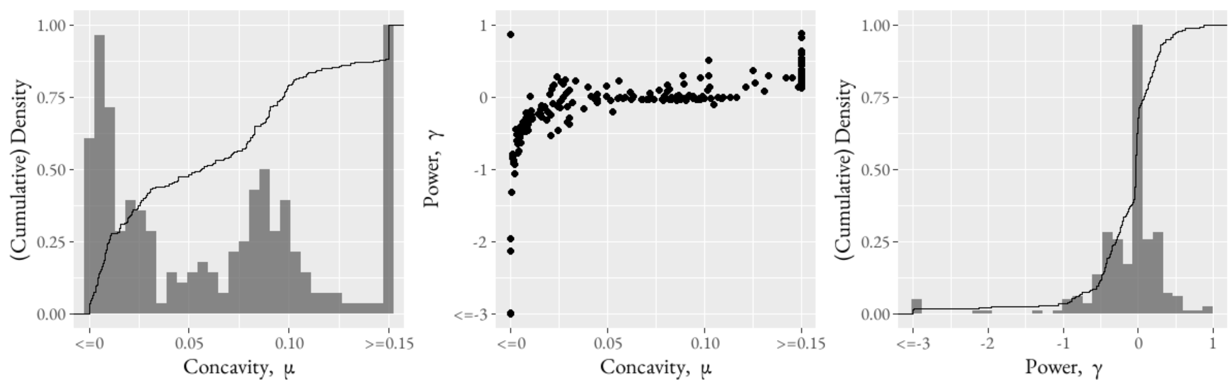

The distribution of γ indicates a lack of consensus in whether participants have increasing-, constant- or decreasing-absolute inequality aversion. By comparing participant AICs with the *prioritarian* model, we do, however, find that 65.65% of participants have preferences in line with constant ($\gamma \approx 0$), whilst 25.6% exhibit increasing ($\gamma < 0$) and only 8.67% exhibit decreasing ($\gamma > 0$) absolute inequality aversion. The relation between μ and γ is shown in the middle panel. Those with lower μ generally exhibit increasing, whilst those with higher μ exhibit decreasing absolute inequality aversion. This additional flexibility enables the *expo-power* model to fit better than the *prioritarian* model (AIC of 113869.8 compared to 114738.7), and whilst the majority 62.6% have constant, there is a substantial minority who exhibit increasing absolute inequality aversion.

D.8 Gains model

Our main models use health, h_i, as an input into the social welfare function. However, we can model health h_i as a function of current health, y_i, and any *gains* to health ($x_i p_i$), caused by an allocation of resources, x_i, times the health production multiplier p_i. A social planner – tasked with allocating resources to determine the health of individuals in a population – may not be concerned with a fair distribution of health outcomes ($h_i = y_i + x_i \pi_i$), but a fair distribution of health gains ($x_i \pi_i$):

$$\omega_i = 1\,, \qquad U\left(x_i \pi_i\right) = \begin{cases} \frac{1-e^{-\mu x_i \pi_i}}{\mu} & \mu \neq 0 \\ x_i \pi_i & \mu = 0. \end{cases} \tag{D1}$$

Figure D11 shows the distribution of concavity parameters, μ, for the gain model. Results show a median μ of 0.051, with less heterogeneity in μ than in the prioritarian model. The gains model, is shown to provide a worse fit, with only 16.96% of participants who focus on gains, rather than outcomes. Moreover, the overall goodness of fit is lower than both the *threshold* and *prioritarian* models, with a log–likelihood of -63089.8 and AICc of 127165.5.

Figure D11: Parameter distribution for gain model

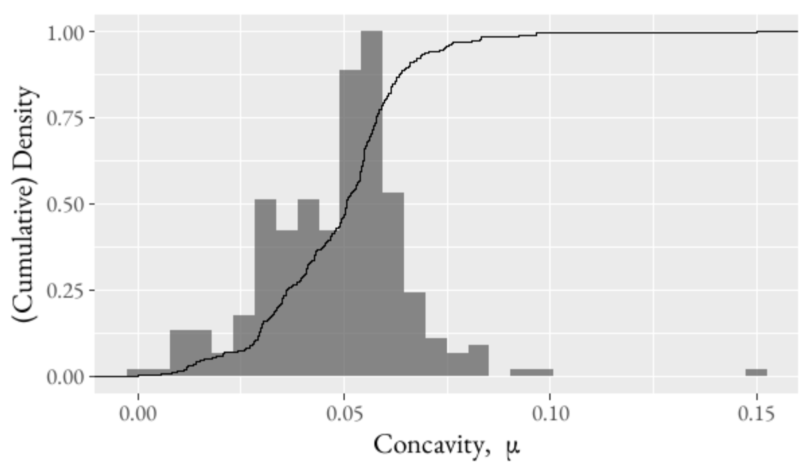

D.9 Poverty-Gap

The simplest version of the threshold approach takes it that social welfare is *only* increased by improvements in the health of those below a health (or age) threshold τ (Harris, 1985). The aim is then to push all individuals to the threshold. This can be represented by a model which is similar to the *poverty-gap* model (see Foster et al., 1984; Clarke and Erreygers, 2020). With welfare linear in QALY's and a health (or age) threshold τ, the weighting scheme can then be modelled as:

$$\omega_i = \begin{cases} 1 & \text{if } h_i < \tau \\ 0 & \text{if } h_i \geq \tau, \end{cases} \qquad U(h_i) = h_i \,. \tag{D2}$$

Figure D12 show the distribution of estimated participant-level τ within the *poverty-gap* model. The median τ is 50 QALEs, but we see little evidence of a homogeneous τ. We find a worse goodness-of-fit, as compared to the *prioritarian* model. The overall LL is -59582.2, compared to -56909.33, and an AIC of 120150.3, compared to 114738.7. At the participant-level, we find that only 12.61% of participants choices are best explained by the *poverty-gap* model, as compared to the *prioritarian* model. In sum, we find limited evidence that the *poverty-gap* model well explains the social preferences of our sample.

Figure D12: Parameter distribution for poverty-gap model

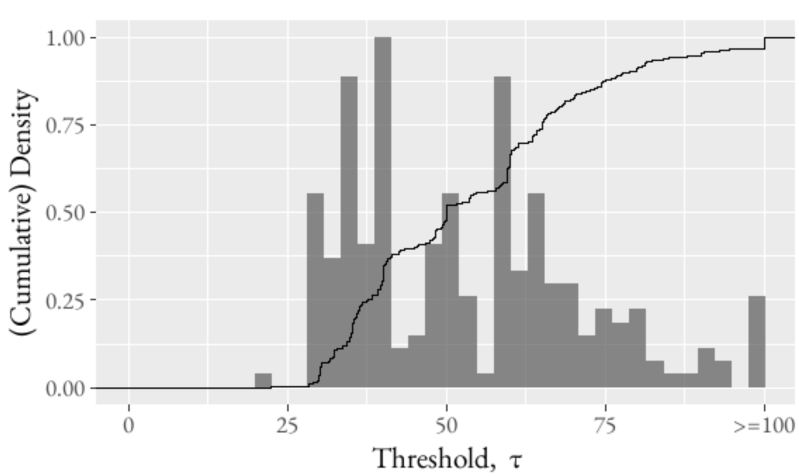

Appendix E Noise and Goodness-of-Fit

E.1 Goodness-of-Fit Overview

Table E6 presents the pooled log-likelihood and AICc values across all alternative models, alongside the % of participants for whom the alternative models provides a ‘better’ fit than the *prioritarian* model. Models are ordered from lowest to highest AIC.

Table E6: Goodness-of-Fit of Alternative Models

	Goodness-of-Fit		
Model	Log-Likelihood	AICc	vs Prioritarian (%)
Expo-Power	-56243.0	114019.5	34.35
Age-Weighted Prioritarian	-56311.4	114156.3	49.13
Discontinuous-Prioritarian	-56344.9	114813.3	26.96
Threshold-Prioritarian	-56553.8	115230.9	27.83
Prioritarian	-56909.3	114804.5	–
Poverty-Gap	-59582.2	120150.3	12.61
Threshold	-60649.9	122833.4	8.70
Gain	-63089.8	127165.5	16.96
Age-Weighted	-68313.2	137612.3	0.00
Efficiency Maximiser	-209417.7	419311.4	0.00

Results show that the *poverty-gap* model performs worse than the *prioritarian* model, but does, however, perform better than the *threshold* model. The *gain* model provides a substantially worse fit than the prioritarian model, with only 16.96% of participants for whom the model is better, this implies that the majority of participants focus on *outcomes*, rather than *gains*. The *expo-power* and *discontinuous-prioritarian* models provide a better fit than the *prioritarian* model, and indeed a lower overall AIC than the *threshold-prioritarian* model. This suggests that these models - seldom utilised in this field - could be investigated further in future research. The *age-weighted* model is shown to have provide the worst fit, from all models compared. We add the *efficiency maximiser* model at the end to show that this model fits substantially worse that all other models. Although this model is assumed for cost-effectiveness analysis, there is not a single participant for whom this model fits best.

E.2 Participant-Level Goodness-of-Fit

Figure E13 provides an overview of goodness-of-fit across the four models, at the participant level. The distribution of log-likelihood (LL) and Akaike information criterion (AICc) are shown in the left and right panels, respectively. The closer to zero these values are, the better the goodness-of-fit. Results show that, compared to the *prioritarian* model, the *threshold* and *age-weighted* models provide a worse fit. The *threshold-prioritarian*, *age-weighted prioritarian*, *discontinuous prioritarian* and *expo-power* models are all found to provide a better, aggregate, fit. The *threshold* and *age-weighted* models are shown to be (weakly) dominated by the other models which include concavity in lifetime health outcomes. Whilst we see the *prioritarian* model being (weakly) dominated by the *age-weighted prioritarian*, *discontinuous prioritarian* and *expo-power* models in the log-likelihoods, when using the AICc to penalise those models for more parameters the gap between cumulative density plots reduces.

Figure E13: Distribution of Participant Level Log-Likelihood and AIC

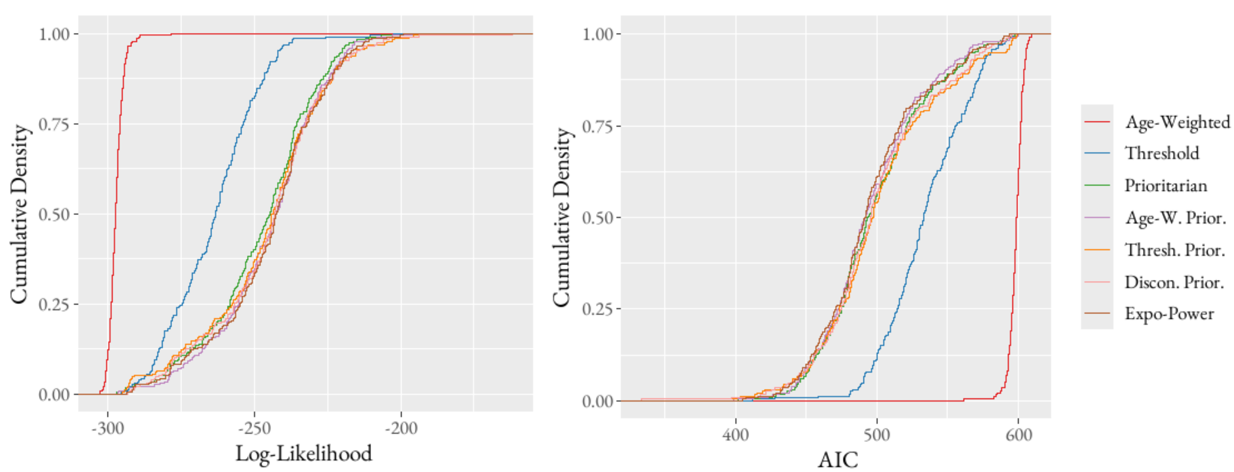

E.3 Noise and Log-Likelihood Distributions

Figure (E14) shows the distribution of participant-level noise parameters, σ, and the relationship between σ and log-likelihood values across all models, in the left and right panel, respectively. Results show heterogeneity amongst participants and between models. We see there is lower noise (lower σ) in the prioritarian models, models which include concavity, compared to threshold and, in particular, age-weighted models. The right panel shows a close relationship between the noise parameters and log-likelihood values: those with lower noise parameters have higher log-likelihood values.

Figure E14: Precision and Log-Likelihood, Across Models

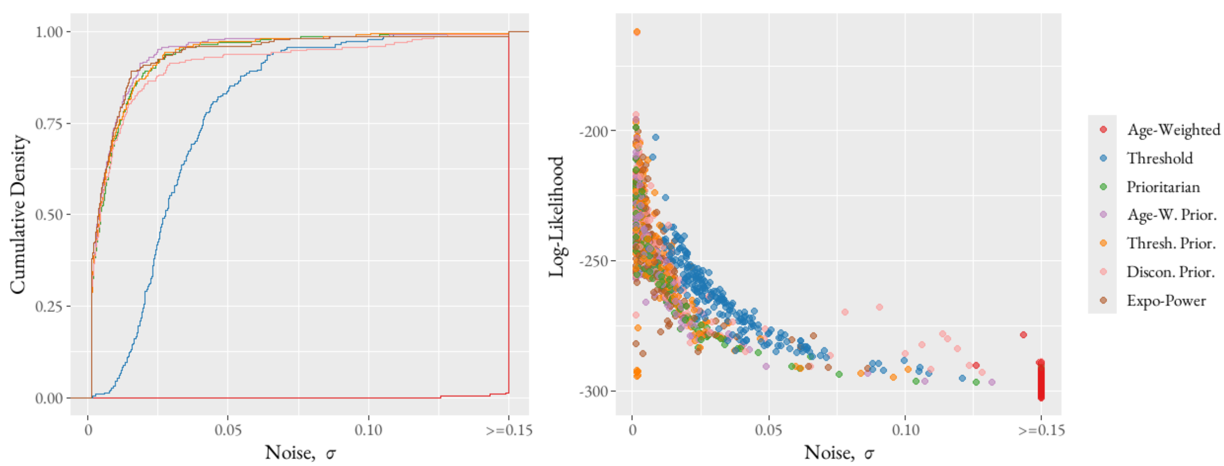

E.4 Sensitivity analysis: exclusion by tutorial questions

Table (E7) estimates differences in log-likelihood, according to whether participants were excluded (Include = 0) or excluded (Include = 1) from the main analysis. Results show that for the *prioritarian*, *age-weighted prioritarian*, *threshold-prioritarian*, *discontinuous-prioritarian* and *expo-power* models those included had higher LLs (i.e. a better fit), whilst there are no significant differences for the *threshold* or *age-weighted* models. As shown in our main analysis, the threshold and age-weighted models provide the worst fit, as compared to the other models. These results were expected. Participants who answer questions to check their understanding of the experiment correctly adhere more closely to predictions of these models, while the choices of those who were excluded are noisier.

Table E7: Sensitivity to Exclusion: Tutorial Questions

	Dependent variable: Log-Likelihood		
	Age-Weighted	Threshold	Prioritarian
	(1)	(2)	(3)
Include	−0.686	6.141	8.248**
	(0.858)	(3.926)	(3.888)
Constant	−296.328***	−269.836***	−255.679***
	(0.840)	(3.800)	(3.685)
Observations	244	244	244
R^2	0.004	0.009	0.011

	Dependent variable: Log-Likelihood			
	Age-Weighted Prioritarian	Threshold Prioritarian	Discontinuous Prioritarian	Expo-Power
	(4)	(5)	(6)	(7)
Include	8.072***	9.681**	11.270***	10.660***
	(0.858)	(3.926)	(3.888)	(3.806)
Constant	−252.904***	−255.567***	−256.248***	−255.194***
	(0.840)	(3.800)	(3.685)	(3.609)
Observations	244	244	244	244
R^2	0.011	0.011	0.017	0.017
Note:				*p<0.1; **p<0.05; ***p<0.01

E.5 Order Effects: Learning and Fatigue

The experiment involves multiple rounds within each treatment, and the (random) order in which participants complete these rounds may influence their decisions and, consequently, the estimates of their preferences. Two key mechanisms through which such order effects may arise are *learning* and *fatigue*. Learning could lead participants' choices to become less noisy over time, improving the goodness-of-fit of our estimated social welfare functions in later rounds, while fatigue could increase noise, causing a decline in goodness-of-fit as the experiment progresses. This appendix presents analyses that test for both of these effects.

We measure goodness-of-fit in each round, t, as the Mean Proportional Likelihood (MPL_t), by taking the average Proportional Likelihood across all participants, $r \in R$, in a given round (see eq.9):

$$MPL_t = \frac{1}{R}\sum_{r}^{R} \frac{L_{rt}}{L_{rt} + L_{rt}^{UNI}}. \tag{E3}$$

The MPL_t is an appropriate measure to test for order effects as a) preference parameters are estimated at the participant-level and are not estimated as a function of round number, and b) the round order (within each treatment) is randomised between participants. Any systematic difference in MPL_t across rounds will, therefore, provide evidence of learning and/or fatigue at the sample level.

Figure E15 plots MPL_t across rounds, with MPL_t coefficients and confidence intervals predicted using non-parametric dummy variables for each round. Here, we estimate preference parameters within the *prioritarian* social welfare function, for each participant. Results show that there is little systematic variation in MPL_t across rounds, within each session, indicating that there is no evidence of learning or fatigue across the rounds of the experiment.

Table E8 shows results of a random effects regression of MPL_t on the round, and round squared, within each session. Results show no significant differences across rounds within each session, confirming the above inference from Figure E15.[32]

[32] We do observe a significant difference between Session 1 (budget treatment) and Session 2 (age-equal and age-varying), with higher fit in Session 1. But note that these are different treatments, with the budget treatment as the less complex, with sessions approximately one week apart. This difference does not, therefore, indicate learning or fatigue within the experiment. We did not randomise these treatments, as more information was introduced in the treatments of Session 2, and so we cannot test for order effects between these treatments.

Figure E15: Mean Proportional Likelihoods by Round

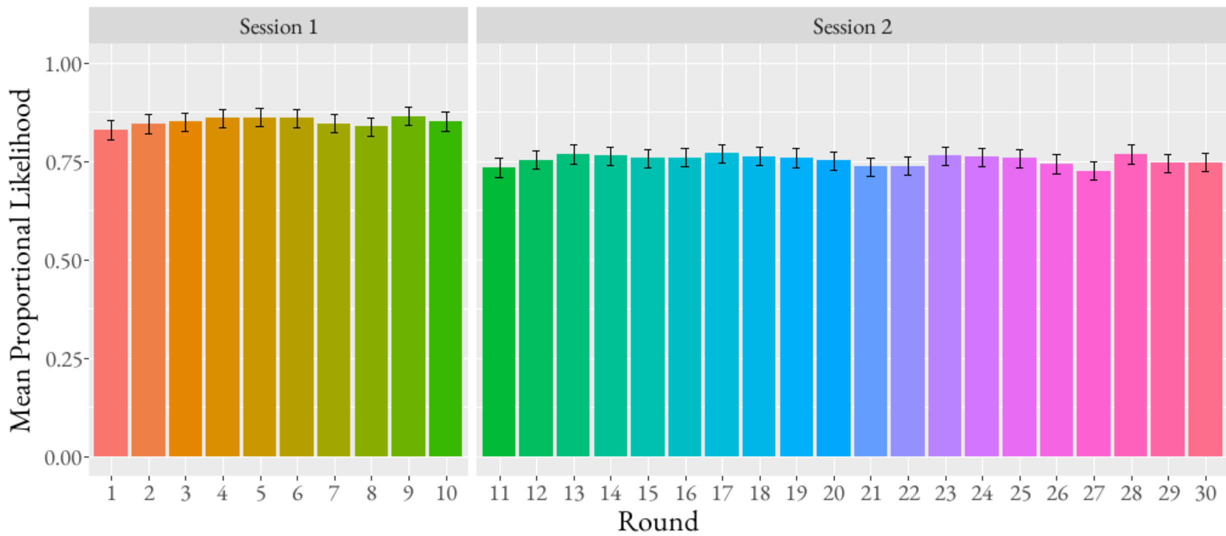

Table E8: Random Effects Regression on Proportional Likelihood

	Dependent variable:
	Proportional Likelihood
Session	−0.075***
	(0.018)
Round S1	0.009
	(0.006)
Round S1 (squared)	−0.001
	(0.001)
Round S2	0.001
	(0.002)
Round S2 (squared)	−0.0001
	(0.0001)
Constant	0.902***
	(0.033)
Observations	6,900
R^2	0.070
Note:	*p<0.1; **p<0.05; ***p<0.01