Health Inequality Aversion and Rank-Dependent Income Weights in the Netherlands

EsCHER Working Paper
No. 2026005
Date: July 2026

Matthew Robson
Matthias Harlass
Chiara Celine Brück

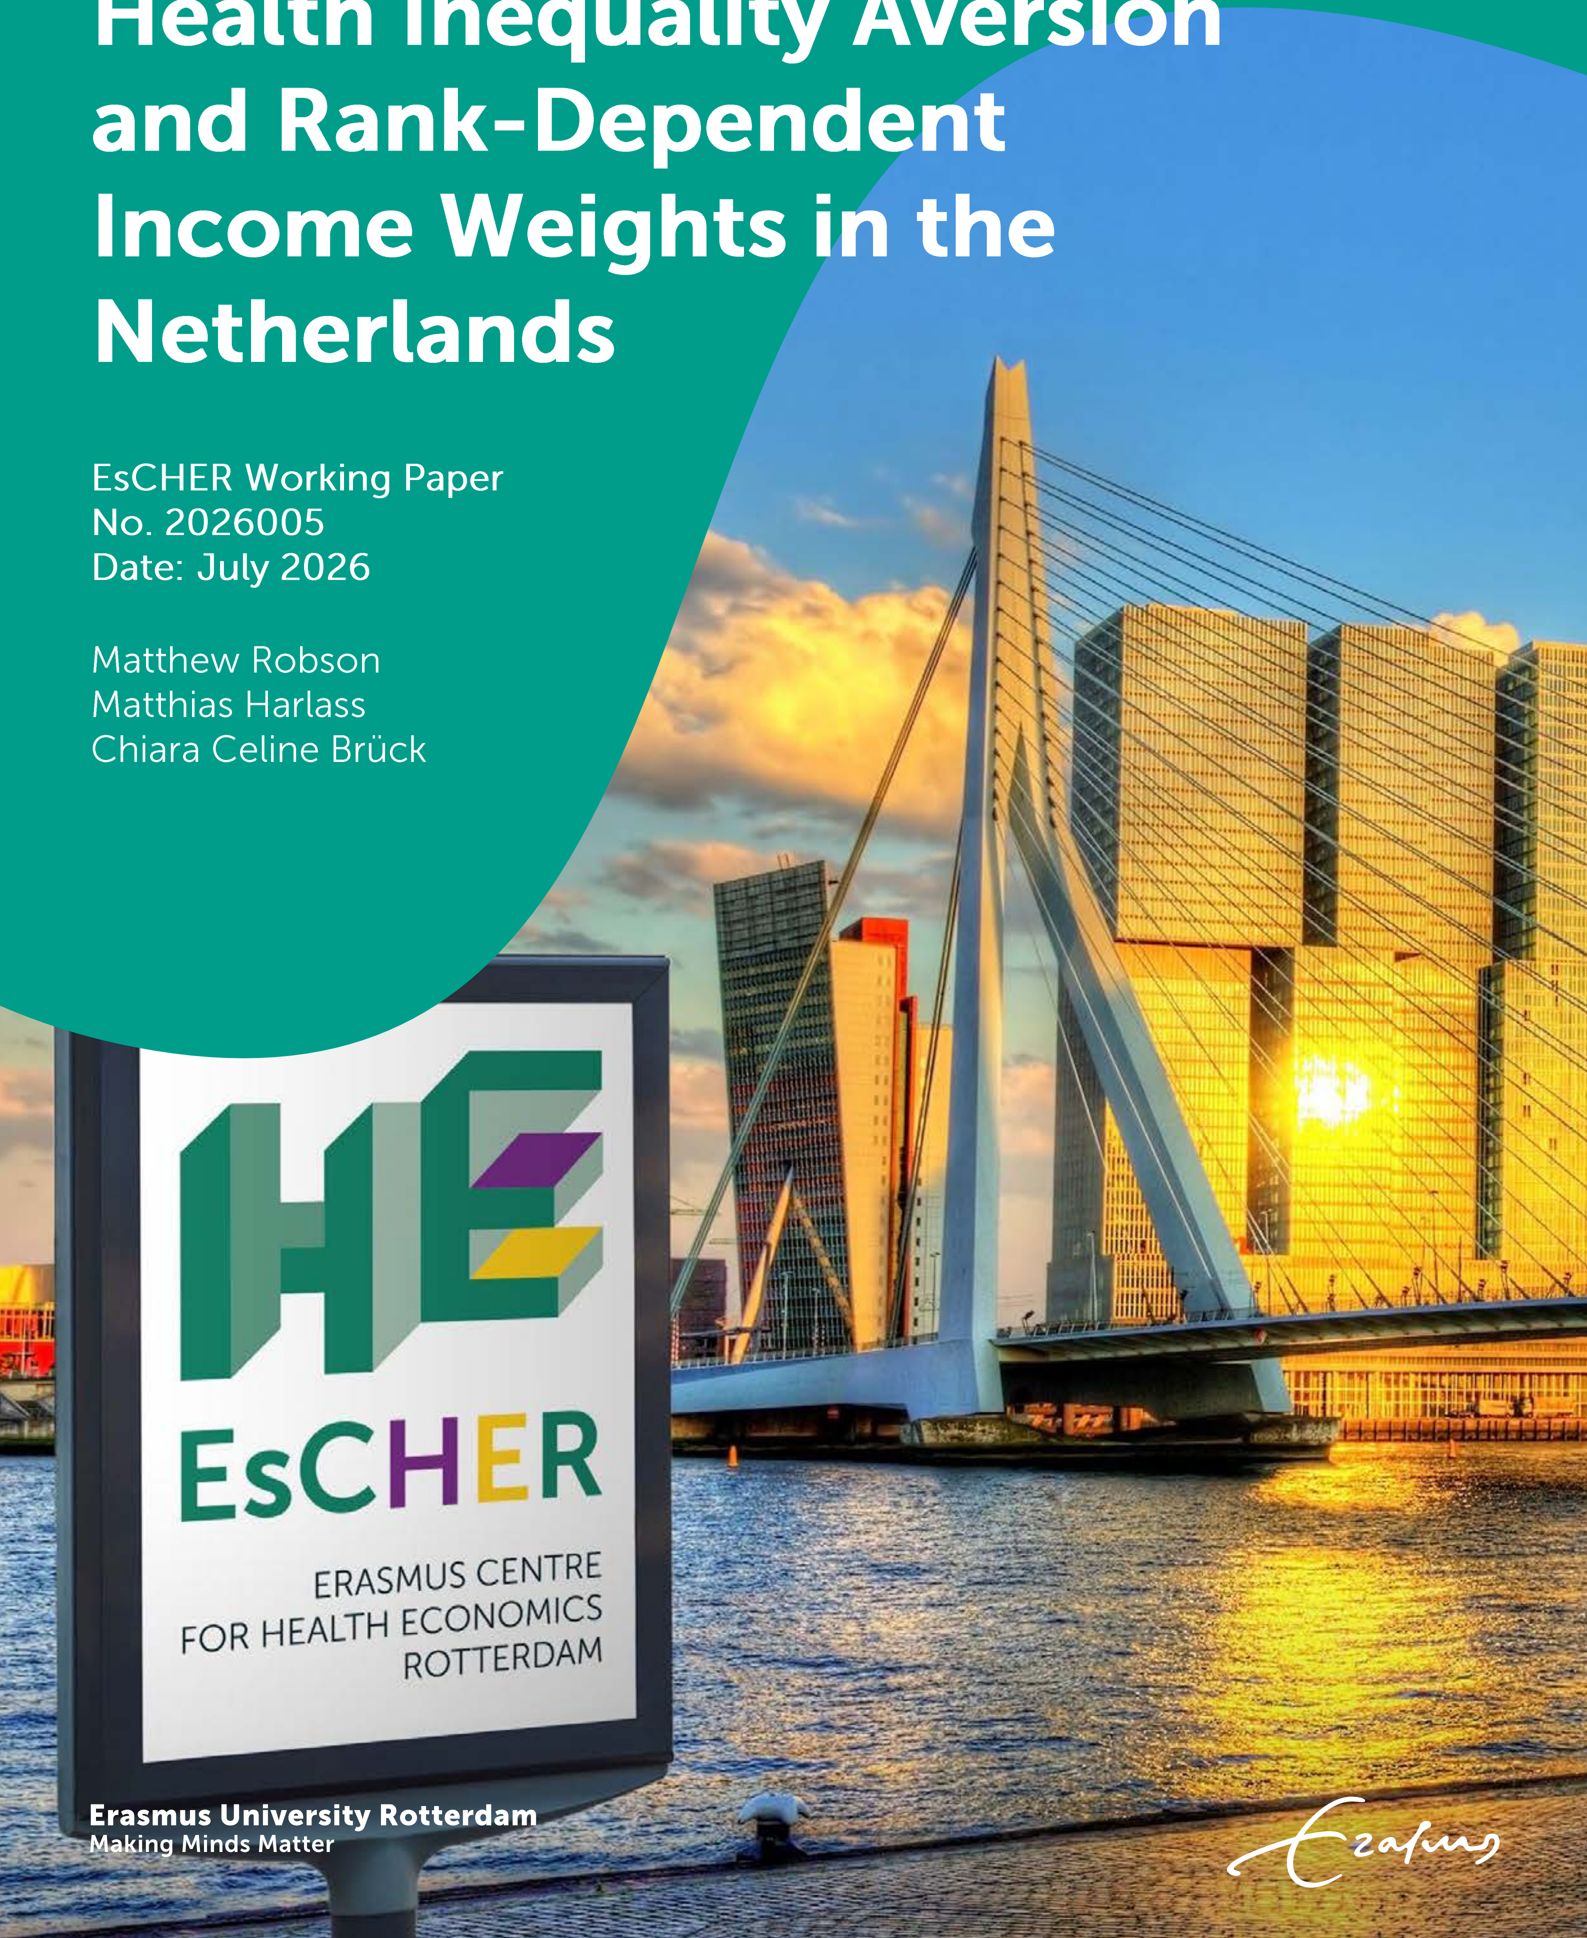

Erasmus University Rotterdam
Making Minds Matter

Title
Health Inequality Aversion and Rank-Dependent Income Weights in the Netherlands

Authors
Matthew Robson
Matthias Harlass
Chiara Celine Brück

Keywords
Social preferences; Benefit trade-off; equity-efficiency trade-off; Atkinson index; Distributional cost-effectiveness analysis.

JEL classification
I10; I18; D63; C91

Cite as
Robson, M., Harlass, M., & Brück, C. C. (2026). Health Inequality Aversion and Rank-Dependent Income Weights in the Netherlands. *EsCHER Working Paper Series* No. 2026005, Erasmus University Rotterdam. Available from: https://www.eur.nl/en/research/research-groups-initiatives/escher/research/working-papers

Erasmus Centre for Health Economics Rotterdam (EsCHER) is part of Erasmus University Rotterdam.
Want to know more about EsCHER? Visit www.eur.nl/escher
Want to contact EsCHER? E-mail escher@eur.nl

Interested in more EsCHER Working Papers? Download from www.eur.nl/escher/research/workingpapers

Health Inequality Aversion and Rank-Dependent Income Weights in the Netherlands

Matthew Robson[†] Matthias Harlass[‡] Chiara Celine Brück[§]

July 6, 2026

Abstract

Objectives: Health-related social welfare functions (HRSWF) quantify how societies value trade-offs between maximizing overall health and reducing health inequalities. We estimated health inequality aversion parameters, income-related health inequality aversion, and rank-dependent income weights in the Dutch population, for use in distributional cost-effectiveness analysis.

Methods: We conducted a two-phase benefit trade-off (BTO) experiment with a representative sample of Dutch adults ($n = 1101$). Participants made equity-efficiency health trade-offs between anonymous groups and income-defined groups. Participant-level parameters were elicited for Atkinson, Kolm–Pollock, and Rank-Dependent HRSWFs. Response-quality checks included graph literacy, conceptual comprehension, monotonicity, and completion time.

Results: Most participants (89%) were willing to sacrifice total health to reduce inequality. The median Atkinson inequality aversion parameter was $\varepsilon = 5.36$ for both univariate and income-related health inequalities, implying marginal health gains to the worse-off fifth were valued 2.6 times more than gains to the better-off fifth. Income-weight preferences were 51.7% neutral, 30.6% pro-poor, and 17.6% pro-rich. Higher-quality responses showed fewer extreme patterns and more categorisable responses.

Conclusions: Most respondents were willing to trade total health gains for greater equity. These parameters support distributional cost-effectiveness analysis in the Netherlands and highlight the importance of response-quality checks in health equity preference studies.

Keywords: Social preferences; Benefit trade-off; equity-efficiency trade-off; Atkinson index; Distributional cost-effectiveness analysis.

JEL: I10; I18; D63; C91.

[†]Erasmus School of Economics, Tinbergen Institute, and Erasmus Centre for Health Economics Rotterdam (EsCHER), Erasmus University Rotterdam, Rotterdam. Email: robson@ese.eur.nl. Website: www.mrobson92.com.

[‡]Department of Public Health, Erasmus MC University Medical Center, Rotterdam. Email: m.harlass@erasmusmc.nl.

[§]Department of Public Health, Erasmus MC University Medical Center, Rotterdam. Email: c.bruck@erasmusmc.nl.

1 Introduction

Health policy affects both the overall level and distribution of population health. However, conventional economic evaluation in health focuses on *average* health gains, implicitly valuing all gains equally and overlooking the distributional consequences of health interventions (Stevens et al., 2003). This health-maximising stance contradicts evidence showing that most people are willing to forgo some efficiency to achieve greater equity (Richardson et al., 2012; Robson et al., 2024). Extensions of cost-effectiveness analysis, such as distributional cost-effectiveness analysis (DCEA), incorporate equity considerations by weighting health gains among different population subgroups (Asaria et al., 2016). Central to this framework are preference parameters embedded in health-related social welfare functions, which capture trade-offs between total health maximisation, inequality reduction, and prioritisation of health gains by income. Although such parameters have been estimated in other countries, no such estimates exist for the Netherlands.

In this study, we estimate health inequality aversion, income-related health inequality aversion, and rank-dependent income weights for the Dutch population using a two-phase benefit trade-off (BTO) design. These parameters can directly inform DCEAs, which evaluate alternative health interventions considering their cost-effectiveness and equity impacts simultaneously. We illustrate this and provide a web tool for empirical applications.

Our study offers three main contributions to the existing literature. First, we provide the first empirical estimates of health inequality aversion parameters amongst a representative sample of Dutch adults, which can be used to inform Distributional Cost-Effectiveness Analyses in the Netherlands. Second, we extend the BTO design of Robson et al. (2017) to allow for the elicitation of health inequality aversion (ε) separately from (rank-dependent) income weights (ω_i). We provide both the theory and experimental design to allow this extension to be readily applied in other countries and contexts, enabling others to disentangle these distinct preferences. Third, we include five data-quality checks — graph literacy, conceptual comprehension, completion time, monotonicity violations and categorisability and document substantial differences in both elicited preference distributions and the share of categorisable responses. These findings highlight the importance of response-quality considerations in BTO designs and their role in improving the reliability of elicited preference distributions for health policy analysis.

2 Methods

In our experiment, participants take the role of social decision makers choosing between alternative policies that affect the distribution of quality-adjusted life expectancy (QALE) across two groups. We implement two within-subject treatments: a univariate treatment with anonymous groups and an income-labelled treatment where groups are identified by income. Using participants' responses, we elicit individual-level parameters capturing health inequality aversion and income-related health inequality aversion within symmetric health-related social welfare functions (Atkinson, Kolm–Pollock, and rank-dependent), alongside rank-dependent income weights in an asymmetric specification. We also include response-quality checks and examine how response quality relates to response categorisability and the distribution of elicited inequality aversion parameters.

2.1 Theory

In our experiment, participants take the role of a social decision maker (SDM) who is assumed to make decisions to maximise social welfare (W): an aggregation of health (h_i; measured as quality-adjusted life expectancy) of individuals (i) in a population of size N. We model this using a weighted health-related social welfare function (HRSWF),

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

where $0 \leq \omega_i \leq 1 \; \forall i$ and $\sum_{i=1}^{N} \omega_i = 1$ are weights assigned to each individual. The social utility function, $U(.)$, is common across individuals and represents the social preference for health of the SDM (Wagstaff, 1991; Dolan, 1998). If it is strictly concave and the weights are constant ($\omega_i = 1/N$, $\forall i$), then there is aversion to health inequality: a transfer of health from a healthier to a less healthy individual increases welfare.

Health Inequality Aversion. In our main analysis, we restrict attention to the iso-elastic function, $U\left(h_i\right) = \left(h_i^{1-\varepsilon} - 1\right) / \left(1 - \varepsilon\right)$ for $\varepsilon \neq 1$, and $U\left(h_i\right) = ln\left(h_i\right)$ for $\varepsilon = 1$, where ε is the Atkinson health inequality aversion parameter (Atkinson, 1970).[1] Here welfare can be measured by the equally distributed equivalent (EDE) level of health,

$$h_{EDE} = \left(\sum_{i=1}^{N} \omega_i \, h_i^{1-\varepsilon}\right)^{\frac{1}{1-\varepsilon}}. \tag{2}$$

With constant weights, ε captures the trade-off the SDM is willing to make between maximising aggregate health and equalising the distribution of health. When $\varepsilon = 0$ social utility is linear in health: the sole objective is to maximise total health ("Health Maximiser"). For $\varepsilon > 0$, the social utility function becomes strictly concave in health ("Prioritarian"), indicating a willingness to forgo health maximisation to reduce health inequality. As ε increases, inequality aversion intensifies and social welfare becomes more sensitive to marginal health improvements for those with lower health. As $\varepsilon \to \infty$, the SDM's preferences approach the Rawlsian maximin — only health improvements experienced by the least healthy raise welfare ("Maximin"). If, on the other hand, $\varepsilon < 0$ then the indifference curve be-

[1] Kolm-Pollak and Rank-Dependent HRSWFs are shown in Appendix B.3 as alternative specifications.

comes strictly convex, indicating a willingness to sacrifice aggregate health to prioritise the better-off, increasing inequalities ("Inequality Seeking").

Income-Weights. We can relax the assumption of constant weights when individuals are identified by income, x_i. Here weights depend on an individual's rank in the income distribution, with a weighting function based on generalised concentration indices (Yitzhaki, 1983; Wagstaff, 2002; Makdissi and Yazbeck, 2016). Let $r(x_i) \in \{1, \ldots, N\}$ denote the rank of individual i, where higher values correspond to higher income (i.e. $x_i > x_j \Rightarrow r(x_i) > r(x_j)$), then income-weights are given by:

$$\omega_i = \left(\frac{N - r(x_i) + 1}{N}\right)^{\beta} - \left(\frac{N - r(x_i)}{N}\right)^{\beta}, \tag{3}$$

where $\beta > 0$ reflects the direction and degree of prioritisation by income.[2] With $\beta = 1$, income weights are constant ($\omega_i = 1/N$, $\forall i$), reflecting income-neutrality and a symmetric HWRSF. If $\beta \neq 1$, the HRSWF is asymmetric. When $\beta > 1$, individuals lower in the income distribution are given higher weights ("Pro-Poor"). Weights used in the standard Concentration Index correspond to $\beta = 2$. As $\beta \to \infty$, the weights approach zero for all but the lowest ranked individual. With $0 < \beta < 1$, the weights *increase* with income rank, giving more priority to those with higher income ("Pro-Rich").

For ease of interpretation, we report the relative weight assigned to the poorest quintile, defined as $\omega_p = \omega_i/(\omega_i + \omega_j)$, where i denotes the poorest quintile and j the richest quintile. Income neutrality is represented by $\omega_p = 0.5$, while $\omega_p > 0.5$ indicates a pro-poor preference and $\omega_p < 0.5$ a pro-rich preference. Figure B1 plots social indifference curves to provide the

[2]When income levels are not unique, individuals with identical income levels are treated as ties. Let $P(H = x_i)$ denote the population share at income level x_i. The difference in rank-based weights, $[P(X \geq x_i)]^{\beta} - [P(X > x_i)]^{\beta}$, is allocated equally across individuals with $X = x_i$. Individual weights are therefore given by $\omega_i = \frac{[P(X \geq x_i)]^{\beta} - [P(X > x_i)]^{\beta}}{P(X = x_i)}$.

intuition of both ε and ω_p. Convexity of social indifference curves is determined by ε and the slope is determined by ω_p.

Indirect Equity Weights. The SDMs marginal rate of substitution between the health of two individuals – i and j – is given by the Indirect Equity Weights (IEW_{ij}),[3]

$$IEW_{ij} = \frac{\partial h_i}{\partial W} \bigg/ \frac{\partial h_j}{\partial W} = \frac{\omega_i}{\omega_j} \left(\frac{h_j}{h_i} \right)^{\varepsilon}. \tag{4}$$

This provides both an intuitive measure of the relative amount of health one individual must gain in order to offset a reduction in the health of another, such that social welfare is constant, and provides a method of weighting marginal gains from alternative policies to evaluate which maximises social welfare. Priority given to individuals, thus, depends on the relative difference in the health, health inequality aversion, ε, and the weights, ω_i.

2.2 Experimental Design

Our experiment used a benefit trade-off (BTO) design, adapted from Robson et al. (2017). Through the BTO design, participant-specific preference parameters were elicited, within health-related social welfare functions (HRSWF), from participant's "indifference points" between health distributions resulting from two alternative programmes.

The experiment comprised of two within-subject treatments: a "univariate" treatment and an "income-labelled" treatment. In each treatment, participants acted as SDMs and had to choose between two alternative programmes that affected the quality-adjusted life expectancy (QALE) of two groups in the population, which differed in their baseline QALE. In the first (univariate) treatment, the group identities were not further characterized and

[3]These weights have different names in alternative literatures, (generalised) social marginal welfare weights (SMWW) (Piketty and Saez, 2013; Saez and Stantcheva, 2016; Robson et al., 2025b) from welfare economics, or indirect equity weights (Cookson et al., 2020) from the Distributional Cost Effectiveness Analysis literature in health economics.

simply labelled "Group H" (higher baseline QALE) and "Group L" (lower baseline QALE). In the second (income-labelled) treatment, the two groups were defined as the highest and lowest income quantiles, labelled "Richest Fifth" and "Poorest Fifth", with the Richest Fifth having a higher baseline QALE. The baseline QALEs and the programme effects were identical across the two treatments. Within each treatment, participants completed seven choice tasks, selecting Programme A, Programme B, or Equally Good.

Figure 1 shows a screenshot of the first choice task in the "univariate" treatment. Across the seven tasks, the QALE gain in Programme A was held constant for both groups, while the gain offered to the worse-off group in Programme B was titrated, from +8 in the first task, by integers, to +2 in the seventh task. This forced participants to make equity-efficiency trade-offs: maximising total gains versus improving the QALE of the worst-off group.

Figure 1: Trade-Off Question

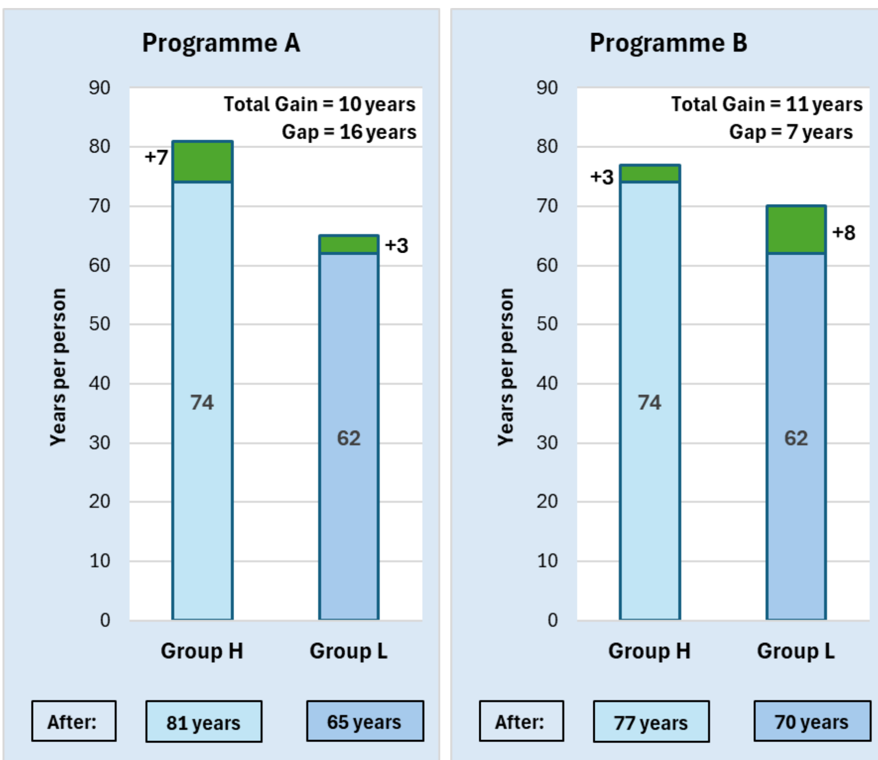

Inequality aversion parameters can be elicited by identifying a "switching point" (i.e. a point of indifference) between two programmes. For a SDM who is indifferent between programmes (A) and (B), the equally distributed equivalent (EDE) level of health is equal across programmes, such that $h^A_{EDE} = h^B_{EDE}$.[4]

Under a *symmetric* Atkinson HRSWF, this condition is equivalent to $\sum_{i=1}^{N} \frac{1}{N}(h_i^A)^{1-\varepsilon} = \sum_{i=1}^{N} \frac{1}{N}(h_i^B)^{1-\varepsilon}$. The health inequality aversion parameter (ε) can then be obtained by numerically solving the above equation, given the health levels associated with the "switching point". Table 1 shows how participant responses can be categorised into Inequality Seeking, Health Maxmiser, Prioritarian, Maximin and Egalitarian standpoints, with the corresponding values of the health inequality aversion parameter ε. The health sacrificed column shows the intuition of their choices: the total QALE gains the participant was willing to sacrifice in order to reduce health inequality. Indirect equity weights, IEW_{ij}, show the marginal rate of social substitution between a worse-off (poorer) group and better-off (richer) group, with QALEs of 62 and 74, respectively. For the health inequality aversion parameters and IEW_{ij} for the Kolm-Pollack and Rank-Dependent HRSWFs see Table B1.

If *symmetry* is relaxed, both an inequality aversion parameter, ε, and rank-dependent income weights can be identified. We first elicit ε in the univariate treatment under the assumption of symmetric weights, as described above. We then hold ε fixed and use each participant's switching point in the income-labelled treatment to infer the rank-dependent income-weight parameter, β, assuming that participants apply the same degree of health inequality aversion across treatments. Under the asymmetric Atkinson HRSWF, β is obtained

[4] Where a participant reports indifference, the corresponding health distributions are used directly. For example, the response pattern B=AAAAA implies indifference between ($h_1^A = 81, h_2^A = 65$) and ($h_1^B = 77, h_2^B = 69$). Where a participant switches between programmes, the indifference point is approximated using the midpoint between the two allocations immediately before and after the switch. For example, the response pattern BBAAAAA implies indifference between programme (A) and the midpoint allocation ($h_1^B = 76.5, h_2^B = 69$), which lies halfway between the adjacent (B) allocations.

Table 1: Response Categorisation, Inequality Aversion, and Indirect Equity Weights

Rank	**Category**	**Response**	**Health Sacrificed**	**Inequality Aversion, ε**	**Indirect Equity Weights**
1	Inequality Seeking 1	AAAAAAA	−1.5	−2.075	0.693
2	Inequality Seeking 2	=AAAAAA	−1.0	−1.419	0.778
3	Inequality Seeking 3	BAAAAAA	−0.5	−0.731	0.879
4	Health Maximiser	B=AAAAA	0.0	0	1.000
5	Prioritarian 1	BBAAAAA	0.5	0.792	1.150
6	Prioritarian 2	BB=AAAA	1.0	1.671	1.344
7	Prioritarian 3	BBBAAAA	1.5	2.673	1.605
8	Prioritarian 4	BBB=AAA	2.0	3.862	1.980
9	Prioritarian 5	BBBBAAA	2.5	5.358	2.581
10	Prioritarian 6	BBBB=AA	3.0	7.430	3.723
11	Prioritarian 7	BBBBBAA	3.5	10.946	6.936
12	Maximin	BBBBB=A	4.0	∞	∞
13	Egalitarian 1	BBBBBBA	4.5	NA	NA
14	Egalitarian 2	BBBBBB=	5.0	NA	NA
15	Egalitarian 3	BBBBBBB	5.5	NA	NA

Notes: Responses categorised by sequences of choices in the benefit trade-off task. "Health Sacrificed" indicates the total QALE gains the participant was willing to sacrifice in order to reduce health inequality. The inequality aversion parameter ε is recovered under a symmetric Atkinson social welfare function; higher values indicate greater aversion to inequality. "Maximin" corresponds to the limiting case of infinite inequality aversion. "NA" indicates response patterns that cannot be represented within the Atkinson framework. Indirect equity weights, IEW_{ij}, show the marginal rate of social substitution between a worse-off (poorer) group and the better-off (richer) group, with QALEs of 62 and 74, respectively.

by numerically solving $\sum_{i=1}^{N} \omega_i^A (h_i^A)^{1-\varepsilon} = \sum_{i=1}^{N} \omega_i^B (h_i^B)^{1-\varepsilon}$, where the income-weights ω_i are determined by β within the rank-dependent weighting function, see eq.(3).

Preferences are "income-neutral" ($\beta = 1 \rightarrow \omega_p = 0.5$) when participants make the same choice in both treatments, "pro-poor" if they sacrifice more total health in the income treatment to benefit the poorer group ($\beta > 1 \rightarrow \omega_p > 0.5$), and "pro-rich" if they sacrifice less total health in the income treatment ($\beta < 1 \rightarrow \omega_p < 0.5$). If participants sacrifice four or more years in either treatment, corresponding to the Maximin and Egalitarian categories (ranks 12-15 in Table 1), their inequality aversion parameter is infinite or undefined, so the weight no longer affects the trade-off and income weights cannot be elicited. See Figure B2 for the elicited ω_p parameters, conditional on the choices in both the "univariate" and "income-labelled" treatments.

Experimental Particulars. The experiment was conducted from June 2nd, 2025 to June 30th, 2025. In the experiment, participants first received instructions, then completed two graph literacy and five conceptual comprehension questions. Then, they completed eight trade-off questions in the "univariate" and "income-labelled" treatment, with the last question testing violations in monotonicity in preferences. Finally, they were asked about their opinion on reducing health inequality, to what extent they identified with the groups mentioned in the trade-off questions (Group H, Group L, Poorest Fifth, Richest Fifth) (Tropp and Wright, 2001), and a question eliciting the degree of altruism of the participant (Falk et al., 2022). The median completion time was 13 minutes. Full experimental instructions in Dutch and English are available on the Open Science Framework (OSF; https://osf.io/h79p8). We ran a pilot of 92 participants from March 29th to June 25th 2025, prior to running the main experiment. Following the pilot we only made minor changes to the language to increase understandability. An overview of the experimental instructions translated into English and details on the pilot can be found in Appendix A.2 and Appendix A.3.

2.3 Sample

We recruited a representative sample of the Dutch adult population from the LISS panel (Longitudinal Internet studies for the Social Sciences).[5] A total of 1142 participants began the experiment, and 1101 (96.4%) completed it. The complete data set with accompanying

[5]The LISS panel is administered and managed by the non-profit research institute Centerdata (Tilburg University, the Netherlands) (Scherpenzeel, 2018). The LISS panel is a representative sample of Dutch individuals who participate in monthly internet surveys. The panel is based on a true probability sample of households drawn from the population register by Statistics Netherlands. Self-registration is not possible, and households that would otherwise be unable to participate are provided with a computer and internet

codebook can be found in the LISS Data Archive (https://dataarchive.lissdata.nl/study-units/view/1657).

For our main analytical sample, we applied five data quality criteria: 1) at least one of two graph literacy questions correct, 2) at least three of five conceptual comprehension questions correct, 3) at least five minutes to complete the survey, 4) no monotonicity violations in either treatment, and 5) categorisable response pattern in both treatments (see Table 1). A total of 629 participants did not meet these criteria, resulting in a main analytical sample of 472 participants. To preserve the representativeness of the sample, we applied inverse probability weighting according to the Dutch age, sex and education distribution in 2024 to the main analytical dataset (Appendix C.2).

2.4 Statistical analysis

The study hypotheses and analysis plan were preregistered on the Open Science Framework (OSF; https://osf.io/a2xgz) prior to data analysis, and the analysis code is publicly available on the OSF project page. The analysis adhered to the preregistered plan with two methodological refinements to ensure robustness and representativeness, which are described in Appendix A.1. All analyses were performed using R version 4.6.0.

Following the theoretical derivation of the inequality aversion parameters (ε) and income weights (ω_i) via numerical optimization, respondents were classified into the corresponding theoretical categories (e.g. Health Maximiser, Prioritarian, Maximin) for each treatment. To compare response distributions beyond summary statistics, we used Kolmogorov–Smirnov (KS) tests on the *health sacrificed* outcome. For comparisons between independent unweighted groups (e.g. data-quality strata), we used the standard two-sample KS test. For weighted comparisons, we compared weighted empirical cumulative distribution func-

connection. It consists of 5,000 households, comprising approximately 7,500 individuals of 16 years and older.

tions. When the same respondents were compared under alternative weighting schemes (IPW vs. unweighted), p-values were obtained from a sign-flip randomisation test with 10,000 draws. We estimated population median parameters and 95% confidence intervals using weighted quantile regression with bootstrapping (10,000 replications). We used Ordinary Least Squares (OLS) regression predicting the magnitude of health sacrificed in both univariate and income-framed contexts to identify the determinants of health inequality aversion. When examining response quality relates to categorisability, we grouped the full sample (n = 1,101, unweighted) into three quality tiers using the first four criteria only — graph literacy, comprehension, timing, and monotonicity — so that categorisability could be analysed as an outcome rather than built into the definition. High-quality respondents answered both graph-literacy questions and all five conceptual-comprehension questions correctly, took longer than 5 minutes, and did not violate monotonicity in either treatment. Medium-quality respondents met the less stringent thresholds used in the main analysis (at least one graph-literacy question correct, at least three conceptual-comprehension questions correct, completion time above 5 minutes, and no monotonicity violation in either treatment) but did not meet the high-quality threshold. Low-quality respondents did not meet the medium-quality threshold. See Appendix C.1 for the composition of the in- and excluded sample.

3 Results

3.1 Health Inequality Aversion

The majority of respondents (88.2%-89.8%) were willing to trade off total health gains to achieve greater health equity, for both univariate and income-related health inequality aversion (Figure 2).[6]

Figure 2: Health Inequality Aversion Distribution: Univariate and Income-Related Treatments

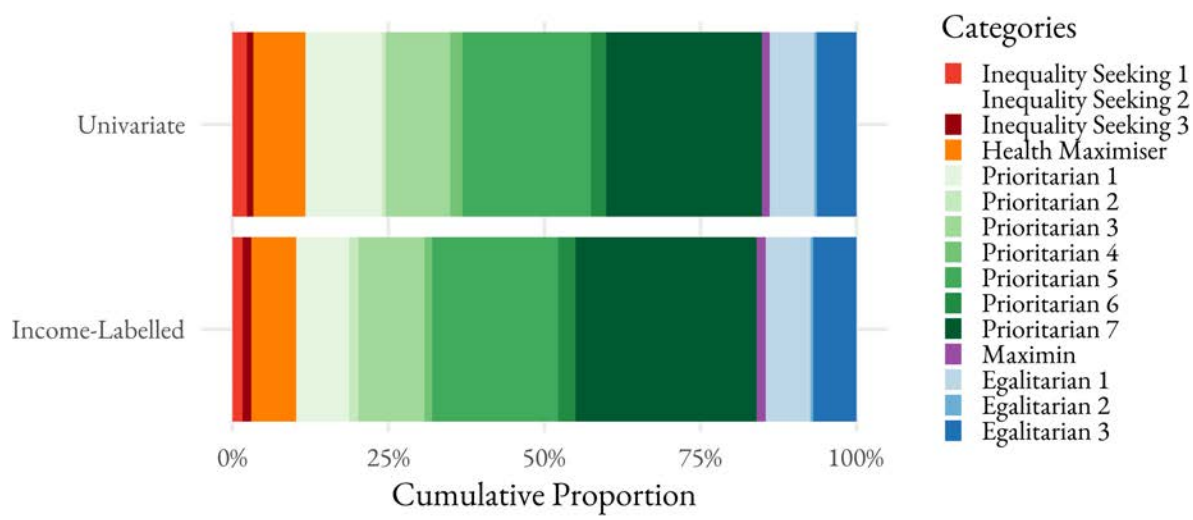

For univariate health inequality aversion, the majority of respondents were classified as *Prioritarian* (73.1%). A minority exhibited more extreme inequality aversion, with 1.3% being classified as *Maximin* and 13.8% as *Egalitarian*. Additionally, there were 8.3% *Health Maximisers* and only 3.3% who were *Inequality Seeking*. The distribution of income-related health inequality aversion was similar. A two-sample weighted Kolmogorov–Smirnov test confirmed that the distributions of health sacrificed did not differ significantly between the two treatments (see Appendix D.2). In comparison, the proportion of *Prioritarian* responses

[6] We find a strong negative association between health sacrificed and the reported importance of reducing health inequalities, confirming that the elicited inequality aversion parameters correspond to the participants reported views on equity and efficiency (see Table D6).

was marginally larger (73.8% vs 73.1%) and more inequality averse, whilst the proportion of *Health Maximisers* was slightly smaller for income-related inequality (7.2% vs 8.3%).

Overall, we estimated a median health inequality aversion parameter, ε, of 5.36 (see Table D4 for bootstrapped 95% confidence intervals), in both the univariate and income-labelled treatments. This indicates that the median respondent displayed substantial, and essentially identical, aversion to univariate and income-related health inequalities.

Translating this into an indirect equity weight, the median participant weighed marginal gains in QALYs for the poorest group 2.581 times that of the richest group. The majority (84.5%-89.2%) gave higher weights to the worst-off group, while few gave equal weight to both groups (5.3%-9.5%) or higher weight to the richer group (3.2%-10.2%).

3.2 Health Inequality Aversion and Income Weights

Figure 3 shows the distribution of participant-level health inequality aversion parameters, ε, and income-weights, ω_p, when we allowed for asymmetric income weights. Health inequality aversion parameters are as described above, with a median ε of 5.36. For the income-weights, we found that the modal group (51.7%) was income-neutral, with $\omega_p = 0.5$. However, we found a substantial fraction (30.6%) who was pro-poor, with $\omega_p > 0.5$ and interestingly, a notable fraction of participants (17.6%) who was pro-rich, with $\omega_p < 0.5$.[7]

3.3 Data Quality

Data quality meaningfully influenced the number and distribution of categorisable responses, underlining the importance of rigorous data quality checks. The proportion of categorisable responses increased dramatically when moving from low (26.0%) to medium (67.6%) to

[7]Note that this is similar to Robson et al. (2024) who found 71.5%, 23.4%, and 5.0%, who were income-neutral, pro-poor, and pro-rich, respectively.

Figure 3: Health Inequality Aversion and Income Weights

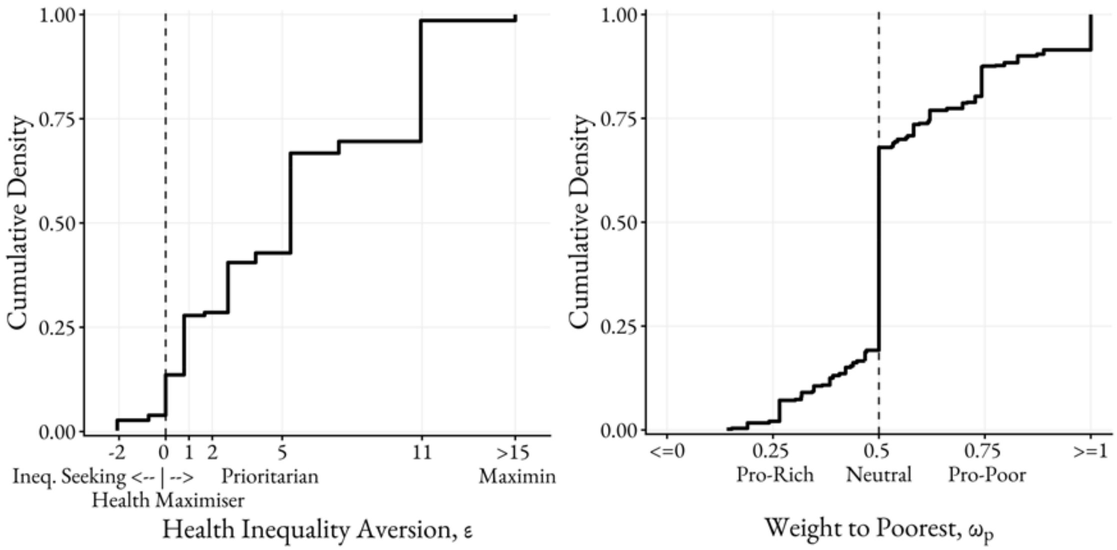

high (85.5%) response quality (Figure 4). Moreover, pairwise Kolmogorov–Smirnov tests (Appendix D.4) confirmed that the distribution of responses of low-quality participants differed significantly from those of high-quality participants ($p < 0.001$), while there was no significant difference between high and medium quality groups ($p = 1.00$). In particular, the prevalence of extreme response types declined: responses classified as *Inequality Seeking 1* decreased from 18.3% to 1.1%, and *Egalitarian 3* responses decreased from 19.8% to 4.9%.

3.4 Explaining Response Heterogeneity

Results in Table D8 show that much of the heterogeneity in health inequality aversion was not explained by observable socio-demographic characteristics.[8] Whilst we found that females were significantly more inequality averse than males (only in the income-labelled treatment), we found no significant differences in responses by age, household composition, rural, household income or education levels.

[8] This finding is consistent with Robson et al. (2025a), who likewise find that socio-demographic characteristics explain little of the variation in health inequality aversion.

Figure 4: Response Distribution by Data Quality

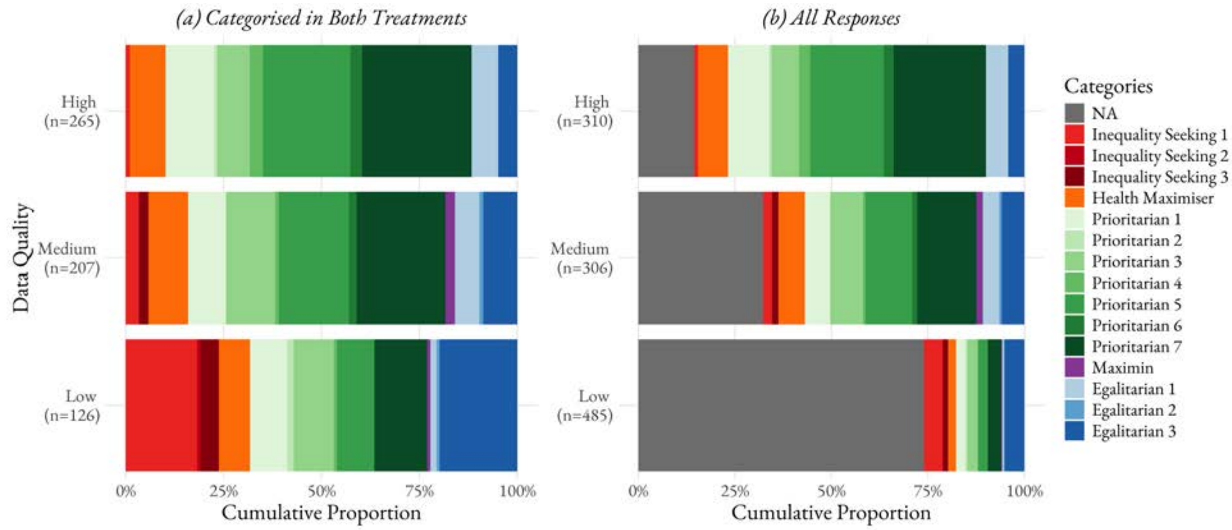

Notes: The left panel shows those responses which could be categorised in both treatments, and the right panel also includes those participants with responses which could not be categorised. This analysis used the whole unweighted sample (n = 1,101).

Interestingly, we found significantly higher health inequality aversion amongst those who reported feeling more connected to the lower baseline group than the higher baseline group (univariate treatment), and to the poorer group than the richer group (income-labelled treatment). However, this association was only present for those classified as "selfish". For those classified as "altruists", relative connectedness and health inequality aversion were not associated. In other words, selfish respondents prioritised their in-group, while altruists did not. This matters in asymmetric designs (where the rich always have more health), as in-group priority raises inequality aversion when the respondent is poor, but lowers it when rich.

Finally, we found that data quality was generally higher among younger, richer, and more educated respondents (see Appendix D.6). Thus, socio-demographics predicted whether responses were categorisable, but not the heterogeneity in the categorisable preferences.

Table 2: Policy Evaluation Example

QALE	Baseline	Policy A	Policy B
Poorer Group	62	63 (+1)	62 (+0)
Richer Group	74	74 (+0)	76 (+2)
Mean	68	68.5	69
h_{EDE}	66.616	67.341	67.158

3.5 Policy Evaluation Example

The parameters we elicit can be used to assess a wide range of policies and interventions that affect population subgroups with unequal baseline levels or unequal treatment effects. We present a policy evaluation example to demonstrate how to apply the elicited parameters in this study and to provide intuition of their meaning. Additionally, we developed a web tool, that allows users to upload their own dataset to compute the equity-adjusted net health benefit of interventions, highlight implied policy recommendations and conduct sensitivity analyses.

Because the results can be applied across many contexts, we illustrate their use with a stylized example. Assume the population is divided in half: the poorer half of the population with a baseline QALE of 62 and the richer half of the population with a baseline QALE of 74. A policymaker must choose between two equally costly policies which have different Net Health Benefits (NHB) across the two groups: Policy A provides a NHB of one QALY only to the poorer half and Policy B provides a NHB of two QALYs only to the richer half. Using the median health inequality aversion parameter of 5.36, we can calculate the EDE of the QALE (h_{EDE}) of the two policies (Table 2).

Without taking the preferences of the population into account, the policy maker would choose Policy B, as the average QALE of the population is higher after Policy B (69) than Policy A (68.5). However, taking the median preference of the population in both univariate and income-related contexts into account, the policy maker would choose Policy A as the h_{EDE} of Policy A (67.341) is higher than the h_{EDE} of Policy B (67.158).[9]

Using the indirect equity weight of 2.581, the policy maker would value the gain of one QALY for the poorer group in Policy A as 2.581 QALYs, while the gain of two QALYs for the richer group in Policy B remains 2 QALYs. Therefore, the policy maker would also choose Policy A (weighted gain of 2.581 QALYs) over Policy B (weighted gain of 2 QALYs), when using the indirect equity weight.

The above presents an example using the median health inequality aversion parameter. The results of this paper can also be applied to individual-level policy evaluation using either symmetric or asymmetric HRSWFs; we demonstrate this application in Appendix D.7. Our example focuses on the equity-adjusted net health benefit of interventions. When policies differ not only in effectiveness but also in costs, the estimated health inequality aversion parameter can additionally be incorporated into a full distributional cost-effectiveness analysis (DCEA), in which policy costs and their corresponding h_{EDE} values are combined to produce inequality-adjusted Incremental Cost-Effectiveness Ratios (ICERs); this lies beyond the scope of the present study, and we refer the reader to Cookson et al. (2026) for its construction and interpretation.

4 Discussion

Using a two-phase BTO experiment, we elicited univariate and income-related health inequality aversion amongst a representative sample of the Dutch public. We found that the majority (88.2%) was inequality averse and a median health inequality aversion parameter ε of 5.36. Translating this to indirect equity weights, the median Dutch respondent would

[9]Note that this is an approximation, indirect equity weights are most appropriate when NHBs only bring marginal changes. Often one also normalises so that they sum to 1; i.e. $IEW_i = IEW_{ij}/(IEW_{ij} + IEW_{ji})$.

weight marginal QALY gains of the worst-off fifth 2.581 times more than QALY gains of the best-off fifth. We found similar distributions of univariate and income-related health inequality aversion, with slightly higher inequality aversion in the income-related treatment. By separately eliciting income-weights ω_i, we found that whilst the modal group (51.7%) was income-neutral, more respondents had pro-poor than pro-rich preferences (30.6% vs 17.6%). The elicited parameters allow for an empirically supported DCEA to be conducted in the Netherlands. We provided a simple illustrative example of the equity-adjusted net health benefit, alongside a complementary web tool.

4.1 Comparison to Literature

Whilst we presented the first estimates of health inequality aversion parameters for the Netherlands,[10] a growing literature provides comparable evidence from other countries. Robson et al. (2017) elicited a median ε of 10.95 in England. Their experimental design, focused on income-related health inequality aversion, was replicated across several countries with elicited median ε of: 27.16 in Australia (Boujaoude et al., 2025), 13.9 in Japan (Shimamoto et al., 2025), and 12.12 in the USA (Slejko et al., 2025). In Canada (Iragorri et al., 2025), the median respondent was found to have *egalitarian* preferences. For univariate health inequality aversion, a somewhat lower median ε of 3.52 was found in Norway (Horn et al., 2025). Earlier BTO designs also found substantial univariate (Edlin et al., 2012) and income-related (Dolan and Tsuchiya, 2011) health inequality aversion in the UK, with *Egalitarian* median respondents in Spain (Abasolo and Tsuchiya, 2004).

[10] A recent Dutch study used a discrete choice experiment to elicit preferences over total health gains and the distribution of those gains across income groups (Meulman et al., 2026). Most respondents preferred an equal distribution of health benefits across income groups, while 29% additionally favoured prioritising lower-income groups. Although not directly comparable to the present study, as equity–efficiency trade-offs in health were not explicitly elicited, these findings similarly suggest that preferences regarding income are predominantly income-neutral or pro-poor rather than pro-rich.

However, cross-country differences should be interpreted cautiously. Published estimates vary not only by setting, but also by framing, identified groups, survey mode, sample composition, and the stringency of data-quality screening. The results of our study show that response-quality filters materially affect the observed distribution of preferences. Therefore, differences across studies may reflect methodological as well as substantive variation. Common to each of these country-contexts is that the majority of respondents are classified as inequality averse, which stands in contrast to the normative assumptions underpinning standard Cost-Effectiveness Analysis.

Whilst the above studies focused on *either* univariate or income-related health inequality aversion, several studies investigated both. Hurley et al. (2020) used a Discrete Choice task to separately estimate health inequality aversion and income-related health inequality aversion (using an extended concentration index) in Canada. Pinho and Botelho (2018) used a two-phase BTO to sequentially elicit health inequality aversion parameters and weights by occupational class, amongst college students in Portugal. Robson et al. (2024) used a constrained budget set experiment to simultaneously estimate health inequality aversion and income-weights in the UK. In each study, income-related health inequality aversion was found to be (slightly) higher than univariate health inequality aversion.

Our study extended the one-phase design of Robson et al. (2017) to a two-phase design, enabling the separate identification of health inequality aversion and income-related weights. Similar approaches have been implemented by Dolan and Tsuchiya (2011) and Pinho and Botelho (2018), which also use univariate and socioeconomic-labelled treatments. However we elicit preferences within a rank-dependent income-weighting framework (as in Robson et al. (2024)), allowing parameters to be extrapolated to alternative numbers of income ranks (e.g. quantiles or deciles). This facilitates applications in distributional cost-effectiveness analysis. Compared to Robson et al. (2024), our design is simpler to implement and estimate, based on static choice tasks rather than interactive sliders, while still allowing for separate identification of health inequality aversion and income-related weighting.

4.2 Data quality

An important contribution of our paper is the incorporation and analysis of extensive data quality checks within the experimental design. Specifically, we implemented checks on (i) graph literacy, (ii) conceptual comprehension, (iii) monotonicity violations, (iv) completion-time and (v) categorisablilty. Imposing increasingly stringent data quality criteria substantially increased the share of categorisable responses, rising from 26.0% to 67.6% and then to 85.5% as we moved from low to medium to high data quality thresholds. Importantly, the composition of categorised responses also changed. In particular, the prevalence of extreme response types declined: responses classified as *Inequality Seeking 1* fell from 18.3% to 1.1%, while *Egalitarian 3* responses decreased from 19.8% to 4.9%. This distinction matters normatively: while extreme responses represent coherent social value judgements, our results suggest that a non-trivial share may instead reflect limited engagement with the task.

Most closely related studies incorporated data quality checks as ex-ante exclusion criteria, such that respondents who failed comprehension or consistency checks did not proceed to the main experiment (e.g. Boujaoude et al. (2025); Slejko et al. (2025)).[11] In contrast, our design retained all respondents and applied data quality criteria ex-post. This allowed us to examine how preference classification and estimated inequality aversion parameters varied with data quality, rather than assuming that excluded responses were irrelevant.

It is, additionally, important to comment on the number of low quality responses in our sample. We classify 485 as "low" quality, and excluded these participants from our analysis. Whilst only 310 were classified as "high" quality, we still retain the "median"

[11] Note, our graph literacy questions were inspired by Boujaoude et al. (2025).

quality respondents. This difference is important to account for in sample size calculations and survey design. Moreover, we found this is non-random, with those who were older, poorer and less educated less likely to be included. While we did apply inverse probability weighting to correct for observable differences in age, sex and education, it cannot address selection on unobserved factors correlated with both response quality and distributive preferences. It is an important consideration in ensuring that the views of the public are accurately represented.

4.3 Limitations

Our study has several limitations. First, we did not randomise the order of treatments. All participants completed the univariate treatment before the income-labelled treatment. Consequently, we cannot rule out order effects: choices in the univariate treatment may have influenced subsequent decisions in the income-labelled treatment, and participants may have learned over time or experienced fatigue. However, the within-subject design was necessary to separately identify ε and ω_p. Following Robson et al. (2024), treatment order was therefore fixed, as the second treatment builds directly upon the first. Had participants been exposed to income labels first, it would arguably have been difficult for them to subsequently view the univariate groups without implicitly assigning richer and poorer labels.

Second, although both the recruited sample and the re-weighted analytical sample were representative of the Dutch adult population with respect to sex, age, and education, respondents were drawn exclusively from the LISS panel. As a result, some groups who are less likely to participate in online panels, such as those with limited digital access or literacy, may be under-represented.

Third, although BTO designs allow preference parameters to be elicited within a particular social welfare function, they do not identify which social welfare function best characterises individual preferences. A functional form (e.g., Atkinson, Kolm–Pollock, or Rank-Dependent) must therefore be assumed *a priori*. More generally, some response patterns (e.g., AAAAAAA or BBBBBBB) yield identifiable parameters but may reflect simple decision heuristics rather than considered preferences. Our findings on response quality provide some reassurance in this regard: as response quality increases, the prevalence of inconsistent responses and extreme response patterns declines substantially.

4.4 Future Research

Our two-phase extension of the Robson et al. (2017) BTO design allows health inequality aversion to be separated from weights attached to other equity-relevant characteristics. While we focus on income, relabelling groups according to other characteristics (e.g., sex, ethnicity, or smoking behaviour) would allow both a health inequality aversion parameter and group-specific priority weights to be identified. Previous applications following Robson et al. (2017) have typically assumed a symmetric Atkinson social welfare function when estimating health inequality aversion between identified groups.

Our findings also highlight the importance of data quality in stated-preference studies of health equity. Both the distribution of preferences and the proportion of categorisable responses changed substantially when analysis was restricted to respondents meeting the response-quality criteria. In particular, inequality-seeking and egalitarian response patterns became markedly less common. While these are valid social value judgements, it remains important to establish that they arise from respondents who have understood the choice task.

Finally, the estimates reported here provide inputs for distributional cost-effectiveness analysis in the Netherlands. The results suggest that most Dutch adults are willing to trade off some total health to reduce health inequalities, while preferences regarding income are predominantly income-neutral or pro-poor. Future applications of these estimates within empirical DCEA could help assess the implications of alternative social value judgements for health policy and resource allocation decisions.

5 Conclusion

This study provides the first empirical estimates of health inequality aversion and income-related equity preferences for the Netherlands, showing that most individuals are willing to sacrifice total health gains to reduce inequality. By introducing a novel two-phase BTO design, it enables the separate identification of inequality aversion and income-based priority weights, improving the precision and interpretability of equity preferences. Importantly, the study demonstrates that rigorous data-quality checks substantially affect both the share of usable responses and the distribution of elicited preferences, underscoring their critical role in stated-preference research. The resulting parameters offer practical, evidence-based inputs for DCEA, supporting more transparent and equity-sensitive health policy decisions.

Acknowledgements

We would like to thank Sem van den Hoogen for excellent research assistance, and Marie-Anne Boujaoude, Richard Cookson, Iris Lansdorp-Vogelaar, Owen O'Donnell, Tom Van Ourti, and Aki Tsuchiya for valuable comments. The LISS panel data were collected by the non-profit research institute Centerdata (Tilburg University, the Netherlands). Funding for the panel's ongoing operations comes from the Domain Plan SSH and ODISSEI since 2019. The initial set-up of the LISS panel in 2007 was funded through the MESS project by the Netherlands Organization for Scientific Research (NWO).

Funding statement

Financial support for this study was provided by the Erasmus Initiative Smarter Choices for Better Health (Erasmus University, Rotterdam, the Netherlands) and the Dutch Institute for Public Health and the Environment (contract number: 4410002642, PI: I. Lansdorp-Vogelaar).

Data availability statement

The complete data set with accompanying codebook can be found in the LISS Data Archive (https://dataarchive.lissdata.nl/study-units/view/1657). The analysis code is publicly available on the Open Science Framework (OSF; https://osf.io/a2xgz).

Ethical Statement

Ethics approval was given by the Internal Review Boards of the Erasmus Research Institute of Management and Erasmus School of Economics [ETH2425-0791]. We have no conflicts of interests to declare.

References

Abasolo, I., & Tsuchiya, A. (2004). Exploring social welfare functions and violation of monotonicity: An example from inequalities in health. *Journal of Health Economics*, *23*(2), 313–329. https://doi.org/10.1016/j.jhealeco.2003.08.003

Aron, A., Aron, E. N., & Smollan, D. (1992). Inclusion of other in the self scale and the structure of interpersonal closeness. *Journal of personality and social psychology*, *63*(4), 596.

Asaria, M., Griffin, S., & Cookson, R. (2016). Distributional Cost-Effectiveness Analysis: A Tutorial. *Medical Decision Making: An International Journal of the Society for Medical Decision Making*, *36*(1), 8–19. https://doi.org/10.1177/0272989X15583266

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

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

Boujaoude, M.-A., Dalziel, K., Cookson, R., Devlin, N., & Carvalho, N. (2025). Aversion to income, ethnic, and geographic related health inequality: Evidence from australia. *Social Science & Medicine*, *364*, 117495.

Cookson, R., Griffin, S., Norheim, O. F., & Culyer, A. J. (2020). *Distributional cost-effectiveness analysis: Quantifying health equity impacts and trade-offs*. Oxford University Press.

Cookson, R., Kaur, G., Skarda, I., Venkatesh, S., Doran, T., Robson, M., Norheim, O. F., O'Donnell, O., & Paulden, M. (2026). The inequality-adjusted incremental cost-effectiveness ratio. *Value in health*.

Dolan, P. (1998). The measurement of individual utility and social welfare. *Journal of Health Economics*, *17*(1), 39–52. https://doi.org/10.1016/S0167-6296(97)00022-2

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

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

Edlin, R., Tsuchiya, A., & Dolan, P. (2012). Public preferences for responsibility versus public preferences for reducing inequalities. *Health Economics*, *21*(12), 1416–1426. https://doi.org/10.1002/hec.1799

Falk, A., Becker, A., Dohmen, T., Huffman, D., & Sunde, U. (2022). The preference survey module: A validated instrument for measuring risk, time, and social preferences. *Management Science*, *69*(4), 1935–1950. https://doi.org/10.1287/mnsc.2022.4455

Gächter, S., Starmer, C., & Tufano, F. (2015). Measuring the closeness of relationships: A comprehensive evaluation of the'inclusion of the other in the self'scale. *PloS one*, *10*(6), e0129478.

Horn, S. A., Norheim, O. F., & Barra, M. (2025). Inequality aversion in health: Eliciting aversion to univariate health inequalities across different health distributions. *Value in Health*.

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

Iragorri, N., Ali, S., Mishra, S., & Sander, B. H. (2025). Measuring aversion to income-related health inequality in canada: An equity-efficiency trade-off experiment. *Value in Health*.

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

Makdissi, P., & Yazbeck, M. (2016). Avoiding blindness to health status in health achievement and health inequality measurement. *Social Science & Medicine*, *171*(1), 39–47. https://doi.org/10.1016/j.socscimed.2016.10.027

Meulman, I., Rotteveel, A., Uiters, E., Cloin, M., Polder, J., & Stadhouders, N. (2026). What value do dutch citizens place on health interventions that provide greater health gains to lower-income groups? a discrete choice experiment. *International Journal of Health Policy and Management*, *15*, 9095.

Piketty, T., & Saez, E. (2013). Chapter 7 - optimal labor income taxation. In A. J. Auerbach, R. Chetty, M. Feldstein, & E. Saez (Eds.), *Handbook of public economics, vol. 5* (pp. 391–474, Vol. 5). Elsevier. https://doi.org/10.1016/B978-0-444-53759-1.00007-8

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

Pollak, R. A. (1971). Additive utility functions and linear engel curves. *The Review of Economic Studies*, *38*(4), 401–414. https://doi.org/10.2307/2296686

Richardson, J., Sinha, K., Iezzi, A., & Maxwell, A. (2012). Maximising health versus sharing: Measuring preferences for the allocation of the health budget. *Social Science & Medicine (1982)*, *75*(8), 1351–1361. https://doi.org/10.1016/j.socscimed.2012.05.036

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

Robson, M., Doran, T., O'Donnell, O., & Ourti, T. V. (2025a). Do social preferences explain health inequality aversion? *The Journal of Economic Inequality*, *23*(3), 933–956. https://doi.org/10.1007/s10888-025-09702-8

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

Robson, M., O'Donnell, O., & Van Ourti, T. (2025b). Responsibility-sensitive welfare weights for health. *Journal of Health Economics*, *102*, 103018. https://doi.org/10.1016/j.jhealeco.2025.103018

Saez, E., & Stantcheva, S. (2016). Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, *106*(1), 24–45. https://doi.org/10.1257/aer.20141362

Scherpenzeel, A. C. (2018). "true" longitudinal and probability-based internet panels: Evidence from the netherlands. In *Social and behavioral research and the internet* (pp. 77–104). Routledge.

Shimamoto, K., Doran, T., & Cookson, R. (2025). Public preferences for reducing health inequality in japan: A national survey. *Value in Health Regional Issues*, 101111.

Slejko, J. F., Ricci, S., Cookson, R., Kowal, S., et al. (2025). Health inequality aversion in the united states. *Value in Health*.

Stevens, A., Milne, R., & Burls, A. (2003). Health technology assessment: History and demand. *Journal of Public Health Medicine*, *25*(2), 98–101. https://doi.org/10.1093/pubmed/fdg022

Tropp, L. R., & Wright, S. C. (2001). Ingroup identification as the inclusion of ingroup in the self. *Personality and Social Psychology Bulletin*, *27*(5), 585–600.

Wagstaff, A. (1991). Qalys and the equity-efficiency trade-off. *Journal of Health Economics*, *10*(1), 21–41. https://doi.org/10.1016/0167-6296(91)90015-F

Wagstaff, A. (2002). Inequality aversion, health inequalities and health achievement. *Journal of Health Economics*, *21*(4), 627–641. https://doi.org/10.1016/S0167-6296(02)00006-1

Yitzhaki, S. (1983). On an extension of the gini inequality index. *International Economic Review*, *24*(3), 617–628. https://doi.org/10.2307/2526814

APPENDICES

A Study Design

A.1 Preregistered Analysis Plan

The study hypotheses and analysis plan were preregistered on the Open Science Framework (OSF; https://osf.io/a2xgz) prior to data analysis. The analysis adhered to the preregistered plan with two methodological refinements to ensure robustness and representativeness. First, we operationalized the preregistered exclusion criterion regarding comprehension checks by restricting the primary analytical sample to respondents who demonstrated graphical literacy (correctly answering ≥ 1 of 2 questions), conceptual literacy (correctly answering ≥ 3 of 5 questions), and who did not speed through the survey (> 5 minutes). We further excluded respondents with intransitive preferences (monotonicity violations). Sensitivity analyses regarding these exclusion criteria are reported in the Appendix. Second, although the study recruited a representative sample, we applied inverse probability weighting (IPW) calibrated to Dutch population data (age, sex, and education) from Statistics Netherlands (CBS) to correct for sampling imbalances after applying data quality filters. We assessed the impact of these weights on the parameter distributions using the weighted Kolmogorov-Smirnov test. Results are reported in Table C3 in Appendix C.

A.2 Experimental Instructions

Introduction

This questionnaire concerns the trade-off between, on the one hand: reducing health inequalities within a population and, on the other hand: improving the health of the population as a whole.

Imagine that you are asked to choose between two large government programmes that will affect the healthy life expectancy of the entire population. Healthy life expectancy refers to the number of years a person is expected to live, combined with the quality of life of those years. For example, imagine someone dies at age 80. This person lived in good health until age 70 and then lived another 10 years with a condition that reduced their quality of life by half. Healthy life expectancy would then be 75 years — a combination of 70 years in full health plus 5 years (half of 10 years), taking into account the reduced quality of life in the last 10 years.

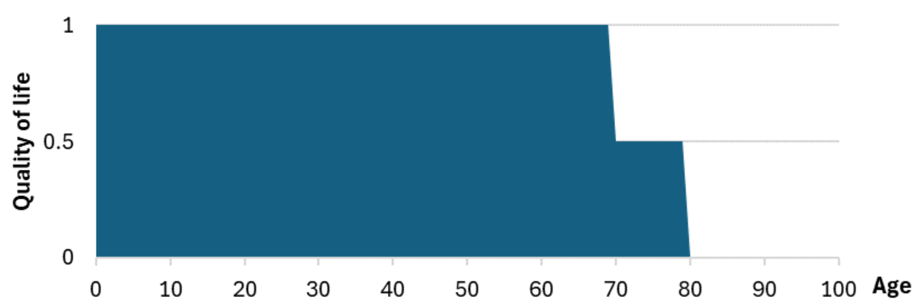

Guidance to help understand the diagrams

In this questionnaire, we ask you to weigh two programmes, as shown in the graphs below. The graphs display the healthy life expectancy for two different groups: Group H and Group L.

- Group H is the top fifth (20%) of the population with the highest (H) healthy life expectancy.
- Group L is the bottom fifth (20%) of the population with the lowest (L) healthy life expectancy.

Each graph shows how a health program can affect these groups differently. The blue bars show the number of healthy life years before a programme is introduced. The green bars on top show the extra healthy life years people gain because of the programme. The total healthy life expectancy after the programme is shown in the boxes below each bar.

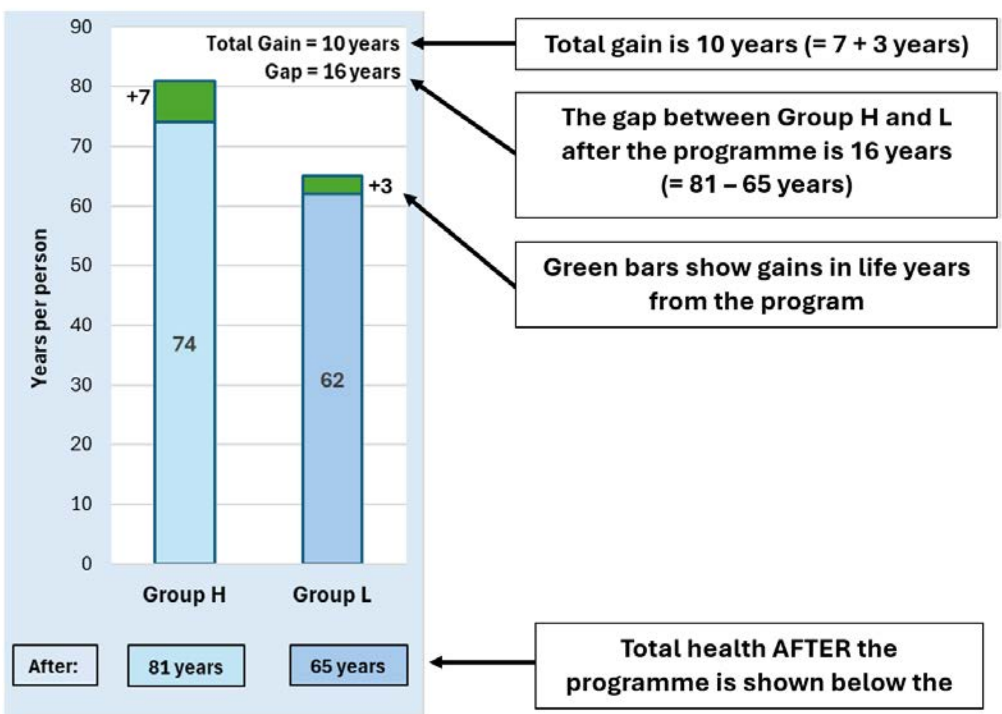

Questions on graph literacy and comprehension

[Correct answers are marked in bold.]

We will now ask you two questions on your understanding of the graphs.

1. Which of the two groups has the highest healthy life expectancy after this programme?
   (a) **Group H**
   (b) Group L
   (c) Both groups live equally long

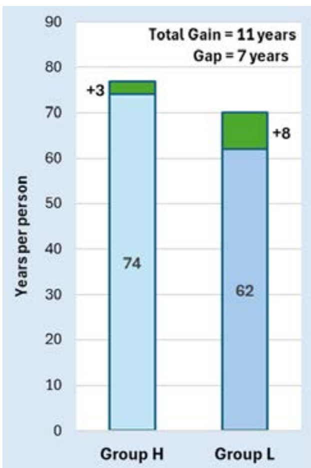

2. Which of the two groups gains more healthy life years due to this programme?
   (a) Group H
   (b) **Group L**
   (c) Both groups equally

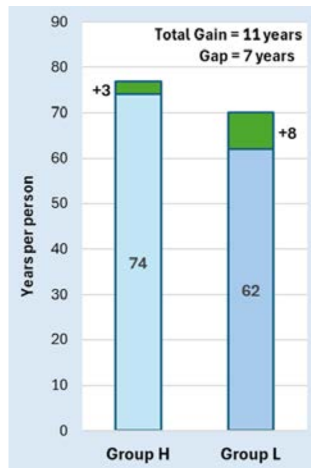

We will now ask you five questions about alternative programmes: Programme A and Programme B. The programmes will change between questions. These questions are to check your understanding of the consequences of these programmes.

3. Which of these two programmes gives the largest gain to Group H?

(a) **Programme A**

(b) Programme A and B give equal gains to Group H

(c) Programme B

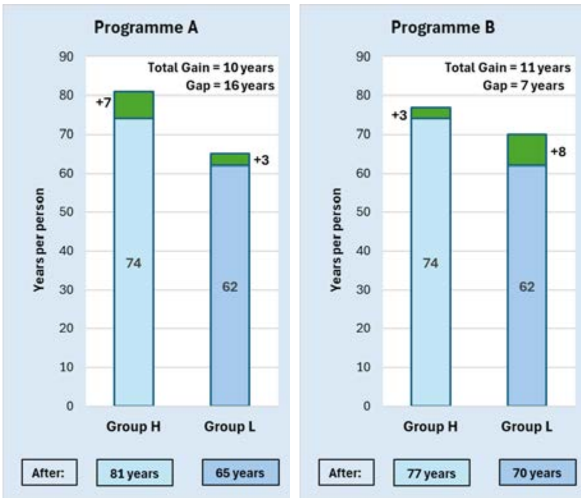

4. Which of these two programmes maximises total population health?

(a) Programme A

(b) **Programme A and B maximise total population health equally**

(c) Programme B

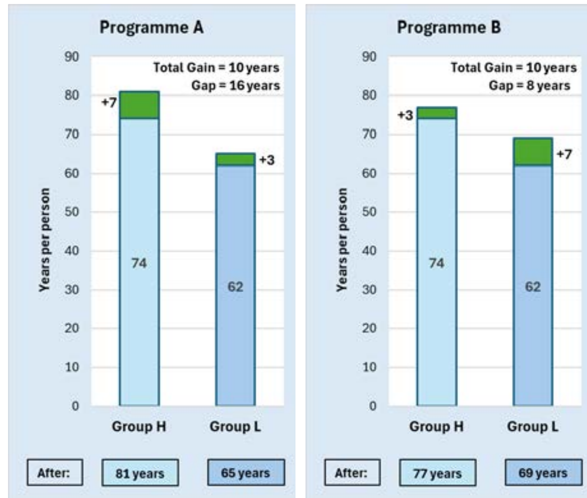

5. Which of these two programmes focuses the most on improving healthy life expectancy for Group L?

(a) Programme A

(b) Programme A and B focus equally on improving healthy life expectancy for Group L

(c) **Programme B**

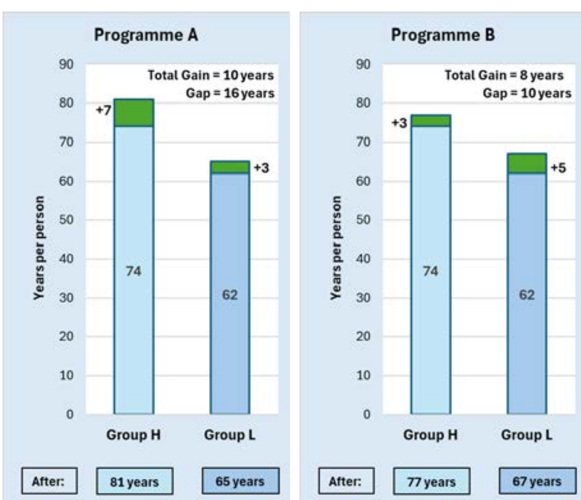

6. Which of these two programmes gives the greatest gain to Group L?

(a) Programme A

(b) **Programme A and B give the same gain to Group L**

(c) Programme B

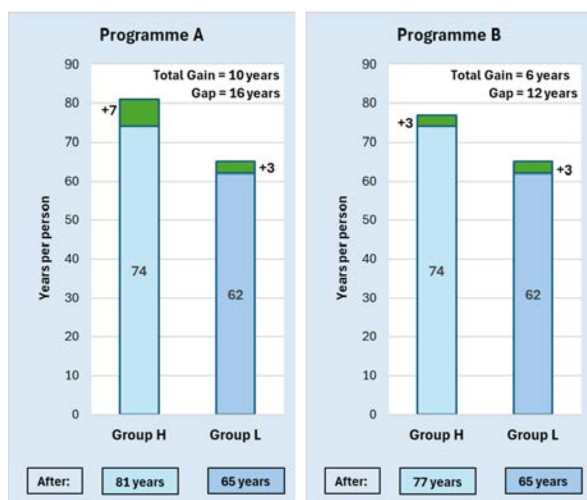

7. Which of these two programmes reduces the gap in healthy life expectancy the most?

(a) Programme A

(b) Programme A and B reduce the gap in healthy life expectancy equally

(c) **Programme B**

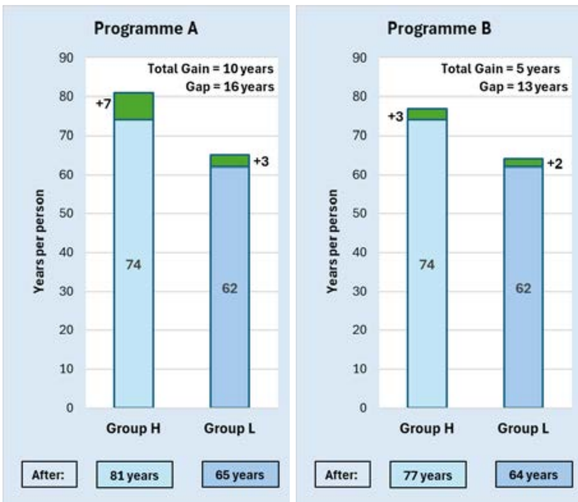

The instructions are now finished. You will move to the trade-off questions.

Trade off questions

Imagine that you are asked to choose between two large government programmes that will influence the healthy life expectancy of the population.

Keep the following in mind when making your choice:

- We cannot pay for both programmes — a **choice must** be made.
- "Equally good" means that it does not matter which programme is chosen.
- Both programmes cost exactly the same.
- The only difference between the programmes is the gain for Group H (the group with the highest healthy life expectancy) and Group L (the group with the lowest healthy life expectancy).
- The middle three fifths of the population are not affected.

8. Which programme should the government choose?

(a) Programme A

(b) Programme A and B are equally good

(c) Programme B

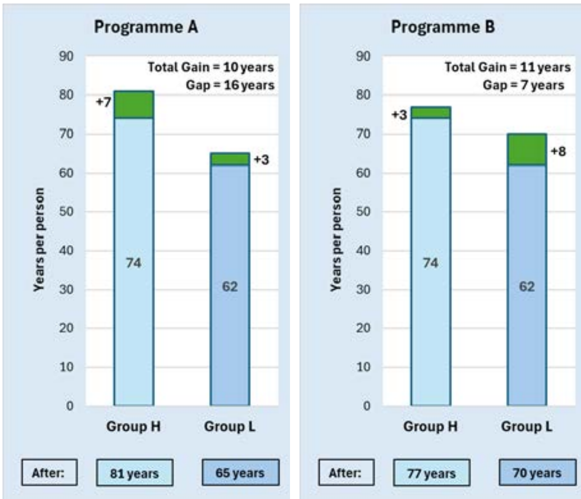

9. Which programme should the government choose?

(a) Programme A

(b) Programme A and B are equally good

(c) Programme B

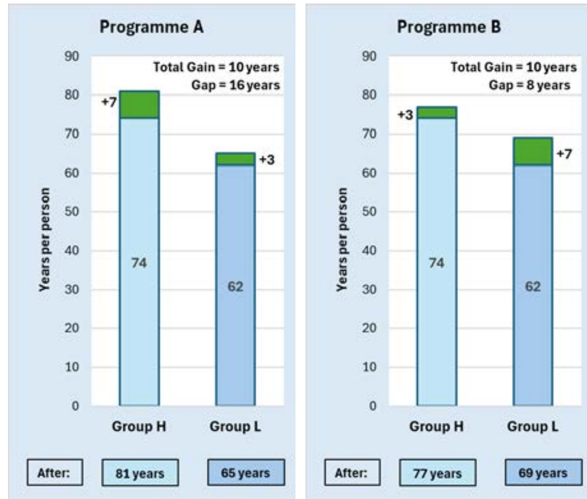

10. Which programme should the government choose?

(a) Programme A

(b) Programme A and B are equally good

(c) Programme B

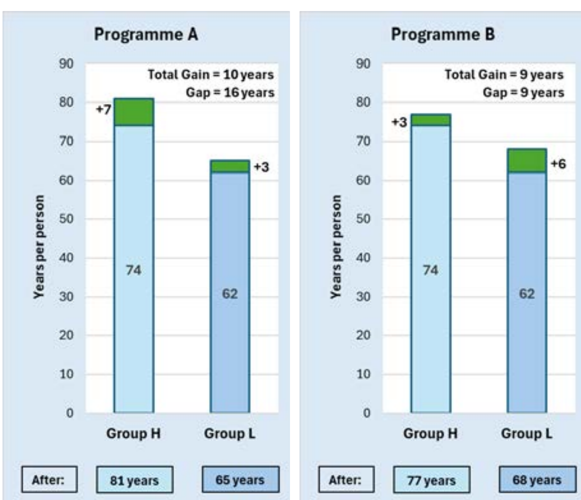

11. Which programme should the government choose?

(a) Programme A

(b) Programme A and B are equally good

(c) Programme B

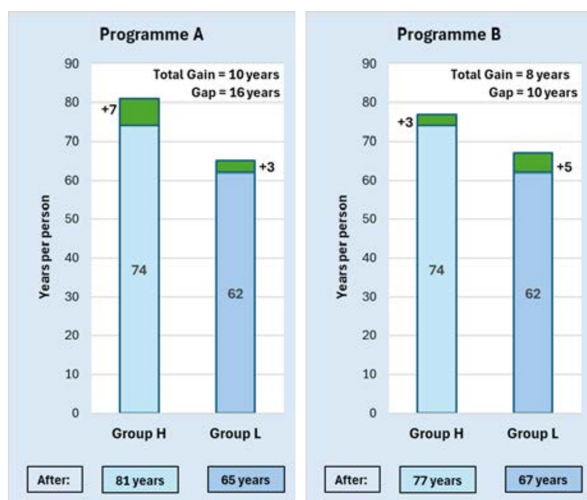

12. Which programme should the government choose?
   (a) Programme A
   (b) Programme A and B are equally good
   (c) Programme B

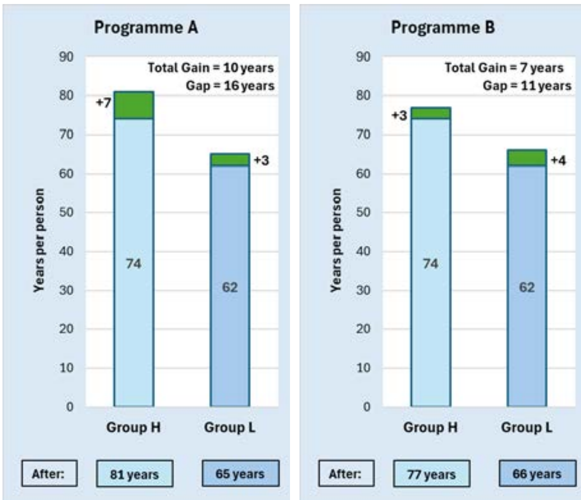

13. Which programme should the government choose?
   (a) Programme A
   (b) Programme A and B are equally good
   (c) Programme B

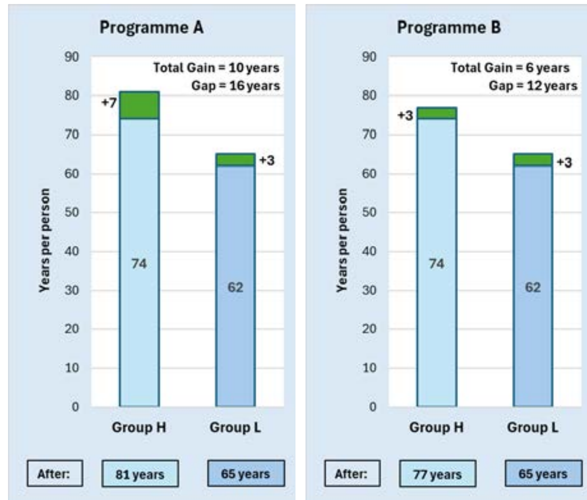

14. Which programme should the government choose?
   (a) Programme A
   (b) Programme A and B are equally good
   (c) Programme B

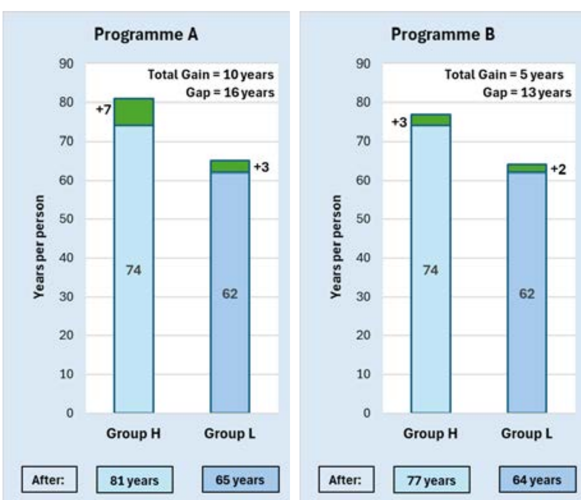

15. Which programme should the government choose?
   (a) Programme A
   (b) Programme A and B are equally good
   (c) Programme B

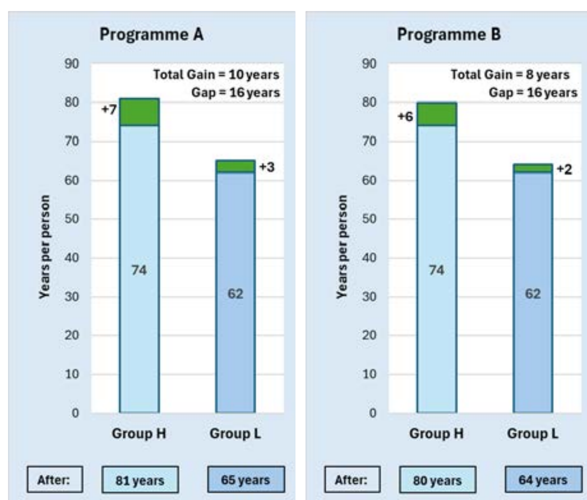

Trade off questions (by income)

Instead of changing the healthy life expectancy of Group H and Group L, the programmes now affect the healthy life expectancy of the **richest** and **poorest fifth** (20%) of the population. The graphs now display healthy life expectancy for two groups:

- **Richest fifth** represents the richest 20% of the population
- **Poorest fifth** represents the poorest 20% of the population

[Questions 16 to 23 repeat questions 8 to 15 with "Richest Fifth" and "Poorest Fifth" instead of "Group H" and "Group L", respectively.]

16. Which programme should the government choose? etc.
   (a) Programme A
   (b) Programme A and B are equally good
   (c) Programme B

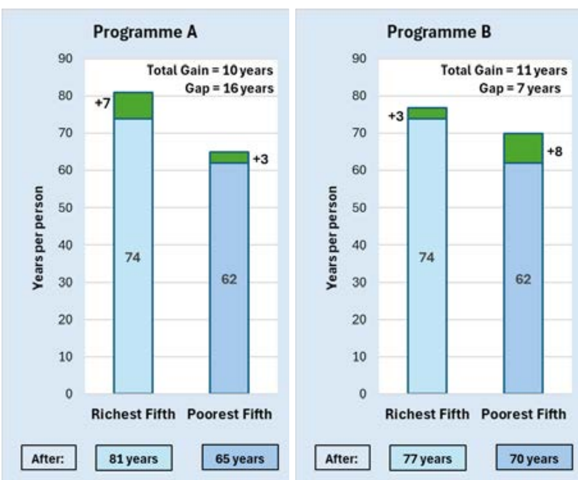

Questions on views, group connectedness and altruism

You are finished with the trade-off questions. We will now ask a few final questions.

To what extent do you agree with the following statement? [Scale from 0 to 10]

24. Improving total population health is more important than reducing inequalities in health.

25. Increasing total national income is more important than reducing income inequality.

26. Improving total population health is more important than reducing inequalities in health between the rich and poor.

27. How strongly do you identify with each of the population groups below? [Seven levels ranging from no overlap to high degree of overlap]

27a Group H (highest healthy life expectancy)

27b Group L (lowest healthy life expectancy)

27c Richest fifth of the population

27d Poorest fifth of the population

28. Imagine the following situation: You won 1,000 euros in a lottery. Considering your current situation, how much would you donate to charity? [Slider from 0 to 1,000]

A.3 Pilot

Prior to running our main experiment, we ran a pilot with a convenience sample of Dutch adults ($n = 92$). We ran the survey online, using Qualtrics, as a test of the experimental design and translation into Dutch. Following the pilot we only made minor changes to the language to increase understandability. Interestingly, we found similar distributions of both univariate and income-related health inequality aversion, with median ε of 5.358.

B Additional Theory

B.1 Social Indifference Curves

Figure B1 plots social indifference curves for different health inequality aversion parameters, ε, and income-weights, ω_i. The left panel illustrates how ε determines the concavity of social indifference curves, and assumes a symmetric HRSWF (i.e. $\omega_i = 1/N$, $\forall i$). The right panel illustrates how verying ω_p alters the slope of social indifference curves, here assuming $\varepsilon = 0$.

Figure B1: Social Indifference Curves

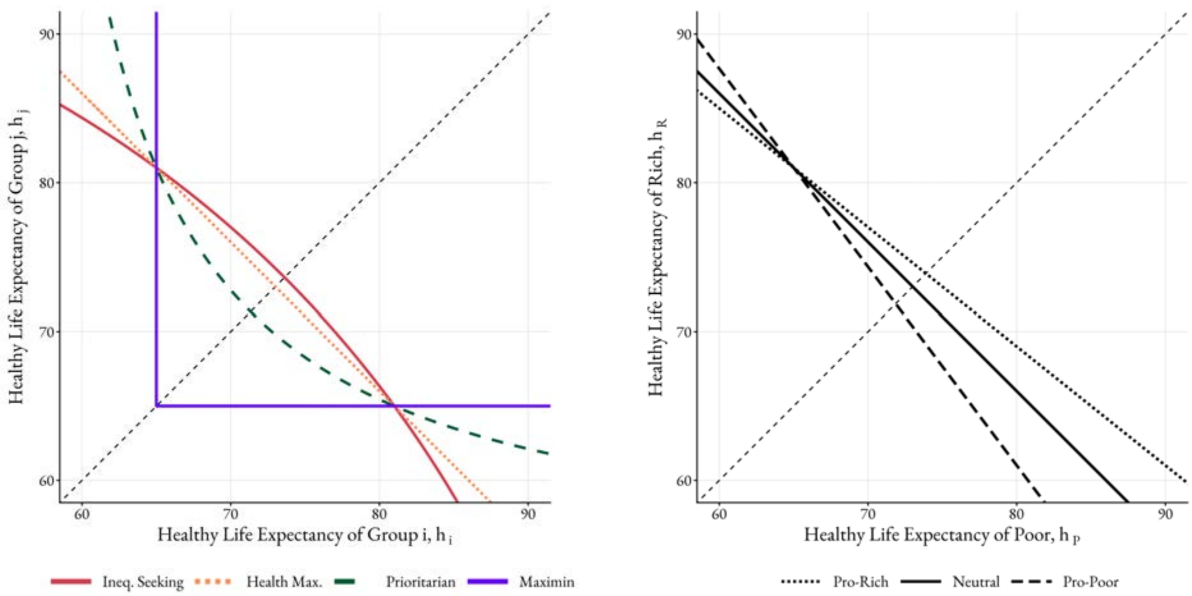

B.2 Income Weights

Figure B2 shows the elicited ω_p parameters, conditional on the choices in both the "univariate" and "income-labelled" treatments, for each level of health sacrificed. Preferences are income neutral ($\omega_p = 0.5$) when participants make the same choice in both treatments (points on the 45° line). If they are willing to sacrifice more total health in the income treatment (to benefit the poorer group), this indicates they are pro-poor ($\omega_p > 0.5$). Conversely, if they sacrifice less total health in the income treatment, participants have pro-rich preferences ($\omega_p < 0.5$). Note that if participants sacrifice four or more years, in either treatment, then income weights cannot be elicited.

Figure B2: Income Weight Elicitation

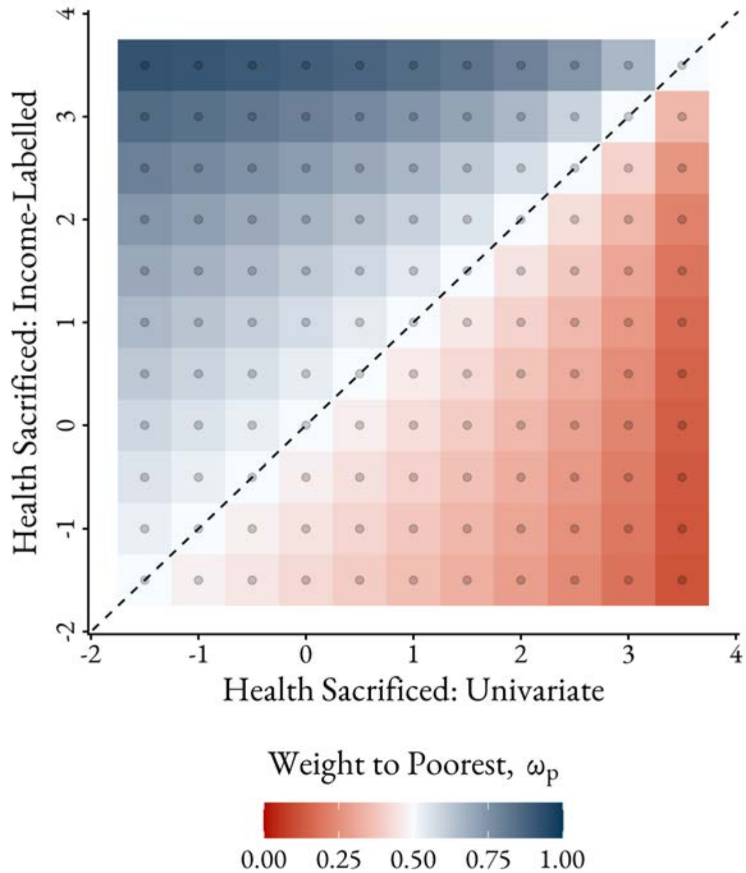

B.3 Alternative Social Welfare Functions

Kolm-Pollak. One alternative to the Atkinson social welfare function is the Kolm-Pollak HRSWF (Pollak, 1971; Kolm, 1976). This function represents a social decision makers concern with *absolute* health inequalities, rather than *relative* health inequalities. This assumes the functional form: $U(h_i) = -exp(-\alpha h_i)$, giving the EDE health of:

$$h^K_{EDE} = -\frac{1}{\alpha} \ln \left(\sum_{i=1}^{N} \omega_i \exp(-\alpha h_i) \right), \tag{B1}$$

where α is the health inequality aversion parameter. Health Maximiser preferences are approximated as $\alpha \to 0$, with health inequality aversion for $\alpha > 0$, and health inequality seeking for $\alpha < 0$. Indirect equity weights are as follows:

$$IEW^K_{ij} = \frac{\omega_i}{\omega_j} \exp\left(-\alpha\left(h_i - h_j\right)\right). \tag{B2}$$

Rank-Dependent. Another alternative is the rank-dependent HRSWF - with a weighting function based on that of the generalised Gini index (Donaldson and Weymark, 1980; Bleichrodt et al., 2004). Social utility is assumed to be linear in health, $U(h_i) = h_i$, but individuals are weighted according to their rank in the health distribution:

$$h^R_{EDE} = \sum_{i=1}^{N} \omega^R_i h_i, \tag{B3}$$

$$where \quad \omega^R_i = \left(\frac{N - r(h_i) + 1}{N}\right)^\gamma - \left(\frac{N - r(h_i)}{N}\right)^\gamma$$

where $r(h_i) \in \{1, \ldots, N\}$ is a strictly increasing ranking of health, i.e. $h_i > h_j \Rightarrow r(h_i) > r(h_j)$ for all $i \neq j$.[12] With $\gamma > 1$, individuals lower in the health distribution are given higher

[12]When health levels are not unique, individuals with identical health are treated as ties. Let $P(H = h_i)$ denote the population share at health level h_i. The difference in rank-based weights, $[P(H \geq h_i)]^\gamma -$

weights. Weights used in the standard Gini coefficient correspond to $\gamma = 2$. As $\gamma \to \infty$, the weights approach zero for all but the lowest ranked individual. With $0 < \gamma < 1$, the weights *increase* with health rank, giving more priority to those with higher health. Indirect equity weights are as follows:

$$IEW_{ij}^{K} = \frac{\omega_i^R}{\omega_j^R}. \tag{B4}$$

Table B1 presents the elicited Kolm health inequality aversion parameters (α), rank-dependent weighting parameters (γ), and Indirect Equity Weights (IEW_{ij}): for each Response Category (see Table 1). [13] The percentage of participants categorised into each response category is shown in the second column.

$[P(H > h_i)]^\gamma$, is allocated equally across individuals with $H = h_i$. Individual weights are therefore given by $\omega_i^R = \frac{[P(H \geq h_i)]^\gamma - [P(H > h_i)]^\gamma}{P(H = h_i)}$.

[13] Note that identification of γ relies on the assumption that the health rank of the worst-off group remains unchanged by Policies A or B. In particular, improvements in their health must not be large enough to move them beyond the second fifth of the population. In the experimental design, participants are informed that the middle three fifths of the population are unaffected. While this allows for the possibility of re-ranking in principle, such changes are unlikely to be salient to participants. We, therefore, highlight this assumption for completeness.

Table B1: Inequality Aversion and Indirect Equity Weights: Kolm and Rank-Dependent

	Responses	**Parameters**		**IEW**$_{ij}$	
Category	**(%)**	Kolm, α	Rank, γ	Kolm	Rank
Inequality Seeking 1	2.3	-0.028	0.876	0.711	0.727
Inequality Seeking 2	0.0	-0.019	0.912	0.792	0.800
Inequality Seeking 3	1.0	-0.010	0.953	0.886	0.889
Health Maximiser	8.3	0.000	1.000	1.000	1.000
Prioritarian 1	12.3	0.011	1.054	1.140	1.143
Prioritarian 2	0.6	0.023	1.118	1.318	1.333
Prioritarian 3	10.3	0.037	1.195	1.557	1.600
Prioritarian 4	2.0	0.053	1.291	1.898	2.000
Prioritarian 5	20.6	0.074	1.420	2.438	2.667
Prioritarian 6	2.4	0.103	1.607	3.450	4.000
Prioritarian 7	24.9	0.152	1.941	6.222	7.999
Maximin	1.3	∞	∞	∞	∞
Egalitarian 1	7.1	NA	NA	NA	NA
Egalitarian 2	0.4	NA	NA	NA	NA
Egalitarian 3	6.3	NA	NA	NA	NA

C Descriptive Statistics

Table C2 presents descriptive statistics for the full sample, alongside the unweighted and weighted analytical sample. For comparison, we show Dutch population data (for age, sex, and education) from Statistics Netherlands (CBS).

Table C2: Descriptive Statistics of the full, unweighted and weighted sample.

	Full sample			Unweighted sample	Weighted sample	Population
Variable	Mean	Min	Max	Mean	Mean	Mean
Sample Size	1101.00			472.00	472.00	
Gender[a]	1.50	1	3	1.47	1.50	1.51
Age	49.45	16	93	46.21	47.90	48.42
Household Size	2.62	1	8	2.70	2.72	
Head Lives with Partner[b]	0.67	0	1	0.72	0.72	
Urbanicity[c]	2.73	1	5	2.76	2.81	
Net Household Income	4280.64	0	23000	4799.64	4646.69	
Education Status (6 levels)[d]	3.72	1	6	4.24	3.85	
Education Status (3 levels)[d]	2.06	1	3	2.35	2.16	2.11
Altruism Score[e]	107.07	0	1000	90.11	90.16	
Connectedness High QALE[f]	4.31	1	7	4.60	4.51	
Connectedness Low QALE[f]	2.88	1	7	2.56	2.57	
Connectedness Rich[f]	3.43	1	7	3.60	3.52	
Connectedness Poor[f]	2.64	1	7	2.19	2.27	

Notes:
[a] 1 = Male, 2 = Female
[b] 1 = Yes, 0 = No
[c] 1 = Most Urban, 5 = Most Rural
[d] Higher values indicate higher education level
[e] Range 0-1000, higher values indicate more altruistic
[f] Range 1-7, higher values indicate more connected to the group

C.1 Reasons for Exclusion and Inclusion

For our main analytical sample, we applied five data quality criteria: 1) at least one of two graph literacy questions correct, 2) at least three of five conceptual comprehension questions correct, 3) at least five minutes to complete the survey, 4) no monotonicity violations in either treatment, and 5) categorisable response pattern in both treatments.

Figure C3: Composition of Excluded Sample

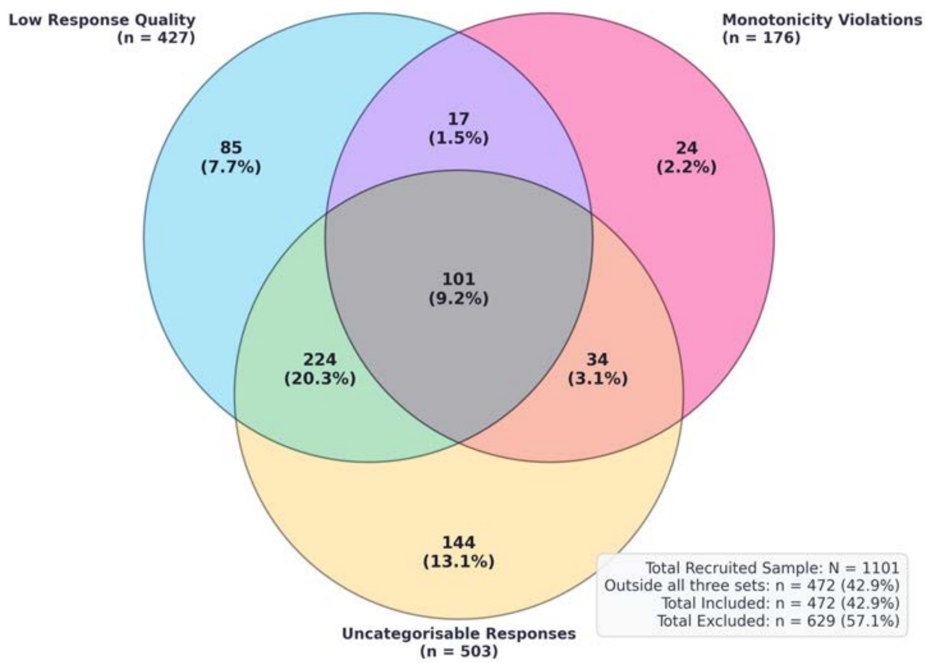

Figure C3 shows the composition of the excluded sample, grouping together those who did not meet the graph literary, comprehension and completion time criteria under "Low Response Quality". Overall, 629 responses (57.1%) are excluded, leaving 472 included. A total of 101 responses (9.2%) simultaneously exhibit low response quality, monotonicity violations, and are uncategorisable. Additional responses fall into overlapping pairs of these categories, with the largest subset being low-quality responses with uncategorisability (224, 20.3%) and being uncategorisable responses only (144, 13.1%). The quality criteria for

Figure C4: Composition of Categorisable Responses

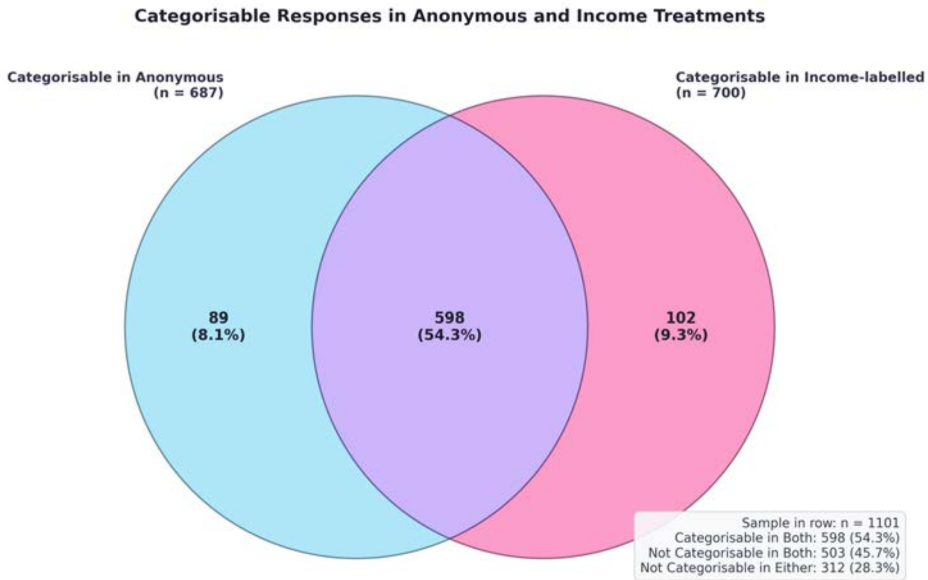

categorisable responses required a categorisable response in both the anonymous and income-labelled treatments. Figure C4 shows the distribution of individuals who had categorisable responses in either of the two treatments. The majority of responses (598, 54.3%) were categorisable in both treatments, indicating substantial overlap between the two conditions. A smaller share is uniquely categorisable in only one condition (89 in anonymous, 102 in income-labelled), while 312 responses (28.3%) cannot be categorised in either.

Figure C5 shows the composition of the included sample. Of the 472 included responses (42.9% of the full sample), all are categorisable, with medium or high response quality and no monotonicity violations. Only small subsets fail one of these criteria individually, with the exception of no monotonicity violations but low quality and/or uncategorisable responses (224, 20.3%).

Figure C5: Composition of Included Sample

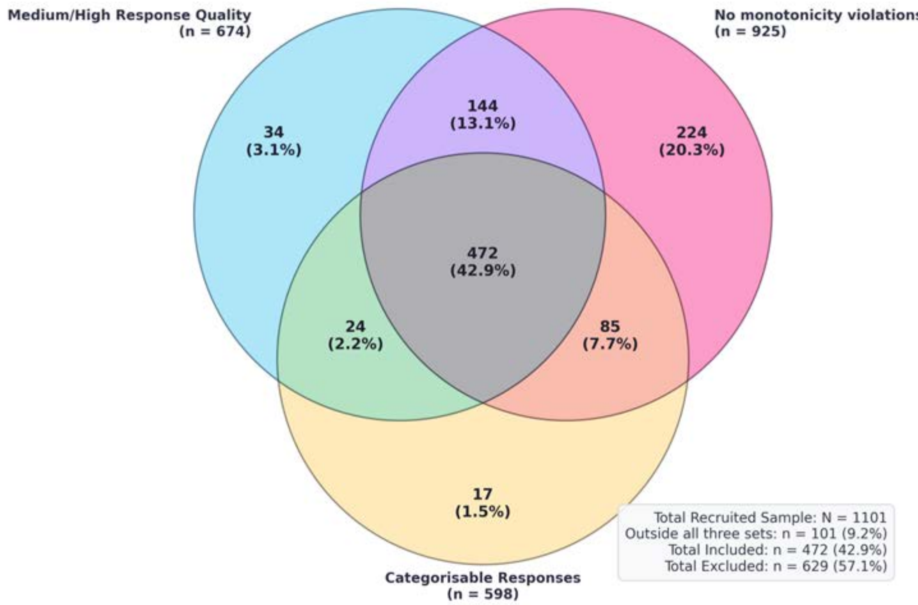

C.2 Inverse Probability Weighting

Figure C6 provides a detailed comparison of the distributions of sex, age, and education across the unweighted and weighted analytical samples, benchmarked against the representative Dutch population.

We compared the inverse-probability weighted and unweighted empirical distributions of the derived inequality aversion and income-weight parameters using weighted KS statistics. Because the same respondents are observed under alternative weight vectors, p-values were obtained from a sign-flip randomisation test with 10,000 draws. We found no statistically significant differences.

Figure C6: Distribution of Sex, Age and Education: Unweighted, Weighted, and Dutch Population

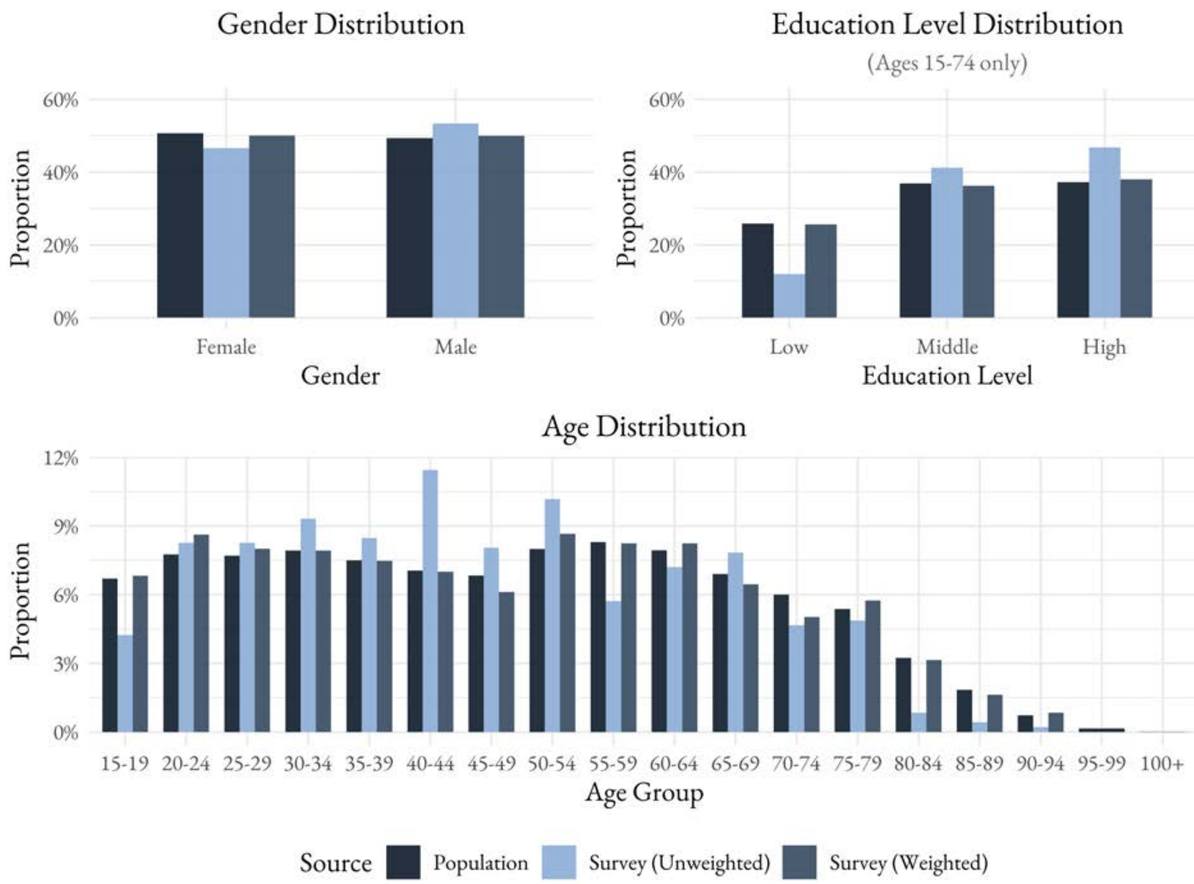

Table C3: Kolmogorov-Smirnov Test: IPW vs. Unweighted Distributions

Parameter	D Statistic	P-value	Observations
Inequality Aversion (Univariate)	0.0131	0.9597	408
Income Weight (Income-Labelled)	0.0247	0.3725	408

Sign-flip randomisation p-values (10,000 draws).

D Additional Results

D.1 Median Health Inequality Aversion

Table D4 shows the estimated median parameters for both univariate and income-related inequality aversion. The median estimates of ε in both experiment treatments were identical and imply that the median respondent is equally averse to pure health inequality as to income-related health inequality. However, the confidence intervals for the income-related aversion parameter are wider.

Table D4: Median Estimates with Bootstrap Percentile Confidence Intervals

	Dependent variable:	
	Univariate, ε	Income-Related, ε
	(1)	(2)
Intercept	5.36***	5.36***
95% CIs	(5.36, 5.36)	(5.36, 10.95)
Observations	472	472
Note:		*p<0.1; **p<0.05; ***p<0.01

D.2 Kolmogorov–Smirnov Test: Univariate vs. Income-Labelled Treatment

Table D5 presents results of a two-sample weighted Kolmogorov–Smirnov test comparing the distributions of health sacrificed in the univariate and income-labelled treatments within the analytical sample (n = 472, IPW-weighted). The test finds no statistically significant difference between the two distributions, consistent with the observation that univariate and income-related health inequality aversion are similar in this sample.

Table D5: Kolmogorov-Smirnov Test: Univariate vs. Income-Labelled Distributions

Comparison	D Statistic	P-value	Observations
Univariate vs Income-Labelled	0.0027	1	472

Two-sample weighted KS test comparing univariate and income-labelled distributions of health sacrificed.

D.3 Sense Check

We ask two questions on participants views on whether reducing health inequalities is more important than improving overall health of the population. The first framed in a univariate context, the second specifically on socioeconomic health inequalities.[14] Table D6 shows results of a regression of these variables on the *health sacrificed* in the univariate and income treatments, respectively. Reassuringly, we find strong and significant negative associations between these variables. The health inequality aversion we elicit in the experiment corresponds to participants reported views on equity and efficiency.

D.4 Data Quality: Literacy, Competency and Monotonicity

Although our main analysis reports response distributions based on aggregated data quality measures, this section presents the distributions separately for each individual data quality indicator. The results show a consistent pattern across measures: respondents who satisfy the data quality checks exhibit a substantially higher proportion of categorisable responses and a lower prevalence of "extreme" responses.

[14]The exact wording of the questions were: 1) Univariate: "*On a scale from 0 to 10, where would you place your opinion? Where: 0 = Reducing health inequalities is more important than improving the overall health of the population 10 = Improving the overall health of the population is more important than reducing health inequalities.*" 2) Income: "*On a scale from 0 to 10, where would you place your opinion? Where: 0 = Reducing health inequality between the rich and poor is more important than improving the overall health of the population 10 = Improving the overall health of the population is more important than reducing health inequalities between the rich and poor.*"

Table D6: Sense Check

	Dependent variable: *Health Sacrificed*	
	Univariate	Income
	(1)	(2)
Reducing Inequalities vs Improving Overall: Univariate	-0.285^{***}	
	(0.026)	
Reducing Inequality vs Improving Overall: Income		-0.279^{***}
		(0.024)
Constant	4.013^{***}	4.079^{***}
	(0.160)	(0.144)
Observations	472	472
R^2	0.203	0.221
Note:	*p<0.1; **p<0.05; ***p<0.01	

Figure D7: Response Distribution by Graph Literacy

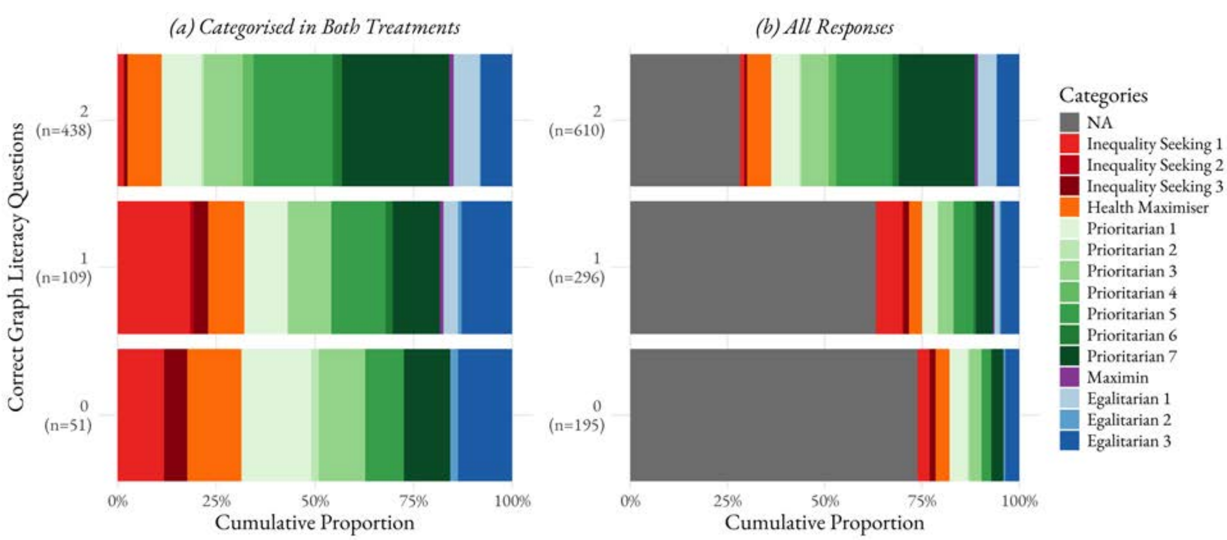

Table D7 presents pairwise Kolmogorov–Smirnov test statistics comparing the distributions of health sacrificed across cumulative data quality groups. We find large and statistically significant differences when comparing high-quality or combined high-and-medium-

Figure D8: Response Distribution by Conceptual Competency

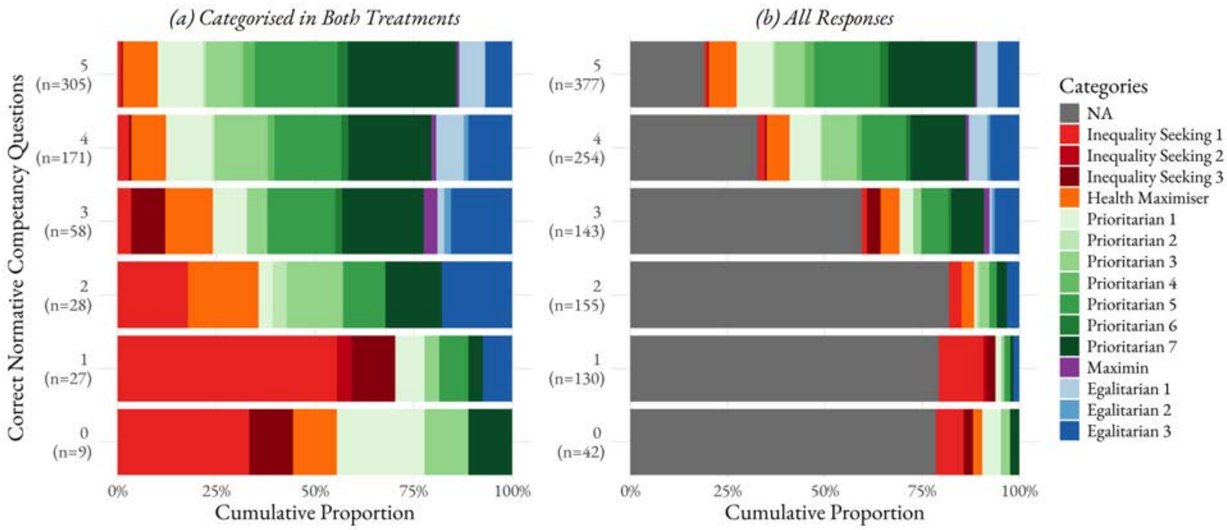

Figure D9: Response Distribution by Timing Violations

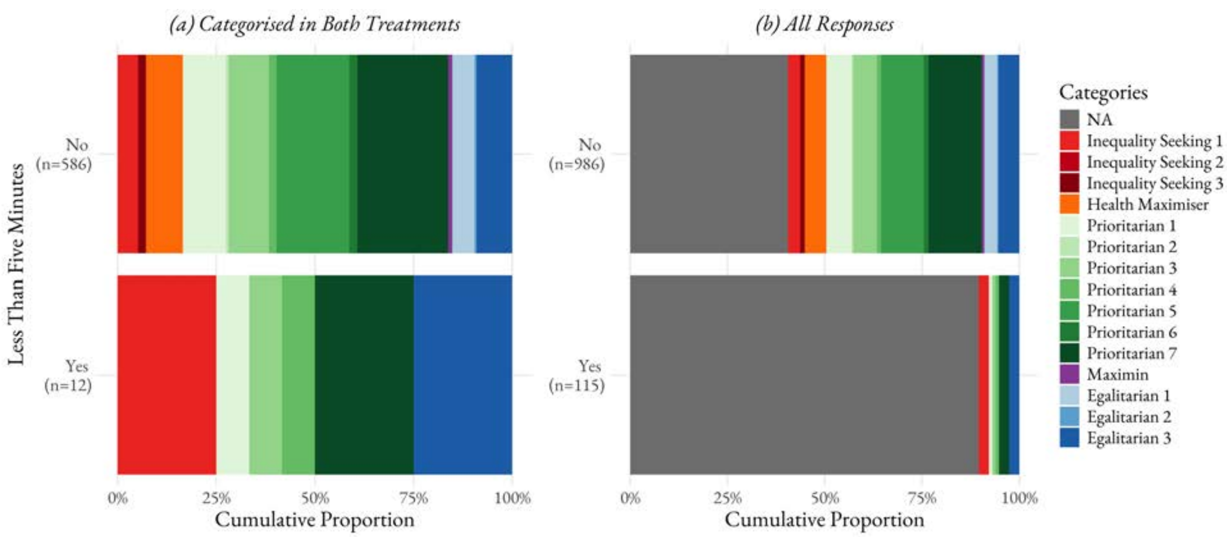

quality responses against low-quality responses ($p < 0.001$). Comparisons between high and medium quality groups, however, do not reach statistical significance, suggesting that excluding low-quality responses is crucial, while combining high and medium quality responses preserves the underlying preference distribution.

Figure D10: Response Distribution by Monotonicity Violations

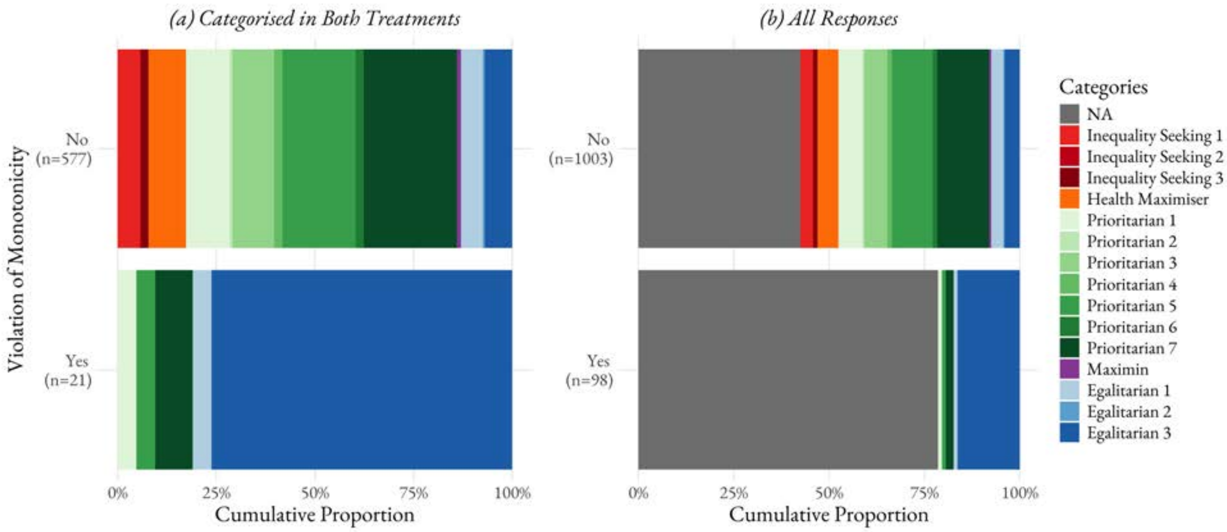

Table D7: Kolmogorov-Smirnov Tests: HealthSacr Distributions by Data Quality Level

Comparison	D Statistic	P-value	N1	N2
Univariate: High vs High+Med	0.0292	0.9987	265	472
Univariate: High vs All	0.0676	0.3704	265	598
Univariate: High+Med vs All	0.0435	0.7012	472	598
Univariate: High vs Low	0.2268	0.0003	265	126
Univariate: High+Med vs Low	0.2063	0.0004	472	126
Income-Labelled: High vs High+Med	0.0358	0.9814	265	472
Income-Labelled: High vs All	0.0720	0.2966	265	598
Income-Labelled: High+Med vs All	0.0362	0.8798	472	598
Income-Labelled: High vs Low	0.2076	0.0013	265	126
Income-Labelled: High+Med vs Low	0.1718	0.0056	472	126

Two-sample KS tests comparing distributions across data quality groups (unweighted).

D.5 Heterogeneity in Health Sacrificed

Table D8 presents regression results for *health sacrificed* on participants' sociodemographic characteristics, as well as measures of altruism and connectedness differences, for both the univariate and income-labelled treatments. Results show that females sacrifice significantly more health than males in the income-labelled treatment; in other words, they have higher levels of income-related health inequality aversion. However, we find no significant differences by age, household composition, rural residence, household income, or education.

We do find significant differences in health sacrificed by connectedness differences when interacted with altruism. Connectedness differences show the difference in reported connectedness to the low (poor) group minus their connectedness to the high (rich) group (Aron et al., 1992; Tropp and Wright, 2001; Gächter et al., 2015). We find that for those who are "selfish", there is a significant association between connectedness differences and health sacrificed. In other words, selfish individuals prioritise their own group. But for "altruists", there is no significant relationship. This indicates that those who are "selfish" prioritise their own group, while those who are "altruists" do not. Figure D11 plots the predicted health sacrificed according to the interaction between altruism and connectedness differences – for the income treatment – to illustrate the intuition behind these interactions. This finding is important to account for as, in socioeconomic-labelled designs (where the rich typically have more health) a priority for the poor is the same as the priority for the worse-off. Meaning that a priority for those in your in-group will either leave higher inequality aversion (if poor) or lower inequality aversion (if rich).

Table D8: Health Sacrificed on Sociodemographics, Connectedness and Altruism

	Dependent variable: *Health Sacrificed*	
	Univariate	Income
	(1)	(2)
Female	0.274	0.534***
	(0.178)	(0.170)
Age	−0.001	−0.0002
	(0.005)	(0.006)
Household size	−0.008	−0.072
	(0.091)	(0.085)
Head lives with partner	−0.006	0.214
	(0.259)	(0.280)
Rural	0.049	0.004
	(0.067)	(0.064)
Household net income	−0.0001	−0.00000
	(0.00005)	(0.00005)
Education category	0.071	0.027
	(0.068)	(0.059)
Altruism	0.218	0.259
	(0.247)	(0.214)
Connectedness: Low-High	0.130**	
	(0.060)	
Connectedness: Low-High × Altruism	−0.099	
	(0.076)	
Connectedness: Poor-Rich		0.158***
		(0.059)
Connectedness: Poor-Rich × Altruism		−0.141*
		(0.076)
Constant	1.940***	1.644***
	(0.504)	(0.468)
Observations	433	433
R^2	0.048	0.069
Note:	*p<0.1; **p<0.05; ***p<0.01	

Figure D11: Predicted Health Sacrificed by Altruism and Connectedness: Poor-Rich

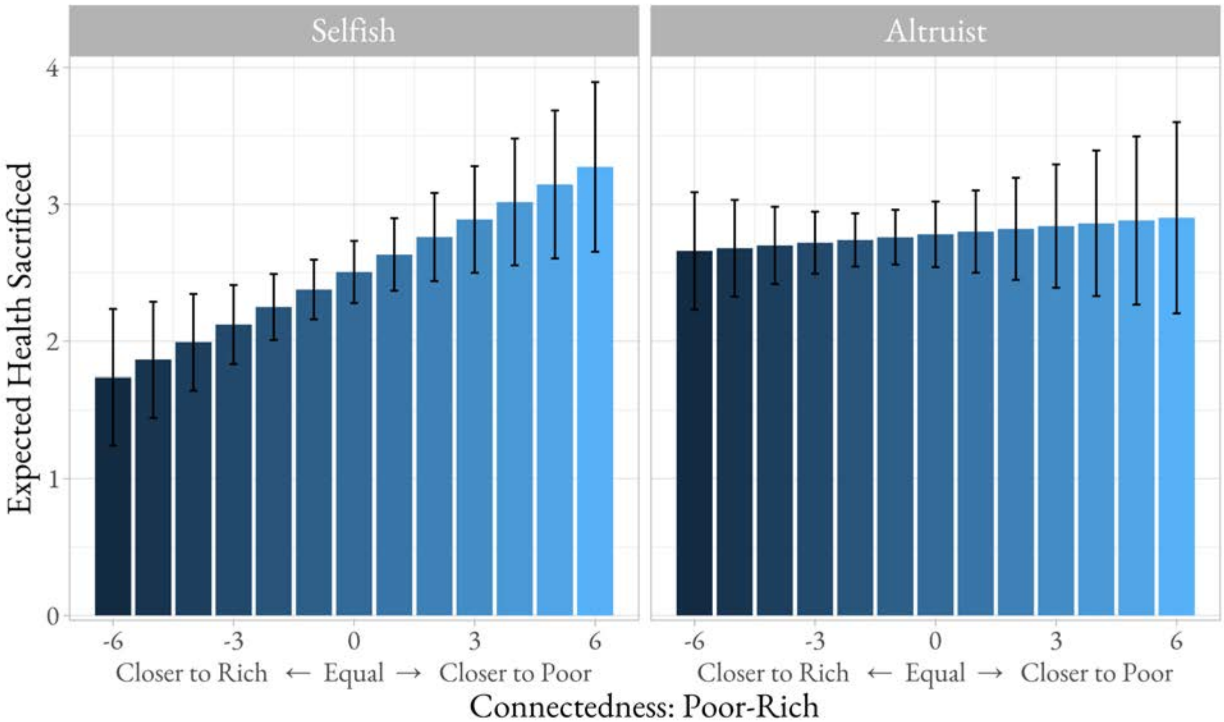

D.6 Heterogeneity in Data Quality

Table D9 presents regression results of sociodemographics on a) whether participants responses could be categorised, b) data quality, and c) inclusion in the analytical sample. Results show that those who are younger, richer and more highly educated are more likely to give categorisable responses, to have higher data quality and are included in our analytical sample.

Table D9: Regressions of Categorised, Quality and Excluded on Sociodemographics

	Dependent variable:		
	Categorised	Data Quality	Included
	(1)	(2)	(3)
Female	−0.039	−0.095*	−0.048
	(0.031)	(0.050)	(0.030)
Age	−0.002*	−0.010***	−0.004***
	(0.001)	(0.001)	(0.001)
Household size	−0.019	−0.043	−0.018
	(0.016)	(0.026)	(0.016)
Head lives with partner	−0.011	0.058	0.052
	(0.044)	(0.067)	(0.042)
Rural	0.001	0.003	0.005
	(0.012)	(0.019)	(0.011)
Household net income	0.00002**	0.00003*	0.00002**
	(0.00001)	(0.00001)	(0.00001)
Education category	0.062***	0.164***	0.082***
	(0.011)	(0.017)	(0.011)
Constant	0.421***	0.872***	0.332***
	(0.098)	(0.159)	(0.097)
Observations	967	967	967
R^2	0.062	0.174	0.117
Note:			*p<0.1; **p<0.05; ***p<0.01

D.7 Participant-Level Policy Evaluation and Indirect Equity Weights

Figure D12 presents participant-level policy evaluations of the simple policy options in Section 3.3. The left panel plots the distribution of participant-level Welfare Differences, $h^A_{EDE} - h^B_{EDE}$, whilst the right panel plots the distribution of Indirect Equity Weights, IEW_{ij}. We plot these distributions for the three HRSWFs: a) univariate-symmetric, ii) income-symmetric and ii) the asymetric HRSWF with income weights.

Figure D12: Policy Evaluation and Indirect Equity Weights

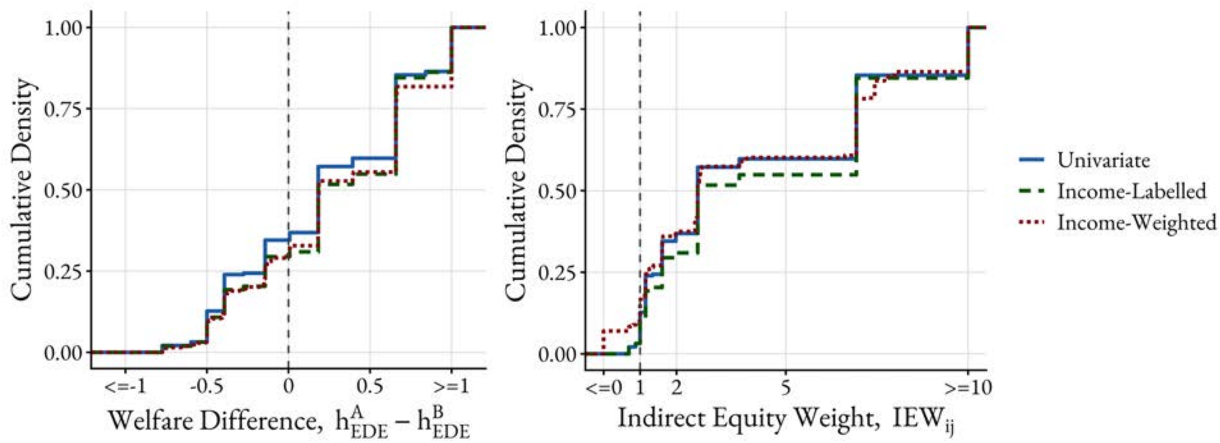

The left panel of Figure D12 shows the distribution of welfare differences between two (equally costly) interventions A and B, which affect two groups: when we compare a poorer group (i) with a QALE of 62, to a richer group (j) with a QALE of 74. Policy A provides a Net Health Benefit of 1 QALE to i (and 0 to j), whilst Policy B provides a Net Health Benefit of 2 QALEs to j (and 0 to i).[15] If $h^A_{EDE} - h^B_{EDE} > 0$, then intervention A is preferred, whilst is $h^A_{EDE} - h^B_{EDE} < 0$ then B is preferred. Results show that the distribution of evaluations are similar, with $65.5\% - 70.6\%$ of participants preferring the less efficient, but inequality reducing, intervention A to the more efficient intervention B. With participant-level estimates, any two interventions can be similarly evaluated, proving a prediction of the

[15]We assume that both health benefits and opportunity costs have been fully accounted for, so that the QALEs entering the HRSWF represent net health consequences: e.g. $h_{iA} = h_i + \Delta h_{iA}$, where Δh_{iA} are health gains from A minus the health foregone from spending elsewhere.

percentage of the public who would support each policy: with majority rule as a decision criteria.[16]

The right panel of Figure D12 shows the distribution of Indirect Equity Weights, IEW_{ij}, when we compare the poorer group (i) to the richer group (j). Results show that irrespective of the SWF or treatment used the median IEW is: 2.581. In other words, the median participant weighs marginal gains in QALYs for poorest group, 2.581 times that of the richest group. We see few (5.3%-9.5%) give equal weight $IEW_{ij} = 1$ to both groups, and fewer (3.2%-10.2%) give higher weight to the richer group $IEW_{ij} = 1$. The majority (84.5%-89.2%) give higher weights to the worst-off group.

[16] For the symmetric HRSWFs, this is equivalent to conducting the above policy evaluation with median health inequality aversion parameters. Here, we obtain $h^A_{EDE} - h^B_{EDE} = 67.341 - 67.158 = 0.183$. According to the median voter, intervention A is preferred, in both univariate and income-related contexts. For asymmetric, two parameter HRSWF, one cannot use median parameters; and should evaluate the policies for each participant, and then aggregate.