**DEMS WORKING PAPER SERIES**

# Moral Preferences over Health-Wealth Trade-offs

**Antonio Filippin, Marco Mantovani**

**No. 531 – March 2024**

**Department of Economics, Management and Statistics**
**University of Milano – Bicocca**
**Piazza Ateneo Nuovo 1 – 2016 Milan, Italy**
**http://dems.unimib.it/**

# Moral Preferences over Health-Wealth Trade-offs

Antonio Filippin,[1*] Marco Mantovani,[2*]

[1]University of Milan, Department of Economics, Via Conservatorio 7, 20122 Milan, Italy
Institute for the Study of Labor (IZA), Schaumburg-Lippe-Str. 5-9, 53113 Bonn, Germany

[2]University of Milan-Bicocca, Department of Economics, Management and Statistics,
Piazza dell'Ateneo Nuovo, 1, 20126 Milan, Italy

[*] Address correspondence to antonio.filippin@unimi.it, marco.mantovani@unimib.it.

**Abstract**

Using a choice experiment we analyze moral preferences over fatalities and jobs losses due to the pandemic in Italy, the UK and the US. A structural estimation displays, surprisingly, aversion to diversification among these two bads. We also find that about 95% of the weight in the participants' utility function goes to health, and that respondents' stable traits (such as political orientation or risk aversion) influence attitudes more than their personal experiences with the consequences of the pandemic. Moreover, policy responses look misaligned with estimated preferences. Italy adopted more stringent containment measures, while Italian respondents display a relatively weaker pro-health attitude.

**Keywords:** Covid-19, Structural estimation, Health-wealth trade-off, Moral preferences.

**Acknowledgements**: The Authors would like to thank Paolo Bertoletti, Giovanna D'Adda, Francesco Guala, Paolo Vineis, and participants in the 7th BEEN workshop, ESA World Meeting 2022 (Boston), ESA European Meeting 2022 (Bologna), the WZB Workshop on Behavioral and Experimental Insights on COVID-19, as well as in the internal seminars at the University of Milan and the University of Milan Bicocca for the useful comments received.

## 1 Introduction

Work and related activities constantly expose people to health risks. In normal times, this dilemma is not perceived. During the Covid-19 pandemic, instead, it rose to the forefront of the political agenda. Urgent choices have been made between conflicting objectives, namely preventing the spread of the virus while at the same time limiting the economic and social impact of containment measures. Different approaches have emerged internationally. Some countries have opted for immediate, strict and prolonged containment strategies. Others have shown to be more concerned about impairing the economic activity and the other aspects of the citizens' daily life. Whatever the choice, this dilemma was solved "behind the scenes," implementing policies that are not necessarily aligned with the citizens' preferences. Ignoring such preferences can also prevent politicians from effectively communicating the reasons for their actions.

This paper studies empirically the citizens' moral preferences with respect to such trade-off between public health and economic activity. The exercise is based on an online questionnaire combined with theoretical principles and statistical methods traditionally used to estimate economic preferences. The main goal is to improve our understanding of how individuals deal with such a moral dilemma, and to highlight differences both across different population groups (e.g. by age, political orientation, or type of occupation) and at the country level. Furthermore, we present a small-scale test of a tool for estimating such preferences, providing a first assessment of their alignment with the governments' actions.

We administer an online questionnaire containing a battery of choices between different combinations of health and economic outcomes. Each choice presents two scenarios, each identified by a total number of Covid-19 fatalities and of jobs lost. We construct a set of 125 binary choices that vary in *a)* how severe the scenarios are along the economic and/or health dimension, *b)* the cost in terms of jobs lost for saving one life embedded in the choice. Individuals' questionnaires select 25 of these choices for each participant. The selection ensures, through a pseudo-random procedure, that each questionnaire gathers information balanced for all types of scenarios and exchange rates.

We explicitly frame choices as scenarios in the participant's country. As such, the individual is asked to solve the trade-off for the society as a whole, rather than for herself. In other words, victims are unidentifiable (Jenni and Loewenstein, 1997; Lee and Feeley, 2016; Small and Loewenstein, 2003). This approach distinguishes our exercise from the related one of investigating how people solves this trade-off at the individual level, often estimating the value of a statistical life from individual behavior or market risk premiums (Ashenfelter, 2006; Belle and Cantarelli, 2022; León and Miguel, 2017; Viscusi and Aldy, 2003).

The analysis proceeds in two parallel ways. First, we adopt linear probability models to estimate the determinants affecting the probability of choosing the scenario with fewer fatalities. Second, we use maximum likelihood estimation methods to represent respondents' preferences over health and economic dimensions. In this way, starting from the information on their local behavior contained in every answer, we retrieve a global representation of the respondents' preferences. This methodology is well-established and has been used for instance to represent preferences under risk (Andersen *et al.*, 2006; Filippin and Crosetto, 2016; Hey and Orme, 1994; Stott, 2006).

The study involves 2490 observations equally distributed across Italy, the United Kingdom and the United States, recruited through the online platform Prolific (Palan and Schitter, 2018). On top of their choices, we collect and relate to the estimated preferences socio-demographics variables (sex and age), economic attitudes (risk aversion and patience), political orientation, and personal experience with the effects of the pandemic.

We expect preferences to be concordant with each country policy response. Since Italy adopted the most stringent policies, and the US the least stringent ones, we expected the strongest pro-health attitude in Italy, and the weakest one in the USA. Furthermore, we expect older, more risk averse and participants with more left-leaning political preferences to have stronger pro-health attitudes. Finally, we expect participants that faced the adverse health contingencies due to the pandemic to have stronger pro-health attitudes, and the converse for participants more exposed to the adverse economic consequences due to the pandemic.

Overall, we find evidence in line with these expectations with respect to persistent individual traits: preferring outcomes with fewer fatalities and more jobs lost is associated with older age, left-leaning political preferences and risk aversion. Conversely, we find no evidence that contingent experiences with the pandemic – having tested positive or lost someone in one's close network, leaving in an area with a high death rate, or belonging to employment categories that were hit most severely by the pandemic from an economic perspective – affect the participants' preferences. Finally, we find that differences in the pro-health attitude across the three countries does not correspond to the strength the governments' responses throughout the pandemic. Relative to the UK, Italian respondents display a weaker pro-health attitude, and US respondents a stronger one. We further leverage data from different sources in order to unveil the mechanics behind this result. The stringency index at the level of US state (from the Oxford COVID-19 Government Response Tracker, Hale *et al.*, 2021) seems to exclude that containment policies have shaped preferences. Data from the bureaus of statistics, Medicare and Medicaid Services, and US Census suggest that cross-country differences in pro-health attitudes do not depend on different costs of the two bads.

Our project contributes to the recent literature that rigorously measures moral preferences in a real situation such as the Covid-19 pandemic. The global characterization of preferences in our study complements those that investigate a trade-off involving more dimensions related to the pandemic, but only locally (Carrieri *et al.*, 2021; Chorus *et al.*, 2020; Lesschaeve *et al.*, 2021; Loría-Rebolledo *et al.*, 2022; Manipis *et al.*, 2021; Oana *et al.*, 2021) and those that estimate the individual welfare costs or economic concerns related to containment measures (Andersson *et al.*, 2021; Codagnone *et al.*, 2020). To our knowledge, our project is the first to propose a direct estimation of moral preferences over the full space of reasonable outcomes. Our exercise belongs to family of contributions investigating moral decision making using hypothetical scenarios. This branch of the literature finds its roots in the seminal contributions on trolley problems (Foot, 1967) and taboo dilemmas (Tetlock, 2003). A similar approach characterizes the growing literature in health economics that elicits the willingness to pay for quality-adjusted life years (see the review by Spencer *et al.*, 2022, and references therein).

A growing literature, bridging economic and epidemiological models, estimates how health and economic outcomes respond to containment policies (Alvarez *et al.*, 2020; Eichenbaum *et al.*, 2021; Favero *et al.*, 2020; Jones *et al.*, 2021; Kaplan *et al.*, 2020). These exercises identify the set of available options to a society, representing outcomes that can or cannot be achieved. Our methodology covers the other side of the choice problem, i.e. the identification of preferences in the same space. In other words, we only focus on what individuals or groups like and do not like (in relative terms). Once taken together, these two sides promise to deliver the optimal combination of health and economic outcomes, meant as the best alternative among those possible given economic and epidemiological constraints.

## 2 Procedures and Data

The study was approved by the IRB of the University of Milan ('Comitato Etico', decision n. 128/21) and pre-registered at Open Science Foundation ($OSF.io/4vznh$). All methods were performed in accordance with the relevant guidelines and regulations, including the Declaration of Helsinki on research involving human subjects. The survey was administered online through Qualtrics on December 23, 2021, and involved samples collected at the same time in the UK, the US and Italy.

**Questionnaire.** The first part of the questionnaire presents the study, and informs the participants that the data collected are anonymous and used exclusively in aggregate form for academic research. Participants are informed that they are free to withdraw from the study at any time. Informed consent is then obtained from all participants.

The second part of the questionnaire contains the task aimed at eliciting participants' moral preferences between economic and health outcomes. The task consists in a battery of 25 binary choices between two hypothetical scenarios presented as 'bads.' In fact, each scenario is characterized by a number of fatalities (health outcome) and a number of jobs lost (economic outcome). Participants must indicate, for every choice, which of the two scenarios they deem relatively preferable. Figure 1 provides an example.

The 25 choices are extracted from a set of 125 alternatives, reported in Figure 2. The $y$-axis represents fatalities, while the $x$-axis represents jobs lost. Every segment in the figure corresponds to one binary choice. The coordinates of the two ends of each segment identify the two scenarios between which the respondent may be called to choose. The

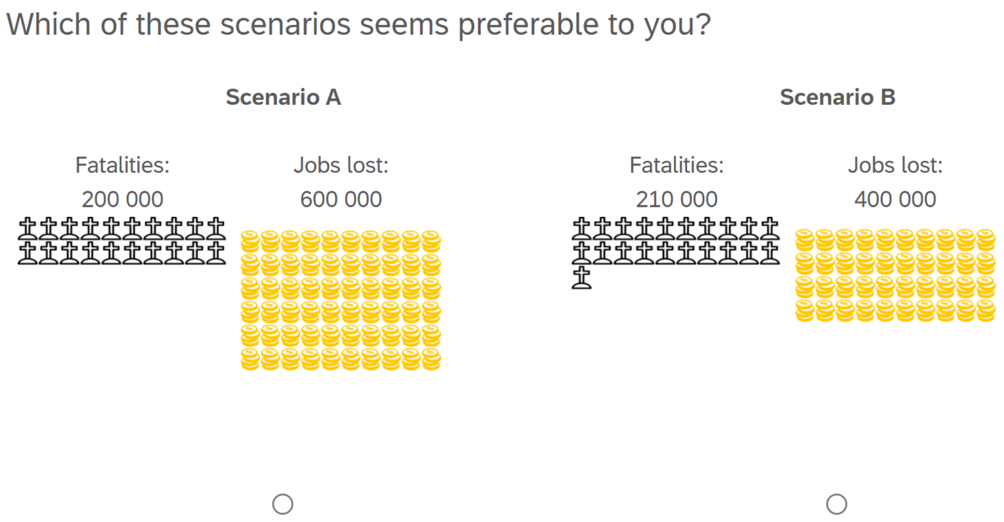


Figure 1: An example of binary choice

*Notes:* Each choice presents two scenarios (called A and B). Every scenario is characterized by a number of fatalities and of jobs lost. The position of each combination is randomized, so that the scenario with fewer fatalities appears as A or B with equal probability. The respondents must indicate which of the two scenarios is relatively preferable.

125 choices can be categorized along two dimensions.

The first is the number of fatalities and jobs lost in one of the two scenarios, i.e. the location of the south-east vertex of each edge in Figure 2. There are 25 locations, spanning a large set of different conditions around the values actually observed.[1]. Doing so we present situations that are severe in one dimension but relatively less in the other one, as well as situations that are severe in both dimensions, and any other combination.

The second is the slope of the segments in each edge, representing the trade-off between these two 'bads' when moving from one scenario to the other. In other words, the slope captures the implicit cost in terms of jobs lost to save one life embedded in each choice. A flat segment represents a trade-off in which a small decrease in fatalities occurs at the cost of a large increase in jobs lost, i.e. a very high cost of one life. Vice-versa, a steep segment implies a large decrease in fatalities at the cost of a relatively smaller increase in jobs lost, i.e. a low cost of one life. There are 5 possible costs: $\{0.25, 5, 10, 20, 100.\}$

The combination of the 25 locations with the 5 exchange rates gives the full set of 125 choices. Each subject is shown a subset of 25 choices extracted from the full set of alternatives in a pseudo-random manner. The pseudo-randomization ensures that, for each participant *i)* the slopes are equally represented (5 choices for every slope) *ii)* all areas in the space are equally represented (5 choices for each letter identifier in Figure

[1]The scenarios reported in Figure 2 are those actually administered in Italy and in the UK. Values are instead multiplied by five in the USA so that both outcomes in all countries are centered approximately around the same rate of incidence in the population

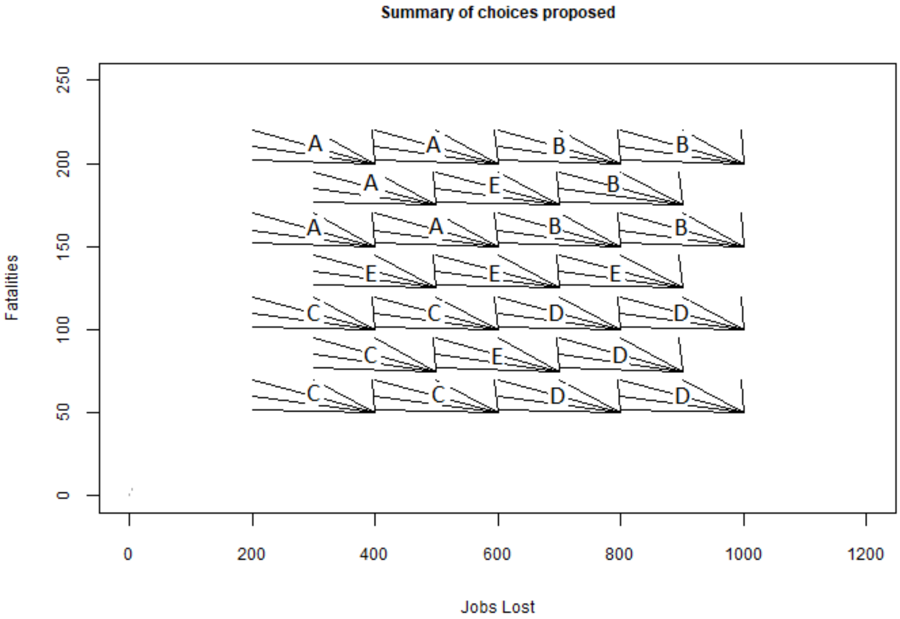


Figure 2: The set of 125 choices

*Notes:* The $y$-axis represents fatalities, while the $x$-axis represents jobs lost. Every segment represent one binary choice between the coordinates at its extremes. The values reported are those administered in the questionnaire in Italy and in the UK, while values are multiplied by 5 in the US. The 125 choices are categorized along two dimensions. First, the number of fatalities and jobs lost in one of the two scenarios, i.e. the location of the south-east vertex of each edge. Second, the slope of each edge, which represents the trade-off between these two 'bads' when moving from one scenario to the other. A flat segment represents a trade-off in which a small decrease in fatalities occurs at the cost of a large increase in jobs lost, i.e. a very high cost of one life. Vice-versa, a steep segment implies a large decrease in fatalities at the cost of a small increase in jobs lost, i.e. a low cost of one life. Letters A-E represent the groups of choices used in the pseudo-randomization to ensure that each respondent receives alternatives evenly distributed in the plan.

2). Specifically, for each area represented by one of the letter identifiers, the procedure selects the five choices with the same slope and picks one at random to be included in the questionnaire. The procedure repeats the random selection for each slope within an area and for all areas. Finally, the procedure randomizes the order in which the 25 choices are presented as well as which of the two alternatives is labelled as Scenario A.

The last part of the questionnaire collects the socio-demographic variables of interest. We ask for employment condition on a seven-category scale. This allows in particular two isolate two categories: self-employed and atypical workers which are most vulnerable to the pandemic shock. We then ask for the participant's political orientation (on a 1-6 scale from left to right) and elicit classic economic attitudes, e.g. risk aversion (scale 1-10) and patience (scale 1-10), using the validated questions from the SOEP German panel (Dohmen *et al.*, 2011). Finally, we collect data on the participants' personal experience with Covid-19: whether they were infected and whether someone in their close network died of Covid-19. The questionnaire also includes a 'dummy' question that we use as attention check to eliminate the respondents who fill the questionnaire randomly for the sole purpose of receiving the remuneration.

**Participants.** Respondents are 2490 adult human participants who voluntarily registered on the Prolific platform (Palan and Schitter, 2018). We recruited an equally sized sample of 830 respondents in Italy, the UK, and the USA. Within each country we stratified by employment status, imposing 70% of the participants among those who registered as employed. This is done to avoid students being over-represented, and to have a sufficiently large sample directly exposed to adverse economic consequences of the pandemic. Within each cell, corresponding to one country and one employment condition, the sub-sample was balanced by sex. Subjects received a fixed amount ($1.20) for the successful completion of the questionnaire, which took on average 6 minutes and 16 seconds to complete.

**Power Analysis.** The sample size of 830 subjects per country is computed as follows. First, we anticipated we may have to discard about 5% of the observations either because respondents do not pass the attention check or because a completion time lower than 2 minutes suggests that responses have been generated by bots or randomly. The remaining 780 observations allow us to reach a statistical power of $1 - \beta = 0.8$ to detect a small effect size (Cohen's $d = 0.2$) in a two-tailed test for the difference between two subgroups within a country. In more detail, we assume:

- a linear probability model for a dummy variable that (without loss of generality) is equal to one when the option involving the lower number of fatalities is chosen;
- two equally sized subgroups within a country;
- a difference in the probability of 0.05 across the two subgroups;
- a standard deviation equal to 0.25 of the measured outcome;
- a confidence level of 95% ($\alpha = 0.05$).

This sample size translates into a statistical power of about 0.97 to detect a similar effect in the comparison between two countries (N=1560).

The power analysis based on a non-parametric rank-sum test of individual preferences delivers similar results. Let us assume that preferences are linear and mostly oriented to contain fatalities, as represented by the distribution of the weight attributed to economic outcomes following a $Beta(2, 18)$, i.e. with mean 0.1 and standard deviation 0.0645. A difference of 0.01 in the average value of the weight between the two groups, corresponding to a small effect size (Cohen's $d = 0.15$), can be detected with a power of 0.8 and a confidence level of 95% ($\alpha = 0.05$) having a numerosity of $N = 668$ per group.

**Additional data.** We merge the dataset resulting from the survey with data from different sources. First, the individual records from Prolific, including sex, age, employment status, nationality, country of origin. Second, the Oxford COVID-19 Government Response Tracker, and in particular the stringency index computed at the state level as well official Covid cases and deaths at the country and at the state level. Third, the geolocation of the respondent provided by Qualtrics that we use to assign respondents to state/regions within each country. Fourth, unemployment rates provided by ISTAT and by the National Bureau of Labor Statistics (overall average for year 2020). Finally, health out-of-pocket expenditure and health spending in the US provided by the US Census, and by the Medicare and Medicaid Services, respectively.

**Exclusion.** We exclude from the analysis observations that do not pass the attention check. Since a correct answer to the attention check can occur by pure chance with 20% probability, we also drop surveys that were completed in less than two minutes, assuming such fast completion times are inconsistent with providing meaningful answers. Overall we exclude 78 observations, or 3.13 percent of our sample. While we consider a good practice to follow these exclusion criteria, our results are unaffected when running the analysis on the full sample.

**Balancedness.** Samples across the three countries are not representative of the corresponding population, and are also not perfectly balanced. While samples do not significantly differ along risk aversion and experience with Covid deaths in one's network, other differences emerge, without a systematic pattern (see Table 1). We address this issue including these variables in our regressions, thereby controlling for such differences. Furthermore, we perform two robustness checks in Section 3.3. First, we include interactions between the country dummies and each of these variables to exclude that country effects are specific to subgroups where differences in the sample are more pronounced. Second, we adopt an inverse probability weighting procedure to account for a possibly different selection on observables of our sample across the three countries.

## 3 Results

A bird's eye view of the data emphasizes a strong pro-health attitude of the respondents. Overall, 67 percent of the choices favors the scenario with fewer fatalities and more jobs lost. This frequency tightly reacts to the cost of saving one life. When this cost is low (0.25 jobs for one life) 92% of the choices minimizes the number of fatalities. When the cost is instead high (100 jobs) only a minority of choices (40%) is pro-health. In between the two extremes, the fraction of pro-health choices decreases monotonically with higher costs. In the US, 73% percent of the choices favors the scenario with fewer fatalities and more jobs lost. This number decreases to 68% in the UK and 60% in Italy. The fraction of participants that always choose to minimize fatalities is 31% in the US, 34% in the UK and 15% in Italy.

Results are further analyzed using two approaches. The first displays the change in the probability of choosing the alternative with fewer fatalities associated to the main

Table 1: **Tests of balancedness across countries.**

| Variable | UK | ITA | US | Type of test | $p$-value | Pairwise comparisons |
|---|---|---|---|---|---|---|
| Age | 40.3 | 28.7 | 34.6 | Kruskall-Wallis | $< 0.001$ | $UK > US > ITA$ |
| Political orientation | 3.4 | 3.0 | 3.1 | Kruskall-Wallis | $< 0.001$ | $UK > ITA \approx USA$ |
| Risk tolerance | 4.6 | 4.8 | 4.6 | Kruskall-Wallis | 0.432 | $UK \approx ITA \approx US$ |
| Patience | 5.9 | 6.5 | 6.1 | Kruskall-Wallis | $< 0.001$ | $ITA > UK \approx US$ |
| Self-employed | 14.4 | 20.9 | 21.6 | Fisher | $< 0.001$ | $US \approx ITA > UK$ |
| Covid deaths | 8.9 | 11.9 | 13.1 | Fisher | 0.147 | $UK \approx ITA \approx US$ |
| Tested positive | 66.6 | 56.8 | 67.4 | Fisher | $< 0.001$ | $UK \approx US > ITA$ |

*Notes*: The table reports the tests of balancedness of groups across countries for age and the variables measured in our questionnaire. *Political orientation* is ordered from left to right-wing, so that $>$ indicates a distribution that is shifted toward a more right-wing orientation. The non-parametric Kruskall-Wallis test is adopted for semi-continuous variables (*Age, Political orientation, Risk tolerance, Patience*), while the non-parametric Fisher test is used for binary variables (*Self-Employed, Covid deaths, Tested positive*). The last column summarizes the pairwise comparisons (Mann-Whitney test for semi-continuous variables, Fisher test for binary variables), with significant differences (at 5%) reported as $>$. $\approx$ is instead used for differences that are not significant. Significance levels are corrected for multiple hypotheses testing (Bonferroni).

explanatory variables. The second reports a structural estimation of respondents' preferences over economic and health outcomes obtained with maximum likelihood methods.

### 3.1 Probability of choosing fewer fatalities

The analysis of the choices is carried out by choosing as dependent variable the probability of opting for the minimization of fatalities (without loss of generality). For ease of interpretation we use a linear probability model, but the results obtained hold unchanged adopting a Probit or Logit specification.[2] The estimated coefficients represent therefore the change in the average probability to minimize fatalities explained by the corresponding explanatory variable.

We estimate how such a probability differs first by country, and then along several dimensions (sex, age, political orientation, risk aversion, patience, employment condition, Covid contagion, Covid death rate at local level), always controlling for the exchange rate between fatalities and jobs lost implicit in each choice. Standard errors are clustered at the individual level and significance is always adjusted for multiple hypotheses testing using the Bonferroni correction. All the results are robust to the coding of the variables (dummy, categorical, continuous, when applicable) and to the inclusion of different sets of controls. Table 2 reports the full set of results, first with the explanatory variables separately (Column 1-7) and then together (Column 8). Figure 3 also plots the estimated coefficients of the linear probability model derived from our richest specifica-

[2]Country marginal effects for these specifications are reported in Table 3. The full results are available upon request.

tion (Column 8 in Table 2).

Table 2: Linear probability model

| | Dependent variable: Choice of the alternative involving fewer fatalities | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| Cost of one life | -0.04183*** | -0.04182*** | -0.04179*** | -0.04182*** | -0.04180*** | -0.04183*** | -0.04183*** | -0.04177*** |
| | (0.000761) | (0.000761) | (0.000761) | (0.000761) | (0.000761) | (0.000761) | (0.000761) | (0.000761) |
| Italy | | -0.0820*** | | | | | | -0.0831*** |
| | | (0.0140) | | | | | | (0.0159) |
| US | | 0.0451** | | | | | | 0.0472** |
| | | (0.0139) | | | | | | (0.0164) |
| Female | | | 0.0118 | | | | | -0.00304 |
| | | | (0.0116) | | | | | (0.0114) |
| Age $\geq$ 50 | | | 0.0457* | | | | | 0.0411* |
| | | | (0.0153) | | | | | (0.0155) |
| Left-wing | | | | 0.0856*** | | | | 0.0965*** |
| | | | | (0.0130) | | | | (0.0131) |
| Risk averse | | | | | 0.0590*** | | | 0.0502*** |
| | | | | | (0.0121) | | | (0.0119) |
| Patient | | | | | 0.00869 | | | 0.0137 |
| | | | | | (0.0119) | | | (0.0116) |
| Self-employed | | | | | | 0.0351 | | 0.0370 |
| | | | | | | (0.0152) | | (0.0149) |
| Covid deaths | | | | | | | 0.0202 | 0.0213 |
| | | | | | | | (0.0195) | (0.0189) |
| Tested positive | | | | | | | 0.0178 | 0.0110 |
| | | | | | | | (0.0123) | (0.0121) |
| High death rate | | | | | | | -0.00101 | -0.00753 |
| | | | | | | | (0.0122) | (0.0141) |
| Constant | 0.784*** | 0.796*** | 0.771*** | 0.724*** | 0.742*** | 0.778*** | 0.658*** | 0.669*** |
| | (0.00546) | (0.00954) | (0.00852) | (0.0111) | (0.0121) | (0.00609) | (0.0104) | (0.0196) |
| *N* | 60270 | 60270 | 60195 | 60270 | 60245 | 60270 | 60270 | 60170 |

*Notes*: Coefficients represent the change in the probability of choosing the alternative with fewer fatalities as compared to the Constant associated to each independent variable. Independent variables: *Cost of one life* is the number of jobs lost necessary to reduce fatalities by one unit implicit in each choice. All the other variables are dummmies taking a value equal to one for a respondent that *(i)* lives in *Italy* or *USA* (UK is omitted variable captured by the Constant); *(ii)* is *Female*, with an *age* $\geq$ 50, *Left-wing, Risk averse, Patient*, *Self-employed*; *(iii) Covid death* means that the respondent reports deaths due to Covid in his/her close network; *(iv) Tested positive* means that the respondent reports that himself or someone in her close network tested positive to Covid-19; *(v) High death rate* captures respondents living in a region/state with a rate of deaths due to Covid over 100,000 inhabitants higher than the median in our sample. Standard errors are reported in parentheses. Significance levels corrected for multiple hypotheses testing (Bonferroni): $^{*} = p-$value$< 0.1$; $^{**} = p-$value$< 0.5$; $^{***} = p-$value$< 0.01$.

Table 2 confirms the sizeable differences emerging across country of residence. The probability that participants in Italy choose the option involving lower fatalities is 8.3 percentage points lower than in the UK (the omitted variable corresponding to the vertical bar in Figure 3 ) and 13 percentage points lower than in the US. All pairwise differences between countries are significant (*t*-test, ITA vs UK: $p$-value $< 0.001$, USA vs UK: $p$-value $= 0.047$, ITA vs US: $p$-value $< 0.001$). In Section 3.3 we provide a robustness analysis of these country differences, further discussing the role played by sample unbalancedness as well as by possible confounding factors.

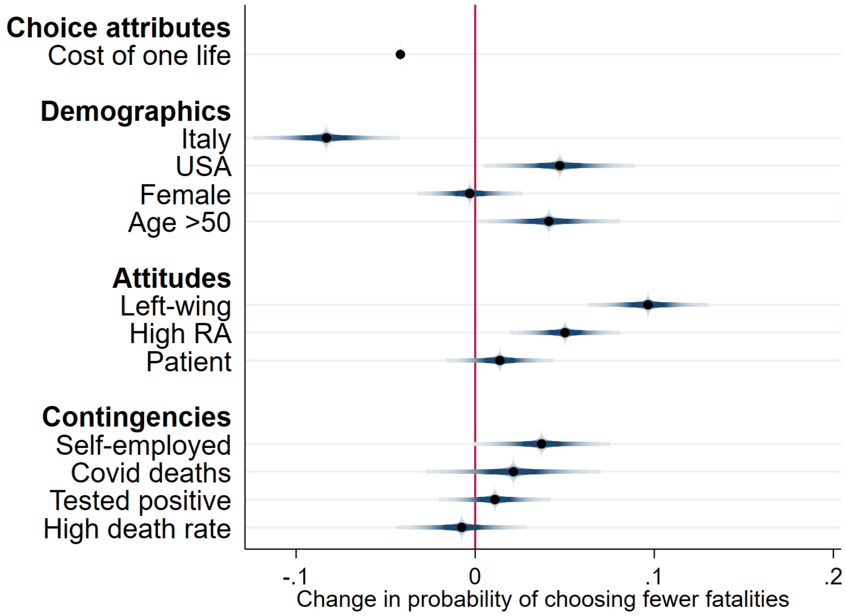


Figure 3: Coefficients for the change in the probability of making pro-health choices

*Notes:* The Figure plots the estimated coefficients of the linear probability model derived from our richest specification (Column 8 in Table 2). For each explanatory variable the Figure reports the change in the probability to opt for the alternative involving fewer fatalities. The vertical bar is the reference value corresponding to the Constant and should be interpreted as the value of the omitted category (e.g. UK for the countries, Male, Age $< 50$, etc.).

In general, what emerges from the data is that the variables that shape respondents' attitude toward this moral dilemma are more their stable traits (demographics and attitudes) rather than contingencies experienced during the pandemic. The coefficient of age approaches traditional significance thresholds ($p$-value $= 0.094$), with a stronger pro-health orientation characterizing older participants, while there is no difference between men and women ($p$-value $= 1.00$).[3] More risk-averse respondents are more inclined toward minimizing fatalities ($p$-value $< 0.001$). As one would expect, the same is true for left-wing participants ($p$-value $< 0.001$). In contrast, patience does not display a significant correlation (p-value = 1.00)

Contingent variables display null effects across the board. We expected that personal experiences with the pandemic significantly shaped the response to the moral dilemma. In particular, we expected a direct experience with Covid-19, such as testing positive or having a fatal loss in one's close network, to induce a more pro-health attitude. One may expect a similar effect also when living in regions/states that were hit more harshly by the pandemic, while a stronger attention to jobs lost was expected within categories that suffered stronger economic consequences, such as self-employed workers. The evidence is not in line with these expectations, as none of these variables shows a coefficient that is

[3]Note that $p$-values equal to 1.00 may mechanically stem from the Bonferroni correction for multiple hypotheses testing

significantly different from zero (Tested positive: $p$-value = 1.00; Covid deaths: $p$-value = 1.00; High death rate: $p$-value = 1.00; Self-employed: $p$-value = 0.157).

When we allow the explanatory variables to show a different impact by country using interaction terms two findings are worth mentioning.[4] First, as shown in the right-hand panel Figure 4, political orientation has a different impact in the countries under investigation, with a larger polarization of preferences along political lines in the US. The difference between left and right-wing respondents is of 15 percentage points in the US, significantly higher than in the UK (7 percentage points) and in Italy (6 percentage points), where the difference is instead not significant (p-value = 0.141). Second, as shown in the left-hand panel Figure 4. participants in different countries behave in a similar way when the cost of saving one life is low. Differences appear and increase monotonically as this cost increases. When saving one more life implies losing 100 jobs respondents in Italy have a probability of minimizing fatalities that is 20 (11) percentage points lower than respondents in the US (UK). The effect of an increase in the cost of saving one life is significantly larger in Italy than in the the other two countries at all possible levels of this cost, while the difference between the UK and the US in this effect is significant only for costs larger than five jobs for one life.

### 3.2 Structural estimation of preferences

Starting from all the choices we estimate with maximum likelihood the shape of the utility function for a representative respondent. To retrieve the underlying moral preferences of our groups of respondents we assume a flexible functional form of their utility function:

$$u(F,J) = \left[\alpha(\overline{J} - J)^{\rho} + (1-\alpha)(\overline{F} - F)^{\rho}\right]^{\frac{1}{\rho}}, \tag{1}$$

where $F$ and $J$ are the level of fatalities and jobs lost, respectively. $\overline{F} > max\{F\}$ and $\overline{J} > max\{J\}$ are used to transform this particular case into a classic utility function over 'goods,' since Equation 1 requires positive arguments. Note, however, that as $F$ and $J$ increase, the utility decreases. Hence, both outcomes are considered 'bads' that the subjects want to minimize.

Preferences are shaped by two parameters that represent intuitively a weight on jobs lost Vs. fatalities (the *distribution* parameter $\alpha$) and the inclination to diversify the outcomes along the two dimensions analyzed (the *curvature* parameter $\rho$). For a given value of $\alpha$ and $\rho$, Equation 1 identifies the indifference curve corresponding to any level of $u(F,J)$.[5]

The subjective trade-off between fatalities and jobs lost is captured by the Marginal Rate of Substitution (MRS) for any combination of the two outcomes:

[4] We return on the whole set of interactions in Section 3.3.

[5] The functional form in Equation 1 encompasses as special cases the most common examples of utility functions according to the values taken by $\alpha$ and $\rho$. Appendix 4 contains a detailed explanation of these functional forms as well as of some technicalities of the estimation procedure.

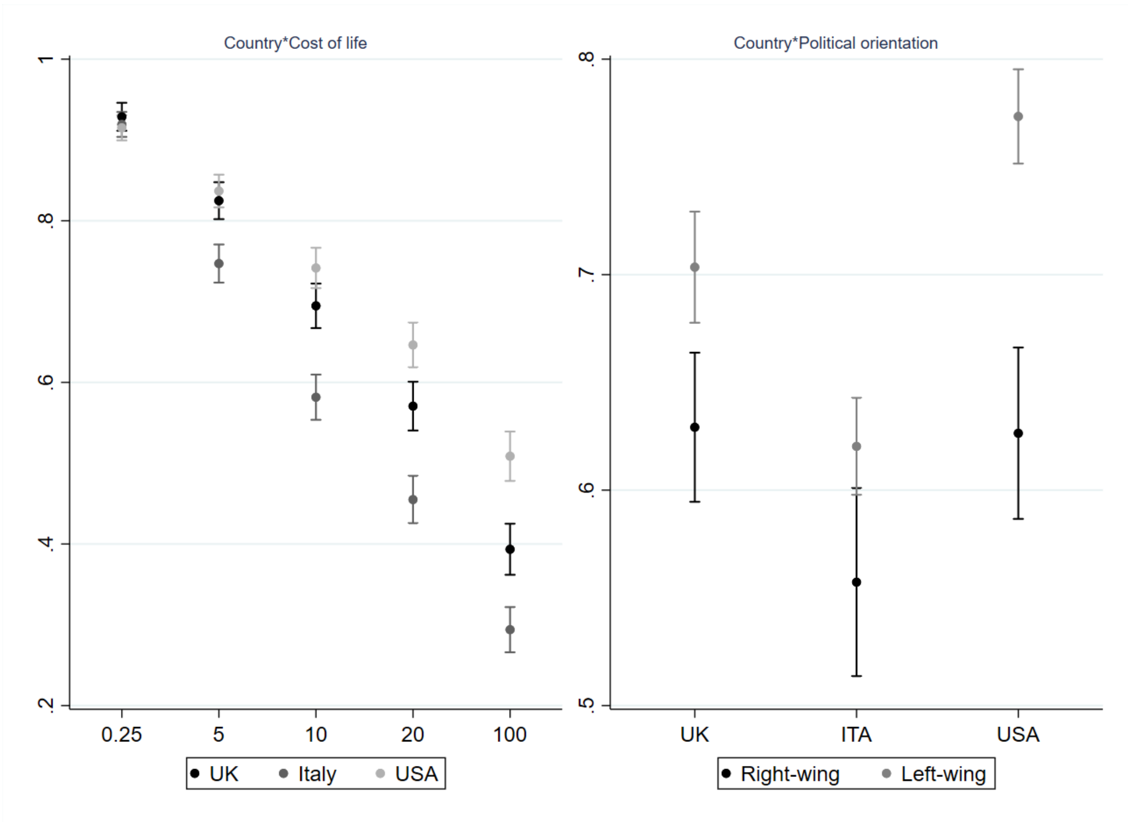


Figure 4: Pro-health choices by country interacted with the cost of saving one life and with political orientation

Notes: The left panel (*Country*Cost of life*) reports the fraction of pro-health choices (vertical axis) by country for different costs of saving one life (horizontal axis). Choices are indistinguishable by country and almost entirely pro-health when the cost in terms of jobs lost to save one life is low. As this cost increases, the fraction of pro-health choices monotonically decreases particularly in Italy and, to a lower extent, in the UK. The right panel (*Country*political orientation*): reports the fraction of pro-health choices (vertical axis) for participants that self-identify as politically left- and right-oriented, in each country. Difference between left and right participants are larger in the US than in the UK and Italy.

$$MRS_{F,J} = \frac{\alpha}{1-\alpha}\left(\frac{\overline{F}-F}{\overline{J}-J}\right)^{1-\rho}. \tag{2}$$

When $\rho = 1$ the MRS is constant, identifying perfect substitutes. When $\rho < 1$ the MRS increases with $J$, showing a preference for diversification over the two 'bads.' This case corresponds to convex preferences and concave indifference curves in a Cartesian plan like Figure 2. Conversely, when $\rho > 1$ the MRS decreases with $J$. In this case the two 'bads' are characterized by decreasing marginal disutility and preferences show aversion to diversification.

Preferences as specified in Equation 1 are estimated using a structural model, which does not impose a deterministic choice between alternatives. The agent is a utility maximizer who can make a zero mean error ($\varepsilon$) in comparing the (dis)utility of the available alternatives. We adopt a Fechner representation of stochastic decisions, meaning that the probability of choosing Scenario B is given by:

$$Pr(B) = Pr(u_B - u_A > \varepsilon). \tag{3}$$

and adopt the error specification proposed by Luce (1959).

The maximum likelihood estimation of this structural model confirms the strong pro-health attitude of the respondents who react to the relative cost of one life. The weight on jobs lost turns out to be $\hat{\alpha} = 0.05$, i.e. $1 - \hat{\alpha} = 0.95$ on fatalities, implying that health outcomes weight sizeably more than economic outcomes.[6] As an intuition, if $\hat{\rho}$ was equal to one the number of jobs lost for saving one life that makes the representative agent indifferent between saving it or not (i.e. the inverse of MRS) would be constant and equal to $(1 - \hat{\alpha})/\hat{\alpha} = 19$.

The curvature parameter is instead significantly larger than one: $\hat{\rho} = 1.21$ (two-sided $t$-test: $p$-value $< 0.001$).[7] An implication of $\hat{\rho} > 1$ is that the rate at which respondents are willing to lose jobs to save one life is not constant in our two-dimensional space. This rate, computed for all the south-east vertexes of the decision problems in Figure 2, ranges between 10.76 and 20.18 (14.78 on average), with lower values observed when the choice is made for relatively worse health situations. In fact, a value of $\hat{\rho} > 1$ implies a diminishing sensitivity to the two 'bads.' For instance, the higher the minimum number of fatalities that has to be accepted in any case, the lower is the value that respondents attach to saving one life. A similar reasoning applies to jobs lost. The implication is an aversion to diversify the two 'bads,' the opposite pattern as compared to what typically observed dealing with 'goods.'

Cross country differences are also confirmed estimating the structural model. The estimates of the parameters for the UK are very similar to the aggregate ones ($\hat{\alpha}_{UK} = 0.05, \hat{\rho}_{UK} = 1.22$). Italy has a larger distribution weight and a smaller curvature parameter, while the converse holds for the US ($\hat{\alpha}_{ITA} = 0.08, \hat{\rho}_{ITA} = 1.06$; $\hat{\alpha}_{US} = 0.03, \hat{\rho}_{US} = 1.47$). Indifference curves are not-significantly different from linear in Italy, while they display diminishing sensitivity in the UK, and even more so in the US. The parameters imply indifference curves that are steeper for Italy and are summarized in Figure 5.[8]

As already mentioned, the inverse of the MRS represents the number of jobs lost for saving one life that makes the agent indifferent between saving it or not. Taking the average over all the south-east vertexes of our decision problems, this number is 10.68 in Italy, 14.61 in the UK and 19.00 in the US. In other words, respondents are on average relatively less willing to lose jobs in order to reduce the number of fatalities in Italy. The

[6]Also including in the analysis the respondents characterized by horizontal or vertical indifference curves the weight on fatalities would be $1 - \hat{\alpha} = 0.98$. Results available upon request.

[7]In what follows $p$-values refer to two-sided $t$-tests when not specified otherwise.

[8]The results reported refer to aggregate estimates by country and are robust (with a comparably lower $\alpha$ in each country) to the inclusion respondents with horizontal and vertical preferences. The comparison across country hold qualitatively unchanged also when based on the average value of the parameters estimated at the individual level, in the sub-sample of respondents for whom the estimation procedure converges. Results available upon request.

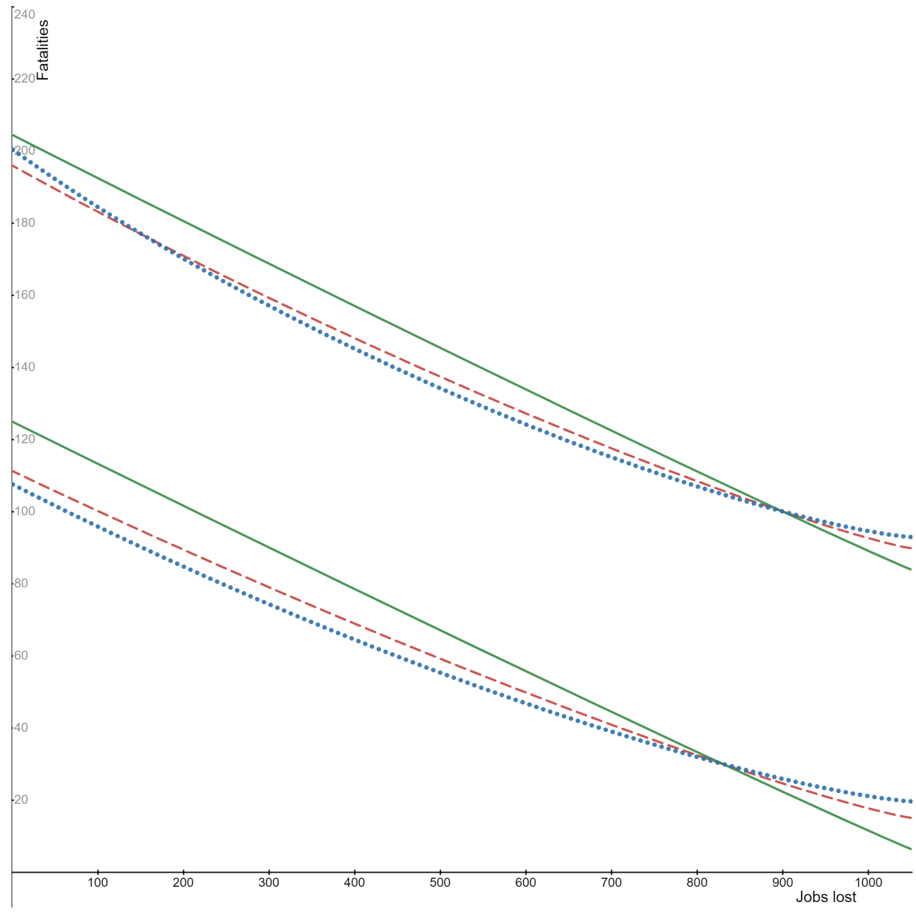


Figure 5: Indifference Curves by country

Solid line: Italy; Dashed line: UK; Dotted line: US.Each line represent an indifference curve. The space below each curve represents the set of combinations that are strictly preferred. Indifference curves that are closer to the origin represent preferred combination of outcomes. The inverse of the slope of the indifference curves captures the amount of jobs lost that the respondent considers equivalent to one life. Hence, a more pronounced pro-health attitude is captured by flatter indifference curves.

differences in the slopes become larger and larger as the economic situation worsens and the health situation improves, so that they are largest in the bottom-right part of the graph. In the scenario with the largest number of jobs lost and the smallest number of fatalities, among those considered in the survey, respondents in the US ask for 37.01 jobs in order to accept one additional fatality, respondents in the UK ask for 20.24, and respondents in Italy for 11.70.

### 3.3 Robustness

In this section, we present additional empirical exercises aimed at excluding competing explanations of the observed cross country differences based on possible artifacts. In particular, we try to connect our findings in Section 3.1 to *a)* differences in the sub-samples, *b)* differences in the fundamentals on both the health and economic side, *c)* effects of the experienced policies. Since we want to be conservative when excluding competing explanations, we also report uncorrected $p$-values for these exercises. The results corroborate the robustness of the observed differences in the solution of the moral dilemma across countries.

**Sample unbalancedness.** A natural comment when observing the results across countries is that the measured moral preferences may be driven by sample unbalancedness rather than reflecting genuine differences. In Section 3.1 the multivariate analysis already partials out the effect of observable characteristics, but in this section we conduct two additional robustness checks.

First, Figure 6 reports the interaction effects between the country and each explanatory variable. The results show that country differences, and in particular the difference between Italy on the one hand and the US and UK on the other hand, are found consistently in all demographic/attitudes/experience subgroups. Note in particular that age is the variable where differences across countries are more pronounced, with Italy being characterized by a significantly younger sample of participants. The upper left panel of Figure 6 shows that differences in the pro-health attitude differs by country also in this sub-sample.

Second, we deal with the heterogeneity across sub-samples replicating the analysis in Section 3.1 using Inverse Probability Weighting (IPW), as in Słoczyński and Wooldridge (2018). IPW is a standard procedure to account for differences across samples. It relies on building a logistic regression model to estimate the individual probability of belonging to one of the three groups, and then using the inverse of such predicted probability as a weight in the linear probability model. Table 3 reports the estimated country effects with and without IPW. The magnitude of the change in the probability of choosing the outcome with fewer fatalities is virtually unaffected by this procedure, suggesting that differential selection on observables across countries is not driving country differences.

**Different cost of the fundamentals.** Another possible explanation for the observed heterogeneity in moral preferences across countries is that the dilemma is perceived differently because the cost attached to either health or economic outcomes is not the same. For instance, labor market conditions vary widely across countries. Unemployment has been historically higher and lasting longer in Italy. It could then be that 'losing one job' is perceived as more costly in Italy than in the other countries. Consequently, apparent differences in moral preferences could actually stem from a higher cost of unemployment.

We exploit within-country differences in labor market conditions to test for this hypothesis. In case of jobs lost we use both Italy and the US, while the concentration of observations in England prevent us from extending the exercise to the UK. We match our dataset with data on unemployment rates in 2020 at the regional level from the Italian National Institute for Statistics (ISTAT), and at the state level in the US from the National Bureau of Labor Statistics. Local labor markets conditions differ widely within

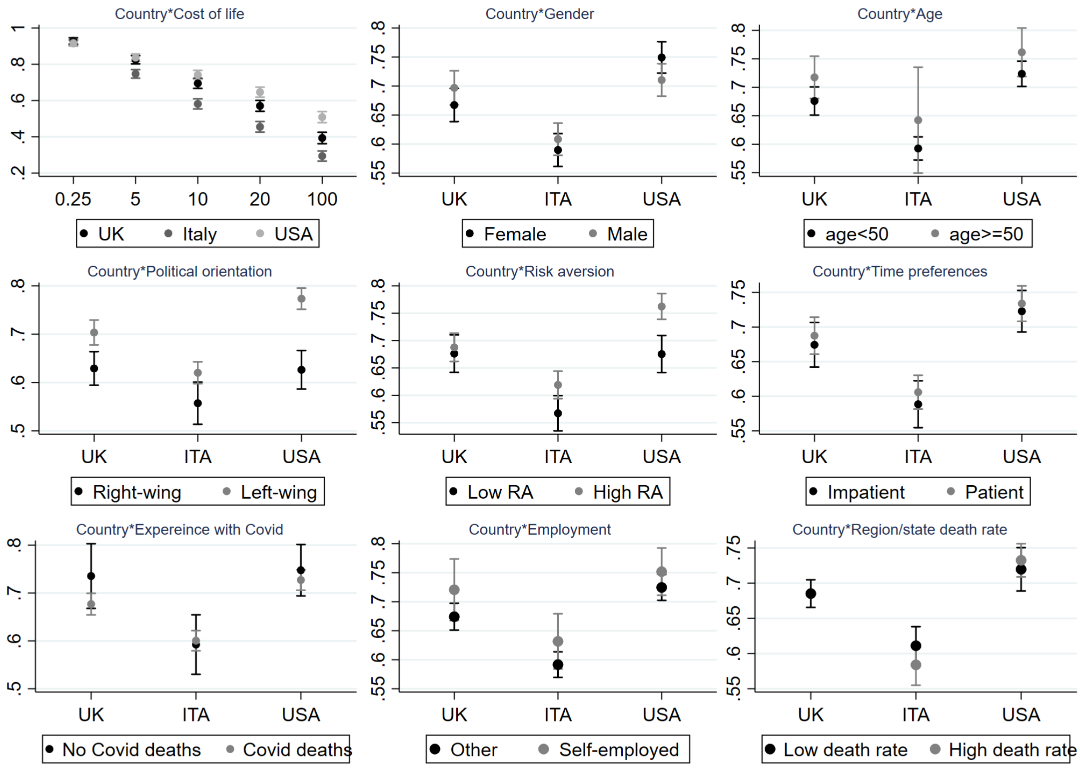


Figure 6: Pro-health choices by country interacted with each explanatory variable

Notes: The figure reports the fraction of pro-health choices (vertical axis) by country for the explanatory variable specified in each panel.

Table 3: Country differences with Inverse Probability Weighting

| | Linear probability | | Probit | | Logit | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| | Baseline | IPW | Baseline | IPW | Baseline | IPW |
| Italy | -0.0831*** | -0.0806*** | -0.0826*** | -0.0805*** | -0.0831*** | -0.0809*** |
| | (0.0159) | (0.0189) | (0.0160) | (0.0189) | (0.0161) | (0.0190) |
| US | 0.0472** | 0.0343* | 0.0468*** | 0.0337* | 0.0471** | 0.0339* |
| | (0.0164) | (0.0176) | (0.0160) | (0.0173) | (0.0162) | (0.0174) |

*Notes*: The table reports the change in the probability of choosing the alternative with fewer fatalities in Italy and the US as compared to the the UK (coefficients for the case of linear probability models and marginal effects for Probit and Logit models). Columns labeled 'IPW' (2, 4, 6) report estimates obtained with Inverse Probability Weighting, that is, after weighting observations with their propensity score of belonging to each country, based on observables. These propensity scores are obtained from a Multinomial Logit that includes all the variables used in the balance tests in Table 1. Columns labeled 'Baseline' (1, 3, 5) report estimates without Inverse Probability Weighting. All regressions include the same independent variables as in Column (8) of Table 2. Standard errors, clustered at the individual level, are reported in parentheses. Significance levels: $^{*}$ $= p-$value$< 0.1$; $^{**}$ $= p-$value$< 0.5$; $^{***}$ $= p-$value$< 0.01$.

Italy, with Northern regions experiencing levels and duration of unemployment similar to those in the US and in the UK, and definitely lower than in Southern regions. The unemployment rate in Italy ranges from 4.5% to 20.1% (mean 9.5 st.dev. 4.6), while in the US from 4.1% to 13.5%. (mean 7.4 st.dev. 1.9).

If the relatively stronger concerns for economic outcomes observed among Italian respondents was explained by higher cost of losing jobs, we should also observe the unemployment rate to significantly correlate with a pro-health behavior of the respondents. As shown in Column 1-3 of Table 4, this indeed is not the case (Column 1: ITA+US, uncorrected $p$-value = 0.261; Column 2: ITA, uncorrected $p$-value = 0.444; Column 3: US, uncorrected $p$-value = 0.939).

A similar reasoning holds for the other dimension, as health costs are higher in the US than in the other two countries. Respondents in the US may indirectly weight the choice to minimize the number of jobs lost against a relatively higher costs in the health domain. We exploit variability across US states in out-of-pocket health expenditure per-capita (source: US Census) and health care spending per-capita (source: Medicare and Medicaid Services) and test whether these variables correlate with a higher probability to minimize fatalities.[9] Also in this case we do not find any significant correlation (Table 4; Column 4: out-of-pocket spending, uncorrected $p$-value = 0.420; Column 5, out-of-pocket spending, uncorrected $p$-value = 0.454; health spending, uncorrected $p$-value = 0.186).

**Policy shapes preferences.** According to the Oxford COVID-19 Government Response Tracker, the stringency index at the country level ranks Italy consistently higher than the UK and the US. This fact clearly emerges using both the average and the maximum score, regardless of the time period considered (e.g. throughout the pandemic rather than in the last 90 days prior to our survey). When observing that policies are not aligned to citizens' preferences, it may be the case that more stringent policies have caused the preferences we measure. For instance, US respondents who have experienced a relatively weaker health protection through milder containment measures may have become more concerned by the health dimension of the dilemma.

To test this alternative hypothesis we exploit the variability of the stringency index across states. Unfortunately, this exercise must be limited to the US, because in Italy the stringency index exists only at the national level while in the UK there are only four states and the large majority of our sample comes from England. The exercise is meaningful nevertheless: the variability of the stringency index within the US is comparable in size to that between Italy and the US as a whole and Italy's average index (67) is relatively close to the highest average index among US states (62). If more stringent policies were at the root of country differences, we should expect respondents from the states characterized by more stringent policies to display a lower pro-health attitude. This is

[9]Insufficient variability within the other countries and non-availability of data prevent us from extending the exercise to the whole sample).

Table 4: Linear probability model: within-country variation in policy and fundamentals

| | Dependent variable: Choice of the alternative involving fewer fatalities | | | | | | |
|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| | ITA+US | ITA | US | US | US | US | US |
| Unemployment | 0.00253 | -0.0176 | 0.000401 | | | | 0.000096 |
| | (0.00226) | (0.0230) | (0.00521) | | | | (0.00646) |
| Long-term unempl. | | 0.0242 | | | | | |
| | | (0.0295) | | | | | |
| Out-of-pocket spending | | | | -0.0000237 | -0.0000220 | | -0.0000302 |
| | | | | (0.0000294) | (0.0000294) | | (0.0000327) |
| Health spending | | | | | 0.0000117 | | 0.0000083 |
| | | | | | (0.00000887) | | (0.0000099) |
| Stringency index | | | | | | 0.00320 | 0.00269 |
| | | | | | | (0.00294) | (0.00337) |
| Str. index (last 90 days) | | | | | | -0.000799 | -0.000434 |
| | | | | | | (0.00209) | (0.00216) |
| US | 0.133*** | | | | | | |
| | (0.0144) | | | | | | |
| Constant | 0.541*** | 0.704*** | 0.606*** | 0.615*** | 0.521*** | 0.480*** | 0.433*** |
| | (0.0337) | (0.0919) | (0.0532) | (0.0353) | (0.0790) | (0.112) | (0.116) |
| Controls | YES | YES | YES | YES | YES | YES | YES |
| *N* | 39882 | 19758 | 20124 | 20124 | 20124 | 20124 | 20124 |

*Notes*: Coefficients represent the change in the probability of choosing the alternative with fewer fatalities (as compared to the Constant) associated to each independent variable. Except for the first model, where the variables are available for both the US and Italy, all other regressions are run within countries (see the top row) and exploit within-country variability. Column 1-3 show that labor market conditions do not affect pro-health attitudes. *Unemployment* is the unemployment rate in 2020 at the state (US) or regional (Italy) level (source: NBLS, ISTAT). *Long-term unemployment* is the rate of unemployment for more than 12 months in 2020 at the regional (Italy) level (source: ISTAT). Coulumns 4-5 show that different cost of health care does not affect the choices. *Out-of-pocket spending* is the estimated median annual out-of-pocket healthcare spending per-capita in each US state (source: US Census). *Health spending* is the average annual healthcare spending per-capita in each US state (source: Medicare and Medicaid Services). Column 6 shows that choices do not correlate with the containment measures experienced by the respondents. *Stringency index* is the mean value of the stringency index throughout the pandemic for each US state, and *Stringency index (last 90 days)* is the same mean computed over the 90 days prior to the survey (source: Oxford COVID-19 Government Response Tracker). All the specifications include the controls of our favorite specification (Table 2, Column 8): Cost of one life, Female, Age $\geq$ 50, Left-wing, Risk averse, Patient, Self-employed, Covid deaths, Tested positive, High death rate. Standard errors, clustered at the individual level, are reported in parentheses. Significance levels corrected for multiple hypotheses testing (Bonferroni): $^{*} = p-$value$< 0.1$; $^{**} = p-$value$< 0.5$; $^{***} = p-$value$< 0.01$.

indeed not the case.

As shown in Column 6 of Table 4, the stringency index does not significantly correlate with the probability of opting for the alternative with the lower number of fatalities (Str. Index overall: uncorrected $p$-value $= 0.276$; Str. Index last 90 days: uncorrected $p$-value $= 0.703$). Equivalent results are obtained with different specifications of the stringency index, i.e. when using the maximum rather than the average level of the stringency score, as well as manipulating the period of time over which the index is computed.

Overall, we do not find any evidence that cross-country differences in the solution of the moral dilemma can be rationalized by unbalanced samples, by the policies experienced during the pandemic, or by different relative costs of the two 'bads' presented in the choice.

# 4 Conclusion

Within our sample, when confronted with a trade-off between fatalities and jobs losses, individuals prioritize saving more lives in most scenarios: overall, our estimation shows that about 95 percent of the weight in the utility function is allocated to fatalities. However, their answers respond to the cost of saving each life: rather than blindly applying a moral imperative to save more lives, most participants carefully weight the consequences of each scenario and choose based also on the severity of the trade-off. Indeed, when 20 jobs or more need to be sacrificed to save one more life, about half of the participants choose to save more jobs.

Decisions are influenced in a predictable way by an individual's age and attitudes, so that older and more left-wing participants put a larger weight on fatalities. We see this as a comforting signal on the quality of our data. We do not find any effect of gender and of any of the variables that capture different experiences throughout the pandemic. We conclude that participants' attitudes toward this moral dilemma are rooted in rather stable traits and preferences, rather than influenced by contingent situations.

A surprising result of our study is that respondents on average show a diminishing sensitivity to the increase of any of these two 'bads.' The higher the number of fatalities that respondents have to accepted in any case, the lower is the value they attach to saving one life. A similar reasoning applies to jobs lost. As a consequence, respondents display an aversion to diversify between the two dimension proposed, the opposite pattern as compared to what typically observed dealing with 'goods.' Interestingly, this opposite finding across 'goods' and 'bads' mirrors in a different framework one of the pillars of Prospect Theory (Kahneman and Tversky, 1979) and may have relevant policy implications.

Another striking result of our study is that the moral preferences over this taboo trade-off look misaligned with the policies adopted during the pandemic across the countries investigated. The stringency indexes rank Italy higher than UK and the US in the strength of containment measures undertaken, but we find instead that US and UK respondent are relatively more pro-health than Italians.

The robustness analysis corroborates the result that the respondents' attitude toward this moral dilemma genuinely differs across countries, and in a relatively different direction as compared to the containment measures undertaken during the pandemic. Our results seems to point at some unobservable cultural reason as the cause for the differences observed. While unfortunately our study is not equipped to identify the ultimate explanation, we believe it sets an interesting avenue for future research.

## References


ALVAREZ, F., ARGENTE, D. and LIPPI, F. (2020). A simple planning problem for covid-19 lockdown, testing, and tracing. *American Economic Review: Insights*.

ANDERSEN, S., HARRISON, G., LAU, M. and RUTSTRÖM, E. (2006). Elicitation using multiple price list formats. *Experimental Economics*, **9**, 383–405, 10.1007/s10683-006-7055-6.

ANDERSSON, O., CAMPOS-MERCADE, P., CARLSSON, F., SCHNEIDER, F. and WENGSTRÖM, E. (2021). The individual welfare costs of stay at home policies. *Scandinavian Journal of Economics*.

ASHENFELTER, O. (2006). Measuring the value of a statistical life: problems and prospects. *The Economic Journal*, **116** (510), C10–C23.

BELLE, N. and CANTARELLI, P. (2022). Your money, your life, or your freedom? a discrete-choice experiment on trade-offs during a public health crisis. *Public Administration Review*.

CARRIERI, V., DE PAOLA, M. and GIOIA, F. (2021). The health-economy trade-off during the covid-19 pandemic: Communication matters. *PloS one*, **16** (9), e0256103.

CHORUS, C., SANDORF, E. D. and MOUTER, N. (2020). Diabolical dilemmas of covid-19: An empirical study into dutch society's trade-offs between health impacts and other effects of the lockdown. *PLoS One*, **15** (9), e0238683.

CODAGNONE, C., BOGLIACINO, F., GÓMEZ, C., CHARRIS, R., MONTEALEGRE, F., LIVA, G., LUPIÁÑEZ-VILLANUEVA, F., FOLKVORD, F. and VELTRI, G. A. (2020). Assessing concerns for the economic consequence of the covid-19 response and mental health problems associated with economic vulnerability and negative economic shock in italy, spain, and the united kingdom. *Plos one*, **15** (10), e0240876.

DOHMEN, T., FALK, A., HUFFMAN, D., SUNDE, U., SCHUPP, J. and WAGNER, G. G. (2011). Individual Risk Attitudes: Measurement, Determinants, And Behavioral Consequences. *Journal of the European Economic Association*, **9** (3), 522–550.

EICHENBAUM, M. S., REBELO, S. and TRABANDT, M. (2021). The macroeconomics of epidemics. *The Review of Financial Studies*, **34** (11), 5149–5187.

FAVERO, C. A., ICHINO, A. and RUSTICHINI, A. (2020). *Restarting the economy while saving lives under Covid-19*. Tech. rep., CEPR Discussion Paper No. DP14664.

FILIPPIN, A. and CROSETTO, P. (2016). A reconsideration of gender differences in risk attitudes. *Management Science*, **62** (11), 3138–3160.

FOOT, P. (1967). The problem of abortion and the doctrine of the double effect. *Oxford review*, **5**.

HALE, T., ANGRIST, N., GOLDSZMIDT, R., KIRA, B., PETHERICK, A., PHILLIPS, T., WEBSTER, S., CAMERON-BLAKE, E., HALLAS, L., MAJUMDAR, S. *et al.* (2021). A global panel database of pandemic policies (oxford covid-19 government response tracker). *Nature Human Behaviour*, **5** (4), 529–538.

HEY, J. D. and ORME, C. (1994). Investigating Generalizations of Expected Utility Theory Using Experimental Data. *Econometrica*, **62** (6), 1291–1326.

JENNI, K. and LOEWENSTEIN, G. (1997). Explaining the identifiable victim effect. *Journal of Risk and uncertainty*, **14** (3), 235–257.

JONES, C., PHILIPPON, T. and VENKATESWARAN, V. (2021). Optimal mitigation policies in a pandemic: Social distancing and working from home. *The Review of Financial Studies*, **34** (11), 5188–5223.

KAHNEMAN, D. and TVERSKY, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, **47** (2), 263–91.

KAPLAN, G., MOLL, B. and VIOLANTE, G. L. (2020). *The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the US*. Tech. rep., National Bureau of Economic Research.

LEE, S. and FEELEY, T. H. (2016). The identifiable victim effect: A meta-analytic review. *Social Influence*, **11** (3), 199–215.

LEÓN, G. and MIGUEL, E. (2017). Risky transportation choices and the value of a statistical life. *American Economic Journal: Applied Economics*, **9** (1), 202–28.

LESSCHAEVE, C., GLAURDIĆ, J. and MOCHTAK, M. (2021). Health versus wealth during the covid-19 pandemic: Saving lives or saving the economy? *Public Opinion Quarterly*, **85** (3), 808–835.

LORÍA-REBOLLEDO, L. E., RYAN, M., WATSON, V., GENIE, M. G., SAKOWSKY, R. A., POWELL, D. and PARANJOTHY, S. (2022). Public acceptability of non-pharmaceutical interventions to control a pandemic in the uk: a discrete choice experiment. *BMJ open*, **12** (3), e054155.

LUCE, R. (1959). *Individual Choice Behavior: A Theoretical Analysis*. Wiley.

MANIPIS, K., STREET, D., CRONIN, P., VINEY, R. and GOODALL, S. (2021). Exploring the trade-off between economic and health outcomes during a pandemic: a discrete choice experiment of lockdown policies in australia. *The Patient-Patient-Centered Outcomes Research*, **14** (3), 359–371.

OANA, I.-E., PELLEGATA, A. and WANG, C. (2021). A cure worse than the disease? exploring the health-economy trade-off during covid-19. *West european politics*, **44** (5-6), 1232–1257.

PALAN, S. and SCHITTER, C. (2018). Prolific.ac—a subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, **17**, 22–27.

SŁOCZYŃSKI, T. and WOOLDRIDGE, J. M. (2018). A general double robustness result for estimating average treatment effects. *Econometric Theory*, **34** (1), 112–133.

SMALL, D. A. and LOEWENSTEIN, G. (2003). Helping a victim or helping the victim: Altruism and identifiability. *Journal of Risk and uncertainty*, **26** (1), 5–16.

SPENCER, A., RIVERO-ARIAS, O., WONG, R., TSUCHIYA, A., BLEICHRODT, H., EDWARDS, R. T., NORMAN, R., LLOYD, A. and CLARKE, P. (2022). The QALY at 50: One story many voices. *Social Science & Medicine*, **296**, 114653.

STOTT, H. (2006). Cumulative prospect theory's functional menagerie. *Journal of Risk and Uncertainty*, **32** (2), 101–130.

TETLOCK, P. E. (2003). Thinking the unthinkable: Sacred values and taboo cognitions. *Trends in cognitive sciences*, **7** (7), 320–324.

VISCUSI, W. K. and ALDY, J. E. (2003). The value of a statistical life: a critical review of market estimates throughout the world. *Journal of risk and uncertainty*, **27** (1), 5–76.

## Structural Estimation

**Utility function.** Equation 1 encompasses several classic utility functions according to the value taken by the parameters $\alpha$ and $\rho$.
$\alpha \in [0,1]$ is the distribution parameter that captures the weight assigned to economic outcomes. The extreme cases are represented by:

- $\alpha = 0$: Horizontal Indifference Curves that characterize respondents for whom life has no price;
- $\alpha = 1$: Vertical Indifference Curves that characterize respondents who only consider the economic dimension of the trade-off.

These two extreme cases end up including also respondents that may instead be characterized by lexicographic preferences, which give strict priority to one of the two outcomes. Lexicographic preferences cannot be given a utility representation because they violate the continuity axiom, but in this study choices made under this type of preferences are observationally equivalent to those made with horizontal/vertical indifference curves.
$\rho$ is the curvature parameter, that identifies several paradigmatic cases according to the value it takes. In particular:

- $\rho < 1$: Constant elasticity of substitution with respect to $(\bar{J} - J)$ and $(\bar{F} - F)$ that express a preference for diversification among the two 'bads'. In this case the parameter $\rho$ also defines the (constant) elasticity of substitution, which is equal to $1/(1-\rho)$. Two well-known sub-cases are:

  1. $\rho \to -\infty$: Preferences represents outcomes that are perfect complements. In this case the decision maker considers a fixed proportion of the two 'bads' as the optimal outcome
  2. $\rho \to 0$: Cobb-Douglas.

- $\rho = 1$: Linear preferences. In this case outcomes are perfect substitutes, with preferences entirely dictated by the distribution parameter $\alpha$.

- $\rho > 1$: This is the case of convex indifference curves over the two 'bads,' that represent aversion to diversification. In other words, the agent prefers unbalanced outcomes over balanced ones.

**Estimation.** The error specification proposed by Luce (1959) implies that:

$$Pr(B) = \frac{u_B^{\frac{1}{\gamma}}}{u_A^{\frac{1}{\gamma}} + u_B^{\frac{1}{\gamma}}},$$

where $\gamma$ controls the stochastic component of the decision. In fact, $Pr(B)$ converges to $\frac{1}{2}$ (random choice) as $\gamma \to \infty$, whilst, as $\gamma \to 0$, it goes to 1 when $u_B > u_A$ and to 0 when $u_A < u_B$ (deterministic choice). The Luce choice model can be given a Fechner representation using the logarithms of utilities. In fact, the last equation can be derived starting from $Pr(B) = Pr[ln(u_B) - ln(u_A) < \varepsilon]$ with the error that follows a logistic distribution $\varepsilon \sim \Lambda(0, \gamma)$.