

Aversion to Health Inequality, Correlation and Causation

Owen O'Donnell [†] Matthew Robson [‡] Tom Van Ourti ^{§¶}

Draft - Please Do Not Cite or Circulate.

December 23, 2021

Abstract

We introduce a novel experiment that disentangles aversion to health inequality from aversion to health correlation, and that allows the latter aversion to be more intense when the correlation is causally determined. The experiment is a modified dictator game in which participants allocate resources that differentially impact recipients' health. In separate treatments, recipients are 1) anonymous, 2) identified by income, and 3) identified by income that causes health. These treatments identify parameters that determine aversion to 1) health inequality, 2) income-related health inequality, and 3) income-caused health inequality, respectively. In a student sample, we find aversion to health inequality that is lower than most previous estimates. On average, there is aversion to positive health-income correlation that intensifies when low income causes worse health. An income-rank-dependent social welfare function that respects relative invariance fits the data slightly better than one that respects absolute invariance, and both fit much better than a model in which health consequences of resource allocations are ignored.

Keywords: Inequality Aversion, Health, Income, Experiment, Social Welfare

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

[‡]Department of Health Sciences and Centre for Health Economics, University of York. Email: matthew.robson@york.ac.uk. Website: www.mrobson92.com.

[§]Erasmus School of Economics, Erasmus University Rotterdam, Tinbergen Institute: vanourti@ese.eur.nl.

[¶]This study was funded by the Erasmus Initiative Smarter Choices for Better Health and the Wellcome Trust (205427/Z/16/Z). The Internal Review Board of the Erasmus Research Institute of Management (ERIM) granted approval to conduct this experiment (2019/10/30-40140tvo). Part of the work was undertaken while Tom Van Ourti was a visiting researcher at the Milken Institute School of Public Health of the George Washington University. We would like to thank Erik Schokkaert, Richard Cookson and Tim Doran, alongside participants at the Applied Health Economics and Policy Evaluation workshop, IRDES; Health Economics Discussion Group seminar, University of York; and Measuring Health: Ethical Perspectives workshop, Aarhus University for useful comments and suggestions. We have no conflicts of interest to disclose. All errors remain our own. Please note that our main experiment concluded at the end of December, so an updated version will be available soon. See conclusion for more details.

1 Introduction

The socially optimal allocation of healthcare resources depends on willingness to sacrifice maximisation of aggregate health for less health inequality. The inequality aversion that determines this trade-off can take various forms. Differences in health can cause disquiet irrespective of who is in good health and who is sick. There may also be aversion to health differences between individuals distinguished by some other characteristic: gender, ethnicity or socioeconomic status, for example. The resource allocation consequences of this aversion to health correlation, or bivariate health inequality, differ from those arising from aversion to univariate health inequality. The latter implies prioritisation of the less healthy, while the former implies prioritisation of the less advantaged in some characteristic that is only imperfectly correlated with health.

Elicitation of social preferences over the distribution of health does not always distinguish aversion to health inequality from aversion to health correlation, and so risks confounding one with the other. We design an experiment to disentangle aversion to health inequality from aversion to health correlation. The experiment is a modified dictator game in which participants play the role of a social planner dividing resources between three individuals. The health impact of the resources varies between these individuals, which represents interpersonal differences in the health benefits derived from medical treatment (Richardson et al., 2012). This implies a trade-off between maximising total health and equalising its distribution. How a participant chooses to allocate the resources identifies their aversion to inequality. In one treatment, the participant is given no information about the identity of the three individuals. Choices under this treatment, which imposes anonymity, identify aversion to univariate health inequality. However, this parameter is potentially confounded if participants project identities onto the three individuals. In a second treatment, the participants are told the income of each individual, which is assigned randomly and is exogenous to the

(health) productivity of resources over multiple rounds. The extent to which the participants take this information on income into account in allocating the resources, as well as the attention paid to heterogeneity in resource productivity, identifies aversion to income-related health inequality separately from aversion to univariate health inequality.

We describe the experimental choices and the implied health inequality and health correlation in a non-parametric way. This avoids confounding by functional form, and therefore has general relevance for a broad set of social welfare functions (SWFs), but precludes identification of underlying preference parameters of particular classes of SWFs. We use the experimental choices to identify the parameters of a rank-dependent SWF consisting of a function that evaluates health and weights that potentially depend on income rank to capture aversion to income-related health inequality. Under constant weights, the model collapses to a utilitarian welfare function that reflects aversion to univariate health inequality through concavity of the health evaluation function.¹ We estimate pooled preference parameters, i.e. the preferences of the representative social planner, as well as the more common respondent-specific preferences (Bleichrodt, Rohde, and Ourti, 2012; Robson et al., 2017; Hurley, Mentzakis, and Walli-Attaei, 2020; Hardardottir, Gerdtham, and Wengström, 2019).²

Aversion to income-related health inequality may derive from a belief that positive correlation between health and income generates greater inequality in welfare than the inequality that exists in health and in income separately. Such a belief may lead to a willingness to allocate more health resources to the poor to compensate for their material disadvantage. But aversion to income-related health inequality could also derive from a belief that lower

¹In appendix A.5, we consider the rank-dependent SWF of Bleichrodt, Diecidue, and Quiggin (2004) which combines a concave utility function with health-rank dependent weights. We find that the choices of the respondents in our experiment are not in line with health-rank dependent weights and more closely resemble the assumption of constant weights. A full test of the health rank-dependent model is outside the scope of our experiment since our setup does not allow non-optimal health level allocations for some individuals in society.

²The references cited and research on the origins of fairness preferences (Cappelen et al., n.d.) have revealed large inter-individual differences in social preferences.

income causes worse health. Those who see injustice in such causality may be willing to allocate more health resources to the poor than they would be prepared to allocate in order to compensate for poverty through health.

Our experiment is designed to separate aversion to health-income correlation from aversion to health-income causation. To do this, a third treatment informs participants not only of the income of each individual but also how that income determines the productivity of resources in producing health. This includes the case where more income leads to better health, but also scenarios where more income leads to worse health. Trade-offs made in this treatment allow us to identify two components of the income-rank weights: one picks up any willingness to compensate low income through allocating more health resources, while the other reflects sensitivity to the causal impact of income on health productivity.

As far as we know, no study has elicited health inequality and health correlation aversion simultaneously. The more common procedure is to estimate them separately, with aversion to health inequality being captured by curvature of the health evaluation function (or the weighting of health by health rank) and aversion to health correlation captured by the weighting of health by income rank. This approach has revealed stronger aversion to health correlation than to health inequality (McNamara et al., 2019; Hurley, Mentzakis, and Walli-Attaei, 2020; Hardardottir, Gerdtham, and Wengström, 2019). Hurley, Mentzakis, and Walli-Attaei (2020) estimate aversion to health correlation and to health inequality separately, and do so by imposing a functional form on the weights (Yitzhaki, 1983; Wagstaff, 2002) that is more restrictive than our specification of beta-distributed weights. They report average aversion parameters in line with the Gini index (Yitzhaki, 1983) and the concentration index (Wagstaff, 2002), which masks the substantial heterogeneity that exists in individual-specific preferences.³

³Hurley, Mentzakis, and Walli-Attaei (2020) identify two groups. One with very strong inequality aversion, and another with mild to inequality-loving attitudes. But there is no group with the moderate aversion

This paper contributes in four ways. First, this is the first elicitation study that distinguishes aversion to univariate health inequality from aversion to bivariate health inequality (health correlation) while allowing the latter to depend on what causes the inequality. Aversion to health correlation that derives from an intrinsic (outcome-based) concern about whether healthy people are poor or rich can be distinguished from a partly procedural concern about the causes of inequality in the outcome (Fleurbaey and Schokkaert, 2009; Schokkaert and Devooght, 2003).

Second, our design is more general than those previously used to identify aversion to bivariate health inequality. In one treatment, we separately identify aversion to (univariate) health inequality from aversion to health correlation, while previous studies have only identified the latter (Hurley, Mentzakis, and Walli-Attai, 2020; Hardardottir, Gerdtham, and Wengström, 2019). This is important because restricting attention to attitudes to health correlation might fail to distinguish a concern for the income poor from concern about inequality in health, particularly when the experimental stimuli assume a positive correlation between income and health. As far as we know, all previous studies that elicit preference parameters present choice problems in which the income poor are also the least healthy.⁴ With this setup, failure to separate out health inequality aversion will lead to overestimation of aversion to health correlation. This might explain the common finding that aversion to health correlation is greater than aversion to health inequality (McNamara et al., 2019; Hurley, Mentzakis, and Walli-Attai, 2020).

Third, to the best of our knowledge, this is the first paper to use an optimal choice framework to jointly estimate parameters for aversion to health inequality and health correlation. Previously, aversion to health inequality (Bleichrodt, Doctor, and Stolk, 2005; Mc-

suggested by their average estimates. Similar heterogeneity in preferences is reported by Hardardottir, Gerdtham, and Wengström (2019)

⁴Bleichrodt, Rohde, and Ourti (2012) do consider both negative and positive health-income correlation, but this study does not elicit preference parameters. It tests axioms underlying health inequality measures.

Namara et al., 2019; Abásolo and Tsuchiya, 2020) and health correlation (Hurley, Mentzakis, and Walli-Attaei, 2020; Hardardottir, Gerdtham, and Wengström, 2019; McNamara et al., 2019) have been elicited by asking participants to choose between health distributions until a switch from one distribution to another reveals the shape of the social indifference curve. An advantage of our approach is that it allows for mistakes participants make, provided that, on average, behaviour is in line with the imposed SWF.⁵

Fourth, we test for the consistency of choices with various SWFs, including whether preferences are best represented by the Atkinson or Kolm-Pollak functional form: do respondents trade off relative or absolute health disparities against average health levels? We also test whether the social planner cares only about the allocation of resources or is a full consequentialist who considers the impact of resources on the distribution of health. And we test whether respondents weight health by the income ranks or the income levels of the recipients of the resources. Rank-based weights are commonly assumed in the measurement of bivariate health inequality (Wagstaff, Paci, and Doorslaer, 1991), but weights that depend on income levels allow for sensitivity to the magnitude of income differences (Erreygers and Kessels, 2017). We further analyze whether the weighting function of the extended Gini (Wagstaff, 2002; Yitzhaki, 1983), used by Hurley, Mentzakis, and Walli-Attaei (2020) and Hardardottir, Gerdtham, and Wengström (2019), is rejected by the data in favor of a more flexible weighting function derived from the beta distribution. This flexibility is required to allow for the symmetric weighting function proposed by Erreygers, Clarke, and Ourti (2012) that puts more weight on the health of those at the extremities of the income distribution than on the health of those in the middle.

⁵The non-parametric approach of Bleichrodt, Doctor, and Stolk (2005), which to some extent accommodates noise in choices, cannot be used in our setting because it exploits the link between health ranks — for the weighting function — and health levels — for the evaluation function. The corresponding link between health levels and income ranks, required for our approach, does not universally exist: it is a feature of social preferences that we want to elicit rather than impose.

We ran a laboratory experiment at Erasmus University Rotterdam, with the participation of 32 students each making choices in 31 rounds.⁶ In each round, participants had to divide a health resource budget between 3 individuals. The resulting within-participant variation is used to identify the aversion parameters. Our non-parametric estimates reveal that resources (and so health) are allocated equally when incomes and the productivity of resources in generating health do not differ between the individuals. On average, the health allocated to an individual (through resources) rises when they are poorer, more productive in turning resources into health, and both poorer and more (health) productive. These average responses hide substantial heterogeneity with a sizeable fraction of the participants penalizing the poor. In this pilot study, the results for treatment 3, in which participants are informed of each recipient’s income and its impact on health productivity, suggest that our current empirical strategy was not able to disentangle health correlation from health causation.

In addition to the non-parametric analysis, we estimate the health inequality aversion parameters and social weights that derive from imposing a rank-dependent SWF with beta distributed weights. Our estimates indicate that constant relative health inequality aversion lies between 0.5 and 1.3. We find that health inequality aversion is lower when simultaneous allowance is made for the weighting of health by income rank, suggesting that earlier studies may have consistently overestimated aversion to (univariate) health inequality. We find that the Atkinson functional form (relative inequality aversion) fits the response patterns better than the Kolm-Pollak functional form (absolute inequality aversion). Participants are consequentialists: they consider the productivity parameters when allocating resources, such that the optimal allocation of the budget depends on the ultimate impact on the distribution of health. Participant-specific parameter estimates reveal substantial heterogeneity.

⁶10 rounds for the first two treatments and 11 for the third treatment. There were also two practice questions.

Some display preferences consistent with “efficiency prioritarianism”, others are “weighted prioritarrians” and the majority of the participants follow a maximin strategy.

In addition to eliciting a number of aversion parameters, we illustrate how they can be used in practice to evaluate distributional impact by simulating several policy scenarios.

2 Theory

2.1 Social welfare function

Social welfare is assumed to be a weighted average of a concave transformation ($u(\cdot)$) of the health (h_i) of individuals (i) in a population of size N ,

$$W = \frac{1}{N} \sum_{i=1}^N \theta_i u(h_i), \quad (1)$$

where the weights, θ_i , can be specified to allow for the possibility that social concern about individuals’ transformed health depends on their non-health characteristics, such as income or another dimension of socioeconomic status (Makdissi and Yazbeck, 2016). Constant weights gives the utilitarian social welfare function (SWF). In this case, inequality aversion derives only from the concavity of $u(\cdot)$. This evaluation function can be thought of as a utility function. Then, with constant weights, aversion to inequality across individuals derives from the assumed diminishing marginal utility of health experienced by each individual. For ease of exposition we will sometimes refer to u as “utility”, but one can simply think of welfare as a social decision maker’s concave aggregation of health levels. We restrict attention to the iso-elastic function, $u(h_i) = \frac{h_i^{1-\varepsilon}}{1-\varepsilon}$, $\varepsilon \geq 0$, $\varepsilon \neq 1$. Welfare can be measured by the equally

distributed equivalent level of health,

$$h_{EDE} = \left(\sum_{i=1}^N \omega_i h_i^{-\eta} \right)^{-\frac{1}{\eta}} \quad (2)$$

where $-\eta = 1 - \varepsilon$, $-1 \leq \eta \leq \infty$ and $\eta \neq 0$, $\omega_i = \frac{\theta_i}{N}$, $\sum_{i=1}^N \omega_i = 1$ and $0 \leq \omega_i \leq 1, \forall i$. Strictly concave utility ($\eta \neq -1$) combined with constant weights implies social preferences that display aversion to health differences between individuals. The inequality aversion parameter, η , captures the trade-off the social decision maker is willing to make between health maximisation and its equal distribution. As η increases social welfare becomes more sensitive to the lowest levels of health. At $\eta = \infty$, preferences are maximin and only health improvements experienced by the least healthy raise welfare.

The individual-specific weights, ω_i , are allowed to be a general function of non-health characteristics that are possibly correlated with health and may even determine health through mechanisms considered to be fair or unfair. In the experiment, we ascertain whether the weights assigned depend on income. If they do, it may be because income is considered to be a causal determinant of health and this is judged to be an unfair source of health inequality. Alternatively, priority may be given to the health of poorer people to compensate for their economic disadvantage.⁷ The experiment is designed to distinguish between these two motivations for (possibly) assigning income-dependent weights. We do not restrict the weights to be decreasing in income. Some may prioritise the health of the economically better off.⁸

⁷This motivation implies that utility — or the social decision maker’s evaluation of individual well-being — is additively separable in health and the other attributes that determine individual well-being such that the socioeconomic distribution of health may be studied in isolation from the socioeconomic distribution of the other attributes. Additive separability is a strong assumption, but it is all that is needed to justify the focus on utility generated by health alone (Makdissi and Yazbeck, 2016), and is sufficiently flexible to evaluate the welfare consequences of interventions/programs that impact health of income groups differently, but not their income (which is exogenously given in our experimental setup).

⁸A possible motivation for such a preference could be a belief that the marginal value of health is increasing in income. This is not inconsistent with the assumed additive separability of utility because the

Preferences are elicited by asking respondents to allocate resources (y_i) and making the impact on the distribution of health evident. The health production function takes the form, $h_i = \pi_i y_i$, where the productivity parameter, π_i , can depend on income, or any other characteristic.⁹ Maximisation of social welfare, given by the equally distributed equivalent health, subject to a budget constraint, $m = \sum_{i=1}^N y_i$, results in the following optimal allocations:

$$y_i^* = \frac{m}{1 + \sum_{j \neq i}^N \left(\frac{\pi_i}{\pi_j} \left(\frac{\omega_j \pi_j}{\omega_i \pi_i} \right)^{\frac{1}{1+\eta}} \right)} \forall i \quad (3)$$

Within this framework, the experiment tests three specifications of the weights, ω_i , that distinguish aversion to pure (univariate) health inequality from aversion to income-related health inequality, and distinguish between two motivations for the latter type of inequality aversion.

2.1.1 Constant weights

If health is considered to be the only characteristic that should be relevant to the allocation of resources to health, then all individuals with the same health will be given equal weight:

$$\omega_i^A = 1/N. \quad (4)$$

change in social welfare from a change in health depends on the marginal utility of health and the social weight which depends on income.

⁹A linear production function is not particularly restrictive with respect to the elicitation of social preference. A concave function would be more realistic but would complicate the task for respondents. Our linear approach coincides with diminishing returns of the resources when the productivity parameters are measured on a concave scaling.

2.1.2 Income-dependent weights

Without distinguishing between motivations for aversion to income-related health inequality, this aversion can be captured by making the weights a general function of income (x_i):

$$\omega_i^B = f(\mathbf{x}_i) \quad (5)$$

The weighting function can be flexible subject to the two constraints stated immediately after eq.2. We choose to specify it as a beta function:

$$\omega_i^B = \int_{r(x_i) - \frac{1}{N}}^{r(x_i)} \left(\frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} (q^{a-1}(1-q)^{b-1}) \right) dq \quad (6)$$

where $r(x_i)$ is the individual's rank in the income distribution, $\Gamma()$ is the gamma function, and a and b are parameters to be estimated. The weight assigned to the health of an individual is given by the integral of the probability density function of the beta distribution over a region surrounding that individual's income rank. The assumption that it is the income rank, and not the income level, that is taken into account when prioritising the health of an individual is consistent with most of the literature on income-related health inequality that predominantly uses rank-dependent measures, such as concentration indices (Wagstaff, Paci, and Doorslaer, 1991; O'Donnell et al., 2008).

The weights can be restricted, and the estimation task made easier, by setting one or other of the beta function parameters to 1. We consider fixing $a = 1$, which rules out a bimodal distribution of weights, and so is inconsistent with symmetric inequality indices (Erreygers, Clarke, and Ourti, 2012), as well as saddle points implying that only functional forms such as those underlying the extended concentration index can be estimated. This rules out preferences that display concern for the poorest but not poorer individuals. It also does not permit greater concern for the health of both the poor and the rich than for the

health of those in the middle. Otherwise, the specification remains reasonably flexible in terms of the preferences that can be accommodated.

2.1.3 Causality-dependent weights

The specification consisting of eq.(2) combined eq.(6) allows aversion to pure health inequality — captured by the parameter η — to be separated from aversion to income-related inequality — captured by $\omega(x_i)$ — but it does not allow the strength of the latter type of inequality aversion to depend on why health is correlated with income. This can be corrected by allowing for the possibility that health is determined by income according to $h_i = \pi_i y_i = \mathbf{x}_i \Pi y_i$ and specifying a weighting function that takes the form

$$\omega_i^C = f(\mathbf{x}_i, c_1; \Pi, c_2) \quad (7)$$

where c_1 is the intrinsic weight given to an individual with \mathbf{x}_i irrespective of the health consequences of that income (i.e. when $\Pi = 0$) and c_2 captures concern about the causal impact of \mathbf{x}_i (i.e. when $\Pi \neq 0$).¹⁰ Specifically, the weights are obtained from the beta function with the restriction that $a = 1$ and $b = b_1 + b_2 \Pi$:

$$\omega_i^C = \int_{r(x_i) - \frac{1}{N}}^{r(x_i)} \left(\frac{\Gamma(1 + (b_1 + b_2 \Pi))}{\Gamma(1) \Gamma(b_1 + b_2 \Pi)} (1 - q)^{b_1 + b_2 \Pi - 1} \right) dq \quad (8)$$

The parameter b_1 captures intrinsic aversion to income-related health inequality irrespective of its source, while b_2 reflects any abhorrence of health being determined by income. The function also gives a counterfactual interpretation to the weights as indicating how much

¹⁰The specification restricts the marginal health product of income to be constant, for given resources. While this linearity may be unrealistic, it is imposed to ensure that the elicitation task is easily understood. Moreover, the linear specification can deal with decreasing returns to additional resources when the productivity parameters are interpreted as measured on a concave scale, essentially meaning that our setup does not allow distinguishing a linear from a concave functional form

more/less weight would have to be given to the income poor if they had a multiplier of x rather than y in order to hold social welfare constant.

2.2 Alternative social welfare functions

2.2.1 Nonconsequentialism

Only consequentialist respondents will allocate resources to maximise the SWF (1). Some may be nonconsequentialist. That is, they may focus purely on the allocation of resources and not consider the consequences for the distribution of health. We can accommodate such ethics by substituting y_i for h_i as the argument of the utility function in (1).¹¹ This resource-focused SWF yields an equally distributed equivalent allocation of resources,

$$y_{EDE} = \left(\sum_{i=1}^N (\omega_i y_i^{-\eta}) \right)^{-\frac{1}{\eta}}. \quad (9)$$

In this case, the optimal allocations are a function of the budget, the weights and (pure) inequality aversion, but not the health production function parameters:

$$y_i^{**} = \frac{m}{1 + \sum_{j \neq i}^N \left(\left(\frac{\omega_j}{\omega_i} \right)^{\frac{1}{1+r}} \right)} \quad (10)$$

2.2.2 Absolute invariance

The social preferences represented by the equally distributed equivalent health (2) allow for the possibility that the decision maker is prepared to sacrifice the maximisation of aggregate health for less *relative* health inequality. Since there is no consensus on whether relative or

¹¹This is not entirely nonconsequentialist since it still presumes that the decision maker evaluates the distribution of resources by applying a concave function, $u(\cdot)$, to them before aggregating to get welfare. Partial consequentialism may be a better description: no consequences for the distribution of health are considered in the allocation of resources.

absolute invariance is the more appropriate principle on which to base the measurement of health inequality, i.e. whether inequality is invariant to equiproportionate or equal absolute changes in health, we also consider the Kolm-Pollak family of SWFs that capture the trade-off between the maximisation of health and reduced *absolute* health inequality. This class of welfare functions, extended to allow the aggregation to depend on non-health characteristics through the weights, can be obtained by specifying $u(h_i) = -\exp(vh_i)$, $v < 0$ in (1). The respective equally distributed equivalent health is

$$h_{KP} = \frac{1}{v} \ln \left(\sum_{i=1}^N \omega_i \exp(vh_i) \right). \quad (11)$$

The allocation of resources that maximises this indicator of welfare is

$$y_i^{***} = \frac{m - \sum_{j \neq i}^N \left(\frac{1}{v\pi_j} \ln \left(\frac{\pi_i \omega_i}{\pi_j \omega_j} \right) \right)}{1 + \sum_{j \neq i}^N \left(\frac{\pi_i}{\pi_j} \right)} \forall_i \quad (12)$$

The nonconsequentialist who is concerned about absolute, not relative, inequality in resources would make an optimal allocation that is given by (12) with all production function parameters (π_i, π_j) set to 1.

2.2.3 Alternative income weighting functions

While restricting the weights, ω_i , to be determined by ranks in the distribution of income is consistent with the predominance of rank-dependent measures of income-related health inequality, some people may pay attention to an individual's level of income in deciding the priority to be given to that individual's health. Such ethics can be partially accommodated by applying the beta distribution to the income share, rather than income rank. Where the income share is each individual's income divided by total income: $\chi_i = \frac{x_i}{\sum_{i=1}^N x_i}$. By taking

the PDF of the beta distribution at χ_i we obtain v_i . These (intermediate) weights, v_i^{B1} are then normalised, to get ω_i^{B1} :

$$\omega_i^{B1} = \frac{v_i^{B1}}{\sum_i^N v_i^{B1}} \quad (13)$$

Where:

$$v_i^{B1} = \frac{\Gamma(1+b)}{\Gamma(1)\Gamma(b)} (\chi_i^{1-1}(1-\chi_i)^{b-1})$$

The weight given to each individual still depends on the full income distribution, through the income share, but now it takes some account of income quantities (not just positions), and so is more cardinal than the rank-dependent weight. This would build in sensitivity to relative income inequality. But it would not capture any attention that may be paid to the income gap between two individuals when prioritising their respective claims on health resources. Allowing for such distributional preferences, which are consistent with the level-dependent inequality indices proposed by Erreygers and Kessels (2017), requires specifying the weights using something other than the beta distribution.

An alternative, is to derive weights from an iso-elastic function:

$$\omega_i^{B2} = \frac{v_i^{B2}}{\sum_i^N v_i^{B2}} \quad (14)$$

Where:

$$v_i^{B2} = \frac{(x_i + 1)^{1-\varepsilon} - 1}{(1-\varepsilon)x_i}$$

This formulation is a re-centred iso-elastic function, divided by x_i . The weighting parameters, $\varepsilon \in \{-\infty, \infty\}, \neq 1$, determines the curvature of the function, where $\varepsilon = 0$ is linear, giving equal weights, $\varepsilon < 0$ is convex, giving higher weights to higher x_i , and $\varepsilon > 0$ is concave, giving lower weights to higher x_i . This allows any level of x_i to be input, and the level-dependant concerns of Erreygers and Kessels (2017) addressed.

3 Experiment Design

3.1 General setup

A stated-preference experiment based on a modified dictator game is used to elicit allocations of health resources from which SWF parameters are estimated. Participants assume the role of a social decision maker who has a budget to distribute amongst three individuals. To identify inequality aversion, η , the participants are forced to make trade-offs between the maximisation of total health and its equal distribution. They are asked to distribute resources, y_i , which when multiplied by individual-specific productivity factors, π_i , determine health, $h_i = \pi_i y_i$. Resources are constrained by a budget, $m = \sum_i^N y_i$. Participants thus face a constrained optimisation problem.¹² We consider three individuals to accommodate potentially nonlinear weighting functions.

Decisions are made through an onscreen interface using R Shiny, which is shown in Figure 1.¹³ In each round, the participant is given a budget¹⁴ and the productivity factors (called “multipliers”). They must allocate resources between the three individuals using sliders at the bottom of the screen.¹⁵ The chosen allocations and the resulting health outcomes are shown by the blue and black bars, respectively, and are shown numerically in the table at the top of the screen. Summary measures are shown to the right. The “Allocation Gap” is the largest absolute difference between individuals in the resources allocated to them. The “Health Gap” gives the equivalent for health. “Total Health” is the sum of the health outcomes. Participants make decisions in 31 rounds, with the budgets, productivity factors and labels of the individuals changing from round to round. In each round, respondents are forced to exhaust all resources, which is consistent with the elicitation of preferences

¹²The resources and resulting health values are cardinal, which allows participants to precisely distribute resources (and health) across the three individuals.


¹³The task is available online here

¹⁴The budget was randomly selected each round from the set 90, 210 and 300.

¹⁵Participants could also press and hold arrow keys to refine their allocations.

from solutions to an optimization problem. For the same reason, we don't show a baseline distribution which has the additional advantage that it avoids redistribution concerns.

Figure 1: Experimental Interface



The experiment is split into three within-subject treatments that are designed to enable identification of a) pure health inequality aversion, η , b) any prioritisation of the health of poorer individuals to compensate for their lower incomes, b_1 in the weight function (8), and c) any aversion to health inequality that is caused by income variation, b_2 in (8). In Treatment 1, participants are given no information on the identity of the three individuals. Anonymity is lifted in Treatment 2 by informing participants of the income of each individual. In Treatment 3, in addition to making incomes known, participants are informed of how these income determine the productivity factors that condition the effect of allocated resources on health.

3.2 Treatment 1: constant weights

In Treatment 1, participants are given no information about the identity of the individuals, who are referred to by the labels A, B and C. The productivity factors, π_i , change across 10 rounds; the combinations are shown in Table 1. The variation in π_i (in absolute and relative terms, $\pi_i / \sum \pi_i$) forces trade-offs between health maximisation and equalisation and so facilitates identification of η . The symmetry across labels through the 10 rounds ensures that comparisons between labels can be made without bias. Both the order of the 10 rounds and the order of labels is randomised across participants.¹⁶

Table 1: Productivity factors - Treatment 1

Round	Absolute, π_i			Relative, $\pi_i / \sum \pi_i$		
	A	B	C	A	B	C
1	1	1	1	0.33	0.33	0.33
2	1	1	0.33	0.43	0.43	0.14
3	1	0.33	1	0.43	0.14	0.43
4	0.33	1	1	0.14	0.43	0.43
5	1	0.33	0.5	0.55	0.18	0.27
6	0.5	1	0.33	0.27	0.55	0.18
7	0.33	0.5	1	0.18	0.27	0.55
8	1	0.5	0.33	0.55	0.27	0.18
9	0.33	1	0.5	0.18	0.55	0.27
10	0.5	0.33	1	0.27	0.18	0.55
Mean	0.7	0.7	0.7	0.33	0.33	0.33

Note: Columns give productivity factors assigned to individuals labeled A, B and C.

3.3 Treatment 2: income-dependent weights

In Treatment 2 anonymity is lifted. Participants are told the income of each individual. Rather than refer to an individual by the label “A”, a label of “5000 Euro” is used, for

¹⁶While theory predicts no order effects (i.e. order is irrelevant to choices), randomising the order of rounds between participants allows this to be tested. The order of the labels is randomised between participants to test and control for two potential order effects: A-B-C and left-middle-right.

example. For each participant, three incomes from the set (500, 1000, 2500, 5000, 10000) are randomly selected — without replacement — to label the three individuals.¹⁷ This allows differentiation between weights based on income ranks and weights based on income shares as different sets of three random draws will always lead to the same ranks, i.e. one sixth, three sixth, and five sixth, while the (differences between) income levels in one set of 3 random draws will differ.¹⁸ The productivity factors are identical to Treatment 1, allowing direct comparison between the treatments.

3.4 Treatment 3: Causality-dependent weights

In Treatment 3, in addition to telling participants the income of each the three individuals, they are also told how that income affects the productivity factor that multiplies the allocated resources to determine the health of each individual. Participants are told that the productivity factor is equal to a *base component* of 0.33 for all individuals plus a *causal component* determined by the income of each individual and a causal multiplier Π .¹⁹ The productivity factor is $\pi_i = 0.33 + f(\Pi, x_i)$, although it is not presented to the participants in this form. When $\Pi > 0$ ($\Pi < 0$), resources generate more (less) health when allocated to higher income individuals. When $\Pi = 0$ there is no relation between income and the productivity of resources. The causal multiplier is made known to the participant and varies between rounds, as shown in Table 2.

Variation in the allocations of resources that are made from round to round as the causal multiplier differs allows “intrinsic weights”, which potentially compensate individuals

¹⁷While this set of 5 incomes is more unequally distributed than the actual income distribution in the Netherlands, it ensures sufficient variation to identify income-dependent weights.

¹⁸The position (left-middle-right) of the income level labels is also randomised between participants.

¹⁹The productivity of resources in generating health can vary with income, or socioeconomic status more generally, in a way that differs across interventions. “Upstream” interventions, e.g. public smoking ban, often give a greater health bang for the buck among the poor, while “downstream” interventions, e.g. screening, are often more efficient among the better off.

for having lower incomes, to be separately identified from “causal weights”, which capture possible aversion to health differences that are caused by income.

While the set of π_i is different from that used in Treatments 1 and 2, the productivity factors are the same across all three treatments in particular rounds (1, 5, 8 and 9 for π_i , and 3, 7 and 10 for $\pi_i/\sum \pi_i$). This allows us to test whether the round-to-round allocations - and so the weights- differ depending on whether or not participants have information about the causes of π_i .²⁰ Table 2 shows how π_i depends on on the base (0.33), x_i and Π .

Table 2: Productivity Factors - Treatment 3

Round	Incomes			Productivity Factors			Π
	x_1	x_2	x_3	π_1	π_2	π_3	
1	X	Y	Z	0.33	0.5	1	1
2	X	Y	Z	0.415	0.5	0.615	0.3
3	X	Y	Z	0.5	0.5	0.5	0
4	X	Y	Z	0.615	0.5	0.415	-0.3
5	X	Y	Z	1	0.5	0.33	-1
6	V	V	W	0.33	0.33	1	1
7	V	V	W	0.5	0.5	0.5	0
8	V	V	W	1	1	0.33	-1
9	V	W	W	0.33	1	1	1
10	V	W	W	0.5	0.5	0.5	0
11	V	W	W	1	0.33	0.33	-1

Note: Left-hand panel indicates the incomes randomly assigned to each of the three individuals. Three random incomes X, Y and Z ($X < Y < Z$) are drawn from the set (500, 1000, 2500, 5000, 10000) and used for the first 5 rounds. Thereafter, V and W are randomly selected, without replacement, from X, Y and Z, such that $V < W$. The right-hand panel shows the productivity factors calculated from the incomes and then causal multiplier parameter, Π , given in the extreme right-hand column. The factors are calculated using incomes of X=?, Y=?, Z=?, V=? and W=?.

²⁰In order to obtain these comparison rounds, it is necessary to change the income levels between rounds. As in Treatment 2, the incomes x_1, x_2 and x_3 are chosen from the set: 500 , 1000 , 2500 , 5000 , 10000. However, after Round 5 only two of the three incomes are used. Table 2 shows how the incomes change across rounds. Three random incomes X , Y and Z ($X < Y < Z$) are drawn and used for the first 5 rounds. Thereafter, V and W are randomly selected, without replacement, from X , Y and Z , such that $V < W$. For example, if X = 1000, Y = 5000 and Z = 10000, then one possibility could be V = 1000 and W = 5000).

All respondents face the same chronology of treatments — first Treatment 1, next 2, and finally 3.²¹

3.5 Experiment particulars

The experiment was run in the ESE-econlab at Erasmus University Rotterdam. It was conducted over two sessions, both on November 28, 2019. The total number of participants was 32, 16 in each session. The sample was a student sample, recruited through the ESE-econlab internal system. Participants were paid a flat rate of €15. The approximate duration of the experiment was 1 hour.

Participants were given instructions, an interactive tutorial and a set of tutorial questions before beginning the experiment. The instructions gave a minimal overview, while the tutorial gave a more extensive description of the experiment to come. To check understanding, a set of questions about the experiment were asked at the end of the tutorial. The experiment had 31 rounds, one practice question and 10 rounds for each of the three treatments. We use the within-variation to gain statistical power and to consider consistency across rounds, but potential order effects between treatments were not addressed. It finished with a questionnaire about demographics and a set of feedback questions. Full details are found in Appendix A.1. The entire experiment was done on a computer interface, using R-Shiny software. Participants made decisions anonymously, within a private booth, which they were randomly allocated to at the start of the experiment. To encourage deliberation, minimum timers were placed on each round. All participants left at the same time.

²¹Presenting Treatment 3 (or 2) first might have made it difficult for the participant to contemplate Treatment 2 (or 1).

4 Estimation

We use a random behavioural model to estimate preference parameters by maximum likelihood. Estimation involves solving for the optimal allocation to each individual, y_i^* (y_i^* & y_i^*), consistent with the participant's choices and estimating parameters, η and those determining the weights, ω .

While the main model provides precise *optimal* allocations, y_i^* , for a particular decision problem and preference set, participants are assumed to err when calculating, or choosing, these allocations. Following Robson et al. (2017), we assume they draw their *actual* allocations, y_i , from the Dirichlet distribution (Dirichlet, 1839) where the expectation of the random variable equals the *optimal* allocation, $E[Y_i] = y_i^*$. The Dirichlet distribution is a multinomial Beta distribution, which allows for N variables, which here correspond to individual allocations (i.e. y_1, y_2, y_3), where $y_i \in (0, 1)$ and $\sum_{i=1}^N y_i = 1$. The Dirichlet distribution can be formulated as a random behavioural model as follows.

The following assumptions are made: (1) $E[Y_i] = y_i^*$, and (2) $Var(Y_i) = \frac{(y_i^*(y_0^* - y_i^*))}{s}$, therefore:

$$E[Y_i] = \frac{a_i}{a_0} = y_i^* \quad (15)$$

$$Var(Y_i) = \frac{(a_i(a_0 - a_i))}{(a_0^2(a_0 + 1))} = \frac{(y_i^*(y_0^* - y_i^*))}{s} \quad (16)$$

Where:

$$a_0 = \sum_{i=1}^N a_i, \quad y_0^* = \sum_{i=1}^N y_i^*$$

It follows that, $\forall i$:

$$y_i^*(s - 1) = a_i \quad (17)$$

The a_i 's determine the shape of the Dirichlet probability density function (*pdf*) and represent the weight given to a particular i . Precision, or how *noisy* decisions are, is represented

by s . Note that the higher the value of s , and therefore the higher α_0 , the lower the variance will be. The flexibility of the Dirichlet distribution is a useful property, and the above derivations allow for easily interpretable parameters to be estimated.

The preference parameters are estimated through maximising the following log-likelihood function. The preference parameters determine the optimal allocations, y_{it}^* , and consequently the shape parameters, a_{it} , in each round $t \in T$. Estimated parameters are those which maximise the log-likelihood function, hence are the ‘most likely’ fit for the observed data.

$$\sum_{t=1}^T \log \left(\frac{\Gamma \left(\sum_{i=1}^N a_{it} \right)}{\prod_{i=1}^N \Gamma(a_{it})} \prod_{i=1}^N y_i^{a_{it}-1} \right) \quad (18)$$

Due to the flexibility of the Dirichlet distribution, if $a_i < 1$, $\forall i$ the PDF is no longer unimodal, leading to singularity at the bounds. A penalty function is used when $a_i < 1$ to ensure parameter estimates exclude this possibility. For sample parameter estimates, the log-likelihood contributions for the decisions of every participant are pooled.

5 Results

Results are presented in two main sections: raw data resulting from the choices made by participants and the parameter estimates inferred from these choices. In the first of these sections, proportional health ($\tilde{h}_i = h_i / \sum h_i$) is used as the main outcome. This shows how participants respond to experimental variables: income and productivity factors.

5.1 Raw data

Table 3 shows the mean proportional health by the respective labels used to discriminate between the three individuals in treatments 1 and 2. In Treatment 1, which uses the labels

A, B and C, health is very close to equally distributed across the three individuals; indeed, no \tilde{h}_i is significantly different from 0.333 at the 10% level. In Treatment 2, however, the “poorest” of the three individuals receives a greater than equal share of health, 0.357, while the “richest” receives a less than equal share, 0.3136. Each of these allocations is significantly different from 0.333 at the 10% level.²²

Table 3: Mean Proportional Health

Label	Treatment 1	Income	Treatment 2
A	0.3296	Poorest	0.3568
B	0.3385	Middle	0.3296
C	0.3319	Richest	0.3136
No. observations	320		320

Note: Table shows the proportion of health ($\tilde{h}_i = h_i / \sum h_i$) allocated to each of the three individuals labeled either A, B or C (Treatment 1) or by income (Treatment 2) averaged across all participants (n=32) and rounds (10). In Treatment 2, the income (in €) attached to each individual varies randomly across participants. We refer to these incomes by the labels “Poorest”, etc. in this table.

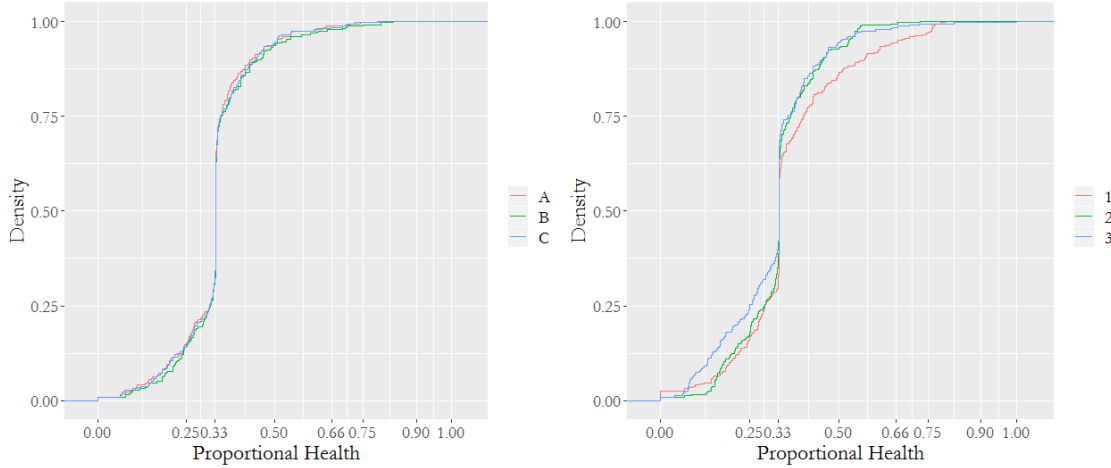
Figure 2 shows the empirical cumulative distribution of proportional health allocated to each of three individual for the Treatment 1 choices (left) and the Treatment 2 choices (right). It demonstrates that the differences between treatments are not restricted to the difference in means observed in Table 3. In Treatment 1, there are no differences between the distributions for individuals identified simply as A, B and C.²³ A large fraction of the choices made in this treatment involve allocating health approximately equally ($\bar{h}_i = 0.33$) across the three individuals. There is also a sizeable fraction allocating health equally in Treatment 2, but this choice is much less prevalent than in Treatment 1. The distribution for the poorest individual clearly lies to the right of that for the richest, indicating that a greater proportion of health is allocated to the poorest than to the richest. In fact, the distribution

²²While there are only 32 respondents in our study, the combination of 10 rounds per treatment with a within-subject design provides sufficient statistical power to reject the null of equal mean allocations at the 10% level.

²³Stochastic dominance tests do not reject the null of equality of distributions for all pairwise comparisons.

for the poorest individual stochastically dominates both of the other two distributions.²⁴ The distribution of the middle-income individual's proportional health stochastically dominates that of the richest individual, which is due the number of observations that allocate less than an equal share of health to the rich individual being greater than the number that give the middle-income individual less than an equal share.

Figure 2: Distributions of Proportional Health: Treatment 1 and 2



Note: Figure shows the empirical cumulative density function of proportional health ($\tilde{h}_i = h_i / \sum h_i$) allocated to each of the three individuals labeled either A, B or C (Treatment 1) or by income (Treatment 2). See notes to Table 3 for explanation of the income labels. Number of observations = 320 (32 participant \times 10 rounds).

Figure 2 pools data across rounds. Since the productivity factors vary from round to round (see Table 1), this does not reveal how allocations vary with those factors. Figure 3 presents empirical cumulative distributions of proportional health conditional on relative productivity, $\tilde{\pi}_i = \pi_i / \sum \pi_i$. A $\tilde{\pi}_i = 0.33$ indicates average productivity.²⁵ Allocating resources to an individual with $\tilde{\pi}_i > 0.33$ has a greater than average impact on health. Allocating

²⁴Dominance is tested using a two-stage procedure (Bennett, 2013) that first tests the null of equality against the alternative of inequality using a two-sided Kolmogorov-Smirnov (KS) test. If the null is rejected, then the minimum of the one-sided KS test statistics is used to distinguish between dominance in each direction and crossing of the distributions. We adopt a conservative approach (with respect to concluding dominance) of using a 1% level of significance at the first stage and a 10% level at the second stage. We adopt the critical values derived by Bennett (2013) for his asymptotic test.

²⁵ $\tilde{\pi}_i = \pi_i / \sum \pi_i = \pi_i / 3\bar{\pi}$, where $\bar{\pi}$ is the mean.

resources to someone with $\tilde{\pi}_i < 0.33$ is inefficient, in the sense that the health gain would be greater if the resources were given to another individual.

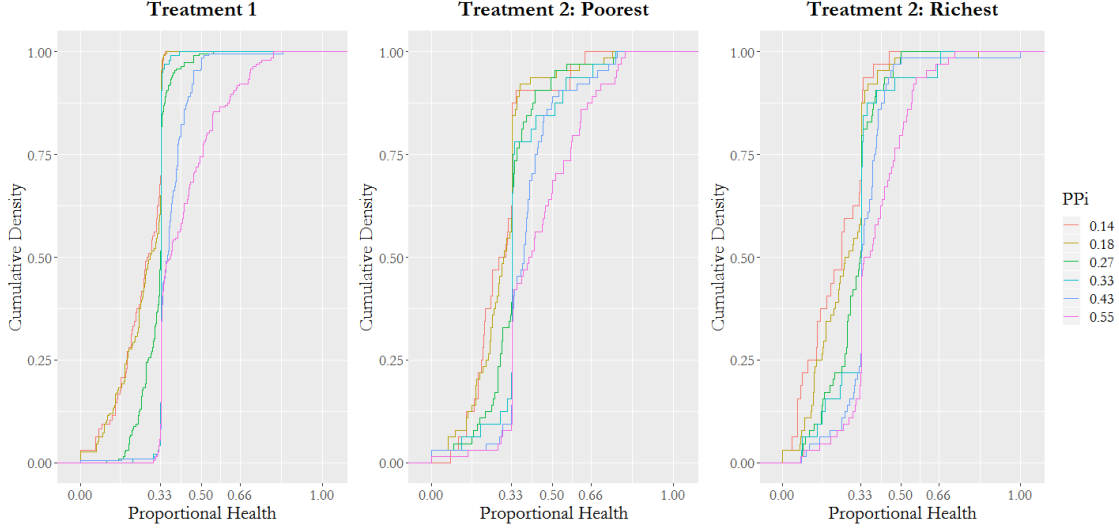
The left panel shows how health allocations change with relative productivity in Treatment 1. When an individual has the average productivity ($\tilde{\pi}_i = 0.33$) they are almost always allocated an equal (one-third) share of health (vertical jump in distribution at $\tilde{h}_i = 0.33$). As an individual's (relative) productivity increases (decreases), a greater (lesser) share of health is allocated to them — the distribution of \tilde{h}_i shifts to right (left) as $\tilde{\pi}_i$ increases (decreases). For given relative productivity, there is heterogeneity across participants in the health allocations made. This suggests that participants trade off efficiency in the production of health for equality in its distribution to varying degrees.

The other two panels of 3 show distributions of proportional health from Treatment 2 conditional on the income of the individual, as well as the relative productivity assigned to them. Unlike in Treatment 1, a significant proportion of participants do not give an equal (one third) share of health to those with average productivity ($\tilde{\pi}_i = 0.33$). Around 25% give the poorest individual with average productivity more than an equal share of health, but about 13% give such an individual less than an equal share of health.²⁶ Less than 25% give the richest individual with average productivity more than an equal share of health, and more than 13% give such an individual less than an equal share. This is evidence that, holding the health impact of resources constant, the number who discriminate in favour of the poor is greater than the number who favour the rich. The impact of productivity on the allocations also differs dependent on whether the recipient is poor or rich. The share of health, \tilde{h}_i , increases more steeply with an increase in relative productivity, $\tilde{\pi}_i$, when the recipient is poorer. And the rich are “penalised” more than the poor through the withdrawal of resources when productivity decreases. The trade off participants are willing to make when

²⁶One participant reasoned (in feedback) that a rich person benefits more from an additional year in good health than a poor person — a bivariate utilitarian, of sorts.

allocating resources between maximising health and equalising its distribution depends on the incomes of the recipients of those resources.

Figure 3: Distributions of Proportional Health by Productivity and Income




Note: Figure shows the empirical cumulative density function of proportional health ($\tilde{h}_i = h_i / \sum h_i$) separately for individuals assigned different relative productivity of health resources ($\tilde{\pi}_i = \pi_i / \sum \pi_i$). The middle and right-hand panels (from Treatment 2) further condition on the assigned income of the recipient of resources. See notes to Table 3 for explanation of the income labels.

Figure 4 combines the information displayed by Figures 2 and 3. It plots the proportional health allocated to one of the three individuals, i , against that given to one of the other two, j . Colours indicate the relative productivity assigned to i , $\tilde{\pi}_i$, and shapes denote that individual's label (Treatment 1 - left panel) or income (Treatment 2 - right panel). The simplex reveals the allocations to all three individuals since $\tilde{h}_k = 1 - \tilde{h}_i - \tilde{h}_j$, where $k \neq i, j$. For Treatment 1, moving from north-west to south-east, the allocations become more lightly coloured, indicating that the proportion of health allocated to i increases as this individual becomes relatively more productive. The shapes appear to be scattered randomly, which confirms that participants pay no attention to the labels A, B and C when allocating resources. The Y shape traced by some of the allocations is due to participants allocating equal health to two individuals with the same relative productivity but varying in

the extent to which they trade efficiency for greater equality in the allocation made to the third individual.²⁷

Figure 4: Relationship between Proportional Health and Relative Productivity



Note: Each figure shows the proportional health ($\tilde{h}_i = h_i / \sum h_i$) of individual i against ($\tilde{h}_j = h_j / \sum h_i$). Colour indicates relative productivity of i ($\tilde{\pi}_i = \pi_i / \sum \pi_i$). In the left panel (Treatment 1) shapes indicate the label (A, B or C) assigned to i . In the right panel (Treatment 2), shapes indicate the income assigned to i . See notes to Table 3 for explanation of the income labels.

The pattern in the right-hand panel, which shows allocations made in Treatment 2, is clearly different from that in the left-hand panel, indicating that participants pay attention to incomes when distributing resources. The Y shape disappears; two individuals who have the same relative productivity can have different incomes, which causes participants to allocate them different health. As in Treatment 1, proportional health increases with relative

²⁷Consider the allocation on the horizontal axis at $\tilde{h}_i = 0.5$. This arises from a participant dividing all the available resources equally between i and k , depriving j , who has lower productivity, from any health. This participant maximises health given the assigned productivities and accepts a great deal of inequality. With the relative productivity of i held constant, moving along the straight line in the north-westerly direction identifies allocations that sacrifice maximisation of the sum of health for greater equality. Similarly, the allocation at $\tilde{h}_i = \tilde{h}_j = 0.5$ is health maximising ($\tilde{\pi}_i = \tilde{\pi}_j > \tilde{\pi}_k$). Moving along the straight line in a south-westerly direction from there traces allocations that generate less total health but greater equality.


productivity (colours are lighter to the south-east). But this relationship appears to weaken somewhat; there are some dark circles to the south-east indicating large shares of health allocated to individuals with low productivity who are poor.

5.1.1 Budget Effects

The budget available to participants to distribute across individuals was randomised between 90, 210 and 300. Behavioural response to this variation in the budget (in addition that in the potential aggregate health created by changes in the multipliers) can be observed. Moreover, it allows for the testing of the Atkinson and Kolm-Pollack welfare functions against one another. The former relies on relative differences, meaning that no changes in \tilde{h}_i should be observed as the budget changes. The latter, on the other hand, is concerned with absolute differences, meaning that changes in \tilde{h}_i should be observed as the budget increases.

Figure 5 shows how the distribution of \tilde{h}_i changes across different budgets. The left, middle and right panels show responses when i has the lowest, mid and highest ranked multipliers, respectively. We do observe some differences when the budget changes. This means that the Atkinson index is not sufficient to explain behaviour for some individuals (if weights are equal), while the Kolm-Pollack could predict such behaviour. Most notably, in the left panel, with the lowest budget $\tilde{h}_i = 0$ when i has the lowest multiplier, for some observations. Intuitively, this means that when the budget is not sufficient to allow for an ‘acceptable minimum’ to be distributed amongst the three individuals, then all resources are taken from the least productive, in order to ensure the other two have sufficient amount of health. When the budget is sufficiently large, positive allocations are given to all participants.

Figure 5: Distributions of Proportional Health by Budget and Rank of Productivity: Treatment 1



Note: Figure shows.....

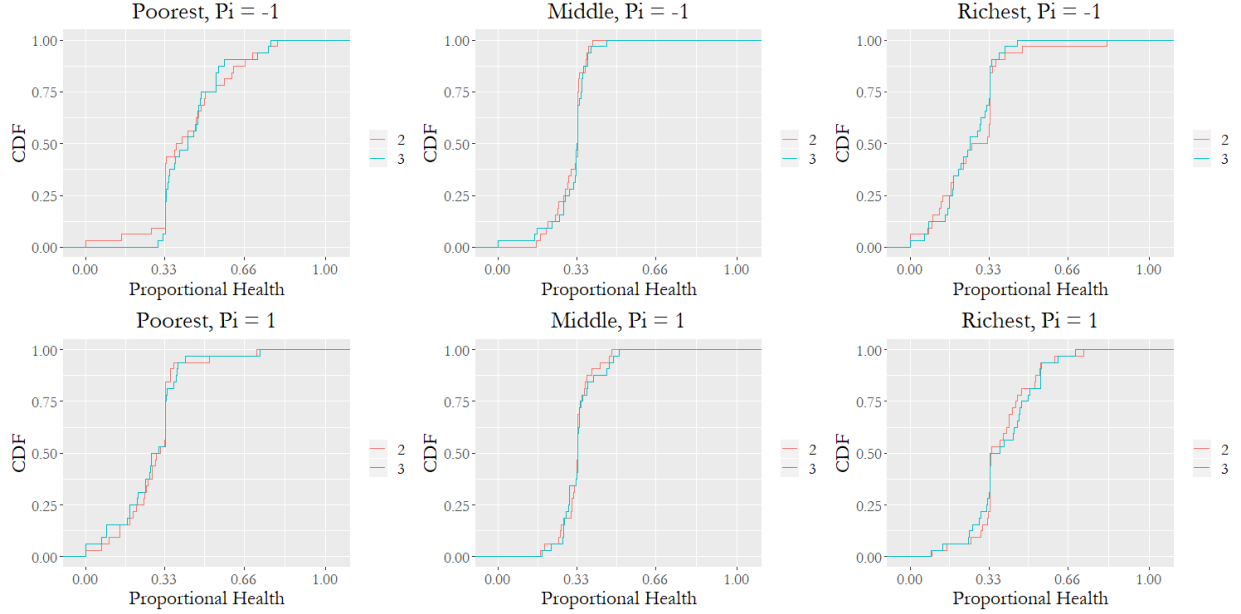
5.1.2 Causal Effects

The allocations made in Treatment 1 show that participants pursue efficiency — they allocate more resources to individuals whose health benefits most from them — but some sacrifice health maximisation for greater equality. Treatment 2 reveals that at least some participants pay attention to incomes in allocating resources. They are averse to income-related health inequality. We now turn to exploratory analysis of whether there are indications in the data that aversion to income-related health inequality depends on how this inequality is generated, We compare allocations made in treatments 2 and 3 for rounds in which the incomes and productivity factors are the same across the treatments. If causality does not matter, then the allocations will be the same in the two treatments.²⁸

Figure 6 shows empirical CDFs of proportional health allocated to the poorest, middle and richest individuals in treatments 2 and 3 (with constant incomes and productivity factors

²⁸Participants may presume that income is the reason for variation in the productivity factors in Treatment 2. Ideally, we would want to keep the overall productivity factors constant, but vary the base level and the casual component to tease out such differences.

Figure 6: Distributions of Proportional Health by Treatment, Income and Income-caused Productivity



Note: Figure shows the empirical cumulative density function of proportional health ($\tilde{h}_i = h_i / \sum h_i$) obtained from treatments 2 and 3 separately by assigned income. See notes to Table 3 for explanation of the income labels. The top panel shows distributions obtained with the strongest *negative* causal effect of income on health productivity in Treatment 3 ($\Pi = -1$). The bottom panel has the strongest *negative* causal effect of income on health productivity in Treatment 3 ($\Pi = 1$). Within each sub-figure, the income level and productivity factor is constant across the two treatments. Number of observations.....

across treatments). The top panel shows the distributions for the most extreme case in which lower income causes higher productivity ($\Pi = -1$). The bottom panel shows the other extreme — higher income causes higher productivity ($\Pi = 1$). There are no clear differences in the distributions arising from the two treatments. Looking closely at the figure in the top left panel, it appears that the poorest individual is less likely to be allocated a low share of health when their productivity is higher *because* they are poor (Treatment 3) as opposed to their productivity simply being higher (Treatment 2). This difference at the bottom of the distributions is not apparent when the poorest individual has lower productivity (bottom left figure) nor when the richest individual has higher productivity. This suggests that there may be less of a tendency to penalise the poor for being less productive in the generation of

health when this is caused by their low income. But the effect, if it exists, would not appear to be strong, at least from this graphical analysis.

5.2 Preferences

5.2.1 Pooled Estimates

Table 4 shows estimates of the preference parameters — health inequality aversion, η , precision, s , income-dependent weight determined by b_1 , and causality-dependent weight determined by b_2 — that are obtained by pooling data from the choices of all participants within each treatment. They can be interpreted as estimates of the preferences of the representative social decision maker. The estimate of η from the Treatment 1 responses indicates a modest degree of aversion to (univariate) health inequality. The point estimate implies that $\hat{\epsilon} = 1.52$, which is lower than most previous estimates of health inequality aversion, some of which approach double figures (refs).

The estimate of health inequality aversion obtained from Treatment 2 is roughly the same as that from Treatment 1. The point estimate of b_1 , which determines the income-dependent weight, is slightly larger than 1, implying that, on average, participants put lower weight on the health of richer individuals.

The choices made in Treatment 3 reveal a higher aversion to univariate health inequality. The estimate of the parameter that determines the non-causal component of the income-dependent weight increases slightly, indicating even greater priority is put on the health of poorer individuals irrespective of whether their poverty impacts on their health. The negative estimate of the parameter, b_2 , implies that when lower income causes the productivity of resources to increase ($\Pi < 0$), on average, participants give even greater weight to the income poor ($b_2\Pi > 0$). However, when effectiveness is positively correlated with income ($\Pi > 0$),

less weight is placed on the health of the poor ($b_2\Pi < 0$). In this case, this causal component offsets the intrinsic aversion to inequality in health that favours the rich.

Table 4: Preference Parameters - Pooled Estimates

Treatment	η	s	b_1	b_2
1	0.507	11.192		
2	0.574	7.827	1.050	
3	1.264	8.484	1.155	-0.226

Note: η is the parameter of the iso-elastic evaluation function $u(\cdot)$, (2). b_1 is determines income-dependent weights irrespective of causality, as in (6) and (8). b_2 determines causality-dependent weights, as (8)

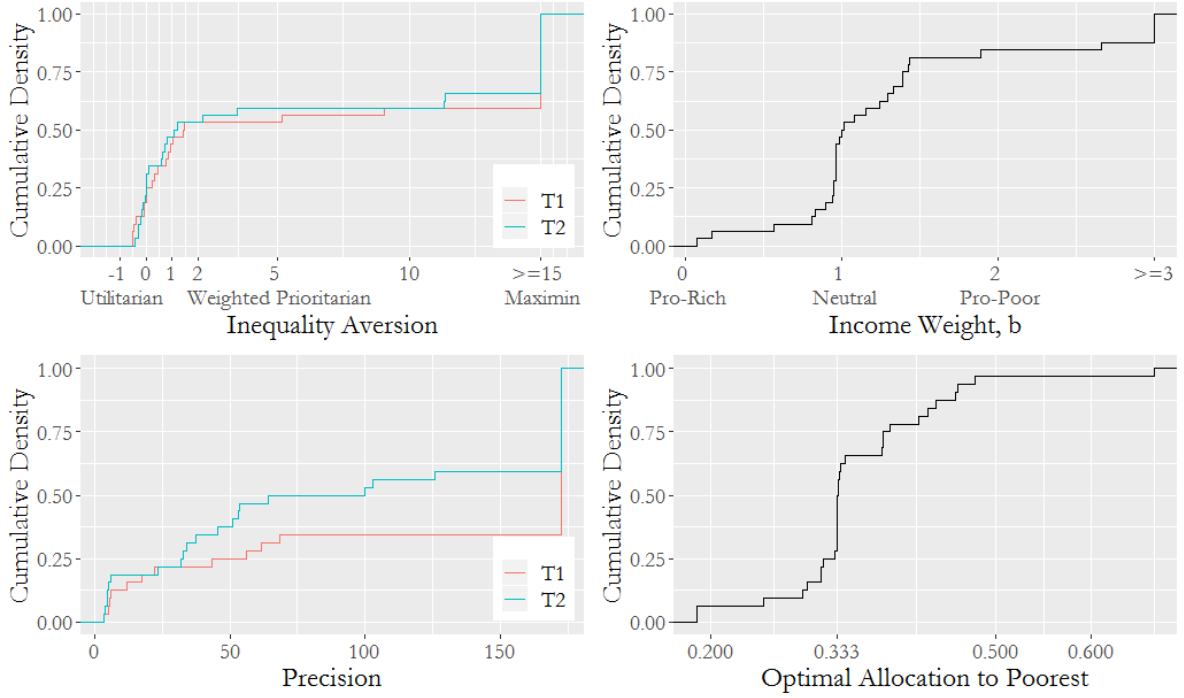
5.2.2 Participant-specific Estimates

The top panel of Figure 7 plots the empirical cumulative distributions of the inequality aversion parameter, η , estimated from treatments 1 and 2. The median η is 1.46. There are 12.5% of respondents who are Efficiency Prioritarians ($\eta \leq -0.1$), 12.5% who have Cobb-Douglas preferences ($-0.1 < r < 0.1$), 34.4% who are Weighted Prioritarians ($0.1 \leq r < 10$) and 40.6% who exhibit Maximin preferences ($r \geq 10$). Many in the sample display a great deal of health inequality aversion, but there is substantial heterogeneity in preferences.

The estimates of health inequality aversion, η , obtained from treatments 1 and 2 differ relatively little. The top panel of the figure shows that the empirical distributions largely overlap one another.

The top right panel of Figure 7 plots the empirical distribution of the income-dependent weight, b_1 , estimated from Treatment 2. The median b is 1.01, which implies a slight prioritisation of poorer individuals. More than a quarter (28.1%) of the sample give higher weight to the health of richer individuals ($b_1 < 0.95$), around a fifth (25%) are approximately neutral in weighting the rich and poor ($0.95 \leq b_1 \leq 1.05$) and more than half (46.9%) give poorer individuals higher weights.

Figure 7: Distribution of Health Inequality Aversion - Treatments 1 & 2



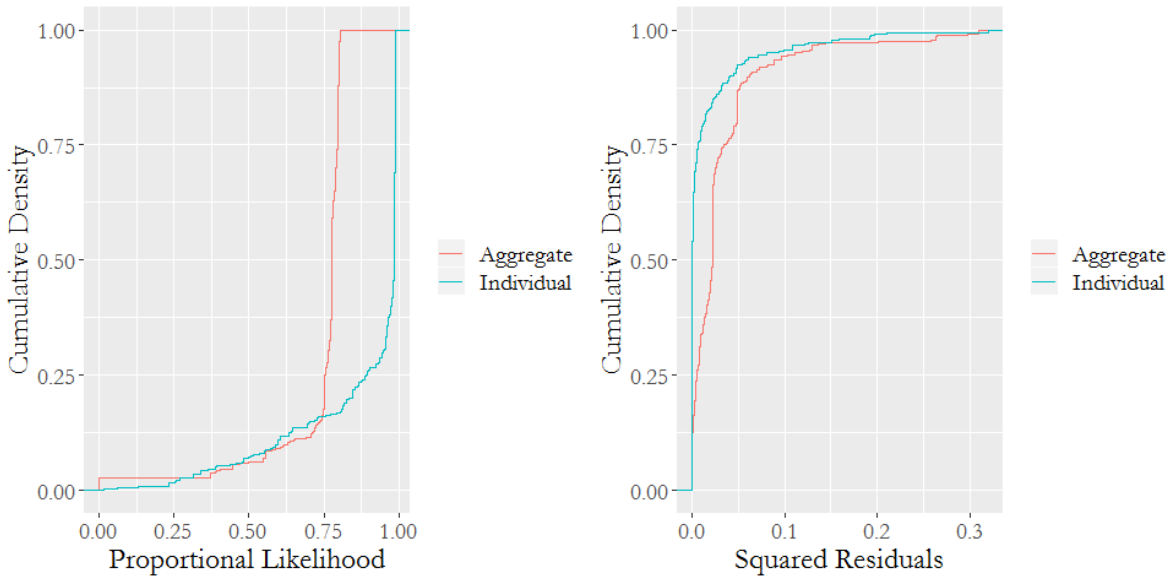
Notes:....

To provide intuition about the implications of these estimated parameters, the bottom-right panel shows the distribution of optimal health allocations, to the individual with the poorest rank, when the $m = 1$ and $pi_i = 1, \forall i$. This combines both inequality aversion and rank income weight parameters, to show how participants would optimally distribute health. Approximately a quarter, 25%, allocation less health to the poorest, ($h_i^* < 0.33$), 37.5% give an equal share to the poorest, ($0.33 \leq h_i^* < 0.34$), while a further 37.5% give proportionately more to the poor ($h_i^* > 0.34$).

The bottom-left panel of Figure 7 shows the distribution of precision, s , in each treatment. The median of s is 172.61 and 64.28 in Treatment 1 and 2, respectively. There are only 12.5% and 17.2% of participants for whom $s < 10$. Comparing this with the aggregate estimates of 11.19 and 7.83 reveals the extent to which allowing for heterogeneity across participants produces predicted allocations that are far more precisely centred around the optimal allocation.

To further highlight this difference between the pooled and participant-specific estimates, Figure 8 shows the distribution of the Proportional Likelihood and Squared Residuals, across all participants and rounds in Treatment 1. Both plots show the clear improvement in fit from estimating preferences at the participant level. The median Proportional Likelihood is 0.775 and 0.985, for pooled and participant-level estimates respectively. The median Squared Residuals are 0.022 and 0.000 respectively.

Figure 8: Distributions of Goodness of Fit - Treatment 1



5.2.3 Social Welfare Functions Comparison

To compare specifications of the social welfare function, data are pooled from treatments 1 and 2. Estimates presented in Table 6 are obtained by further pooling across participants to give representative decision maker preferences, and corresponding goodness of fit statistics, for three specifications of the welfare function: Atkinson (2), Kolm-Pollak (11) and non-consequentialist (9). The estimates of the health inequality aversion parameter, η , and the income-dependent weight parameter, b_1 , for the Atkinson function are very similar to the corresponding estimates from treatments 1 and 2 shown in Table 4. The negative estimate of

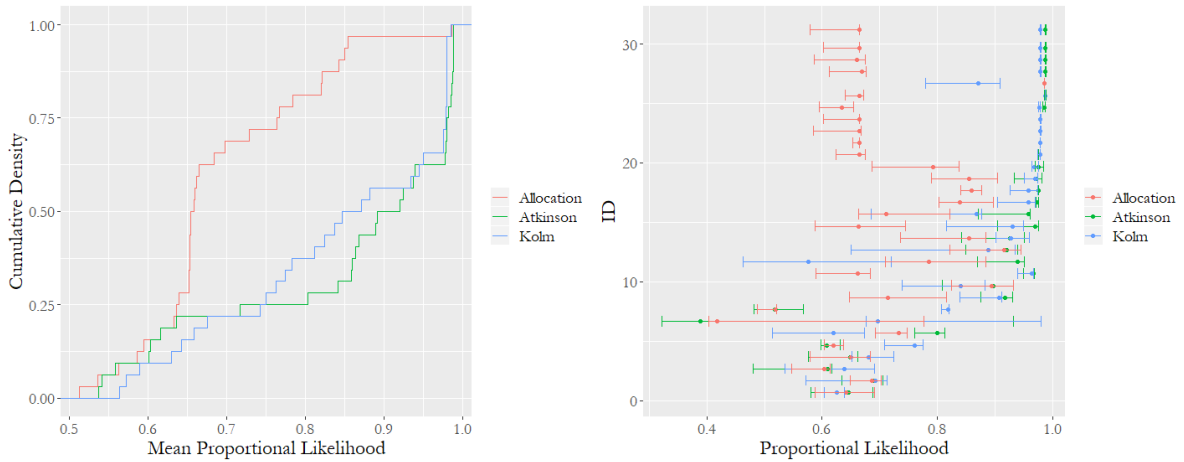
the Kolm-Pollak v parameter also indicates aversion to health inequality. This function and the nonconsequentialist function detect slightly greater income-dependent weights in favour of poorer individuals than the Atkinson function, although the differences are not substantial. More striking is that aversion to inequality in resource allocations estimated using the nonconsequentialist function is considerably greater than aversion to health inequality obtained from the Atkinson function.

Table 5: Social Welfare Functions - Pooled Estimates and Goodness of Fit

Treatment	η, v	s	b_1	MPL	MSE
Atkinson	0.541	9.150	1.050	0.691	0.040
Kolm	-0.046	9.831	1.100	0.682	0.041
Allocation	3.021	8.230	1.127	0.637	0.056
No. observations	?	?	?	?	?

Note: Estimates and goodness of fit statistics from pooling data from treatments 1 and 2 across all participants and rounds. η is the inequality aversion parameter of the iso-elastic evaluation function $u(\cdot)$ from the Atkinson (2) or nonconsequentialist (9) SWF. v is inequality aversion parameter from the Kolm-Pollak SWF (11). b_1 determines income-dependent weights in (6). MPL and MSE are the Mean Proportional Likelihood and Mean Square Error explained in the text.

Figure 9: Distribution of Proportional Likelihoods



We use the Mean Proportional Likelihood (MPL) and the Mean Squared Error (MSE) to compare how well the models fit the data. The MPL compares the likelihood value for

the observed data given the estimated model parameters with an equivalent likelihood for a random draw from a uniform distribution. The Proportional Likelihood in round t , is $P_{\iota_t} = \iota_t / (\iota_t + U_{\iota_t})$, where ι_t is the likelihood value for the data and U_{ι_t} is likelihood value for the uniform distribution draw. The MPL is $1/T \sum_t (P_{\iota_t})$. This can be thought of as a likelihood equivalent to the R^2 . As $MPL \rightarrow 1$ the observations have an increasingly high likelihood, while as $MPL \rightarrow 0$ the uniform random draws become relatively more likely. The MSE is the expected difference between the optimal allocation and the observed allocation. Both measures show that the Atkinson SWF best explains behaviour, closely followed by the Kolm-Pollak SWF, with the nonconsequentialist SWF the least good model.

5.2.4 Weighting Functions Comparison

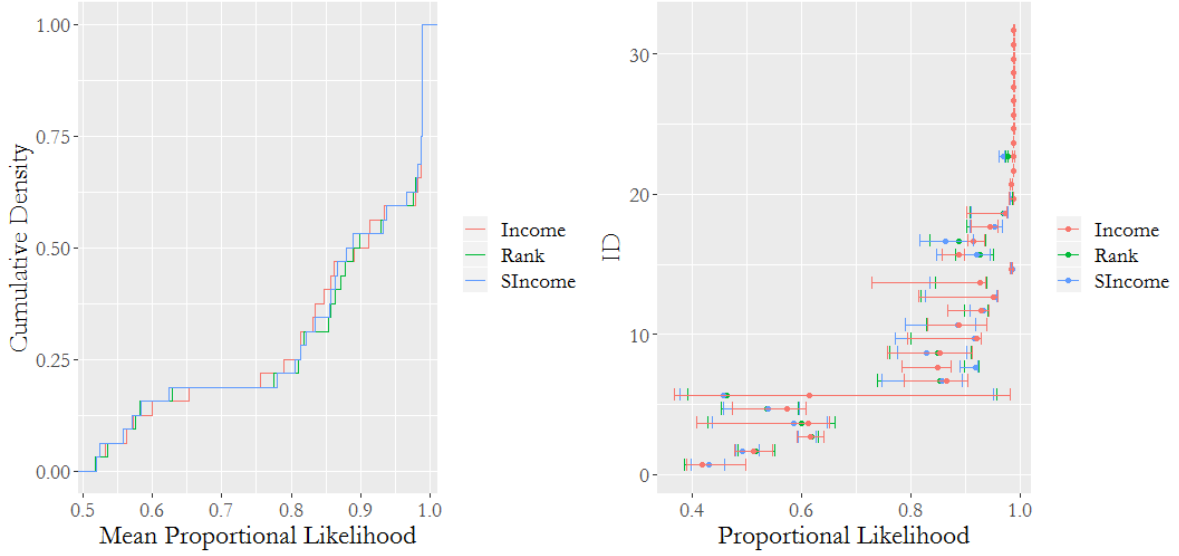
As above, to compare alternative weighting functions we use pooled estimates and generate goodness-of-fit statistics. Table 6 shows these results. Preference parameters for inequality aversion and precision are very close. The goodness of fit for each of the three alternative weighting functions is very similar. This implies that the specification of the weighting function is, here, not as important as the specification of the social welfare function.

Table 6: Alternative Weighting Functions - Pooled Estimates and Goodness of Fit

Weighting	η	s	b_1	MPL	MSE
Rank	0.574	7.827	1.050	0.6545	0.0477
Income Share	0.576	7.844	1.107	0.6547	0.0477
Income	0.576	7.841	0.051	0.6549	0.0474

Note: Estimates and goodness of fit statistics from pooling data from treatments 2 across all participants and rounds. η is the inequality aversion parameter of the iso-elastic evaluation function $u(\cdot)$ from the Atkinson (2). b_1 determines income-dependent weights in (6) across rank income, income share and income level. MPL and MSE are the Mean Proportional Likelihood and Mean Square Error explained in the text.

Figure 10: Distribution of Proportional Likelihoods



6 Social Choice and Policy Implications

Once behaviour is observed and preferences are estimated these can be used to evaluate policy. This could, however be done in multiple ways. A) we could see for a given set of external parameters; budget, multipliers and the size of the population we could estimate the optimal distribution of health that would be given. B) we could show a set of different possible policies and evaluate which policy is preferred. But, this can be done on multiple levels, using: individual preferences or social ones. Then the question emerges, how to use individual or social preferences?

6.1 Stylised Policy Evaluation

Let there be three individuals in society. The three individuals have income, x_i ; where $x_1 = 5000$, $x_2 = 20000$ and $x_3 = 50000$. Two policies, A and B, determine the health, h_i , of the individuals. The social planner has to choose from the two policies; shown in Table 7.

Table 7: Policy Outcomes

Policy	h_1	h_2	h_3	Mean	Gap
A	64	66	70	66.67	6
B	60	67	75	67.33	15

If the social planner is willing to incorporate the results from this experiment in their choice of policy, there are several decisions which need to be made. 1) Should they use aggregate or individual-level preference parameters? 2) Is income a relevant criterion to incorporate?²⁹ The combinations of the decisions made here are shown in Table 8. The choice of policy depends on the decision made.

Table 8: Preference Parameters and Optimal Policies

	Preferences		Optimal Policy	
	r	b	h_i	h_i & x_i
Representative	0.541	1.050	B	B
Median	1.344	1.000	A	A
Individual	[-0.413...213.1]	[0.019...32.05]	B (44%)	A (59%)

To make this example more realistic, these three “individuals” could easily be groups of individuals. The health and income levels could be average levels for those groups. The weighting parameters within the welfare functions would simply be the (weighted) proportion of the population within that group. If the estimation of potential health outcomes, across different groups, is difficult and therefore only group averages can be estimated, this approach is appropriate.

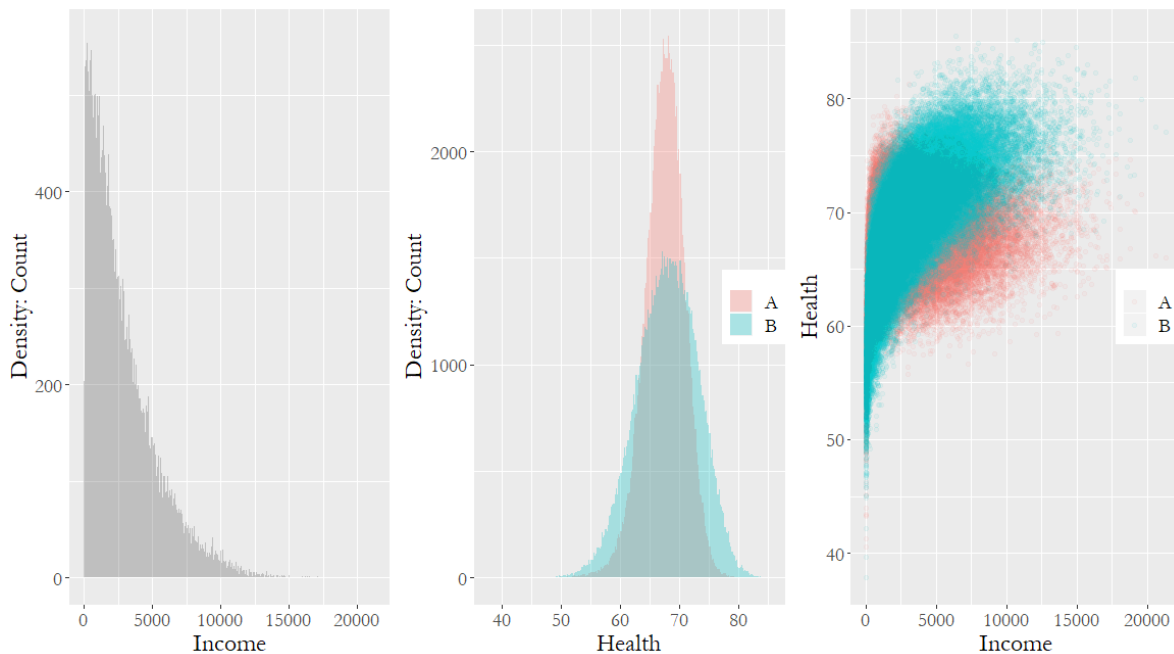
²⁹**Individual-level** - Pros: 1) much more accurate at representing what people actually prefer, 2) allows for percentage of vote. Cons: 1) Difficult to estimate individual-level, 2) more than 2 policy options voting issues. **Aggregate** - Pros: 1) Less individual-level data needed to estimate, 2) can easily rank any number of policies. Cons: 1) Coarse measure for a diverse set of preferences. **Median** - Pros: 1) For one dimensional preferences is equivalent to individual-level vote, 2) able to rank many policies. Cons: 1) when more than 1 preference parameter Frankenstein’s monster, 2) still need data to estimate individual-level preferences.

6.2 Individualised Policy Evaluation

However, with the advent of big data and development of more sophisticated econometric methods more disaggregated variables are observed and individualised potential health outcomes more easily estimated. In those situations the methods proposed here are even more useful and insightful.

Let there be a society with 100,000 individuals. Individuals have income, x_i . Health is a function of a constant, income and an error term, where those with higher income more productively transform income into health. The policies provide resources to individuals, these resources are again transformed into health via a productivity factor, which is a function of income. In Policy A the poorer an individual the more resources they receive, while in Policy B the richer individuals receive more resources.

Figure 11: Distribution of Income and Health: Resulting from Policy



As before the welfare of each policy can be estimated and the optimal policy chosen. For the case of the individual-level (no income) 40.6% prefer A, so B would be the optimal

policy. However, for individual-level when incorporating income it is 59.4% who prefer A to B, so policy A would be chosen. This mean that, the gain in average health of B over A is sufficiently high that individuals are willing to accept an increase in inequality, when ignoring the dimension of income. However, when accounting for the concentration of bad health amongst the income poor, the weights that individuals give to the income poor ensure that some change their mind and now prefer A, the more equal policy.

Although the experiment only uses three individuals, the preferences estimated within the utility functions and income-weighting functions proposed allows for the extrapolation of the analysis to an N dimensional society.³⁰

7 Discussion [Early draft]

7.1 Comparison to Previous Papers

Across the majority of experimental papers seeking to estimate health inequality aversion, with socioeconomic framing, there is a lack of symmetry. Decision problems are presented where the income poor have the lowest health. One issue with this experimental design is the inability to distinguish between weighting for the income poor and a concern for health equality.

In papers which elicit parameters by observing switching points (or indifference), a unique inequality aversion parameter can be elicited, but only if equal income-related weights are assumed. However, if weights were not assumed to be equal, then a unique parameter can not be elicited. Instead, preferences could be drawn from possible sets of weighting parameters and inequality aversion parameters (i.e. $b, \varepsilon = \{0.33, 9\} \dots \{0.34, 8\} \dots \{0.4, 2\}$). Generally speaking, with higher weights for low income groups, we would observe lower

³⁰Indeed, the reason to use three individuals within the experiment is to allow for the estimation of a non-linear income-weighting function; whether that be rank or level of income.

levels of inequality aversion. So the higher inequality aversion parameters found in income-specific vs neutral groups is partially explained by this.³¹

A contribution of our experimental design is to remove this framing of the poor income groups with lower health. The symmetry present allows for us to estimate health inequality aversion and income-related weights separately.

7.2 Rank-dependent versus prioritarian

We have considered social welfare functions that are both rank-dependent (i.e. they depend on a_i) and prioritarian (i.e. inequality aversion is derived from the r and v parameters). While rank-dependence (i.e. the a_i 's are not equal to $1/N$) has many advantages (refer to work of Han), one crucial consequence is that interventions cannot be evaluated by only considering the affected population when the ranks of the unaffected population are also affected (even though their health levels are not). We can not test this assumption in our experiment because our estimation procedure does not allow for this, i.e. it exploits the optimal allocation for every individual, and rules out cases where only part of the population are affected. Nevertheless, by varying the income levels and the productivity factors in table 1, we have considered cases with very different average income and health levels. Nevertheless, our estimates suggest that the inequality aversion parameter and a_i 's do not depend on these average income and health levels.

³¹Indeed, if individuals did actually prefer giving health to the poor, then it must be that their health inequality aversion is not approaching maximin. Otherwise they would always allocate equally, and never give priority to the poorest.

8 Conclusion

Through running a novel experiment we were able to disentangle aversion to health inequality from aversion to health correlation. Amongst a student sample we found lower levels of health inequality aversion, than in most previous studies; this still, however, reflected a significant degree of weighted prioritarianism. On average, participants are shown to have an aversion to positive health-income correlation that intensifies when low income causes worse health. This highlights the importance of separately accounting for aversion to univariate health inequalities and bivariate income-related health inequalities, particularly when low income causes these inequalities.

While this experiment was run with a small student pilot, we have only just completed our main experiment, at the end of December 2021. This experiment was run online and participants were recruited through Prolific. We have a representative sample of the UK population of 401 participants. The design of Treatments 1 and 2 are similar, but Treatment 3 has been substantially reworked, due to issues in the pilot. We are currently conducting the analysis, and will update this paper very soon. Initial results are promising and largely consistent with the findings of the pilot.

References

- Abásolo, Ignacio and Aki Tsuchiya (2020). “Comparing aversions to outcome inequality and social risk in health and income: An empirical analysis using hypothetical scenarios with losses”. In: *Health Economics* 29(1), pp. 85–97. DOI: [10.1002/hec.3974](https://doi.org/10.1002/hec.3974).
- Bennett, Christopher (2013). “Inference for dominance relations”. In: *International Economic Review* 54, pp. 1309–1328. DOI: [10.1111/iere.12038](https://doi.org/10.1111/iere.12038).
- Bleichrodt, Han, Enrico Diecidue, and John Quiggin (2004). “Equity weights in the allocation of health care: the rank-dependent QALY model”. In: *Journal of Health Economics* 23(1), pp. 157–171. DOI: [10.1016/j.jhealeco.2003.08.002](https://doi.org/10.1016/j.jhealeco.2003.08.002).
- Bleichrodt, Han, Jason Doctor, and Elly Stolk (2005). “A nonparametric elicitation of the equity-efficiency trade-off in cost-utility analysis”. In: *Journal of Health Economics* 24(4), pp. 655–678. DOI: <https://doi.org/10.1016/j.jhealeco.2004.10.001>.
- Bleichrodt, Han, Kirsten I.M. Rohde, and Tom Van Ourti (2012). “An experimental test of the concentration index”. In: *Journal of Health Economics* 31(1), pp. 86–98. DOI: <https://doi.org/10.1016/j.jhealeco.2011.12.003>.
- Cappelen, Alexander W. et al. (n.d.). “The Effect of Early Childhood Education on Social Preferences”. In: *Journal of Political Economy* 0(0), p. 0. DOI: <https://doi.org/10.1086/706858>.
- Dirichlet, Peter Gustav Lejeune (1839). “Sur une nouvelle méthode pour la détermination des intégrales multiples”. In: *Journal de Mathématiques Pures et Appliquées* 4, pp. 164–168.
- Erreygers, Guido, Philip Clarke, and Tom Van Ourti (2012). ““Mirror, mirror, on the wall, who in this land is fairest of all?”—Distributional sensitivity in the measurement of socioeconomic inequality of health”. In: *Journal of Health Economics* 31(1), pp. 257–270. DOI: <https://doi.org/10.1016/j.jhealeco.2011.10.009>.

- Erreygers, Guido and Roselinde Kessels (2017). “Socioeconomic Status and Health: A New Approach to the Measurement of Bivariate Inequality”. In: *International Journal of Environmental Research and Public Health* 14(7). DOI: [10.3390/ijerph14070673](https://doi.org/10.3390/ijerph14070673).
- Fleurbaey, Marc and Erik Schokkaert (2009). “Unfair inequalities in health and health care”. In: *Journal of Health Economics* 28(1), pp. 73–90. DOI: <https://doi.org/10.1016/j.jhealeco.2008.07.016>.
- Hardardottir, Hjördis, Ulf-G. Gerdtham, and Erik Wengström (2019). “What kind of inequality do you prefer? Evaluating measures of income and health inequality using choice experiments”. In: *Lund University Department of Economics Working paper* 2019:7.
- Hurley, Jeremiah, Emmanouil Mentzakis, and Marjan Walli-Attaei (2020). “Inequality aversion in income, health, and income-related health”. In: *Journal of Health Economics*, p. 102276. DOI: <https://doi.org/10.1016/j.jhealeco.2019.102276>.
- Makdissi, Paul and Myra Yazbeck (2016). “Avoiding blindness to health status in health achievement and health inequality measurement”. In: *Social Science & Medicine* 171(1), pp. 39–47. DOI: [10.1016/j.socscimed.2016.10.027](https://doi.org/10.1016/j.socscimed.2016.10.027).
- McNamara, Simon et al. (2019). “How averse are the UK general public to inequalities in health between socioeconomic groups? A systematic review”. In: *European Journal of Health Economics*. URL: <https://doi.org/10.1007/s10198-019-01126-2>.
- O’Donnell, Owen et al. (2008). “Analyzing Health Equity Using Household Survey Data : A Guide to Techniques and Their Implementation”. In: URL: <https://openknowledge.worldbank.org/handle/10986/6896>.
- Richardson, Jeff et al. (2012). “Maximising health versus sharing: Measuring preferences for the allocation of the health budget”. In: *Social Science Medicine* 75(8), pp. 1351–1361. DOI: <https://doi.org/10.1016/j.socscimed.2012.05.036>.
- Robson, Matthew et al. (2017). “Inequality Aversion, Self-Interest and Oneness: A Ugandan Lab-in-the-Field Experiment”. In: *DERS Discussion Papers* 17/12.

- Schokkaert, Erik and Kurt Devooght (2003). “Responsibility-sensitive fair compensation in different cultures”. In: *Social Choice and Welfare* 21, pp. 207–242. DOI: <https://doi.org/10.1007/s00355-003-0257-3>.
- Wagstaff, Adam (2002). “Inequality aversion, health inequalities and health achievement”. In: *Journal of Health Economics* 21(4), pp. 627–641. DOI: [https://doi.org/10.1016/S0167-6296\(02\)00006-1](https://doi.org/10.1016/S0167-6296(02)00006-1).
- Wagstaff, Adam, Pierella Paci, and Eddy van Doorslaer (1991). “On the measurement of inequalities in health”. In: *Social Science Medicine* 33(5), pp. 545–557. ISSN: 0277-9536. DOI: [https://doi.org/10.1016/0277-9536\(91\)90212-U](https://doi.org/10.1016/0277-9536(91)90212-U). URL: <http://www.sciencedirect.com/science/article/pii/027795369190212U>.
- Yitzhaki, Shlomo (1983). “On an Extension of the Gini Inequality Index”. In: *International Economic Review* 24(3), pp. 617–628. URL: <http://www.jstor.org/stable/2648789>.

A Appendix

A.1 Instructions, Tutorial Script and Tutorial Questions

The instructions, tutorial script and tutorial questions below are shown to all participants 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 understanding.

A.1.1 Instructions

Welcome. Thank you for taking part today.

Please Read These Instructions Carefully.

You will be asked to allocate resources between three individuals. The allocations you choose will determine the health of the individuals. You will make these allocation decisions on the screen.

The health an individual gains from an allocation of resources is determined by what we call a "multiplier". The lower the multiplier, the less health is gained from each additional resource allocated to an individual. We define Health as the number of years an individual will live in full health.

A budget determines the total amount of resources available. If you allocate more resources to one individual, there will be less available for the others. You must use all of the budget. You will be shown on the screen the gap in resources and in health between individuals, as well as the total amount of health. These will vary according to the allocations you choose.

You will be asked to choose allocations a number of times under different scenarios. The allocations you choose in one round will have no effect on subsequent rounds.

There are no right or wrong answers. We are interested in the decisions that you make.

After the experiment you will be asked to fill out a questionnaire. Your responses from the questionnaire, and from the entire experiment, will be kept strictly confidential.

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

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

Please click Next to continue.

A.1.2 Tutorial 1

This tutorial will show you how to use the on-screen interface. The choices you make during the tutorial will not be recorded.

At the bottom of the next screen is a horizontal slider. By dragging the slider to the right, you allocate more resources to Individual A.

The amount of resources you allocate to the individual is shown by the height of the blue bar in the chart above the slider. The allocation is also indicated by the number to the left of the bar and 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 allocation. Press the arrow key for a change of 0.1 and hold the arrow key for changes of 1.

The resources you allocate to an individual will be taken from the Budget, which is shown on the left of the screen. As you increase the allocation, the Remaining Budget will decrease. You must always Allocate all of the Budget, so that the Remaining Budget is zero. When there is only one individual, this means dragging the slider to the extreme right. Later you will have to divide the budget between individuals.

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

A.1.3 Tutorial 2

The Allocation of resources you give to an individual determines their Health. Health is equal to the Allocation multiplied by a number we call the Multiplier. Health is defined as the number of years an individual will live in full health.

The Multiplier is shown in the table at the top of the screen. When you make an Allocation by moving the slider, the resulting Health is shown by the height of the grey bar. The number to the right of this bar also indicates Health, and this number is also shown in the table at the top of the page.

In this first scenario, the multiplier is 1. So, when you allocate the total budget of 100 to the individual, their Health is 100. They are expected to live 100 years in full health.

Press Next and see how Health changes as you adjust the Allocation. When you are done, allocate all of the Budget and press Next.

A.1.4 Tutorial 3

The Multipliers will change across the scenarios you are presented with. In the previous scenario the Multiplier was 1. Now the Multiplier is 0.5.

Press Next and see how Health changes as you adjust the Allocation. When you are done, Allocate all of the Budget and press Next.

A.1.5 Tutorial 4

Now all three Individuals are shown. There is a slider for each individual at the bottom of the screen that you use to allocate resources to each individual. Each Allocation is multiplied by the individual's Multiplier, which is shown in the table at the top, to give the resulting Health of that individual.

The Allocations and Health of the three Individuals are shown by the bars and numbers in the three columns.

Now you must allocate the Budget across the three individuals. In doing so, you determine the Health of the three individuals. You must use up all of the Budget, so that the Remaining Budget (shown on the left) equals zero. You cannot Allocate more than the Budget.

Press Next and then make Allocations across the three individuals. Remember: you can use the left and right arrow keys to make small changes to an Allocation. When you have allocated all of the Budget, press Next.

A.1.6 Tutorial 5

The Multipliers will change across scenarios. In the previous scenario, they were all equal to 1. But they can differ. As they do in the next scenario.

Move the sliders and notice how the Health resulting from an Allocation varies depending on the Multiplier. If you make the same Allocation to Individuals that have different Multipliers, they will get different levels of Health.

Press Next and then make Allocations across the three individuals. When you have Allocated all of the Budget, press Next.

A.2 Tutorial 6

On the right of the screen you are shown further information on the consequences of the allocations you choose.

“Allocation Gap” is the gap in resources between the Individuals resulting from the Allocations you choose (largest Allocation - smallest Allocation).

“Total Health” is the total amount Health across the three individuals resulting from your Allocation (Health to A + Health to B + Health to C).

“Health Gap” is the gap in Health between the individuals resulting from the Allocation you choose (largest Health - smallest Health).

Also notice that if you have Allocated the whole Budget, you will not be able to move any slider to the right. If you want to increase the Allocation to one Individual, you will need to reduce the Allocation to another Individual first.

Press Next and then make Allocations across the three individuals. When you have Allocated all of the Budget, press Next.

A.2.1 Tutorial Questions

Following the tutorial participants were presented with 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. Amongst how many Individuals will you be Allocating resources? - Options: 2; **3**; 4; Not Sure.
2. You can make and adjust Allocations by (tick all that apply): - Options: **Clicking and Dragging the Allocation Sliders**; **Using the Arrow Keys**; Moving the Vertical Bar; Not Sure.
3. If the Allocation to an Individual is 100 and their Multiplier is 1, then their Health will be? - Options: 25; 50; **100**; Not Sure.
4. If the Allocation to an Individual is 100 and their Multiplier is 0.5, then their Health will be? - Options: 25; **50**; 100; Not Sure.
5. Once you have finished making Allocations you proceed to the next scenario by: - Options: Clicking Next; **Ensuring the Remaining Budget = 0, then Clicking Next**; Waiting; Not Sure.

A.3 Regression Results

Table 9

<i>Dependent variable:</i>	
Proportional Health	
Position Mid	0.008 (0.009)
Position Right	-0.004 (0.008)
Label B	0.009 (0.008)
Label C	0.001 (0.008)
Round	-0.000 (0.001)
Constant	0.328*** (0.010)
Observations	960
<i>Note:</i> *p<0.1; **p<0.05; ***p<0.01	

Table 10

	<i>Dependent variable:</i>		
	Proportional Health		
	(1)	(2)	(3)
Income Middle	-0.027*** (0.010)		
Income Richest	-0.043*** (0.010)		
Income		-0.059*** (0.016)	-0.230*** (0.071)
Income Squared			0.212** (0.087)
Round	-0.000 (0.001)	-0.000 (0.001)	0.000 (0.001)
Constant	0.357*** (0.024)	0.353*** (0.023)	0.371*** (0.025)
Observations	960	960	960
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01		

Table 11

	<i>Dependent variable:</i>		
	Proportional Health		
	(1)	(2)	(3)
Prop. Pi	0.457*** (0.020)	0.412*** (0.026)	0.374*** (0.027)
Prop. Income		-0.059*** (0.014)	-0.169* (0.097)
Prop. Causal Pi			-0.004 (0.111)
P.Pi X P.Income		-0.043 (0.099)	-0.053 (0.096)
P.Causal Pi X P.Income			-0.518 (0.490)
Constant	0.333*** (0.003)	0.333*** (0.004)	0.333*** (0.022)
Observations	960	960	1,056
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01		

A.4 Bootstrapped Aggregate Preference Parameters

Table 12: Preference Parameters - Bootstrapped Estimates

Treatment	η	s	b_1	b_2
1	0.5281 (0.242 - 0.928)	11.715 (8.332 - 17.508)		
2	0.5943 (0.317 - 0.992)	8.086 (6.125 - 11.043)	1.0480 (0.938 - 1.171)	
3	1.378 (0.311-2.755)	8.793 (6.709-12.205)	1.2099 (1.029-1.547)	-0.2638 (-0.748 0.215)

Note: These are bootstrapped parameter estimates. The parameter estimate is the mean of 1000 bootstrapped iterations. The confidence intervals are the values at the 2.5th and 97.5th percentiles.

A.5 Alternative constant weighting functions

While the parameter η captures aversion to pure health inequality, our approach of specifying constant weights is subject to the criticism of Bleichrodt, Diecidue, and Quiggin (2004): that without health-rank-dependent weights no distinction is made between inequality aversion and the valuation of health captured by the concavity of utility. Indeed, other models allow weights to vary with the rank in the distribution of health, as in the the rank-dependent QALY model (Bleichrodt, Diecidue, and Quiggin, 2004; Bleichrodt, Doctor, and Stolk, 2005). However, in our experimental setup, participants will be choosing optimal health levels, which will themselves determine the health ranks. If the weights are a function of health ranks, these optimals will be intractable as the two are endogenous: the health rank depends on the health levels chosen. An alternative, which addresses these concerns, while maintaining tractability, is to have the weights be a function of the rank of productivity factor. Reflecting the “potential” health rank an individual would have. Which, as before, we choose to specify as a beta function:

$$\omega_i^P = f(\pi_i) = \int_{r(\pi_i) - \frac{1}{N}}^{r(\pi_i)} \left(\frac{\Gamma(1+b)}{\Gamma(1)\Gamma(b)} (q^{1-1}(1-q)^{b-1}) \right) dq \quad (19)$$

Using this alternative weighting we can estimate sample level preference parameters, pooling data from treatment 1 across all participants and rounds. The preferences estimated and goodness of fit statistics which compare this alternative weighting function to the standard equal weighting are shown in Table 13.

Results show that b is positive, meaning that participants give a higher weight to individuals with lower productivity factors. The inclusion of b also appears to slightly reduce the estimate of η . However, what we observe from the goodness of fit statistics (a lower MPL and higher MSE) is that the addition of b actually provides a marginally worse fit

Table 13: Alternative Constant Weighting Function - Pooled Estimates and Goodness of Fit

Weighting	η	s	b	MPL	MSE
Equal	0.507	11.192	.	0.7343	0.0319
Productivity	0.408	11.193	1.045	0.7341	0.0320

Note: Estimates and goodness of fit statistics from pooling data from treatments 1 across all participants and rounds. η is the inequality aversion parameter of the iso-elastic evaluation function $u(.)$ from the Atkinson (2). b_1 determines rank productivity factor dependent weights in (19). MPL and MSE are the Mean Propotional Likelihood and Mean Square Error explained in the text.

to the data. Therefore, within the main analysis we will maintain the assumption of equal weighting when individuals are anonymous.