Are distributional preferences for safety stable? A longitudinal analysis before and after the COVID-19 outbreak

Abstract

Policy makers aim to respect public preferences when making trade-offs between policies, yet most estimates of the value of safety neglect individuals' preferences over how safety is distributed. Incorporating these preferences into policy first requires measuring them. Arroyos-Calvera et al. (2019) documented that people cared most about efficiency, but that equity followed closely, and self-interest mattered too, but not enough to override preferences for efficiency and equity. Early 2020 saw the outbreak of the COVID-19 pandemic. This event would impose major changes in how people perceived and experienced risk to life, creating an opportunity to test whether safety-related preferences are stable and robust to important contextual changes. Further developing Arroyos-Calvera et al.'s methodology and re-inviting an international general population sample of participants that had taken part in pre-pandemic online surveys in 2017 and 2018, we collected an April 2020 wave of the survey and showed that overall preferences for efficiency, equity and self-interest were remarkably stable before and after the pandemic outbreak. We hope this offers policy makers reassurance that once these preferences have been elicited from a representative sample of the population, they need not be re-estimated after important contextual changes.

Keywords: efficiency, equity, self-interest, risk, preference stability

Introduction

Improving people's safety and preventing injury, illness, and death is an important goal of many government policies. Measures to reduce the risk of such negative outcomes should be cost-effective,

raising the issue of how to value health and safety. When priority setting, policymakers must also navigate difficult trade-offs between efficiency (maximizing the total health benefit from interventions) and distributional or equity issues (minimizing differences amongst population groups). Policymakers' allocation decisions, and the values underpinning them, ought to reflect the preferences of the general public so it is essential to understand how the public prioritises these principles.

Evidence from the experimental laboratory shows that measures of preferences parameters outside of the health and safety context are likely to depend in part on the elicitation context (e.g., Chuang and Schechter, 2015). If public preferences over health and safety outcomes are similarly influenced by changing contexts, it may be necessary to conduct context-specific valuation exercises, whereas if preferences are relatively stable then policy values may be portable across contexts. As such, the stability of preferences regarding health and safety is a matter of significant importance for policymaking. However, no studies have addressed whether stated preferences for health and safety outcomes are context-dependent, and specifically whether they are impacted by external shocks. To investigate whether these preferences are stable, we make use of the outbreak of the COVID-19 pandemic, which provides an external shock influencing the context in which health and safety allocation decisions are made.

We report a longitudinal study comparing preferences for efficiency, distributional concerns and self-interest elicited in the first wave of the COVID-19 pandemic (i.e., April 2020) with preferences elicited from the same sample well beforehand (i.e., January-February 2017 or September 2018). This longitudinal approach allows a direct, like-for-like comparison of their consciously expressed attitudes to competing principles, as well as the attitudes implicit in their choices between policy options. We found that, overall, preferences were remarkably stable. We contribute within-subjects evidence of the stability of preferences in the health and safety context, aimed to reassure policy-makers that re-estimation of preferences after major external shocks may not be required.

Background

Trade-offs between efficiency and distributional preferences

Prioritizing safety policies on the basis of efficiency assumes that the policy objective should be to reduce risks as much as possible, generating the greatest gain in terms of lives saved or illness cases prevented. A growing body of evidence questions this assumption both normatively and empirically (e.g., Arroyos-Calvera, Covey, Loomes and McDonald, 2019; Arrow et al, 1996; Bobinac et al., 2012; Ubel et al., 1996; 2000). Distributional principles also matter, with people often expressing a preference for an equal distribution of policy benefits (Ubel et al. (1996, 2000). Overall, basing rationing decisions solely on maximization may fail to respect public opinion.

Efficiency-equity trade-offs were also demonstrated in a policy-framed experiment (Arroyos-Calvera et al., 2019). Participants chose between options that differed in efficiency (the expected number of fatalities or cases of ill health prevented) and equity (how evenly the prevented fatalities or cases of ill health were distributed across two sections of the population). They found an overall preference for more efficient policies, although people were willing to sacrifice some efficiency to reduce risks across both population groups.

However, preferences for equality in the distribution of benefits might only apply to contexts where everyone's needs are equal, and hence where equal distributions of benefits do not lead to unequal distributions of outcomes. People are known to be averse to inequalities in health and many people are unwilling to compromise outcome equality for gain equality (Tsuchiya & Dolan, 2009). Lindholm et al. (1998) found that people prioritized the allocation of safety resources on the basis of need and helping those most at risk.

Another attribute that has received some (albeit rather limited) attention in research on the general public's preferences is the extent to which people will forgo benefits to themselves for the benefit of others. Most of this research has been conducted in laboratory settings using economic games such as the dictator game (see 2011 meta-review by Engel). These experiments demonstrate that self-

interest alone cannot predict most people's behaviour. People are averse to inequities and desire fair distributions of payoffs between themselves and other agents (e.g., Camerer, 1997). Whether insights from experimental economics apply to real-world decision-making is an open question, although the limited research that exists suggests that people are also not entirely selfish when distributing health and safety resources. For example, Arroyos-Calvera et al. (2019) found that although people preferred policies that benefited them personally, their concern for self-interest did not override their concerns for efficiency or equity.

Stability of Distributional Preferences

This research investigates the stability of trade-offs between the four competing concerns (i.e., efficiency, equity, helping those most at risk, and self-interest), using the outbreak of COVID-19 as an external shock to test this stability. Unlike other research conducted in response to COVID-19, this project does not examine specific interventions such as lockdown measures e.g., Genie et al. (2020); Blayac et al. (2020). Its focus is on general allocation principles and specifically whether the weight of each principle in guiding policy preferences changed in the wake of the pandemic. In what follows, we discuss possible reasons why the relative importance of the principles might change in this context. We do so to highlight that change is plausible. If preferences did change, policy approaches may need to adapt to take account of public preferences in this and future pandemics and global crises.

Efficiency

During the first COVID-19 wave, in many countries such as the UK and USA the central message around implementing restrictions was to "save lives". For example, the slogan "stay home, protect the NHS, save lives" was used in daily UK government press conferences through March – May 2020. Coming from the government, the message that we ought to save lives (i.e., focus on efficiency) may have garnered increased support – as Dolan and colleagues' (2012) *messenger effect* would suggest. In the area of consumer choice, individuals are known to attach disproportionate weight to the most salient attributes of the choice (Bordalo et al., 2013). We could thereby expect that in cases where

the number of (avoided) fatalities is salient, people would prefer policies offering the highest efficiency.

Equity

In relation to equity, COVID-19 has highlighted issues relating to the uneven distribution of resources both nationally and globally. Inequalities driven by economics, gender, ethnicity and professional influence were emerging in the provision of personal protective equipment (PPE) (Shelton et al., 2021). Emerging divisions between the PPE 'haves' and 'have nots' may have been salient enough considerations during the early part of the first wave of the pandemic to increase weight on equity principles.

Self-interest versus helping others (Most at Risk)

Self-preservation might also have been prioritised, particularly in the first wave when there was considerable uncertainty about how the pandemic might progress and people's fears for their physical well-being were raised (Jungmann & Witthöft, 2020). Many people may have overestimated the probability that they would catch or die from COVID-19 (Attema et al., 2021), which might have led them to focus more on self-interest over helping others. Early in the pandemic, people's fears about shortages of basic supplies led many towards selfishly motivated hoarding behaviours (Baddeley, 2020).

Somewhat paradoxically, the unknown and potentially catastrophic consequences of COVID-19 during the study period could have contributed to an altruistic rather than egoistic perspective. According to terror management theory, mortality salience should intensify the desire to adhere to cultural norms behaving prosocially (Pyszczynski et al., 2021; Solomon et al., 1991). By acting altruistically and generously, people can bolster the belief that they are valuable contributors to society (Jonas et al., 2002). Recent research by Grimalda et al. (2021) is consistent with these theories, finding that direct personal exposure to COVID-19 was associated with greater charitable giving to pandemic relief efforts. Luo et al. (2021) found that as the pandemic worsened during the spring of 2020 people

made more altruistic choices in a dictator game. An altruistic rather than egoistic perspective could also have been bolstered by much of the messaging around COVID-19 policy and prevention behaviours which in the first wave targeted cooperative/ prosocial ("don't spread it") rather than self-interested ("don't get it") motivations, as exemplified by the aforementioned UK government advice to "stay home, protect the NHS, save lives".   COVID-19 policy decisions like the "stay home" requirement relied on many people incurring personal costs for the benefit of others in society. We might therefore expect the experience of self-sacrifice during COVID-19 to increase the weight placed on helping others, particularly those at higher risk.

Taken together, this evidence provides our main hypothesis: that the outbreak of COVID-19 changed people's preferences and trade-offs between the four principles. The importance placed on saving lives in COVID-19 policy and messaging would be expected to increase weight on efficiency. The weight on self-interest could increase or decrease – more weight through self-preservation motives induced by COVID-19 raising fears for personal well-being; less weight through mortality salience promoting the value of prosocial behaviour. We might expect increased weight on helping those most at risk in response to COVID-19 policies aimed at protecting high risk groups.

In the sections that follow we provide details about our study design and data collection, report our findings, and finish with a discussion.

Methods

We compare responses collected prior to the COVID-19 pandemic to responses collected at the height of the first wave to ascertain whether trade-offs between efficiency, equity, helping yourself, and helping those most at risk changed. There were two pre-pandemic data collection phases: phase 1 (January and February 2017), reported in detail in Arroyos-Calvera et al, 2019), and phase 2 (September 2018). Participants from the pre-pandemic phases were invited back to complete phase 3 which took place during the pandemic (April 2020).

There were two main tasks: the principles task and the preference task, described in detail below. Phase 1 participants completed the principles task, plus a version of the preference task with a slightly different design and purpose (reported in full in Arroyos-Calvera et al., 2019 and not included in this paper). Phase 2 participants completed the preference task but not the principles task, and phase 3 participants did both tasks.

We refer to the phase 3 participants who had completed the principles task before (phase 1) as Principles Returners; and to those who completed the preference task before (phase 2) as Preferences Returners.

This resulted in a sample of participants for each task with within-respondent repeated observations before and during the pandemic. Table 1 provides a summary of the phases and tasks.

Table 1. Phases and tasks

Phase	Date	Pandemic Outbreak	Task	Relevant Participants
1	January/February 2017	Before	Principles	Principles returners
2	September 2018	Before	Preferences	Preferences returners
3	April 2020	After	Both	Both returner subsamples

Scenario and treatments

The methods we employ closely follow those reported in Arroyos-Calvera et al. (2019). We described a city with two zones, East and West, each home to 100,000 people. Residents in each zone faced a risk of dying due to bacterial contamination of their water supply, and different policy options that would reduce the risk to residents of each zone by different amounts were described.

In the scenario, the bacteria in the water led to different baseline risks of dying in the East and West zones: 24 in 100,000 for the lower risk zone and 36 in 100,000 for the higher risk zone. Between-participants, we varied whether the East zone or the West zone had the higher initial risk. This is a key difference compared to the methods in Arroyos-Calvera et al. (2019), where the effect of the concern for those most at risk and those getting the larger risk reduction were conflated because the larger

risk reduction was always given to those most at risk. This variation allowed us to identify whether people prioritise those most at risk, all else held equal. It also avoids any confounding with pre-conceptions of differences between East and West in real cities, for instance, on the basis of wealth. In all other respects, the two zones were identical.

Half of participants were informed that they lived in the East zone (the *self-interest* condition), the other half were not told that they lived in the city (the *impartial* condition). The impartial participants would not personally benefit more from risk reductions in one zone compared to the other, whereas those who were told they lived in the East zone would benefit from East Zone risk reductions. This variation allows us to explore the influence of self-interest.

Finally, through the timing of the data collection, we have within-participants variation in whether the preference elicitation occurred before or during the COVID-19 pandemic.

The four policy options are summarised in Table 2. In expectation, policy EO (East Only) would save 10 lives, all in the East zone; Policy WO (West Only) would save 20 lives, all in the West zone; Policy BE (Both Equal) saves 8 lives in each zone (16 lives in total); and policy BD (Both Different) saves 7 lives in total, 3 in the East and 4 in the West. Participants saw shorter policy names (i.e., "E" instead of "EO", "W" instead of "WO", "B" instead of "BE", and "X" instead of "BD") which, to aid recall, were always associated with a specific colour.

The policies are designed to differ along four attributes: efficiency, equity, how much they benefit the decider (benefit to self), and how much they benefit those most at risk. The efficiency of a policy is defined simply as the total number of lives it would save. Policy WO is the most efficient (20 lives), and policy BD is the least efficient (7 lives). Equity is defined as the difference between the number of lives saved in the zone where fewer lives are saved, and in the zone where more lives are saved. Policy BE scores 0 for equity since it is completely equitable by this definition (i.e., 8 lives saved in each zone), and all other policies score negatively for equity, with policy WO the least equitable scoring -20 (i.e., 20 lives saved in the West zone, 0 lives saved in the East zone). The scoring for the

remaining two attributes depends upon the between-participant condition or treatment: in the *self-interest* condition, *benefit to self* records the number of lives saved in the East zone, but in the *impartial* condition the *benefit to self* is always 0 for all policies. When the East zone faces a higher baseline risk (*higher risk East* condition), the score for *benefit to most at risk* records the number of lives saved in the East zone, but when the West zone faces a higher baseline risk (*higher risk West condition*), the score for *benefit to most at risk* records the lives saved in the West zone instead. Table 2 includes the scores for each policy on each attribute.

Table 2. Policy options and their scores on the efficiency, equity, benefit to self, and benefit to most at risk attributes

				Attribute (*Label*)					
		Lives saved in East	Lives saved in West	Efficiency (*Effic*)	Equity (*Equity*)	Benefit to self (*BenSelf*)		Benefit to most at risk (*BenMAR*)	
				All conditions		Self-interest condition	Impartial condition	Higher risk East condition	Higher risk West condition
Policy	EO	10	0	10	-10	10	0	10	0
	WO	0	20	20	-20	0	0	0	20
	BE	8	8	16	0	8	0	8	8
	BD	3	4	7	-1	3	0	3	4

Tasks

The experiment consisted of four types of tasks: the rating/comprehension tasks (in all phases), principles task (in phases 1 and 3); preference task (in phases 2 and 3); and personal characteristics questions (all phases). Screenshots of the interface of phase 3, which included all of the tasks, are provided in Online Appendix A1.

Rating and comprehension

The rating and comprehension tasks introduced participants to the policy scenarios. Each scenario was presented as a paragraph describing the policy, plus a condensed summary setting out the baseline risk, risk reduction and resulting risk in each zone, plus a residency reminder where applicable. Figure 1 shows an example.

Figure 1. Example of policy description: Option BE in the *self-interest - higher risk West zone* condition

Some bacteria can be eliminated from the water supply in the East zone and the West zone. 16 people's lives would therefore be saved over the next year – 8 lives from each zone. This means that those who live in the East zone (like you do) would have their risk of dying from the bacteria in the water supply reduced from 24 in 100,000 to 16 in 100,000. Those who live in the West zone would have their risk of dying from the bacteria in the water supply reduced from 36 in 100,000 to 28 in 100,000.

East Zone (where you live)
– 8 lives saved
– Risk reduced from 24 in 100,000 to 16 in 100,000
West Zone
– 8 lives saved
– Risk reduced from 36 in 100,000 to 28 in 100,000

Participants rated each policy from "very poor" to "very good". The comprehension questions presented participants with pairs of policies. They had to identify the policy in each comparison that fitted a particular criterion. For example, we asked which of policy EO vs policy WO "would make the water safer for the people who [were] most at risk from the bacteria". The policy descriptions remained on the screen while participants deliberated. We provided instant, tailored feedback. A table with the comprehension questions and the correct answer by condition can be found in Online Appendix A2. The rating and comprehension tasks were included to ensure each policy had been individually considered, and that all participants understood the key attributes of each policy.

Principles task

The principles task posed direct trade-offs between the four concerns of interest: efficiency, equity, benefit to self, and helping those most at risk. It was completed by participants in phase 1 (reported in Arroyos-Calvera et al., 2019) and phase 3.

The principles questions in phase 3 were the same as in phase 1, with one exception. The phase 1 question that delved deeper into equity concerns (by asking participants to trade off the belief that a policy should "… make the water a little safer for everyone" versus the belief that it should "… make

the water a lot safer for only some of the people in the city") was substituted for a question that completed the list of all possible pairwise comparisons: Benefit to self vs. Equity.

Each question presented two statements, representing opposite opinions about the relative importance of two concerns. For example, the question trading off helping those most at risk against efficiency elicited the degree to which participants favoured statement A "The chosen option should make the water safer for the people who are most at risk from the bacteria even if that means that fewer lives are saved" over statement B "The chosen option should save the most lives even if that means that the water is not made safer for the people who are most at risk from the bacteria", or B over A. A table showing the six questions that arise from pairing each of the four concerns can be found in Online Appendix A3. Responses were elicited on a seven-point scale with points labelled as "Strongly favour A over B", "Moderately favour A over B", "Slightly favour A over B", "Equally favour A over B", "Slightly favour B over A", etc. This pairwise presentation was designed to ensure participants engaged with the trade-offs between the principles, as opposed to stating agreement with all of them. To control for potential bias towards the statement presented on the left or right the assignment of each statement to the labels statement A or B was randomised between participants.

Preference task

The preference task involved all six pairwise comparisons of policies EO ("East Only"), WO ("West Only"), BE ("Both Equal") and BD ("Both Different"). It was completed by respondents in phases 2 and 3.

Participants were told that resource constraints meant the government could not implement all the policies to treat the bacteria in the water, so they would have to choose between them. In each question, participants used a seven-point scale to indicate the degree to which they would prefer for the government to choose one option out of the pair.

For example, in the comparison between policies EO and BE, the scale points were "Strongly prefer option EO, Moderately prefer option EO, Slightly prefer option EO, Equally prefer options EO and BE,

11

slightly prefer option BE" etc. For each comparison, we displayed the short form description of the policies, with the option to revisit the longer description if desired. An example preference task is shown in Figure 2. To control for potential bias towards the option presented on the left or right of the screen, the assignment of each option to the sides of the screen (and the corresponding left and right ends of the rating scale) was randomised for each choice. To control for potential order effects, the presentation order of the questions within each set of tasks was also randomised at the participant level.

Figure 2. Example of preference task question

Option X	Option E
East Zone (where you live) – 3 lives saved – Risk reduced from 24 in 100,000 to 21 in 100,000 West Zone – 4 lives saved – Risk reduced from 36 in 100,000 to 32 in 100,000	East Zone (where you live) – 10 lives saved – Risk reduced from 24 in 100,000 to 14 in 100,000 West Zone – No lives saved – Risk unchanged at 36 in 100,000

VIEW FULL DESCRIPTION BELOW
○

Which option would you prefer the government to choose?

Strongly prefer option X	Moderately prefer option X	Slightly prefer option X	Equally prefer option X and option E	Slightly prefer option E	Moderately prefer option E	Strongly prefer option E
○	○	○	○	○	○	○

Note: Option X corresponds to BD "Both Different" in this paper, Option E corresponds to EO "East Only". To aid recall, each option's name on the heading of the table was always presented using the same colour.

Estimation

Our aim is to understand how the COVID-19 pandemic influenced the importance of the four attributes, both in terms of choice and expressed principles.

Our analysis of the preference task involves within-participant modelling of choices between the policies across phases 2 and 3, involving the Preferences returner subsample. We employ an alternative specific conditional logit (McFadden's choice) model implemented with the *asclogit* command in Stata 17. The regression assumes a random utility model (McFadden, 1973); i.e., that out of each pair of alternatives, individuals select the policy option that gives them the highest utility, subject to some random error in their decision making.

We assume a latent continuous variable P, capturing preference for the alternative (equation 1):

$$P = \beta_0 + Attributes\ \beta + \alpha\ Pandemic + (Pandemic \times Attributes)\ \gamma + Demographics\ \delta + Treatments\ \theta \tag{1}$$

To estimate the model, the dependent variable (*Chosen*) was coded 1 for the chosen alternative and 0 for the non-chosen alternative in each choice. Cases where participants stated indifference between the two alternatives (3-5% of the total) were dropped from the model. Each alternative was defined by a set of *attributes* which varied as part of the experimental design and summarised in Table 2: scores measuring each alternative's efficiency (*Effic*), equity (*Equity*), benefit to those most at risk (*BenMAR*), and benefit to self (*BenSelf*).

The model also captures the effect of the predictors that vary by individual on the odds of choosing the alternative to the base alternative, which in our model was the option that had been randomly assigned to be presented on the left-hand side of the screen.

Alternative specific conditional logit models use data that vary at an individual level but not at an alternative level to explain the choice of each of the alternatives presented. In our setup, there are 4 alternatives overall (policies EO, WO, BE and BD), but participants encounter only two in every question, so rather than predicting the choice of one of the four alternatives, we can only predict the choice of one of the two presented. Because of this, we cannot predict the choice of, for example, Policy EO, but we need to find a way to distinguish the alternatives within each pair. One way is to classify the alternatives that we present according to an alternative-specific factor such as whether it

is the most efficient in the pair. However, all alternative-specific factors are being used as explanatory variables in the regression. Therefore we use a classification criterion that is independent of the alternative attribute, whether the option was presented on the left (vs. the right), which was determined at random.

The case-specific predictors (i.e., *Demographics, Treatments, Pandemic)* can therefore be considered control variables and the coefficients are of no particular interest aside from providing insight into whether choices favoured options presented on the left or right of the screen. The *Demographics* included participants' age and gender. The *Treatments* included dummy variables for self-interest (vs. impartial) perspective (*selfint*); higher baseline risk presented in the West (vs. East) zone of the city (*highriskWest*) and the *Pandemic* variable denoting whether the data was collected in the pre-pandemic (Pandemic = 0) or pandemic wave (Pandemic = 1).

To explore how the importance of attributes on choice might have changed after the start of the pandemic we tested two-way interactions between the *Pandemic* case-specific variable and each of the *Attributes* alternative-specific variables.

Our analysis of the principles task involves between- and within-participant comparisons of the distributions of responses for each principles trade-off shown in Table 2. This analysis compares phase 1 and phase 3 responses, involving the Principles returners subsample. We present summary statistics for each pair of principles, with a negative (positive) score indicating preference for the first (second) principle in the pair, and with the absolute value indicating the strength of preference. We conduct Kolmogorov-Smirnoff (KS) tests comparing the distributions of answers. We also compute the fraction of occasions on which participants favoured each principle over its alternative in the principle trade-off questions and, by interacting these scores with the matching attributes in regression analyses, we test whether favouring a principle significantly increases the odds of choosing the policy alternative that offers an increase in the matching attribute.

Results

Participants

The studies were distributed using Prolific, a participant recruitment website, and completed online. The median participant took 13 to 15.5 minutes to finish (depending on the phase) and was paid £2.50 (in phase 1) or £2 (in phases 2 and 3).

Table 3 shows the sample sizes and the age and gender composition. The Returners subsamples consist of participants that took part in two of the data collection phases, one before and one during the pandemic. Participants in these samples were members of the general population. In phase 1, participation was restricted to non-students residing in the UK; in phase 2 no restrictions were imposed. The minimum number of preference returners required in phase 3 to test the two-way interactions between the Pandemic case-specific variable and the Attributes alternative-specific variables was 167 based on the approach proposed by Orme (1998). As shown in Table 3 the sample of preferences returners (phases 2 & 3) exceeded this minimum. Table 3 also shows the percentage of comprehension questions that participants successfully answered. We consider this the lower bound of participant understanding and attention because we provided comprehensive feedback.

Table 3. Sample sizes and demographics

Sample	N	Age Mean (SD)	Female N (%)	Correct Test Questions
Full sample (phase 1)[a]	161	37.83 (11.07)	74 (46%)	87%
Full sample (phase 2)[b]	732	32.18 (11.55)[c]	357 (48%)[c]	83%
Full sample (phase 3)	391	39.26 (12.67)	192 (49%)	86%
Principle returners subsample (phases 1 & 3)	81	42.72 (12.45)	36 (44%)	91%[d]
Preference returners subsample (phases 2 & 3)	226	35.03 (12.30)	116 (51%)	89%[d]

Notes: [a] In phase 1 two versions of the policy options were administered on a between-participant basis – one that focussed on mortality risks and another on morbidity risks. Only participants allocated to the mortality risk condition were included. [b] In phase 2 two versions of the policy options were administered – one that focussed on mortality risks and another on a non-health risk. Only participants allocated to the mortality risk condition were included. [c]

Three participants did not disclose their age and two did not disclose their gender therefore the statistics for these characteristics exclude them. [d] These were the same in both phases, pre- and post-pandemic outbreak.

Did preferences change?

In this section, we explore how the effect on policy choices of concerns for efficiency, equity, helping those most at risk, and helping oneself changed from a pre-pandemic to a pandemic context. In short, preferences did not change significantly.

The first set of models in Table 4 showcases the impact of the *attributes* on choice in the pre-pandemic phase of the study: model 1 with the whole sample of respondents from phase 2, and model 2 with the subsample of respondents that also completed the study during the pandemic (i.e., Preference returners subsample). We feature both the Whole and Returner sample models to offer reassurance that attrition did not substantially affect our results. The second set of models are for the pandemic wave: model 3 for the whole sample from phase 3 and model 4 for the subset of respondents that had also taken part in the pre-pandemic wave (i.e., Preference returners subsample). In the third set of models, we compare respondents who took part in both phases before and during the pandemic (i.e., Preference returners subsample), and present two different specifications: model 5 uses the *Attributes* as alternative-specific explanatory variables, and model 6 adds two-way interactions between the *Attributes* and the *Pandemic* dummy. All models include case-specific control variables (i.e., *Demographics* and *Treatment* dummies), as described in the Estimation section. Full output is available in Online Appendix A4.

The top part of Table 4 shows the effect on choice of the alternative-specific predictors that change across policy options. For example, the odds-ratio (OR) for *Effic* shows how the odds of choosing an option are increased (OR > 1) or reduced (OR < 1) for a 1-unit difference in efficiency across the two policy options in each question. The bottom part of the table reports the effect of the case-specific predictors that do not vary across policy options but vary across individuals. An example is the *Pandemic* dummy, which indicates whether the response was collected before the pandemic (*Pandemic* = 0) or during (*Pandemic* = 1). As explained, the base alternative was set as the option that

was presented on the left-hand side of the screen (which was randomly assigned for each choice). The relative-risk ratio (RR) for *Pandemic* therefore tells us how the probability of choosing the option on the left-hand side of the screen is increased (RR > 1) or reduced (RR < 1) after the start of the pandemic. The constant reflects the baseline probability of choosing the alternative to the base alternative.

Table 4. Regression output on pre-pandemic and pandemic concerns and choice.

	BEFORE PANDEMIC OUTBREAK (Phase 2)		AFTER PANDEMIC OUTBREAK (Phase 3)[a]		COMPARISON (Phase 2 vs. 3)	
	Whole sample	Returner subsample	Whole sample	Returner subsample	Returner subsample	
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Alternative-specific variables CHOOSING ANY OPTION						
Effic	1.242***	1.258***	1.289***	1.310***	1.281***	1.259***
	(0.00974)	(0.0181)	(0.0151)	(0.0214)	(0.0145)	(0.0181)
Equity	1.042***	1.047***	1.051***	1.055***	1.051***	1.048***
	(0.00390)	(0.00743)	(0.00584)	(0.00736)	(0.00568)	(0.00748)
BenMAR	1.024***	1.015	1.026***	1.024*	1.020**	1.016
	(0.00584)	(0.0103)	(0.00782)	(0.0106)	(0.00721)	(0.0103)
BenSelf	1.117***	1.094***	1.072***	1.070***	1.082***	1.093***
	(0.0108)	(0.0174)	(0.0135)	(0.0180)	(0.0132)	(0.0174)
Effic x Pandemic						1.039+
						(0.0214)
Equity x Pandemic						1.008
						(0.00921)
BenMAR x Pandemic						1.009
						(0.0151)
BenSelf x Pandemic						0.980
						(0.0217)
Case-specific variables CHOOSING OPTION PRESENTED ON THE LEFT SIDE						
Pandemic						0.951
						(0.0965)
Constant	1.112	0.944	0.975	1.059	1.005	1.010
	(0.148)	(0.219)	(0.209)	(0.275)	(0.180)	(0.187)
Age, Gender controls	Yes	Yes	Yes	Yes	Yes	Yes
Treatment dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8,260	2,578	4,534	2,660	5,182	5,182
LL	-1969	-605.6	-985.3	-556.6	-1157	-1153
Cases	4130	1289	2267	1330	2591	2591
Clusters	719[b]	224[c]	390[d]	225[d]	225[d]	225[d]

Notes: Robust standard errors in parentheses. *** p<0.001, ** p<0.01, * p<0.05, + p<0.1.

Overall, the significance of the *attributes* coefficients indicates that they influence choice and, as expected, this influence is positive. The exception is *BenMAR*, which captures the benefit to those most at risk and in some cases does not significantly influence choice, including in our preferred specification (model 6).

Pre-pandemic (model 1 in Table 4), a one unit increase in *Effic* (efficiency) had the largest effect on choice, with every additional fatality prevented increasing the odds of participants choosing an option by 24.2% (OR = 1.242, p < .001). *BenSelf* (benefit to oneself) had the second largest effect, with every additional life saved in the zone of the city that respondents were told they lived in increasing the odds of choosing an option by 11.7% (OR = 1.117, p < .001). *Equity* (the concern for equity) followed, with an increase of 4.2% (OR = 1.042, p < .001) in the odds of choosing the policy that increased equity in the distribution of the expected life saved by 1 unit across the two areas of the city. The concern that influenced choice the least was *BenMAR* (benefit to those most at risk), with every additional fatality prevented in the group with the highest baseline risk increasing the odds of choosing that option by 2.4% (OR = 1.024, p<.001).

In the returner subsample (model 2), we see some small differences in point estimates compared to the pre-pandemic sample, but none is significantly difference (see Tables A4.2 and A4.3 in the Online Appendix for an interacted model and tests comparing the coefficients between the returner and non-returner subsamples).

For the returner subsample, as indicated by Wald tests, the odds ratios for *Effic* and *BenSelf* are significantly different to each other (p<.001), and so are those of BenSelf and Equity (p=.025), and *Equity* and *BenMAR* (p=.012). Therefore, the ranking of the concerns according to their effect on choice is *Effic* (efficiency) ≻ *BenSelf* (benefit to self) ≻ *Equity* (equity) ≻ *BenMAR* (benefit to those most at risk).

The results in the pandemic wave (models 3 and 4 in Table 4) were very similar: the increase in the odds of choosing a policy option was more affected by *Effic* and *Equity* than it was before the pandemic, similarly affected by the *BenMAR* at most risk, but less affected by *BenSelf*. This pattern repeats with both the whole and the returner sample. The concern preference ranking remained the same as before for the whole sample, but in the returner sample, equity now had a significantly higher impact on choice than benefit to self (p = .041).

We compare the model 4 attributes odds ratios using Wald tests (efficiency vs. benefit to self p<.001; benefit to self vs. equity p=.492; equity vs benefit to those most at risk, p=.015) and obtain a very similar transitive preference ordering as we did from the pre-pandemic data (model 2): efficiency ≻ benefit to self ∼ equity ≻ benefit to those most at risk. Whilst before the pandemic benefit to self was prioritised over equity, during the pandemic these two were equally influential in choice.

It is therefore not surprising that the models pooling the responses of the returner subsample from before and during the pandemic (models 5 and 6) show no major differences compared to models that capture each of the two periods separately (models 1 to 4).

Model 6 tests for the effect of the pandemic, which we find is negligible. The pandemic main effect is not significant (p=.619), meaning there is no change in preference for the left-hand alternative between waves. Turning to the interactions between the attributes and pandemic dummy, the only interaction close to significant at the 5% level is that with *Effic* (i.e., OR = 1.039, p=.066). Efficiency increased the odds of choosing a policy by 25.9% by unit pre-pandemic (OR_{Effic}= 1.259, p<.001), but in the pandemic this effect increased slightly to 30.8% (i.e., OR_{Effic} x $OR_{Effic\ x\ Pandemic}$ =1.308). That is, the

only change in the influence of the attributes in choice from before vs. the start of the pandemic was a 5 percentage-point increase in the odds of choosing the policy alternative that offers an additional unit of efficiency.

Did principles change?

In this section, we turn to stated principles to answer the question of whether they were different at the start of the COVID-19 pandemic than prior to it. We elicited these principles from binary trade-offs where respondents indicated which of two attributes they favoured over the other, and how strongly.

Figure 3 shows the distribution of answers in each trade-off for the principles returner subsample. For the first three trade-offs, n=81; for the fourth and fifth trade-offs, n=20 because only participants who were in the *self-interest* treatment were asked these questions and we only included participants who were in the same *self-interest* treatment in both phases; and in the last trade-off, n=42, which is the number of (returner) pandemic wave respondents allocated to the *self-interest* condition.

The lighter grey columns represent choices in the pre-pandemic wave (phase 1). The darker grey represents answers in the pandemic wave (phase 3). On the extremes of the scale, we have strong preference for each concern over its alternative. For example, in Trade-off 1 (Equity vs. Efficiency), about 1% of participants reported strongly favouring equity over efficiency in the pre-pandemic wave, and about 6% did in the pandemic wave. The middle bar captures indifference.

Figure 3. (Dis)agreement with statements of principle

No clear pattern emerges with regards to a shift between survey waves. In trade-offs involving self-interest, some participants did shift towards favouring the alternative principle over self-interest, but these differences were small.

Table 5 provides sample means and standard deviations (SD) for each trade-off: a negative (positive) score indicates preference for the first (second) principle in the pair over the second (first). The absolute value indicates the strength of preference, either slight (1), moderate (2) or strong (3). A score of 0 indicates indifference between the two principles. The table also reports p-values from Kolmogorov-Smirnoff (KS) tests comparing the distributions of answers.

Table 5. Principle trade-offs scores and comparisons before vs. after the pandemic outbreak

Principle Trade-off	BEFORE PANDEMIC OUTBREAK (Phase 1)	AFTER PANDEMIC OUTBREAK (Phase 3)	Comparison (1 vs. 3)
	Mean (SD)	Mean (SD)	KS p-value

1	Equity vs. Efficiency	0.89 (1.62)	0.27 (1.82)	.085
2	Helping those most at risk vs. Efficiency	0.69 (1.50)	0.65 (1.53)	>.999
3	Equity vs. Helping those most at risk	1.06 (1.47)	0.59 (1.62)	.248
4	Benefit to self vs. Helping those most at risk	1.20 (1.70)	0.90 (1.68)	>.999
5	Benefit to self vs. Efficiency	1.15 (1.69)	1.30 (1.66)	.978
6	Benefit to self vs. Equity		1.38 (1.59)	

Note: Sample sizes are the same as in Figure 3.

The trade-off between equity and efficiency, is the only one that differed between the waves, although the slight shift is only statistically significant at a 10% significance level. The tests for the remaining trade-offs support what the figures suggest: the principles stated by participants do not significantly differ before and after the pandemic outbreak.

If we consider the mean score in each trade-off to represent the direction of support between the two principles (e.g., the 0.89 in the pre-pandemic trade-off 1 would indicate an overall preference for efficiency over equity), we can aggregate the responses of each of the phases into the same transitive ranking order whereby efficiency ≻ benefit those most at risk ≻ equity ≻ benefit to oneself.

One question that Arroyos-Calvera et al. (2019) addressed was whether moral principles were aligned with choices in the preference task. In this study, we provide further evidence to answer this question. As described in the Methods section, we computed scores to capture the fraction of times participants favoured each principle over its alternative in the principle trade-off questions interacted these scores with the matching attributes (e.g., *Effic x favourEf*), testing whether the odds of choosing a policy offering an increase in an attribute are increased when the relevant principle is favoured.

We replicated Arroyos-Calvera et al.'s (2019) finding that endorsing efficiency in principle is associated with greater importance of the efficiency attribute in policy choices. Our methodology also allowed us to investigate this relationship for the rest of the principle-attribute pairs. Overall, we

found that endorsing the attributes in the principles tasks increased the effect that each additional

unit of the given attribute had on choice for efficiency, equity and benefit to self (when applicable),

but not for benefit to those most at risk. This multiplicative effect aligns with the relative importance

of the attribute on choice in the case of efficiency and equity (i.e., the higher the impact of the

attribute on choice, the higher the multiplicative effect of endorsing the principle), but less so in the

case of the benefit to self and not at all in the case of benefit to those most at risk. The full regression

output and detailed discussion of these results can be found in Online Appendix A5.

Discussion

We presented a study investigating the stability of preferences over the efficiency and distribution of

safety risk reductions. We captured preferences in two ways: by directly eliciting the principles by

which participants thought that allocations ought to be made, and by inferring preferences from

choices between policy scenarios. We examined whether these preferences changed following the

onset of the COVID-19 pandemic.

Our main finding is that preferences over factors such as self-interest, efficiency, helping those most

at risk, and spreading benefits equally between population groups are remarkably stable despite the

external shock of the pandemic, despite there being plausible reasons to believe that the importance

of these factors may have changed because of the changing context.

The results provide reassurance to policymakers and researchers in this area because it appears that

policy values, and the principles guiding their application in practice, may not need to be re-estimated

even in the wake of significant external shocks.

The stability of the measured importance of the four attributes we studied is somewhat surprising

given the plausibility of changes to their importance as we outlined in the introduction. Studies of

specific preferences for relevant goods and services, such as for transport (Shamshiripour et al., 2020)

and grocery shopping (Truong and Truong, 2022) did tend to show shifts in preferences during the

pandemic. More directly relevant, the apparent increase in self-interest evidenced by hoarding behaviour (Baddeley, 2020), and the increased altruistic behaviour noted by Grimalda et al. (2021) and Luo et al. (2021), did not manifest in increased importance of self-interest or preference for helping those most at risk in our experiments. It is possible that whilst these preferences expressed through or about real-world behaviours changed, the underlying attitudes to the four concerns we explored have remained stable.

To further set our results in context, there exists substantial evidence about the stability of distributional preferences from outside the health and safety context, and not related to COVID-19. Much of this evidence is summarised in Fisman et al., (2022) and in Bruhin et al. (2019). Fisman et al., (2022) conduct a study of Americans and find stability in fair-mindedness and in equality-efficiency orientation between 2013 and 2016, which is in line with our findings of preference stability. Bruhin et al. (2019) conduct a test-retest study of the stability of distributional preferences over a three-month time period and again find evidence consistent with reasonable stability of these preferences, which supports our findings. However, neither study evaluated the influence of an external shock.

Our results highlight a disparity between the stated importance of the principles governing policy allocation, versus the importance of the same principles as implied by people's choices. This discrepancy is stable across the phases of our study. Specifically, the policy choice questions imply that the most important principle governing choice is efficiency, followed by self-interest and equity (which are given equal priority in choice tasks), and then helping those most at risk. By contrast, the importance ranking implied by the stated principles is efficiency, then benefiting those most at risk, then equity and finally self-interest. This difference constitutes evidence compatible with the notion that the different tasks tap into different ways of thinking about prioritisation. The principles task may reveal a degree of virtuous responding, whereby people state that their own self-interest should not guide policy and helping those most at risk should be an important guiding factor, yet their choices suggest otherwise. Another possibility is that participants endorsed the principle of helping those most at risk, but the risks presented in the preference task were not different enough to set apart

those most at risk from those with the lower risk. We note that, whilst Arroyos-Calvera et al. (2019) found consistency between principles and preferences, their study did not allow for the inclusion of the benefit to those most at risk, which is the principle that appears to vary most between the elicitation modes in the current study. Despite these differences between elicitation frames, we find no difference within frames across the pre- and post- treatments. That is, our main empirical result about the stability of these rankings over time is unaffected.

Despite an overall consistent picture between the phases, we observe some modest changes in the strength of preference over some of the factors. In the Principles task, the preference for efficiency over equity observed before the pandemic loses significance. In the Preferences task, the preference for benefiting oneself over equity loses significance in the pandemic phase. These slight differences tend to be in favour of equity, suggesting that the pandemic may have heightened the salience or perceived importance of equity in health and safety resource allocation. However, the differences are not sufficient to overturn the implied order of importance of the concerns in either task type.

While we believe our study makes a significant contribution by highlighting the stability of safety preferences over time and despite changing contexts, we acknowledge some limitations. First, whilst the rate of return between study phases was good (50% for phase 1 and 31% for phase 2) considering that there was a gap of one and a half to three years between the studies, we cannot rule out selection effects in the attrition, which could mean that our sample returners were more consistent than the general sample. Restricting our analyses to returning participants also meant that some sub-group sample sizes were small, notably in comparisons involving Self-Interest. However, as shown in Table 4, repeating the regression analyses with the full sample and the restricted (returner) sample results in no qualitative difference between the results. Second, we only sampled UK residents for this survey. Different countries experienced the pandemic differently, and it is plausible that stability may be different in different study populations. Third, our results relate to the COVID-19 pandemic, and it is not clear whether or how they would generalise to different types of shock. The latter two concerns are both outside of the scope of this study, but may be fruitful avenues for future research.

Overall, we have made both methodological and empirical contributions. We demonstrate the feasibility of conducting multi-phase studies with lengthy time horizons using online panels (in this case, Prolific), and have further developed and refined the experimental approach introduced in Arroyos-Calvera et al. (2019). We also hope to have provided useful empirical evidence both on the stability and on the nature of the preferences of members of the public over the principles that could be used to guide public spending on health and safety policies. Our main message is one of reassurance – we find no evidence that preferences fundamentally shift in response to a large external contextual shift, which suggests that policy values elicited in one context may well be portable even in the event of large external shocks.

References

Arrow, K. J., Cropper, M. L., Eads, G. C., Hahn, R. W., Lave, L. B., Noll, R. G., Portney, P. R., Russell, M., Schmalensee, R., Smith, V. K., & Stavins, R. N. (1996). Is there a role for benefit-cost analysis in environmental, health, and safety regulation?. *Science*, 272(5259), 221-222.

Arroyos-Calvera, D., Covey, J., Loomes, G., & McDonald, R. (2019). The efficiency-equity trade-off, self-interest, and moral principles in health and safety valuation. *Social Science & Medicine*, *238*, 112477.

Attema, A. E., L'haridon, O., Raude, J., Seror, V., & Coconel Group. (2021). Beliefs and risk perceptions about COVID-19: evidence from two successive French representative surveys during lockdown. *Frontiers in Psychology*, *12*, 619145.

Baddeley, M. (2020). Hoarding in the age of COVID-19. *Journal of Behavioral Economics for Policy*.

Blayac, T., Dubois, D., Duchêne, S., Nguyen-Van, P., Ventelou, B., & Willinger, M. (2021). Population preferences for inclusive COVID-19 policy responses. *The Lancet Public Health*, *6*(1), e9.

Bobinac, A., van Exel, N. J. A., Rutten, F. F., & Brouwer, W. B. (2012). Inquiry into the relationship between equity weights and the value of the QALY. *Value in Health*, 15(8), 1119-1126.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, *121*(5), 803-843.

Bruhin, A., Fehr, E., & Schunk, D. (2019). The many faces of human sociality: Uncovering the distribution and stability of social preferences. *Journal of the European Economic Association*, 17(4), 1025-1069.

Camerer, C. F. (1997). Progress in behavioral game theory. *Journal of economic perspectives*, 11(4), 167-188.

Chuang, Y., & Schechter, L. (2015). Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results. *Journal of development economics*, 117, 151-170.

Dolan, P., Hallsworth, M., Halpern, D., King, D., Metcalfe, R., & Vlaev, I. (2012). Influencing behaviour: The mindspace way. *Journal of Economic Psychology*, 33(1), 264-277.

Engel, C. (2011). Dictator games: A meta study. *Experimental Economics*, *14*(4), 583-610.

Fisman, R., Jakiela, P., Kariv, S., & Vannutelli, S. (2022). The distributional preferences of Americans, 2013–2016. *Experimental Economics* (forthcoming).

Genie, M. G., Loría-Rebolledo, L. E., Paranjothy, S., Powell, D., Ryan, M., Sakowsky, R. A., & Watson, V. (2020). Understanding public preferences and trade-offs for government responses during a pandemic: a protocol for a discrete choice experiment in the UK. *BMJ open*, *10*(11), e043477.

Grimalda, G., Buchan, N. R., Ozturk, O. D., Pinate, A. C., Urso, G., & Brewer, M. B. (2021). Exposure to COVID-19 is associated with increased altruism, particularly at the local level. *Scientific Reports*, *11*(1), 1-14.

Guiso, L., & Paiella, M. (2008). Risk aversion, wealth, and background risk. *Journal of the European Economic association*, *6*(6), 1109-1150.

Jonas, E., Schimel, J., Greenberg, J., & Pyszczynski, T. (2002). The Scrooge effect: Evidence that mortality salience increases prosocial attitudes and behavior. *Personality and Social Psychology Bulletin*, *28*(10), 1342-1353.

Jungmann, S. M., & Witthöft, M. (2020). Health anxiety, cyberchondria, and coping in the current COVID-19 pandemic: Which factors are related to coronavirus anxiety?. *Journal of Anxiety Disorders*, *73*, 102239.

Lindholm, L., Rosén, M., & Emmelin, M. (1998). How many lives is equity worth? A proposal for equity adjusted years of life saved. *Journal of Epidemiology & Community Health*, *52*(12), 808-811.

Loewenstein, G., & O'Donoghue, T. (2004). Animal spirits: Affective and deliberative processes in economic behavior. *Available at SSRN 539843*.

Loomes, G. (2006). (How) Can we value health, safety and the environment?. *Journal of Economic Psychology*, *27*(6), 713-736.

Luo, Y., Shuster, A., Chung, D., O'Brien, M., Heflin, M., Perl, O., ... & Gu, X. (2021). Altruism in a time of crisis: Dissociable social valuation and perception during COVID-19 in the United States. PsyArXiv preprint. Available from: https://psyarxiv.com/kb9tu/

McFadden, D. (1973). Conditional logit analysis of qualitative choice behaviour, in Zarembka, P. (ed.) Frontiers in Econometrics, Academic Press, New York, pp. 105-142.

Orme, B. (1998). Sample size issues for conjoint analysis studies. *Sequim: Sawtooth Software Technical Paper*.

Pyszczynski, T., Lockett, M., Greenberg, J., & Solomon, S. (2021). Terror management theory and the COVID-19 pandemic. *Journal of Humanistic Psychology*, *61*(2), 173-189.

Shamshiripour, A., Rahimi, E., Shabanpour, R., & Mohammadian, A. K. (2020). How is COVID-19 reshaping activity-travel behavior? Evidence from a comprehensive survey in Chicago. *Transportation research interdisciplinary perspectives*, 7, 100216.

Shelton, C., El-Boghdadly, K., & Appleby, J. B. (2021). The 'haves' and 'have-nots' of personal protective equipment during the COVID-19 pandemic: the ethics of emerging inequalities amongst healthcare workers. *Journal of Medical Ethics*.

Solomon, S., Greenberg, J., & Pyszczynski, T. (1991). A terror management theory of social behavior: The psychological functions of self-esteem and cultural worldviews. In M. P. Zanna (Ed.), Advances in experimental social psychology (Vol. 24, pp. 91-159). San Diego, CA: Academic Press.

Truong, D., & Truong, M. D. (2022). How do customers change their purchasing behaviors during the COVID-19 pandemic?. *Journal of Retailing and Consumer Services*, 67, 102963.

Tsuchiya, A., & Dolan, P. (2009). Equality of what in health? Distinguishing between outcome egalitarianism and gain egalitarianism. Health economics, 18(2), 147-159.

Ubel, P. A., DeKay, M. L., Baron, J., & Asch, D. A. (1996). Cost-effectiveness analysis in a setting of budget constraints—is it equitable? *New England Journal of Medicine*, 334(18), 1174-117

Ubel, P. A., Baron, J., Nash, B., & Asch, D. A. (2000). Are preferences for equity over efficiency in health care allocation" all or nothing"?. *Medical Care*, 38(4), 366-373.