Supplementary Material: Policy Discounting Across and Beyond the Lifespan

Contents

1 Ethics

The study received ethical approval from University of Essex Ethics Sub Committee 3 (ETH2223-0950).

This follow-up survey received ethical approval from University of Essex Ethics Sub Committee 3 (ETH2223-1501).

2 Data Availability

Replication data and code are available at https://osf.io/rge8v/.

3 Pre-registration

3.1 Pre-registration details

This study was pre-registered on OSF prior to data collection on 13th April 2023: https://osf.io/n6kcm.

3.2 Deviations from pre-registered procedure

The hypothesis stated in my pre-registration was "People discount future policy benefits that will be realised after the end of their own lives more than those that will be realised within their own lifetime (precise wording may differ)." In the manuscript, I decompose this hypothesis into a 'weak' and 'strong' version and report how the results correspond to each of these. Both retain the meaning of the original pre-registered version.

The pre-registered analyses focus on estimating average marginal component effects (AMCEs). In the manuscript, I opt instead to report marginal means in most cases as they speak more directly to my hypotheses. However, I also report corresponding AMCEs for differences between these marginal means, and there is a direct correspondence between the two measures regardless.

The follow-up study measuring respondents' parental status and subjective life expectancy was not pre-registered. The pre-registration makes no mention of measuring effects across parental subgroups or using subjective measures of life expectancy to measure the lifespan. These were suggested by peers who read and commented on an early version of the manuscript that did not feature these additional analyses. Partly for this reason, the analyses using these measures are only reported in the Supplementary Material. They do not contradict the results reported in the main text.

4 Sample

4.1 Age group surveys

The survey was fielded simultaneously on 17th April 2023 to three separate age group samples: 20-34-year-olds, 35-54-year-olds, and over 55s. In the 20-34-year-old age group (N = 802) the mean age was 28.05, with a standard deviation of 4.01. Men constituted 45.51% of this sample, women constituted 53.37%. White respondents constituted 79.05, Asian 9.73, black 4.49, and mixed 4.49 of this sample. In the 35-54-year-old age group (N = 803), the mean age was 43.7, with a standard deviation of 5.84. Men constituted 50.31% of this sample, women constituted 49.19%. White respondents constituted 89.29%, Asian 5.11%, black 1.87%, and mixed 1.87% of this sample. In the over 55s age group (N = 800), the mean age was 62.46%, with a standard deviation of 5.99%. Men constituted 49.12% of this sample,

Table 1: Distribution of ethnicity in sample and population

Ethnicity	Census	Survey
Asian	8.46%	5.5%
Black	3.6%	2.69%
Mixed	1.81%	2.31%
White	84.17%	88.91%
Other	1.96%	0.59%

women constituted 50.75%. White respondents constituted 95.75%, Asian 1.5%, black 1.62%, and mixed 0.5% of this sample.

When combined, the total sample from the age group surveys is approximately representative of the wider UK population in terms of gender and ethnicity. Men comprise 48.6% of the total sample, women 51.4%, almost exactly reflecting the population distribution of 49% to 51%. Table 1 shows that the percentages of each ethnic group in the sample are close to their presence in the wider over-20 population, slightly under-representing the Asian and Black population while correspondingly over-representing the White population.

4.2 Follow-up survey

Following peer feedback on the original study, respondents were later recontacted via Prolific to complete a short follow-up survey, 18-24th July 2023. I received complete responses from 1995 respondents of the total of 2405 who originally participated – approximately 83%. Of this recontacted group, 29.22% came from the original 20-34 group, 34.59% came from the original 35-54 group, and 36.19% came from the original 55+ group. Men constituted 49.07% of the follow-up sample, and women constituted 50.53%.

5 Power Analysis

Following the approach to conjoint power analysis recommended by Stefanelli and Lukac (2020), the 'quantitative' conjoint (1237 respondents, 8 tasks, 5 levels of relevant feature) has a 98% chance of correctly detecting statistically significant AMCEs of 0.05 or larger, and the 'qualitative' conjoint (1168 respondents, 8 tasks, 3 levels) has approximately a 100% chance. The weakest-powered age group analysis (20-34 year olds, 412 respondents, 8 tasks, 5 levels) has around an 83% chance of detecting significant AMCEs of 0.05 or larger. The weakest-powered covariate subgroup analysis (right-wing ideology, 558 respondents, 8 tasks, 3 levels) has a 97% chance of detecting subgroup AMCEs of 0.05 or larger.

6 Covariates

6.1 Impartial beneficence

Impartial beneficence was measured using the impartial beneficence subscale of the Oxford Utilitarianism Scale. Respondents report agreement with the following items (order randomised) on a five-point scale (strongly disagree, somewhat disagree, neither agree nor disagree, somewhat agree, strongly agree):

- From a moral perspective, people should care about the well-being of all human beings on the planet equally; they should not favour the well-being of people who are especially close to them either physically or emotionally.
- From a moral point of view, we should feel obliged to give one of our kidneys to a person with kidney failure since we don't need two kidneys to survive, but really only one to be healthy.

- If the only way to save another person's life during an emergency is to sacrifice one's own leg, then one is morally required to make this sacrifice.
- It is just as wrong to fail to help someone as it is to actively harm them yourself.
- It is morally wrong to keep money that one doesn't really need if one can donate it to causes that provide effective help to those who will benefit a great deal.

A score was constructed by normalising summed responses to all items. Figure 1 displays the distribution of these scores across each of the three age groups. These normalised scores were then split into three tercile groups for the purpose of subgroup analysis.

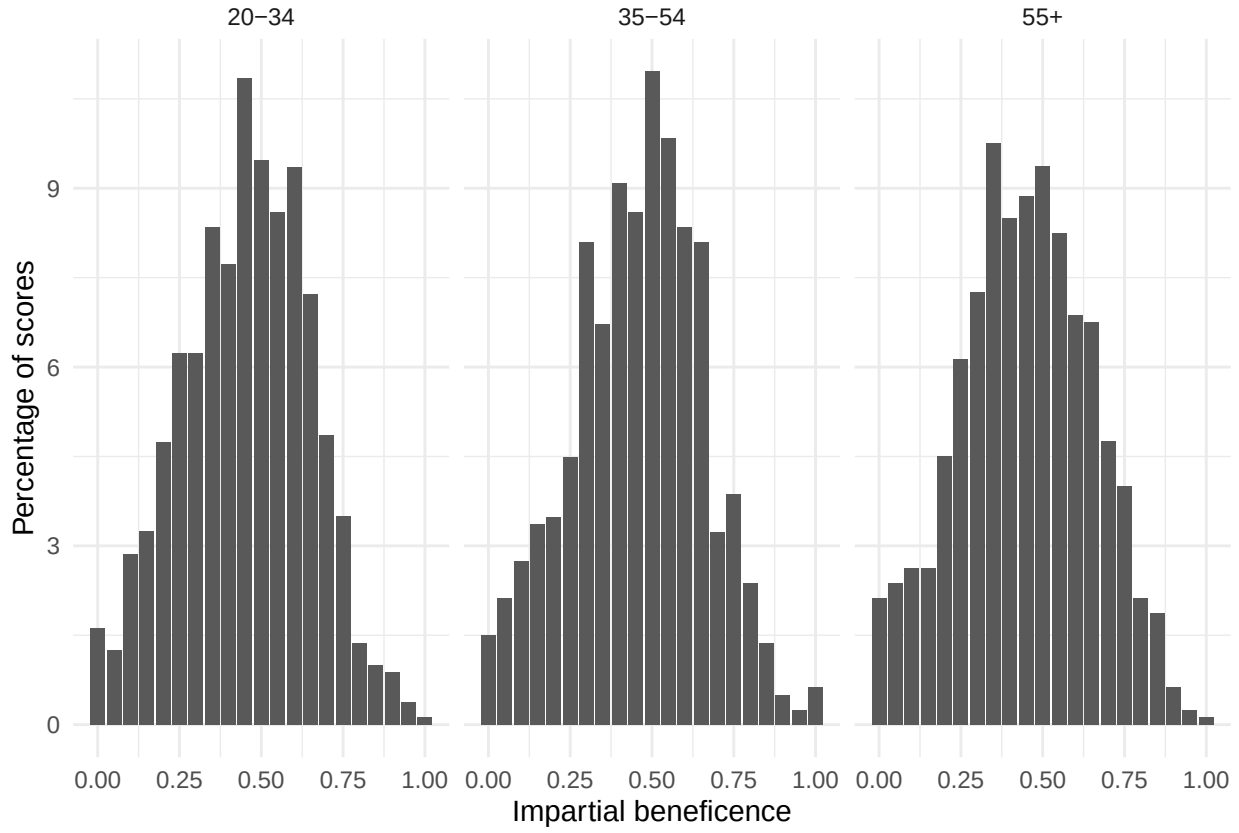

Figure 1: Distribution of normalised impartial beneficence scores. Larger values indicate greater endorsement of the principle of impartial beneficence.

6.2 Temporal focus

Future focus was measured using a reduced version of the Temporal Focus scale. Respondents report agreement with the following items (order randomised) on a five-point scale (strongly disagree, somewhat disagree, neither agree nor disagree, somewhat agree, strongly agree):

- I live my life in the present. (Present focus.)
- I think about what the future has in store. (Future focus.)
- I replay memories of the past in my mind. (Past focus.)
- I imagine what tomorrow will bring for me. (Future focus.)
- My mind is on the here and now. (Present focus.)
- I reflect on what has happened in my life. (Past focus.)
- I think my life is going well. (Filler.)
- I reflect on other people's well-being. (Filler.)

Future focus was constructed by normalising summed responses to future focus items and reverse-coded responses to past focus items. Figure 2 displays the distribution of these scores across each of the three age groups. These normalised scores were then split into three tercile groups for the purpose of subgroup analysis.

6.3 Political ideology

Left-right political ideology was measured on the following 11-point scale:

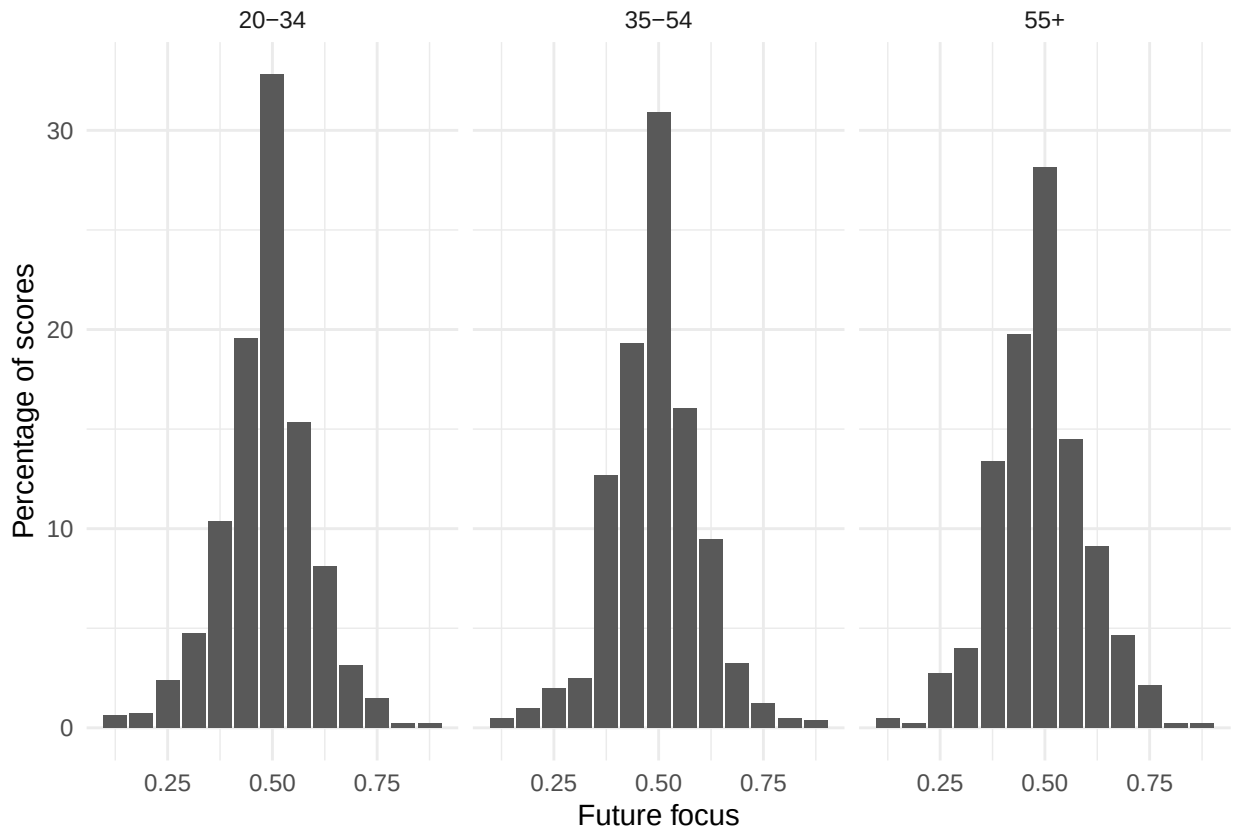

Figure 2: Distribution of normalised future focus scores. Larger values indicate greater psychological emphasis on the future.

- In politics people sometimes talk of left and right. Where would you place yourself on the following scale, where 0 indicates that you are very left-wing and 10 indicates that you are very right-wing?

Figure 3 displays the distribution of normalised responses across each of the three age groups. Leftist, centrist, and rightist groups were constructed by splitting normalised responses to this item into tercile groups.

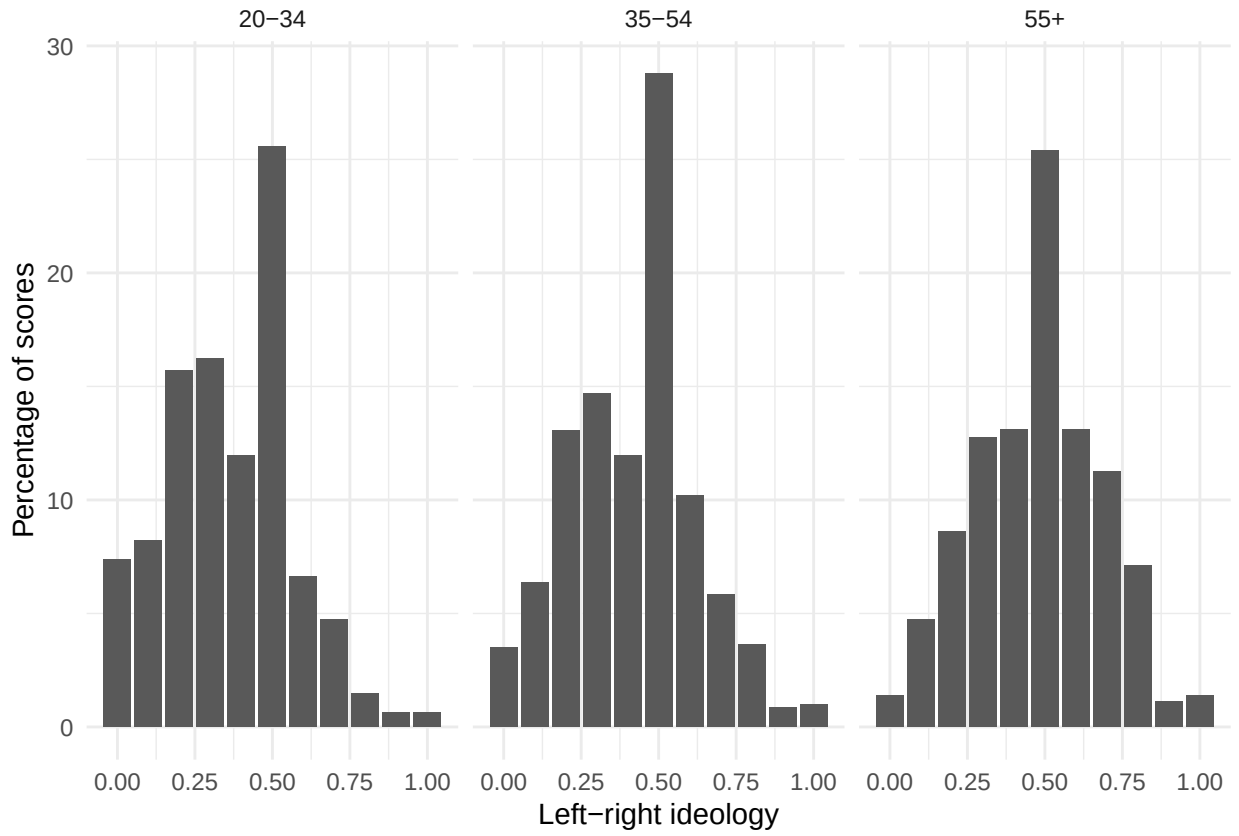

Figure 3: Distribution of normalised left-right ideology scores. Larger values indicate more right-wing political ideology.

7 Experiment

7.1 Preamble

Prior to completing the conjoint choice tasks, respondents read the following text:

Imagine that the UK government is considering adopting a new policy focused on saving lives. The total cost of the policy would be paid out immediately, but the benefits would not be felt until later. You will now be provided information about different pairs of proposals for such a program and will be asked to choose which you prefer.

7.2 Example screenshot

Figure 4 displays an example of the task presented to respondents in both the quantitative and qualitative condition (also featured in main text).

8 Additional analyses

8.1 Merged and unmerged

Figures 5, 6 and 7 display the marginal means (MMs) from the unmerged qualitative, unmerged recoded quantitative, and merged, analyses respectively. Figure 5 repeats the qualitative MMs reported in the main text. Figure 6 displays the MMs of the quantitative conjoint when the payoff delays are recoded into the qualitative categories used in the qualitative conjoint, following the pre-registered schema detailed in the

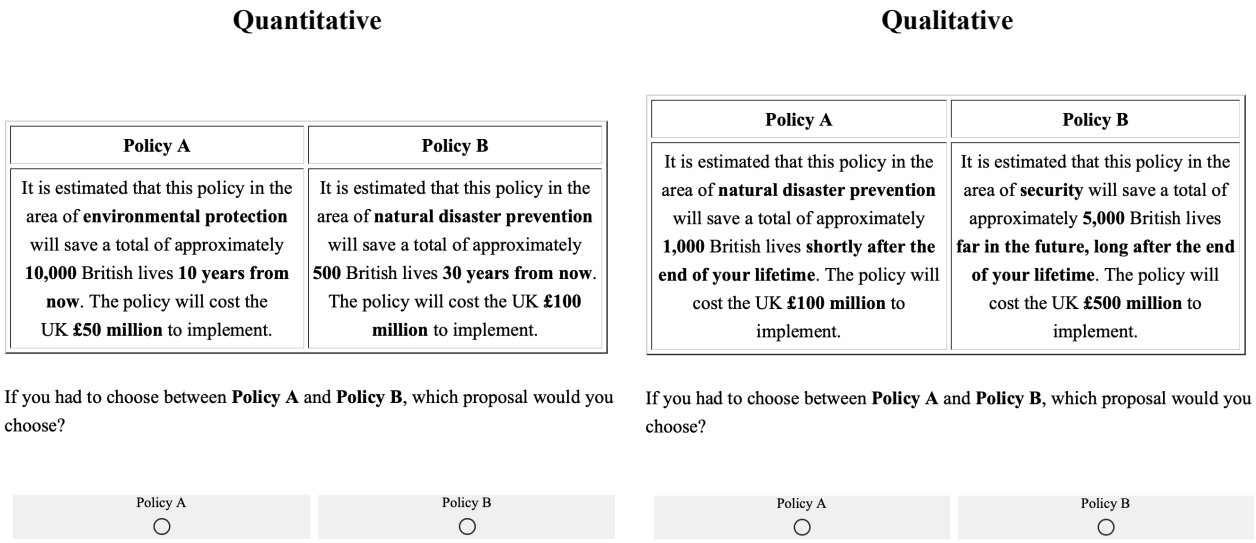

Figure 4: Example conjoint tasks.

main text. Figure 7 displays the MMs of the two merged. The results are very similar across all three versions.

8.2 AMCEs

In the main text, I opted to report marginal means for their correspondence to my hypotheses. Figures 8 and 9 instead display the AMCEs from the qualitative (and merged) and quantitative conjoints, respectively. AMCEs are displayed pooling across age groups.

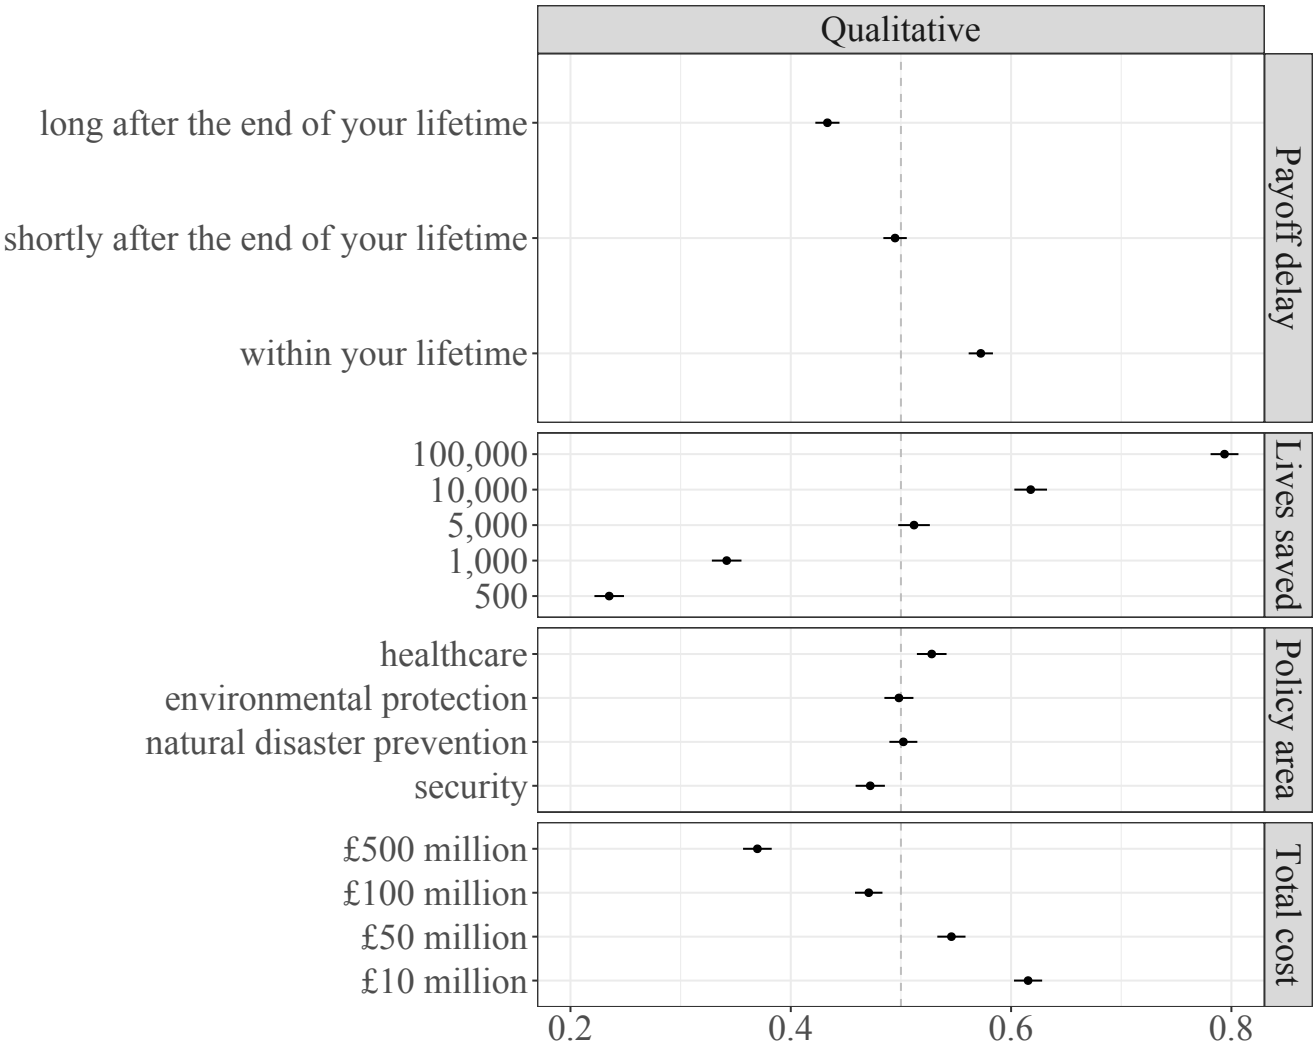

Figure 5: Marginal means from qualitative conjoint.

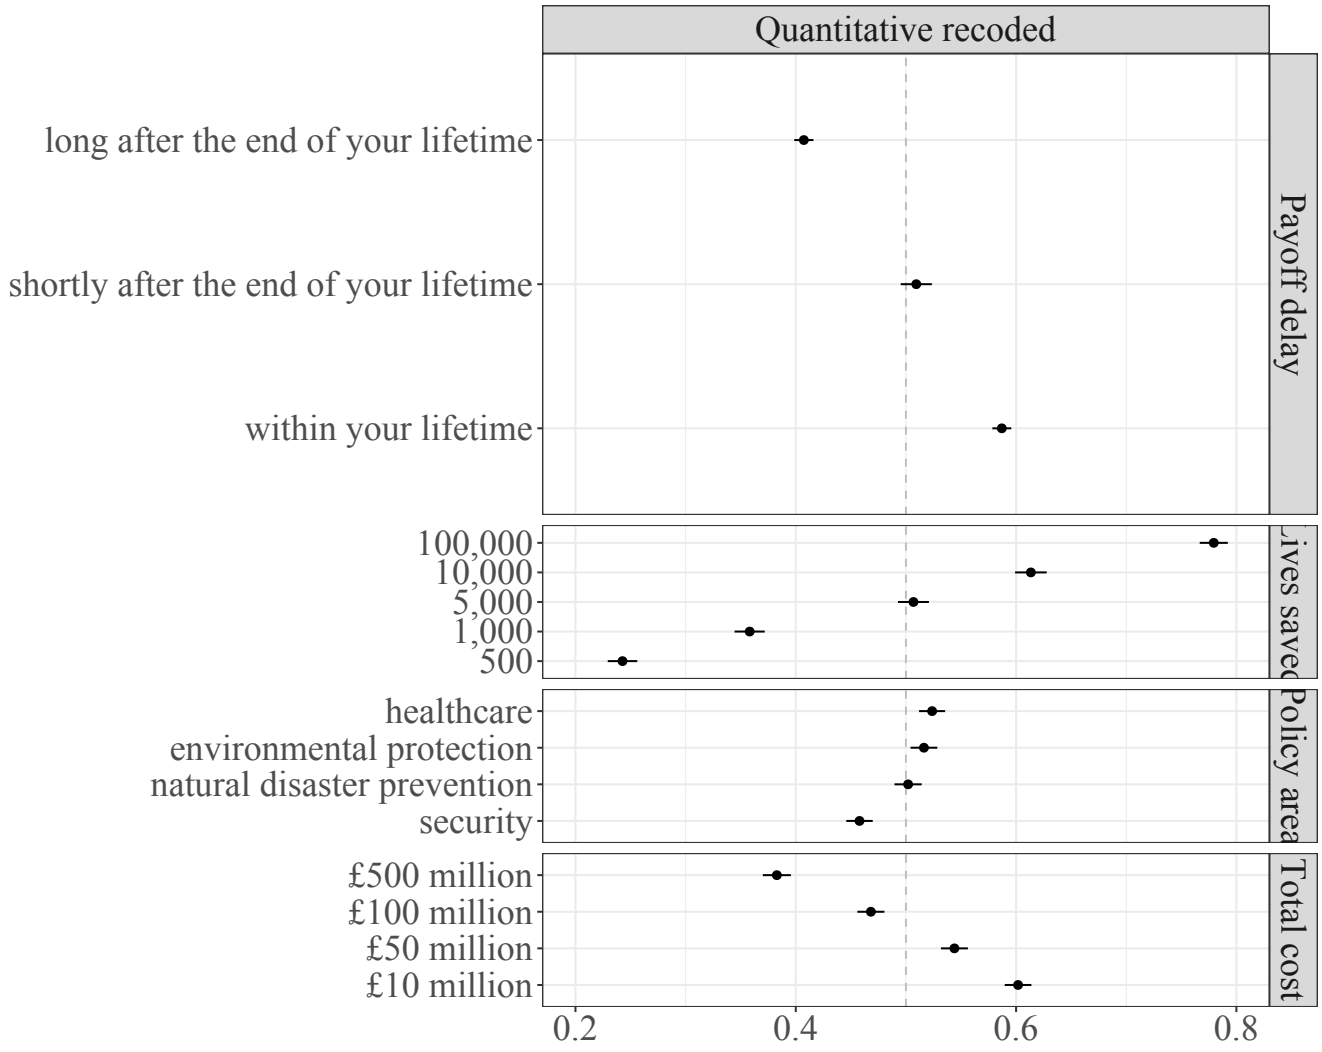

Figure 6: Marginal means from recoded quantitative conjoint.

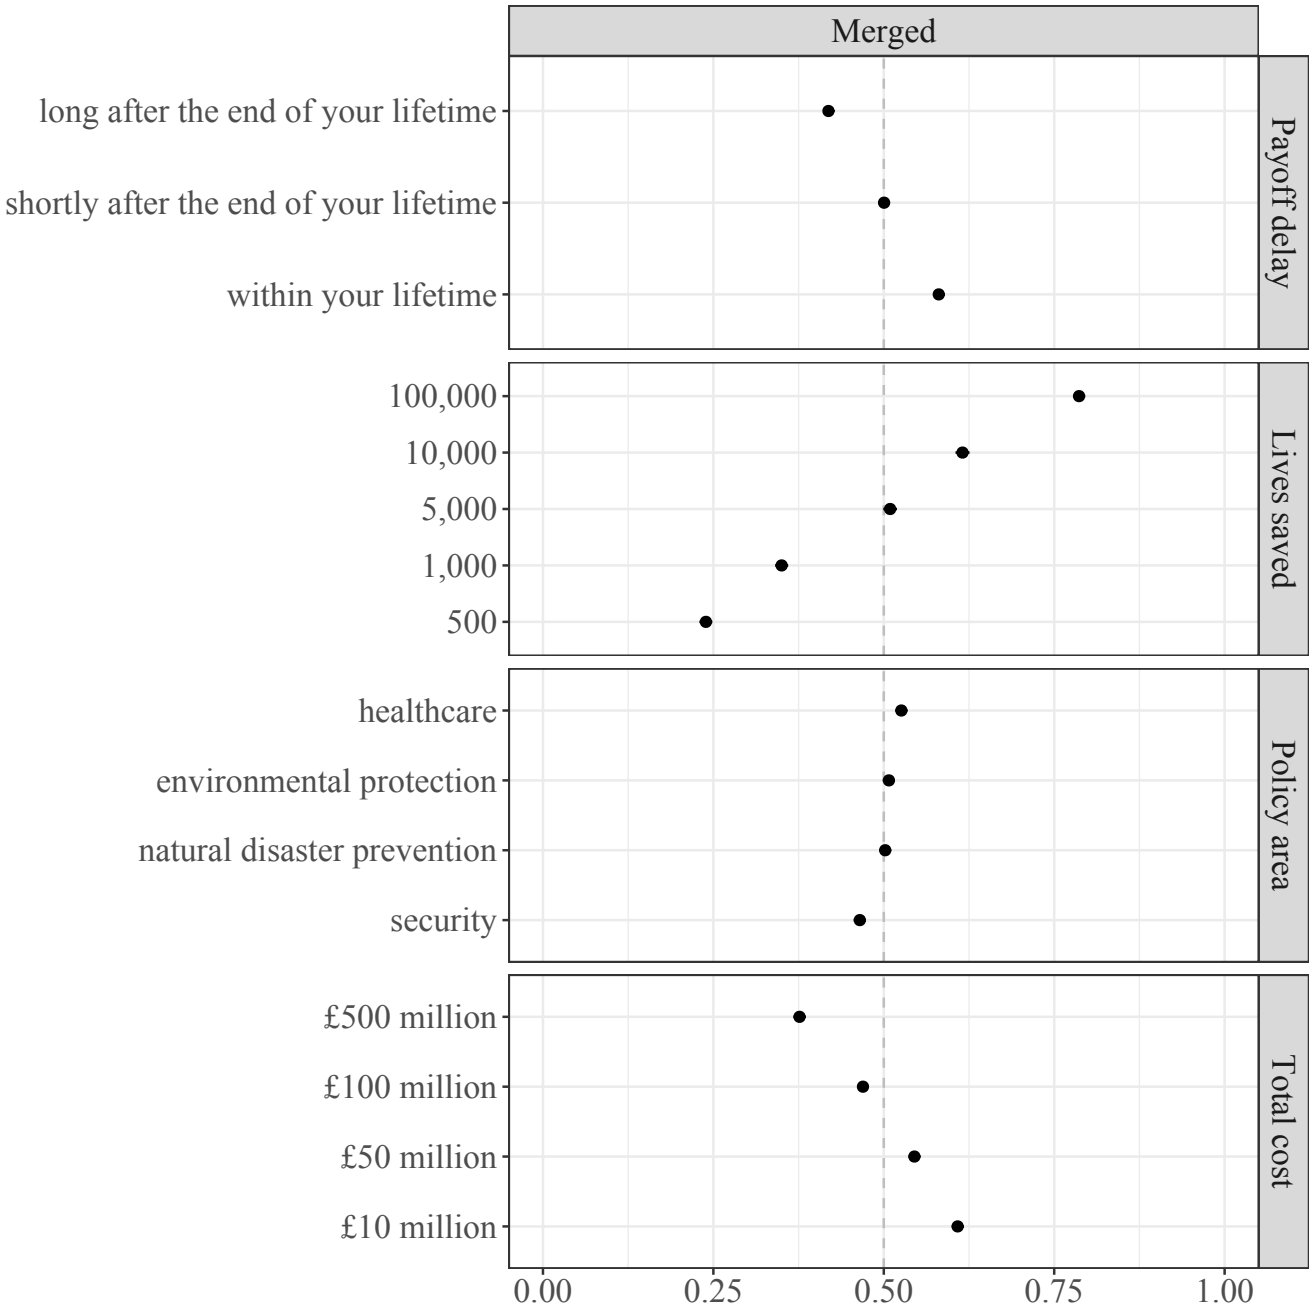

Figure 7: Marginal means from qualitative conjoint and recoded quantitative conjoint merged.

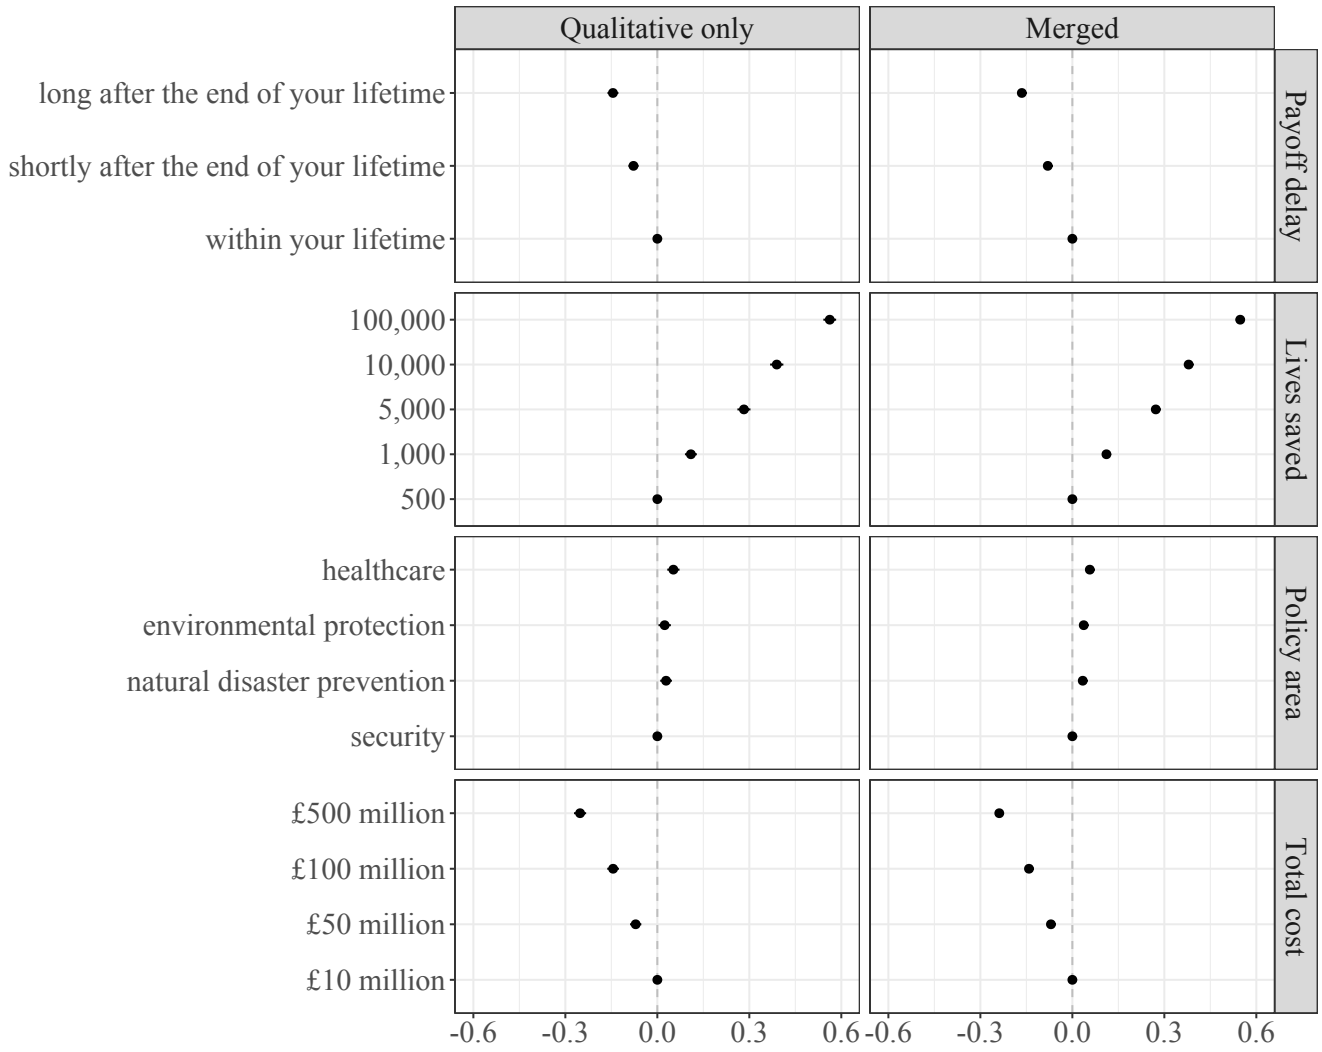

Figure 8: AMCEs from qualitative and merged conjoints (merged version contains data from qualitative and recoded quantitative versions.

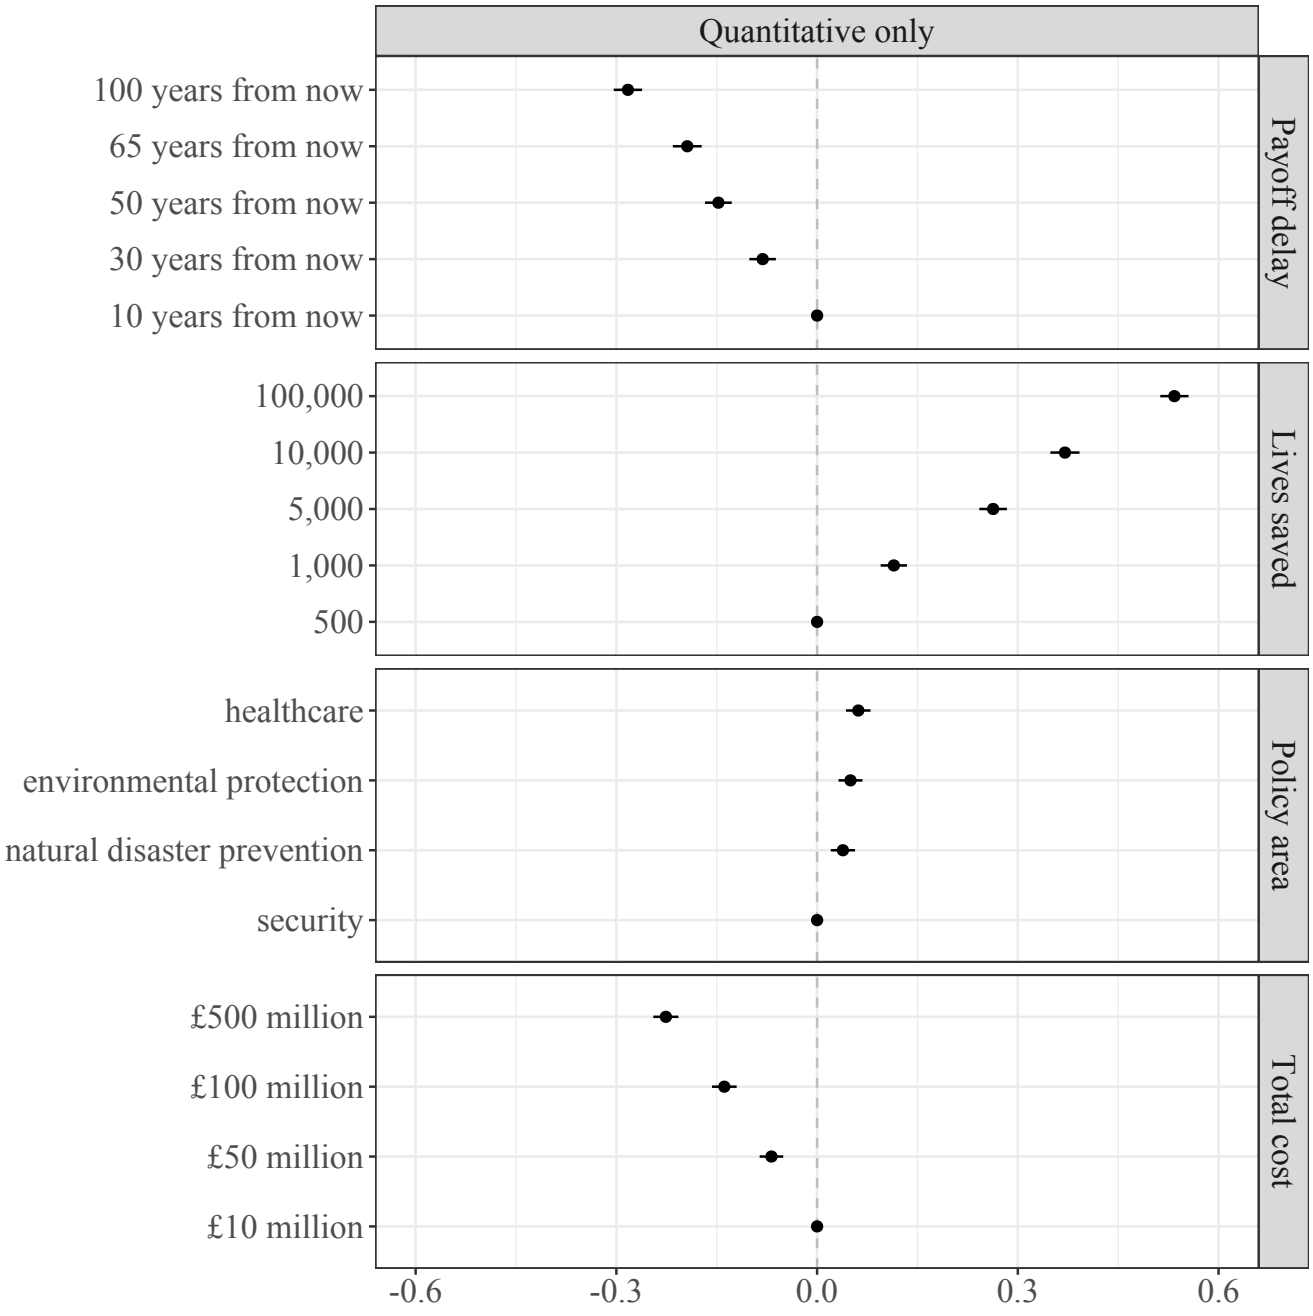

Figure 9: AMCES from quantitative conjoint, pooled across age groups.

8.3 Bayesian monotonic effects

Figure 10 plots the results of Bayesian monotonic effects models. This approach estimates the effect of an ordinal predictor broken down into a total, average linear effect, and the proportion of that effect accounted for by each increase in the level of the ordinal predictor. For example, the average slope from the leftmost point (10 years) to the rightmost point (100 years) in the bottom-left panel of 10 captures the overall average effect of policy payoff delay in the 20-34 sample, while the change in slope between each two successive points captures how much of that total effect each increase in the number of years accounts for. This approach has the advantage that H_{1b} maps directly onto the distribution of monotonic effects: to confirm H_{1b}, we would need to observe an especially steep slope between a payoff delay that falls within a respondent's likely lifespan and one that falls beyond that lifespan. However, in the bottom row showing the results from the quantitative conjoint by age group, none of the groups exhibit this pattern. providing further evidence that the lifespan does not majorly structure temporal discounting. To reinforce this suggestion, the top two panels plot the Bayesian monotonic effects of the payoff delay feature in the qualitative conjoint, and in the qualitative and quantitative merged version. H_{1b} would predict that the slope between the first two points would be steeper than that between the second and third points. However, the slopes are almost identical in both cases. These patterns are, again, inconsistent with H_{1b} and suggest that the lifespan does not structure temporal discounting.

8.4 Parental status

As an additional exploratory hypothesis, I consider whether parents and non-parents differ in their policy discounting behaviours. Parents may be more willing to accept intertemporal policy tradeoffs that pay off

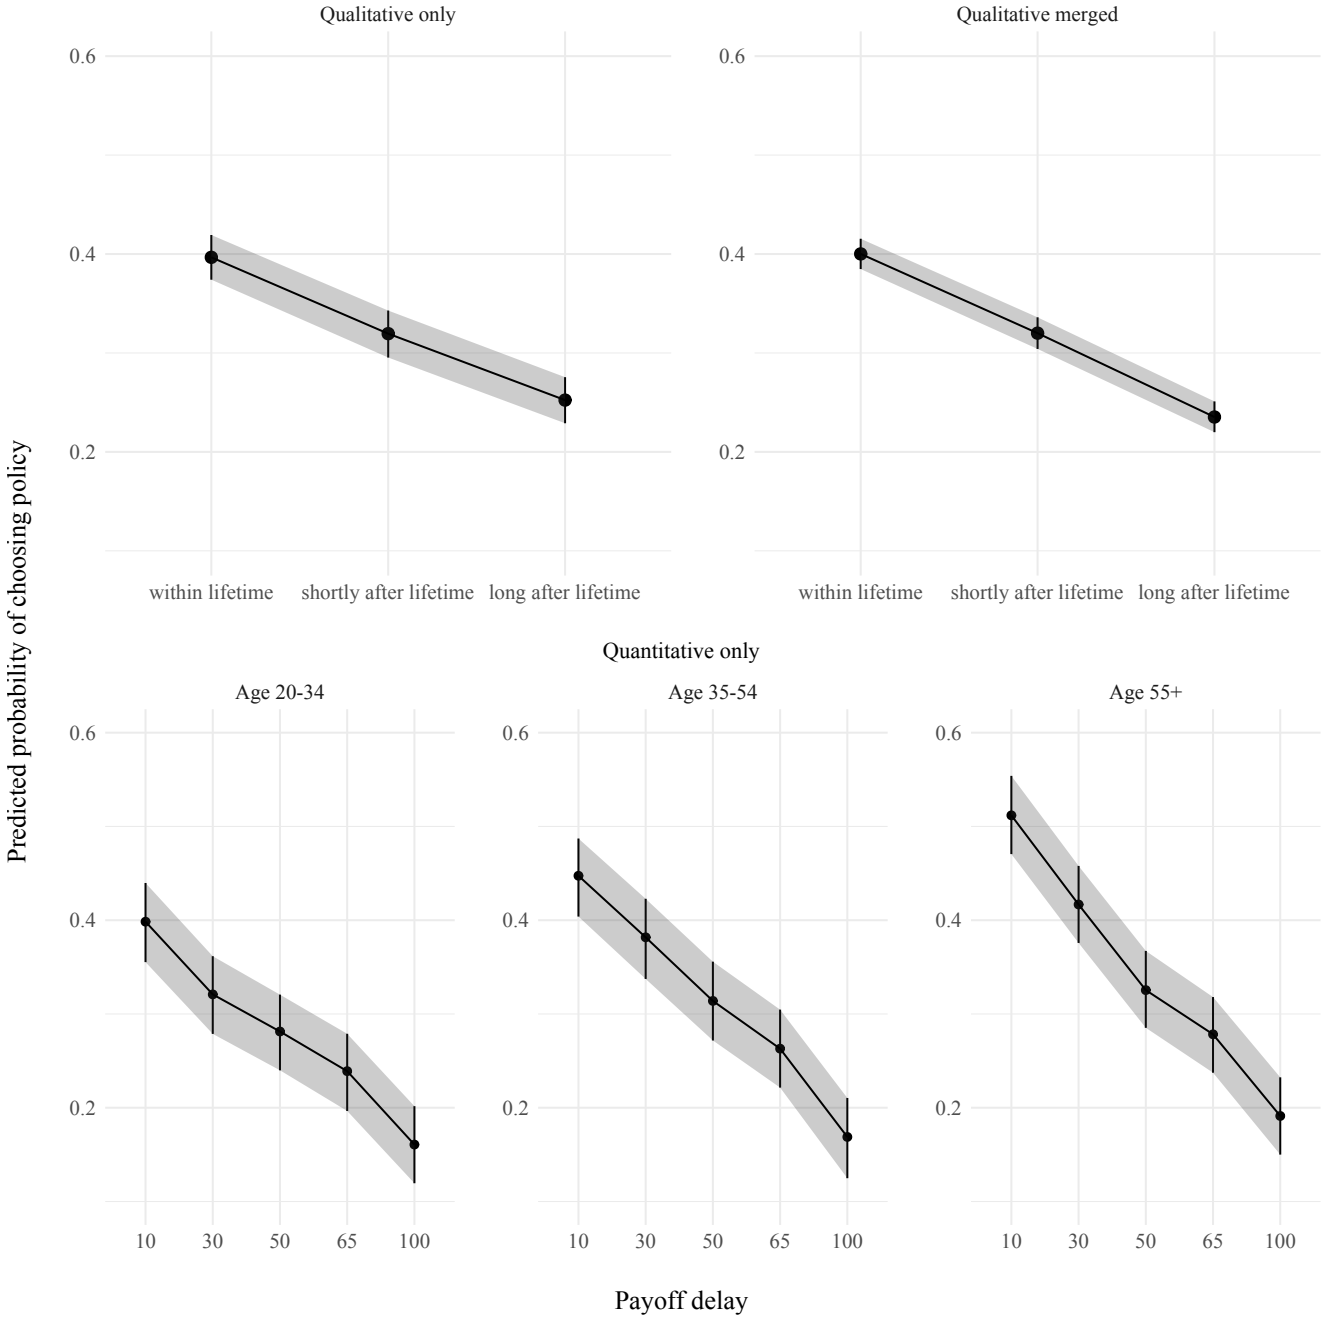

Figure 10: Bayesian monotonic effects. Points represent the predicted probability of choosing a policy with a given feature-level. Bars and shaded region represent 95% Bayesian credible intervals. Slope of line varies between points representing the proportion of the overall effect accounted for by each one-unit change in payoff delay. Top left panel displays monotonic effect of qualitative payoff delay. Top right panel displays monotonic effect of qualitative payoff delay merged with recoded quantitative delay. Bottom three panels show monotonic effect of quantitative payoff delay in each age group.

after the end of their *own*, but not their *children's* lifetimes, whereas non-parents may not have such a strong affective attachment to the next generation who may live to see the benefits that they themselves will not see.

However, Figure 11 shows that preferences for policies paying off within, after, and long after the lifespan are very consistent between parents and non-parents. Non-parents have around a 0.57 probability of choosing a policy that pays off within their lifetime, around a 0.49 probability of choosing one that pays off shortly after, and around a 0.41 probability of choosing one that pays off long after. For parents, those probabilities are 0.59, 0.50, and 0.43. Again, both parents and non-parents evaluate policies paying off shortly after the end of the lifespan significantly worse than those paying off within the lifespan, and evaluate those paying off long after the end of the lifespan significantly worse still.

8.5 Subjective life expectancy

It is possible that the patterns observed in the analyses are partly due to a mismatch between the life expectancies assumed in my pre-registered schema and respondents' own subjective life expectancies. For example, if respondents think they will live longer than they are objectively likely to live, then they might see some policy payoffs as likely to materialise within their own lifetimes, or closer to the end of their life, than assumed in my analyses. Therefore, in the follow-up survey, respondents were asked to report their subjective life expectancy. Specifically, I asked respondents:

- If you had to guess, approximately how long would you expect to continue to live?

The response options matched the possible quantitative payoff delays: Up to 10 more years, Up to 30 more years, Up to 50 more years, Up to 65 more years, Up to 100 more years, Longer than 100 more

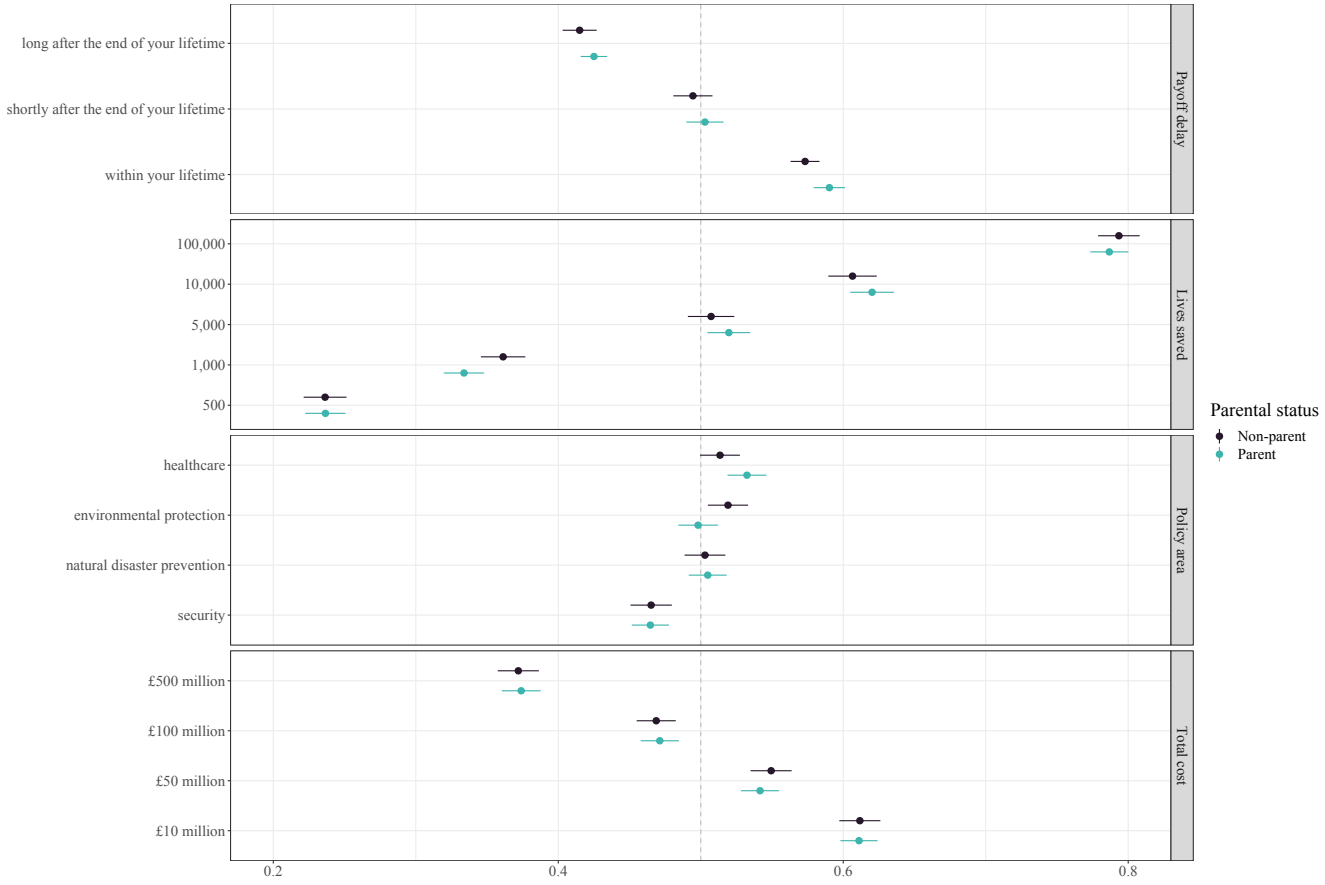

Figure 11: Marginal means among parents and non-parents. Points represent the probability that a policy is chosen when it has a given feature-level, amongst parents and non-parents. Horizontal bars represent 95% confidence intervals.

years, I would rather not say.

These responses were then matched to the quantitative payoff delay features to give, for each respondent, a subjective indication of whether that policy was expected to pay off within their lifetime. For example, if a respondent who answered 'up to 30 more years' saw a policy expected to pay off in 10 years, this observation is coded as 'within your lifetime'. The cutoff between 'shortly after' and 'long after the end of your lifetime' for these combinations is set at 20 years. (Varying this cutoff to reasonable alternative levels barely changes the results.) For example, if a respondent who expected to live 'up to 50 more years' saw a policy expected to pay off in 65 years (implying it will pay off approximately 15 years after the end of their life), this observation was coded as 'shortly after'.

Figure 12 displays the MMs of the quantitative payoff delays recoded, following this procedure, into qualitative categories in combination with these subjective life expectations. The results are substantively the same as those found in the main, pre-registered models.

8.6 Effects in first and last task

To check for the possibility of respondent fatigue over the repeated conjoint task affecting the results of the experiment, Figure 13 compares the marginal mean of each conjoint level from only responses to the first conjoint task (left column) and the last conjoint task (right column). The results reveal that there were minimal differences in preferences expressed in the first and last conjoint task. Most importantly, the pattern of preferences along the lifespan (top panels) remains almost exactly the same.

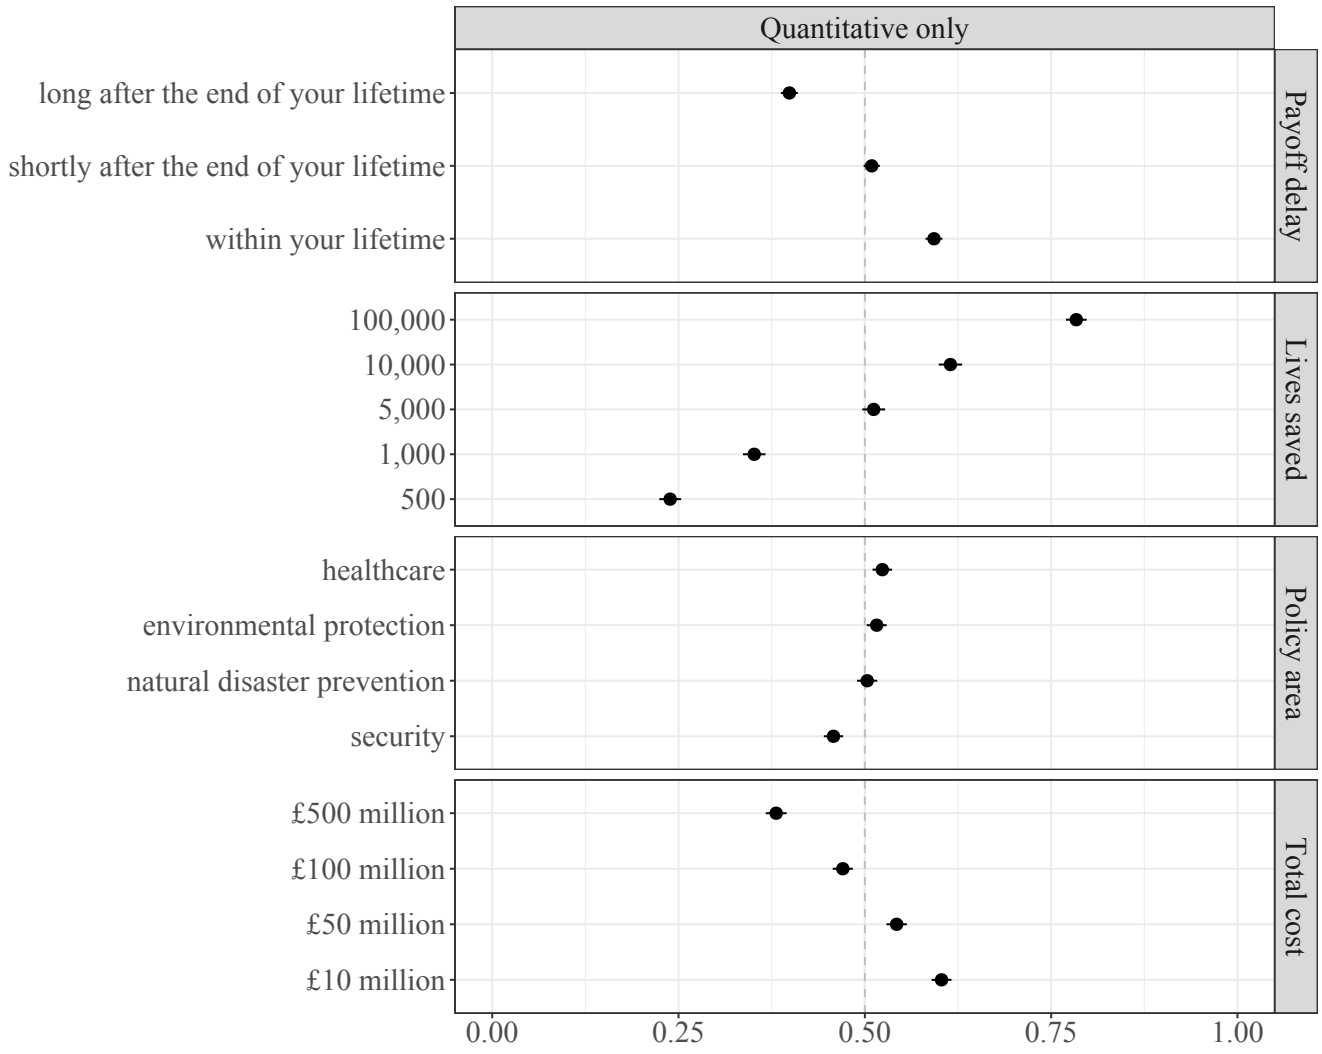

Figure 12: Marginal means for quantitative payoff delays recoded into qualitative categories, using subjective life expectancy and a shortly after cutoff of 20 years.

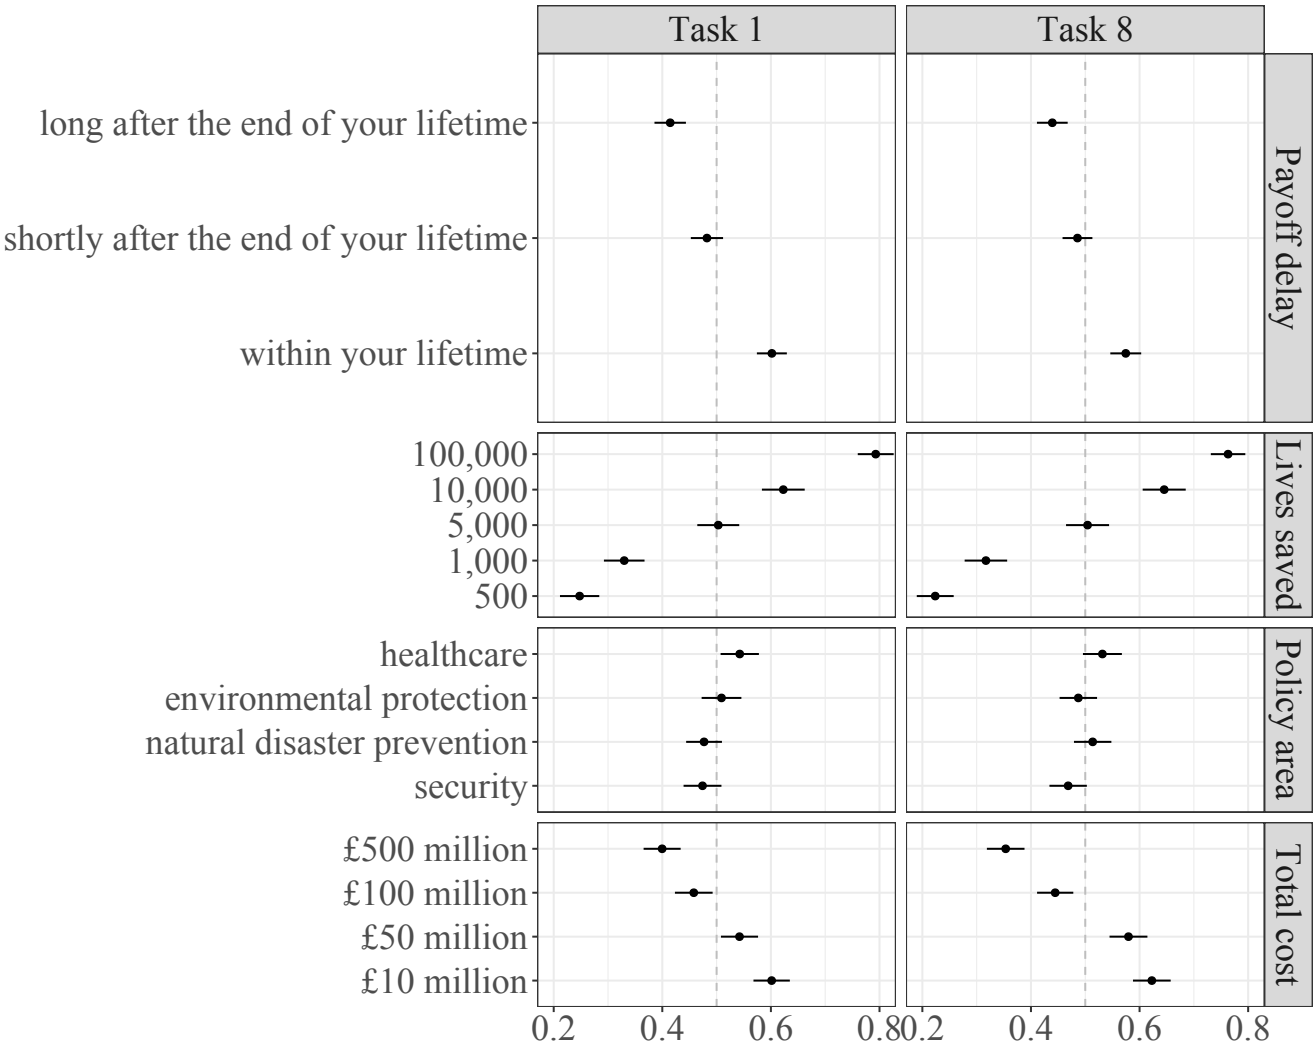

Figure 13: Marginal means for qualitative conjoint only first (left column) and last (right column) conjoint tasks.