Journal of Economic Behavior and Organization 235 (2025) 107033

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Research paper

Moral preferences over health-wealth trade-offs

Antonio Filippin[a,b,*], Marco Mantovani[c]

[a] University of Milan, Department of Economics, Via Conservatorio 7, 20122 Milan, Italy
[b] Institute for the Study of Labor (IZA), Schaumburg-Lippe-Str. 5-9, 53113 Bonn, Germany
[c] University of Milan-Bicocca, Department of Economics, Management and Statistics, Piazza dell'Ateneo Nuovo, 1, 20126 Milan, Italy

ARTICLE INFO

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

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. We find that the participants' utility function is mostly driven by health outcomes, 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. A structural estimation also displays, surprisingly, aversion to diversification among these two bads. Moreover, policy responses look misaligned with estimated preferences. Italy adopted more stringent containment measures, while Italian respondents display a relatively weaker pro-health attitude.

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 across population groups. In particular, differences across countries that implemented different policy responses and across individuals with different personal experiences during the pandemic. Furthermore, we present a small-scale test of a tool for estimating such preferences.

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 choose this framing because the moral dilemma induced by the pandemic is naturally rooted in the bads domain. 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 contain 25 of these choices for each participant.

* Corresponding author at: University of Milan, Department of Economics, Via Conservatorio 7, 20122 Milan, Italy.
E-mail addresses: antonio.filippin@unimi.it (A. Filippin), marco.mantovani@unimib.it (M. Mantovani).

https://doi.org/10.1016/j.jebo.2025.107033
Received 14 May 2024; Received in revised form 10 April 2025; Accepted 12 April 2025
Available online 8 May 2025

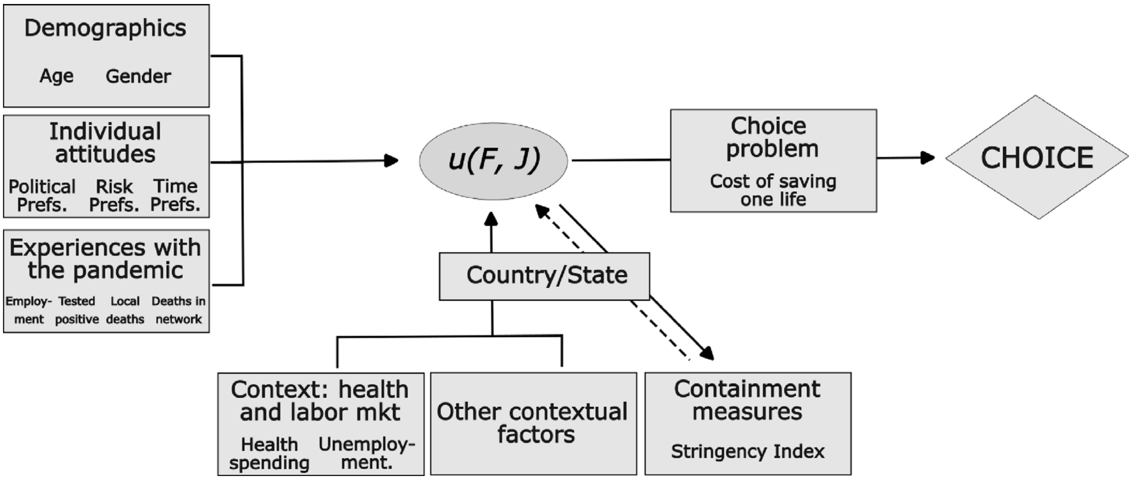

Fig. 1. Conceptual map.

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 (Viscusi and Aldy, 2003; Ashenfelter, 2006; Belle and Cantarelli, 2022; León and Miguel, 2017).

The analysis proceeds in two parallel ways. First, we adopt different regression models to estimate the determinants that affect 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 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 (Hey and Orme, 1994; Andersen et al., 2006; Stott, 2006; Filippin and Crosetto, 2016).

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 in the moral dilemma, 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.

Fig. 1 summarizes the conceptual structure of how these determinants are expected to shape the respondents' choices in the dilemma through their preferences.

One central tenet of representative democracies is that policies are responsive to the preferences of the voters. Evidence of the responsiveness of policies has been found in multiple domains (Page and Shapiro, 1983; Kang and Powell, 2010; Reynolds, 2013). We expect preferences to be concordant with each country policy response also in the case of the pandemic. 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. However, it is worth noting that containment policies may have shaped preferences besides (or rather than) reflecting them. For instance, more stringent policies may have weakened pro-health attitudes. Moreover, the relation between the country of origin and individual preferences may reflect other factors. For instance, countries differ in labor market conditions and healthcare systems. This may make economic and health consequences of the pandemic more or less salient and costly to individuals. We leverage within-country variations along all these dimensions to explore these alternative explanation. In particular, we exploit variation in the stringency index across US states (from the Oxford COVID-19 Government Response Tracker, Hale et al., 2021), variation in labor market conditions across US states and Italian regions (data from Italian and US bureaus of statistics) and variation in healthcare costs across US states (data from Medicare and Medicaid Services, and US Census).

Individual decision-making is often found to be affected by personal experiences in a wide range of domains (Andre et al., 2022; Conrad et al., 2022; Malmendier, 2021). Lifetime experiences with inflation (Malmendier and Nagel, 2016) as well as daily grocery shopping experiences (D'Acunto et al., 2021) shape inflation expectations. Unemployment has persistent effects on consumption (Malmendier and Shen, 2024). Experience shapes economic preferences as well. Risk preferences are affected by past exposure to different levels of risk, returns and ambiguity both in the lab and in the field (Hoffmann and Post, 2017; Mengel et al., 2016; Ert and Haruvy, 2017; Charness et al., 2023). In general, experiences seem to affect the belief updating process. Individuals act as if past outcomes that they experienced were overly likely to occur again, emphasizing their relevance. We therefore expect that participants who faced adverse health (economic) contingencies due to the pandemic display stronger pro-health (pro-economy) attitudes.

Preferences may also vary depending on demographic characteristics and individual traits. Women and men often show different preferences, with respect, for instance, to economic and social preferences, to policy preferences and in health behavior. We then ask whether women and men differ in the way they solve the health-wealth trade-off.

On top of these (pre-registered) hypotheses, other individual characteristics and traits may affect the preferences we elicit: for instance, older participants, for whom adverse health outcomes are more relevant than economic ones, more risk-averse participants and those with more left-leaning political orientation may display stronger pro-health attitudes.

We find that differences in the pro-health attitude across the three countries do not correspond to the strength of the governments' responses throughout the pandemic. Relative to the UK, Italian respondents display a weaker pro-health attitude, and US respondents a stronger one. At least within countries, different labor market conditions or healthcare systems do not seem to affect preferences over the health-wealth trade-off, nor do more stringent containment measures weaken pro-health attitudes. Similarly, 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. While preferences are similar between men and women, other persistent individual traits and demographics are associated with preferences as expected: preferring outcomes with fewer fatalities and more jobs lost is associated with older age, left-leaning political preferences and risk aversion.

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 (Chorus et al., 2020; Carrieri et al., 2021; Lesschaeve et al., 2021; Manipis et al., 2021; Oana et al., 2021; Loría-Rebolledo et al., 2022) 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 recent literature, bridging economic and epidemiological models, estimates how health and economic outcomes respond to containment policies (Eichenbaum et al., 2021; Jones et al., 2021; Alvarez et al., 2020; Kaplan et al., 2020; Favero 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.

The structure of the paper is as follows. Section 2 describe the task representing the moral dilemma and the procedures of our study. Results of both the regression models and of the structural estimation are reported in Section 3. Section 4 concludes.

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.

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. Fig. 2 provides an example.

The 25 choices are extracted from a set of 125 alternatives, reported in Fig. 3. 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 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 Fig. 3. 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

[1] The scenarios reported in Fig. 3 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.

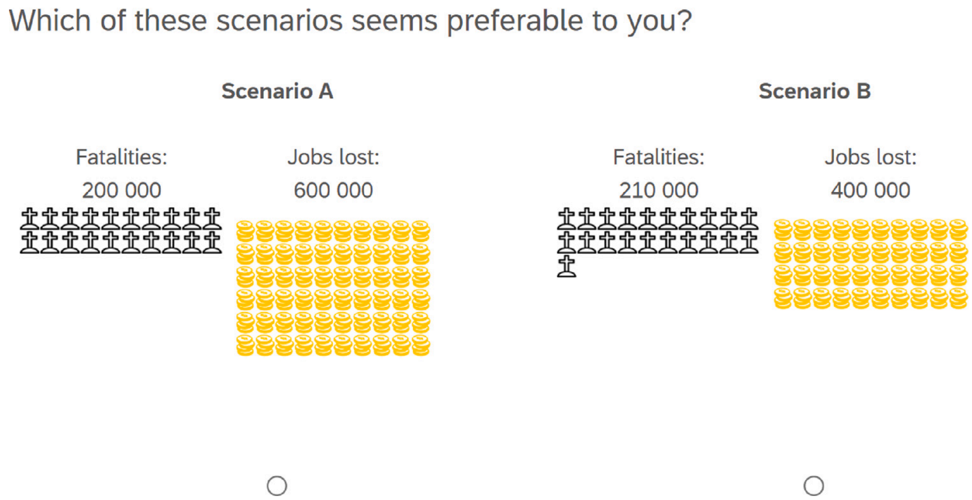

Fig. 2. An example of binary choice
Notes: Each choice presents two scenarios (called A and B). Every scenario is characterized by a number of fatalities and of jobs lost. The position of each combination is randomized, so that the scenario with fewer fatalities appears as A or B with equal probability. The respondents must indicate which of the two scenarios is relatively preferable.

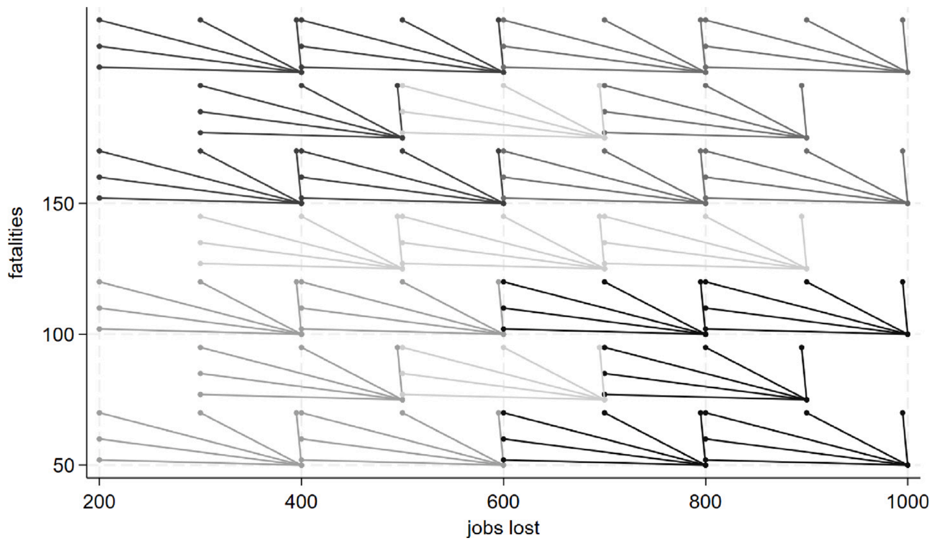

Fig. 3. The set of 125 choices
Notes: The *y*-axis represents fatalities, while the *x*-axis represents jobs lost. Every segment represent one binary choice between the coordinates at its extremes. The values reported are those administered in the questionnaire in Italy and in the UK, while values are 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. Different shades of gray identify the groups of choices used in the pseudo-randomization to ensure that each respondent receives alternatives evenly distributed in the plan.

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 of the five shaded areas in Fig. 3). Specifically, for each shaded area 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 labeled as Scenario A.

Choices must be hypothetical given their nature. It is therefore natural to wonder whether the lack of incentives may bias the results. Hypothetical bias (HB) is a well-known concern in contingent valuation, stated preference studies, and willingness-to-pay elicitation. Our design, in principle, is less prone to HB because it does not rely on a direct elicitation of monetary values or individual stated preferences but instead presents binary choices between pre-defined alternatives. Nevertheless, the two domains (health and wealth) may be affected differently by social desirability bias, potentially influencing responses. For instance, respondents might overweight fatalities in a hypothetical setting, assigning greater importance to health than they would in real-life decisions. This caveat suggests to avoid a cardinal interpretation of the relative weight attributed to each domain and to interpret cross-country differences with caution, in line with findings from (Penn et al., 2024; Ehmke et al., 2008). Socio-economic variables have also been explored as potential predictors of HB in past research, but the evidence remains inconclusive. For instance, (Mjelde et al., 2012) find that older and more educated respondents exhibit a smaller bias, yet they acknowledge that most of the variation in HB remains unexplained. We will discuss the potential impact of HB on our results in Section 3.1.

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 to isolate two categories: self-employed and atypical workers which are most vulnerable to the pandemic shock (Graeber et al., 2021; Mindes and Lewin, 2021). 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.

Invited participants. Invited 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 number has been chosen to avoid students being over-represented, and to have a sample directly exposed to adverse economic consequences of the pandemic in the same order of magnitude of the general population. Within each cell, corresponding to one country and one employment condition, the sub-sample was balanced by sex.

Procedures. 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. All subjects received a fixed amount ($1.20) for the successful completion of the questionnaire, which took on average 6 min and 16 s 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 min 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).[2]

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.[3] 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.

[2] 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.

[3] Geolocation data from Qualtrics can be imprecise because respondents answer to the questionnaire from a different state/region than where they live. While possible in principle, this imprecision seems unlikely to affect our results. In fact, the breakdown of responses by state/region (available upon request) tracks closely the distribution of the corresponding populations.

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 >. ≈ is instead used for differences that are not significant. Significance levels are corrected for multiple hypotheses testing (Bonferroni).

Exclusion criteria. 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 (3.13%). The sample used for the analysis is equal to 814, 793, and 805 in the US, UK, and Italy, respectively. While we consider a good practice to follow these pre-registered 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). One should keep these limitations in mind when interpreting our results.

Nevertheless, we cautiously design a number of exercises so as to make our exploratory evidence as informative as possible. First, all variables are included in the main regressions, so that each coefficient is obtained after controlling for the other variables. In Appendix A we further check the robustness of our set of results interacting the country dummies with each of these variables. This measures the various effects within each country sample, and also excludes that country effects are specific to subgroups where differences in the sample are more pronounced. For instance, Italy has the oldest population among the three countries, but the youngest one in the sample. The interaction allows to assess if age has similar effects across countries and whether young/old Italians have different preferences than young/old Americans in the sample. Furthermore, we adopt an inverse probability weighting procedure to exclude 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 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). 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 along several dimensions: country, 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.

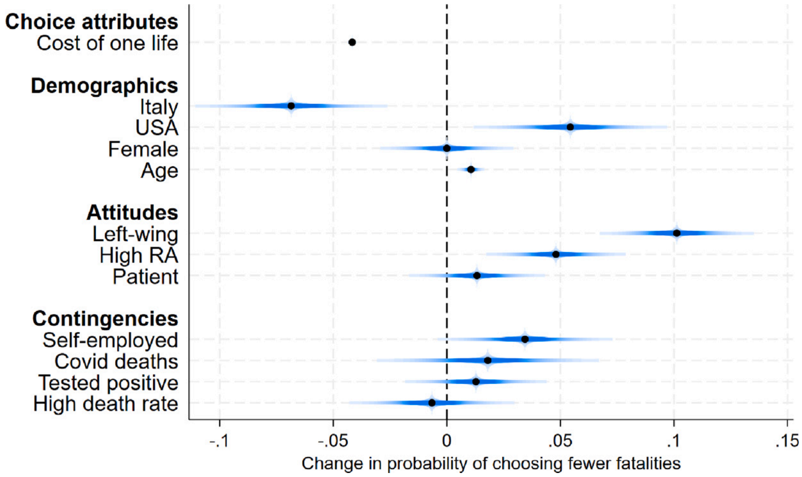

Fig. 4. Coefficients for the change in the probability of making pro-health choices
Notes: The Figure plots the estimated coefficients of the linear probability model derived from our richest specification (Column 1 in Table 2). For each explanatory variable the Figure reports the change in the probability 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 country, Male, etc.).

Columns (1)–(4) in Table 2 reports different regression models. The first column reports a standard linear probability model. The second takes into account the multilevel structure of the data by allowing for both individual and country random effects. Column (3) reports the marginal effects from a Probit specification. Finally, Column (4) deals with the heterogeneity across country sub-samples replicating the analysis in Column (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 countries based on the observables, and then using the inverse of such predicted probability as a weight in the linear probability model. In other words, these propensity scores are obtained from a Multinomial Logit with country as dependent variable, and including all the variables used in the balance tests in Table 1 as explanatory variables. In all models, standard error are clustered at the individual level and the significance of the regression coefficients is always adjusted for multiple hypotheses testing using the Bonferroni correction. Fig. 4 also plots the estimated coefficients of the linear probability model in Column 1.

The results are robust across the different specifications (as well as to the coding of the variables – dummy, categorical, continuous – when applicable). We expected the participants' choices to be aligned with the response to the pandemic in their country of origin, i.e. to observe Italian respondents showing the stronger pro-health orientation. Table 2 shows that sizeable differences emerging across country of residence have instead the opposite sign. The probability that participants in Italy choose the option involving lower fatalities is 6.9 percentage points lower than in the UK (the omitted variable corresponding to the vertical bar in Fig. 4) and 12.3 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). Our samples are not representative of the corresponding populations. Therefore, one cannot generalize this finding to the overall population of these countries. Nevertheless, the stability of the coefficients when using IPW suggests that a possible differential selection of the sample across countries on the variables that we observe do not explain the country differences.[4]

Contrarily to our expectations, variables that capture contingent experiences with the pandemic display null effects across the board. 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 in the moral dilemma. 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 significantly different from zero in the baseline specification (Tested positive: p-value $= 1.00$; Covid deaths: p-value $= 1.00$; High death rate: p-value $= 1.00$; Self-employed: p-value $= 0.157$, correction for multiple hypotheses testing applied).[5]

While we observe no difference between men and women (p-value $= 1.00$),[6] other stable trais and demographics are the variables that shape respondents' attitude toward this moral dilemma. Older participants show a stronger pro-health orientation: being 5 years older implies a 1 percent increase in the probability of choosing the option with fewer fatalities (p-value < 0.001). More risk-averse

[4] The analysis in Section 3.2 also tends to exclude an explanation based on reverse causality and on a different cost of the two bads.

[5] These results are robust across all the specifications, with the exception of Self-employed that turns out to be marginally significant in the IPW model (Column 4), but with the opposite sign as that we expected.

[6] Note that p-values equal to 1.00 may mechanically derive from the Bonferroni correction.

Table 2
Linear probability model.

	(1) lpm baseline	(2) lpm multilevel	(3) probit (margins)	(4) lpm ipw	(5) lpm always health	(6) lpm trade-off
Cost of one life	−0.0417***	−0.0417***	−0.0368***	−0.0415***		−0.0545***
	(0.000762)	(0.00428)	(0.000648)	(0.000887)		(0.000782)
Italy	−0.0687***	−0.0687***	−0.0672***	−0.0819***	−0.0854***	−0.0389**
	(0.0164)	(0.00982)	(0.0166)	(0.0194)	(0.0236)	(0.0164)
USA	0.0544**	0.0544***	0.0543***	0.0346*	0.0714**	0.0369*
	(0.0165)	(0.0126)	(0.0163)	(0.0177)	(0.0260)	(0.0171)
Female	0.0000136	0.0000135	−0.000471	−0.00452	−0.0273	0.0156
	(0.0114)	(0.0184)	(0.0114)	(0.0136)	(0.0173)	(0.0116)
Age	0.0106***	0.0106***	0.0108***	0.00767*	0.00843*	0.00931***
	(0.00232)	(0.00203)	(0.00240)	(0.00348)	(0.00368)	(0.00240)
Left-wing	0.101***	0.101***	0.102***	0.0933***	0.0883***	0.0834***
	(0.0132)	(0.0243)	(0.0132)	(0.0156)	(0.0186)	(0.0130)
Risk averse	0.0480***	0.0480**	0.0478***	0.0594***	0.0453**	0.0377**
	(0.0119)	(0.0162)	(0.0119)	(0.0147)	(0.0173)	(0.0121)
Patient	0.0133	0.0133***	0.0133	0.0158	0.00351	0.0154
	(0.0116)	(0.00206)	(0.0116)	(0.0133)	(0.0175)	(0.0118)
Self-employed	0.0342	0.0342**	0.0334	0.0502*	0.0665*	0.00525
	(0.0149)	(0.00477)	(0.0148)	(0.0166)	(0.0232)	(0.0158)
Covid deaths	0.0179	0.0179	0.0175	0.0257	0.0634	−0.0151
	(0.0190)	(0.0133)	(0.0189)	(0.0212)	(0.0294)	(0.0202)
Tested positive	0.0128	0.0128	0.0136	0.00242	0.00348	0.0161
	(0.0121)	(0.0180)	(0.0120)	(0.0141)	(0.0179)	(0.0122)
High death rate	−0.00651	−0.00651	−0.00641	−0.00292	0.00383	−0.00801
	(0.0141)	(0.0177)	(0.0140)	(0.0167)	(0.0210)	(0.0143)
_cons	0.623***	0.623***	0.667***	0.689***	0.104***	0.581***
	(0.0230)	(0.0497)	(0.0189)	(0.0123)	(0.0338)	(0.0232)
N	59 870	59 870	59 870	59 870	59 870	45 774

Notes: Except for model (4), coefficients represent the change in the probability of choosing the alternative with fewer fatalities as compared to the Constant associated to each independent variable. Model (1): linear probability model, random effects at the individual level. Model (2): linear probability model, random effects at the country and individual level. Model (3): probit model (marginal effects). Model (4): Inverse probability weighting. Model (5): linear probability model, dependent variable: choosing always the pro-health option. Model (6): linear probability model on subsample that does not always choose the pro-health option. Independent variables: *Cost of one life* is the number of jobs lost necessary to reduce fatalities by one unit implicit in each choice. *Age* is semi-continuous (5 year intervals). All the other variables are dummies 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*, *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, 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.

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 $= 1.00$).

As mentioned at the beginning of this section, the fraction of participants that always choose to minimize fatalities is twice as large in the US and in the UK than it is in Italy. Country differences in Columns (1)-(4) may be driven by the choices of participants with such extreme preferences. More generally, these extreme preferences are connected to strong moral norms, such as the idea that no economic benefit is worth a human life. The factors associated with adhering to these norms may differ from those that affect the choices provided that one considers the possibility of not minimizing fatalities. Columns (5) and (6) in Table 2 inspect this possibility.

Columns (5) reports a linear probability model on the probability of *always* choosing to minimize fatalities. Column (6) restrict the analysis in Column (1) on the subsample of participants that do *not* minimize fatalities at least once. Looking at Column (5), we observe that the variables associated with extreme pro-health attitudes are generally the same associated with a higher probability of minimizing fatalities provided that sometimes one does not (Column (6)). One notable exception is that having experienced some deaths from Covid in one's close network makes respondents significantly more likely to always choose to minimize fatalities.

When we allow the explanatory variables to show a different impact by country using interaction terms two findings are worth mentioning (see the full set of interactions in Fig. A.1 in Appendix A). First, as shown in the right-hand panel of Fig. 5, 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 Fig. 5. 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 other two countries at all

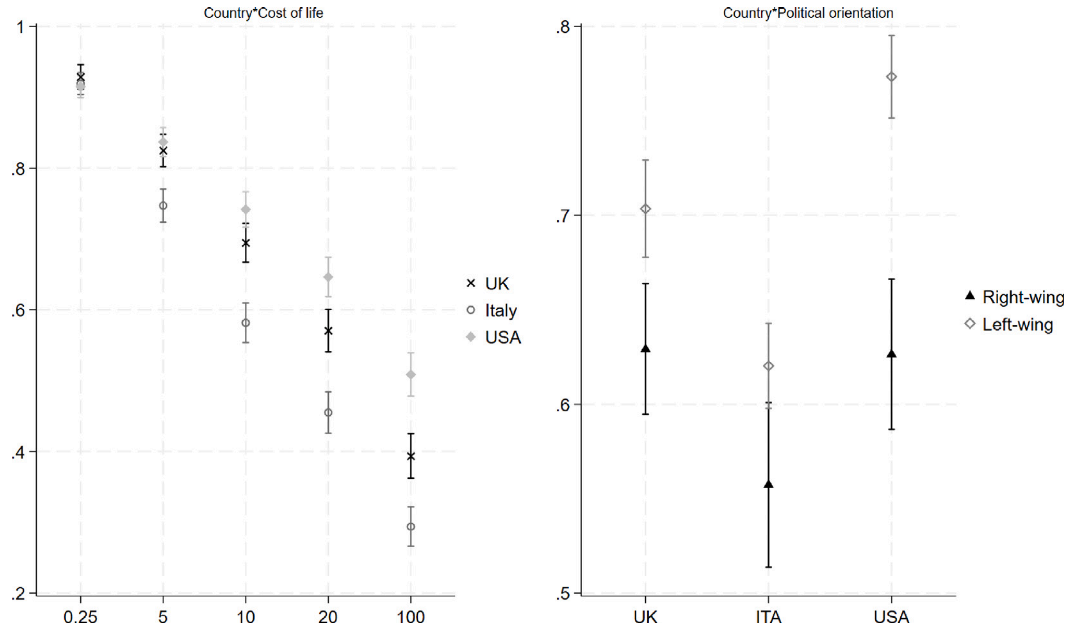

Fig. 5. Choices by country interacted with the cost of one life and political orientation
Notes: The left panel (*Country*Cost of life*) reports the fraction of pro-health choices (vertical axis) by country for different costs of saving one life (horizontal axis). Choices are indistinguishable by country 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.

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.

Several considerations suggest that our data remains valid and meaningful despite choices have been made in hypothetical scenarios:

1. Choices respond consistently to the cost of saving a life, suggesting that they are not primarily driven by social desirability;
2. As long as HB is more likely to manifest through extreme responses (e.g., always prioritizing fatalities) rather than through marginal shifts in trade-off weights, our findings in Table 2 are reassuring: The determinants behave similarly when restricting the sample to respondents who react to the trade-off;
3. There is evidence that some significant correlations with pro-health preferences are genuine rather than HB-driven. For example, left-wing ideology aligns with revealed preferences in policy choices, and age trends in the opposite direction from (Mjelde et al., 2012), suggesting that if HB exists, our estimates may represent a lower bound rather than an overstatement.

3.2. Within-country exploratory analyses on competing determinants

The evidence in the previous section does not support many of the expectations we had. In particular, contingent experiences with the pandemic do not seem to explain the choices of the respondents. Moreover, their preferences are not aligned with the strength of containment policies at the country level. In this section we present an exploratory analysis scrutinizing other potential factors that may possibly explain or qualify some of our results. In particular, we explore the association between the participants' choices and *(a)* the cost of the fundamentals on both the health and economic side, *(b)* the containment policies faced. As mentioned in the Introduction, containment measures may not only reflect preferences, but may influence them. Labor market conditions and the functioning of the healthcare system may shape the cost of the fundamentals and therefore influence the participants' choices. As these contextual factors cannot be distinguished from other determinants that vary at the country level, we proceed by leveraging within-country variations.

Different cost of the fundamentals. Labor market conditions differ both across and within countries. For instance, unemployment has been historically higher and more persistent in Italy, where also regional differences are particularly pronounced. Northern regions experience levels and duration of unemployment similar to those in the US and in the UK, and definitely lower than those in the Southern regions. It could then be that 'losing one job' is perceived as more costly where the cost of unemployment is higher (e.g. in the South of Italy). If this is the case, apparent differences in moral preferences could actually be driven by different labor market

Table 3
Linear probability model: within-country variation in policy and fundamentals.

	Dependent variable: Choice of the alternative involving fewer fatalities						
	(1) ITA+US	(2) ITA	(3) US	(4) US	(5) US	(6) US	(7) US
Unemployment	0.00253 (0.00226)	−0.0176 (0.0230)	0.000401 (0.00521)				0.000096 (0.00646)
Long-term unempl.		0.0242 (0.0295)					
Out-of-pocket spending				−0.0000237 (0.0000294)	−0.0000220 (0.0000294)		−0.0000302 (0.0000327)
Health spending					0.0000117 (0.00000887)		0.0000083 (0.0000099)
Stringency index						0.00320 (0.00294)	0.00269 (0.00337)
Str. index (last 90 days)						−0.000799 (0.00209)	−0.000434 (0.00216)
US	0.133*** (0.0144)						
Constant	0.541*** (0.0337)	0.704*** (0.0919)	0.606*** (0.0532)	0.615*** (0.0353)	0.521*** (0.0790)	0.480*** (0.112)	0.433*** (0.116)
Controls	YES	YES	YES	YES	YES	YES	YES
N	39 882	19 758	20 124	20 124	20 124	20 124	20 124

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 healthcare 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 1): Cost of one life, Female, Age ≥ 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.

conditions. We exploit within-country differences in labor market conditions both in Italy and in the US to test this hypothesis. Unfortunately, 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.

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 was explained by a higher cost of losing jobs, we should observe the unemployment rate to significantly correlate with a pro-health behavior of the respondents. As shown in Column (1)–(3) of Table 3, this is not the case. The coefficient of unemplyment is fairly close to zero and never significant (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, i.e. fatalities, since health costs can also differ. For instance, 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 than the other two countries. We can exploit the variability across US states in out-of-pocket health expenditure per-capita (source: US Census) and healthcare spending per-capita (source: Medicare and Medicaid Services) and test whether these variables correlate with a higher probability to minimize fatalities.[7] Also in this case we do not find any significant correlation (Table 3; 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. The preferences that we measure are not exogenous to the containment policies faced. In contrast, more or less stringent policies may have affected the choices in the dilemma. For instance, US respondents who have experienced a relatively

[7] Insufficient variability within the other countries and non-availability of data prevent us from extending the exercise to Italy and UK.

weaker health protection through milder containment measures may have become more concerned by the health dimension of the dilemma. 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 over time, regardless of the time window considered (e.g. throughout the pandemic rather than in the last 90 days prior to our survey).

To test the hypothesis that pro-health attitudes are recorded where containment measures are more stringent, we exploit the variability of the stringency index across states in the US. Unfortunately, this exercise cannot be extended to the other countries, 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. Italy's average index (67) is relatively close to the highest average index among US states (62). If more stringent policies affect the choice in the dilemma, we should expect respondents from the states characterized by more stringent policies to display a lower pro-health attitude. This is indeed not the case.

As shown in Column (6) of Table 3, 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.

The results above hold unchanged in a specification in which all these determinants are included at the same time (Column (7)). Overall, we do not find any evidence that the solution of the moral dilemma can be rationalized by the policies experienced during the pandemic, or by different relative costs of the two 'bads' presented in the choice.

3.3. Structural estimation of preferences

In this section we estimate with maximum likelihood the shape of the utility function for a representative respondent starting from all the choices in the task. 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 Eq. (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 α) and the inclination to diversify the outcomes along the two dimensions analyzed (the *curvature* parameter ρ). For a given value of α and ρ, Eq. (1) identifies the indifference curve corresponding to any level of $u(F,J)$.[8]

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:

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

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'. In fact, this case corresponds to convex preferences in a Cartesian plan like Fig. 3, in which the origin belongs to the upper contour set. 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.

Note that the distribution parameter α determines the slope of the indifference curves regardless of the value of ρ when $\overline{F}-F = \overline{J}-J$, i.e. along the 45-degree line where fatalities and job losses are exchanged on a 1:1 basis. Therefore, the absolute level of α is not scale invariant, although the different levels would represent the same preferences.

Preferences as specified in Eq. (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 (ε) 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. 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 sizably more than economic outcomes.[9] As an intuition, on the 45 degree line, where there are the same number of jobs lost and fatalities, the extra

[8] The functional form in Eq. (1) encompasses as special cases the most common examples of utility functions according to the values taken by α and ρ. Appendix B contains a detailed explanation of these functional forms as well as of some technicalities of the estimation procedure.

[9] This result refers to respondents who react to the relative cost of one life. 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.

Table 4
The structural parameters ρ and α with all the covariates.

	ρ		α	
	Coefficient	(SE)	Coefficient	(SE)
Italy	-0.140^{**}	(0.0573)	0.0256^{***}	(0.00755)
USA	0.157^{**}	(0.0700)	-0.0146^{**}	(0.00606)
Left-wing	0.0753	(0.0621)	-0.0208^{***}	(0.00690)
Risk averse	−0.0320	(0.0555)	−0.00271	(0.00554)
Patient	−0.00909	(0.0479)	−0.00276	(0.00465)
Age	0.00184	(0.00158)	-0.000387^{***}	(0.000135)
Female	−0.0426	(0.0446)	0.00155	(0.00429)
Self-employed	−0.0611	(0.0560)	0.00351	(0.00575)
Covid deaths	0.0506	(0.0801)	−0.00207	(0.00665)
Tested positive	0.0902	(0.0469)	-0.0116^{**}	(0.00517)
High death rate	0.0163	(0.0352)	0.000558	(0.00379)
Constant	1.169^{***}	(0.0649)	0.0740^{***}	(0.00742)
N	45000			

Notes: Coefficients represent the change in the structural parameters ρ and α associated to each independent variable. *Age* is semi-continuous (5-year intervals). All the other variables are dummies taking a value equal to one for a respondent that *(i)* lives in *Italy* or *USA* (UK is captured by the Constant); *(ii)* is *Female, Left-wing, Risk averse, Patient, Self-employed*; *(iii) Covid deaths* indicates that the respondent reports deaths due to Covid in his/her close network; *(iv) Tested positive* means that the respondent or someone in his/her close network tested positive for Covid-19; *(v) High death rate* captures respondents living in a region/state with a Covid death rate over 100,000 inhabitants above the sample median. Standard errors, clustered at the individual level, are reported in parentheses. * = p-value < 0.1; ** = p-value < 0.05; *** = p-value < 0.01.

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 $(1-\hat{\alpha})/\hat{\alpha} = 19$. If $\hat{\rho}$ was equal to one, the same would also hold at any other allocation.[10]

The curvature parameter is instead significantly larger than one: $\hat{\rho} = 1.21$ (two-sided t-test: p-value < 0.001).[11] 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 Fig. 3, 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. 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 Fig. 6.[12]

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 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.

Interacting the parameters with all the main explanatory variables at the same time delivers results that are qualitatively similar to those described in Section 3.1. Table 4 reports the coefficients.

4. Conclusion

Within our sample, when confronted with a trade-off between fatalities and jobs losses, individuals often prioritize saving more lives. 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.

[10] This shows that the weight parameter, even if standardized between 0 and 1, is not scale invariant.

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

[12] The results reported refer to aggregate estimates by country. The comparison across countries 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.

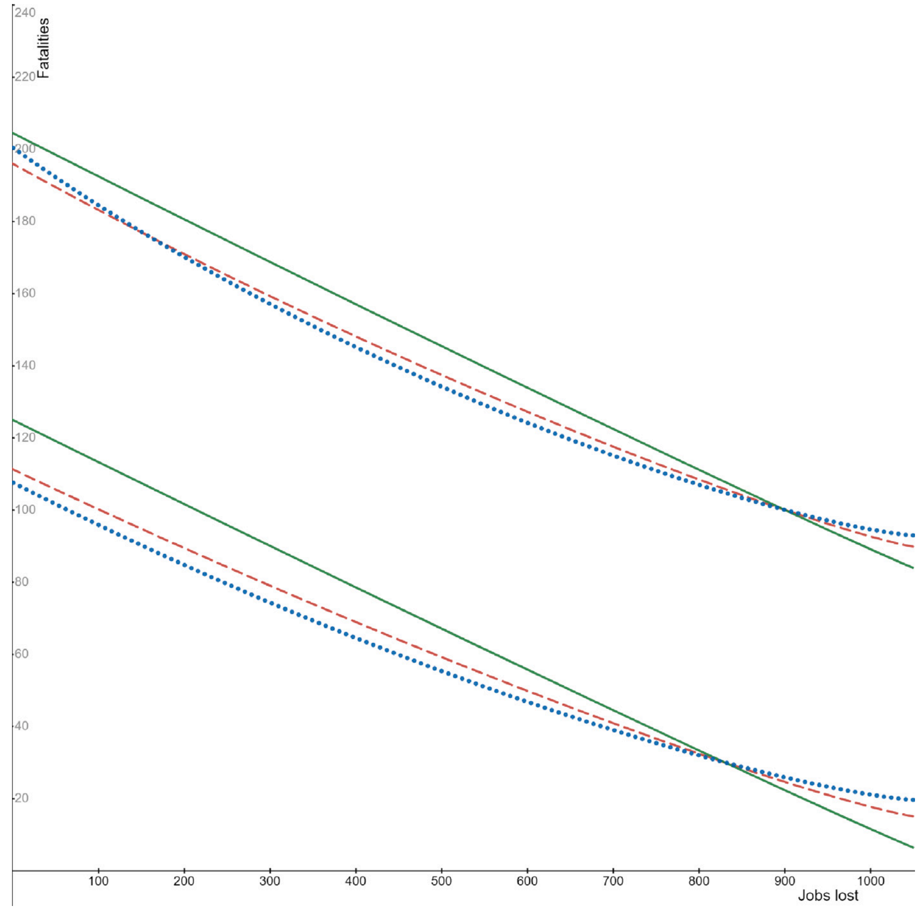

Fig. 6. Indifference Curves by country.
Solid line: Italy; Dashed line: UK; Dotted line: US.Each line represent an indifference curve. The space below each curve represents the set of combinations that are strictly preferred. Indifference curves that are closer to the origin represent preferred combination 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.

Decisions are influenced in a predictable way by an individual's age and attitudes, so that older, more left-wing and more risk-averse participants put a larger weight on fatalities. We do not find consistent effects 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.

Moreover, 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.

Different contextual factors vary across country and may influence preferences. The most directly connected to our problem are labor market conditions and healthcare costs. Furthermore, more stringent policies may have caused weaker pro-health attitudes. Leveraging on within-country variations we do not find evidence that any of these competing explanations is connected to the participants' choices. The country differences that we observe must then be related to unobservables, either among our samples or among the corresponding populations. A possible candidate are cultural differences regarding both the health-wealth tradeoff and the attitudes toward state intervention. While unfortunately our study is not equipped to identify the ultimate explanation, we believe it sets an interesting avenue for future research.

Finally, our 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, resulting in an aversion to diversify between the two dimension proposed. This contradicts a basic characteristic of economic preferences over '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.

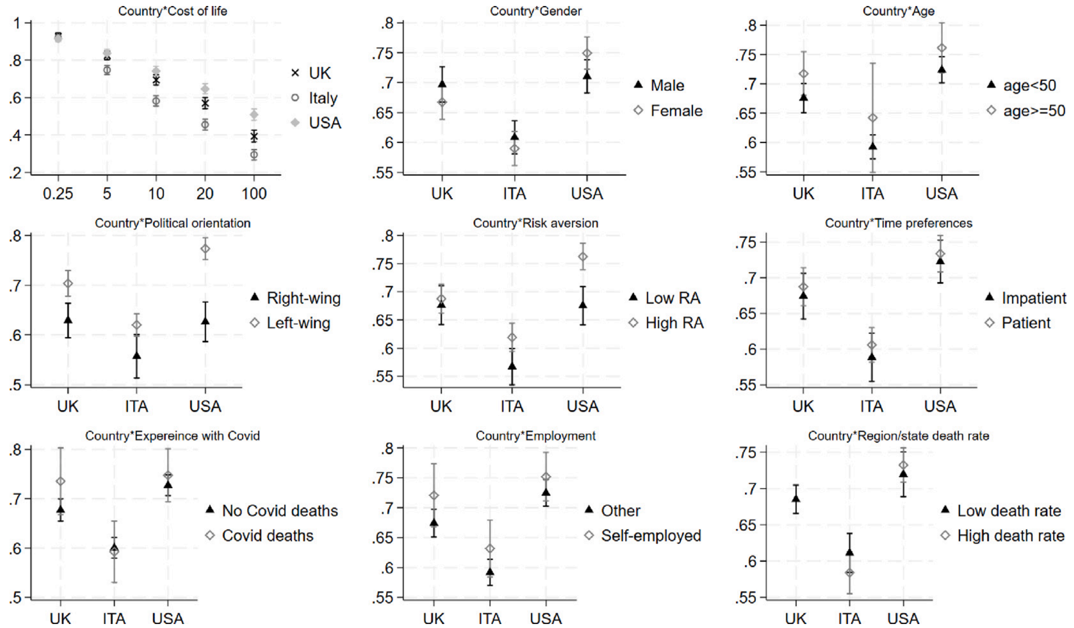

Fig. A.1. Pro-health choices by country interacted with each explanatory variable.
Notes: The figure reports the fraction of pro-health choices (vertical axis) by country for the explanatory variable specified in each panel.

Funding

A.F. was supported by the University of Milan. M.M. was supported by the University of Milan Bicocca, Italy through institutional grant 2021-ATE-0486. The research leading to these results has received funding from project 'PRIN 2022-NAZ-0601' funded by the Italian Ministry of University and Research.

Declaration of competing interest

There is no conflict of interest to disclose that could influence or bias our work I did not receive any financial support from interested third parties, and neither did a close relative or partner. I did not hold any paid or unpaid positions as officer, director, or board member of relevant non- profit organizations or profit-making entities held by me, or by a close relative or partner.

Acknowledgments

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.

Appendix A. Robustness

In Section 3.1 the multivariate analysis already partials out the effect of observable characteristics, but in this section we conduct an additional robustness check. Fig. A.1 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 right panel of Fig. A.1 shows that differences in the pro-health attitude differs by country also in each sub-sample.

Appendix B. Structural estimation

Utility function. Eq. (1) encompasses several classic utility functions according to the value taken by the parameters α and ρ.

$\alpha \in [0, 1]$ 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.

ρ 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 $(\overline{J} - J)$ and $(\overline{F} - F)$ that express a preference for diversification among the two 'bads'. In this case the parameter ρ 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 α.
- $\rho > 1$: 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 γ 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 $\varepsilon \sim \Lambda(0, \gamma)$.

Data availability

Data will be made available on request.

References

Alvarez, F., Argente, D., Lippi, F., 2020. A simple planning problem for COVID-19 lockdown, testing, and tracing. Am. Econ. Rev.: Insights 3 (3), 367–382.
Andersen, S., Harrison, G., Lau, M., Rutström, E., 2006. Elicitation using multiple price list formats. Exp. Econ. 9, 383–405. http://dx.doi.org/10.1007/s10683-006-7055-6.
Andersson, O., Campos-Mercade, P., Carlsson, F., Schneider, F., Wengström, E., 2021. The individual welfare costs of stay at home policies. Scand. J. Econ. 124 (2), 340–362.
Andre, P., Pizzinelli, C., Roth, C., Wohlfart, J., 2022. Subjective models of the macroeconomy: Evidence from experts and representative samples. Rev. Econ. Stud. 89 (6), 2958–2991.
Ashenfelter, O., 2006. Measuring the value of a statistical life: problems and prospects. Econ. J. 116 (510), C10–C23.
Belle, N., Cantarelli, P., 2022. Your money, your life, or your freedom? A discrete-choice experiment on trade-offs during a public health crisis. Public Adm. Rev. 82 (1), 59–68.
Carrieri, V., De Paola, M., Gioia, F., 2021. The health-economy trade-off during the Covid-19 pandemic: Communication matters. PloS One 16 (9), e0256103.
Charness, G., Chemaya, N., Trujano-Ochoa, D., 2023. Learning your own risk preferences. J. Risk Uncertain. 67 (1), 1–19.
Chorus, C., Sandorf, E.D., 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., 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.
Conrad, C., Enders, Z., Glas, A., 2022. The role of information and experience for households' inflation expectations. Eur. Econ. Rev. 143, 104015.
D'Acunto, F., Malmendier, U., Ospina, J., Weber, M., 2021. Exposure to grocery prices and inflation expectations. J. Political Econ. 129 (5), 1615–1639.
Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., Wagner, G.G., 2011. Individual risk attitudes: Measurement, determinants, and behavioral consequences. J. Eur. Econ. Assoc. 9 (3), 522–550.
Ehmke, M.D., Lusk, J.L., List, J.A., 2008. Is hypothetical bias a universal phenomenon? A multinational investigation. Land Econ. 84 (3), 489–500.
Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2021. The macroeconomics of epidemics. Rev. Financ. Stud. 34 (11), 5149–5187.
Ert, E., Haruvy, E., 2017. Revisiting risk aversion: Can risk preferences change with experience? Econom. Lett. 151, 91–95.
Favero, C.A., Ichino, A., Rustichini, A., 2020. Restarting the economy while saving lives under Covid-19. Tech. rep, CEPR Discussion Paper No. DP14664.

Filippin, A., Crosetto, P., 2016. A reconsideration of gender differences in risk attitudes. Manag. Sci. 62 (11), 3138–3160.
Foot, P., 1967. The problem of abortion and the doctrine of the double effect. Oxf. Rev. 5.
Graeber, D., Kritikos, A.S., Seebauer, J., 2021. COVID-19: a crisis of the female self-employed. J. Popul. Econ. 34, 1141–1187.
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). Nat. Hum. Behav. 5 (4), 529–538.
Hey, J.D., Orme, C., 1994. Investigating generalizations of expected utility theory using experimental data. Econometrica 62 (6), 1291–1326.
Hoffmann, A.O., Post, T., 2017. How return and risk experiences shape investor beliefs and preferences. Account. Financ. 57 (3), 759–788.
Jenni, K., Loewenstein, G., 1997. Explaining the identifiable victim effect. J. Risk Uncertain. 14 (3), 235–257.
Jones, C., Philippon, T., Venkateswaran, V., 2021. Optimal mitigation policies in a pandemic: Social distancing and working from home. Rev. Financ. Stud. 34 (11), 5188–5223.
Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. Econometrica 47 (2), 263–291.
Kang, S.-G., Powell, Jr., G.B., 2010. Representation and policy responsiveness: the median voter, election rules, and redistributive welfare spending. J. Politics 72 (4), 1014–1028.
Kaplan, G., Moll, B., 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., Feeley, T.H., 2016. The identifiable victim effect: A meta-analytic review. Soc. Influ. 11 (3), 199–215.
León, G., Miguel, E., 2017. Risky transportation choices and the value of a statistical life. Am. Econ. J.: Appl. Econ. 9 (1), 202–228.
Lesschaeve, C., GlaurdiĆ, J., Mochtak, M., 2021. Health versus wealth during the Covid-19 pandemic: Saving lives or saving the economy? Public Opin. Q. 85 (3), 808–835.
Loría-Rebolledo, L.E., Ryan, M., Watson, V., Genie, M.G., Sakowsky, R.A., Powell, D., 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.
Malmendier, U., 2021. Exposure, experience, and expertise: Why personal histories matter in economics. J. Eur. Econ. Assoc. 19 (6), 2857–2894.
Malmendier, U., Nagel, S., 2016. Learning from inflation experiences. Q. J. Econ. 131 (1), 53–87.
Malmendier, U., Shen, L.S., 2024. Scarred consumption. Am. Econ. J.: Macroecon. 16 (1), 322–355.
Manipis, K., Street, D., Cronin, P., Viney, R., 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 Res. 14 (3), 359–371.
Mengel, F., Tsakas, E., Vostroknutov, A., 2016. Past experience of uncertainty affects risk aversion. Exp. Econ. 19, 151–176.
Mindes, S.C., Lewin, P., 2021. Self-employment through the COVID-19 pandemic: An analysis of linked monthly CPS data. J. Bus. Ventur. Insights 16, e00280. http://dx.doi.org/10.1016/j.jbvi.2021.e00280.
Mjelde, J.W., Jin, Y.H., Lee, C.-K., Kim, T.-K., Han, S.-Y., 2012. Development of a bias ratio to examine factors influencing hypothetical bias. J. Environ. Manag. 95 (1), 39–48.
Oana, I.-E., Pellegata, A., Wang, C., 2021. A cure worse than the disease? Exploring the health-economy trade-off during COVID-19. West Eur. Politics 44 (5–6), 1232–1257.
Page, B.I., Shapiro, R.Y., 1983. Effects of public opinion on policy. Am. Political Sci. Rev. 77 (1), 175–190.
Palan, S., Schitter, C., 2018. Prolific.ac—a subject pool for online experiments. J. Behav. Exp. Financ. 17, 22–27.
Penn, J., Hu, W., Ye, T., 2024. Efficacy of hypothetical bias mitigation techniques: A cross-country comparison. J. Environ. Econ. Manag. 125, 102989.
Reynolds, A., 2013. Representation and rights: The impact of LGBT legislators in comparative perspective. Am. Political Sci. Rev. 107 (2), 259–274.
Słoczyński, T., Wooldridge, J.M., 2018. A general double robustness result for estimating average treatment effects. Econometric Theory 34 (1), 112–133.
Small, D.A., Loewenstein, G., 2003. Helping a victim or helping the victim: Altruism and identifiability. J. Risk Uncertain. 26 (1), 5–16.
Spencer, A., Rivero-Arias, O., Wong, R., Tsuchiya, A., Bleichrodt, H., Edwards, R.T., Norman, R., Lloyd, A., Clarke, P., 2022. The QALY at 50: One story many voices. Soc. Sci. Med. 296, 114653.
Stott, H., 2006. Cumulative prospect theory's functional menagerie. J. Risk Uncertain. 32 (2), 101–130.
Tetlock, P.E., 2003. Thinking the unthinkable: Sacred values and taboo cognitions. Trends Cogn. Sci. 7 (7), 320–324.
Viscusi, W.K., Aldy, J.E., 2003. The value of a statistical life: a critical review of market estimates throughout the world. J. Risk Uncertain. 27 (1), 5–76.