

Experimental Immigration Ethics

Mollie Gerver*, Patrick Lown[†] and Dominik Duell[‡]

Thursday 4th January, 2024

Abstract

Experiments serve many functions in moral and political philosophy. One is to establish public opinion: experiments can establish whether certain policies are popular, or can be made popular when certain facts are made salient. In some contexts, this helps establish what policymakers ought morally to do. Experiments can further establish whether encouraging individuals to try and follow certain moral principles results in individuals expressing wrongful biases against certain groups, sometimes creating reasons for individuals to avoid trying to follow these principles. We present an example of an experiment we conducted which fulfils the above two functions. The experiment aimed to establish whether the public supports immigration policies consistent with a principle of liability, defined as a principle which holds that enforcement against migrants is wrong when they either pose no threat or are not responsible for the threat they pose. We demonstrate that the public in both the US and UK is sensitive to this principle, and generally avoids any bias in the application of the principle. This, in turn, has implications for what policymakers ought to do, and what individuals ought to try to follow.

Keywords:

*King's College London. Email: Mollie.Gerver@kcl.ac.uk

[†]University of Essex. Email: plown@essex.ac.uk

[‡]University of Innsbruck. Email: dominik.duell@uibk.ac.at

1 Introduction

At age 11 Agnesa Murselaj fled war in Kosovo, arrived in the United Kingdom, and went to school in Glasgow. At age fifteen she was woken at dawn by a border official wearing a bullet proof vest, accompanied by ten other officials who brought her to a detention centre in England. Years later she recalled how she felt: "A stranger coming into your room when you're asleep...it just takes me back to Kosovo and what happened there." She was ultimately freed from detention after her friends organised a national campaign in her aid, which additionally aimed to end dawn raids against all immigrants. The campaign did not succeed in ending dawn raids against all migrants, but it sparked media focus on immigration enforcement, and in particular on the detention of refugees (BBC 2019).

The idea that it is wrong to detain refugees is fairly uncontroversial amongst political philosophers (Gibney 2004; Miller 2005; Brock 2020; Hidalgo 2018; Hosein 2019; Mendoza 2015; Lister 2020; Silverman 2014). In contrast, the use of detention is more controversial amongst voters in a range of countries (Koos and Seibel 2019; Bruneau, Kteily, and Laustsen 2018; Betts 2001). In this article, we argue that the opinions of the public should be of concern to and inform philosophers. More specifically, experiments aiming to understand the opinions of the public are important for establishing what voters and policymakers morally ought to do.

In Sections 1 and 2, we present two reasons why experiments matter for ethics in general. First, experiments can provide evidence of what is popular, and so evidence of what policymakers are able to do. A policymaker attempting to pass unpopular legislation concerning immigration - such as legislation to stop detention of refugees - may be unable to do so if voted out of office before the legislative process is complete. If she cannot pass such legislation due to popular opposition, and ought implies can, she has no duty to pass such legislation.

Second, experiments can establish what individual voters ought to try to do. Sometimes voters trying to follow a given principle hold implicit biases that they struggle to counteract. For example, imagine voters try to apply a principle of liability to immigration policies, and hold that it is wrong to harm migrants not liable to harm, including refugees like Agnesa who do not pose any threat. Imagine, also, that voters trying to follow this principle end up

applying implicit biases, viewing migrants from certain regions as liable to more harm. Given the effects of trying to follow the principle, these voters have one reason to not try and follow this principle, and activists have one reason to not encourage voters to try and follow this principle. While this reason needn't be decisive, it is one consideration, and experiments can provide evidence of when this consideration arises.

In Sections 3-4, we present a factorial vignette-based experiment we conducted which fulfilled the above two functions. The experiment, conducted in the US and UK, sought to establish evidence for whether voters support policies consistent with a principle of liability in immigration ethics, and whether they demonstrate implicit bias in following this principle. We found that subjects supported policies in a way that was generally consistent with a principle of liability, and that subjects generally (with one potential exception) did not demonstrate bias when sensitive to liability. Based on this suggestive pattern of results, we argue in Section 5 that there is evidence that (a) from the standpoint of public opinion, policies sensitive to liability are possible for policymakers to implement, and (b) encouraging voters to be sensitive to a principle of liability would not necessarily result in voters holding a biased application of this principle. If so, organizations ought not refrain from appealing to liability out of fear that this will create bias amongst voters.

2 What policymakers can do

We presume an elected policymaker has no duty to pass a given piece of legislation if she cannot do so. An elected policymaker cannot pass legislation if she will lose an election before the legislative process is complete. Such is the case when a policymaker starts the legislative process, causing them to be voted out of office before the process is complete, resulting in the next elected policymaker not completing the legislative process. For example, passing a law to grant permanent residency to all unauthorised migrants might take several years of persuading co-parliamentarians to vote for this law, or simply take several years for the bill to navigate parliamentary procedures. If a policymaker begins the process, and beginning the process results in her being voted out of office before the process is complete, she cannot have an obligation to complete the process.

Though a policymaker has no duty to pass a law if she cannot do so, sometimes a policymaker can pass a law by persuading voters to support the law, ensuring that she can complete the legislative process. Experiments can examine how malleable public opinion is on a particular issue, and whether policymakers will likely be able to sway public attitudes. For example, imagine a policymaker seeks to establish if she has public support to pass legislation which prevents the deportation of those who would face risks in their home countries. An experiment can establish if she is likely to obtain public support if she successfully communicates key information to the public. The experimenter might ask voters to assess whether they think a given migrant should be deported, then randomly assign some participants to be told about the risks the migrant will face if deported (i.e. treatment group) and others to receive no information (i.e. control group). If those in the treatment group are more likely to oppose the deportation, this provides evidence that policymakers can increase public support for legislation banning certain deportations by communicating risks certain migrants face. If the information is persuasive enough, policymakers cannot claim that passing such a policy is not possible, and so cannot claim they lack a duty to do so based on such an impossibility.

It is worth noting that the above claims about the value of experiments are consistent with a broad range of theories of the relationship between public opinion and philosophical truths. In particular, our claims are consistent with moral realism, the claim that at least some moral facts are real whether people think they are or not (Enoch 2011; Parfit 2011). For example, we presume the moral fact "ought implies can" is real whether people think it is or not; it is just that, if ought implies can, experiments can help establish what legislation policymakers may have no duty to pass, on account of being unable to pass such legislation. And they can establish whether policymakers may indeed be able to pass legislation, if they present voters with certain information.

Of course, if it turned out that voters do not support a given piece of legislation even with certain information, and this indicates that policymakers cannot (and therefore have no duty) to pass legislation, the state as a whole may still be committing a wrong. If states have certain obligations, and states are constitutive of both voters and policymakers, then states act wrongly in instituting policies which violate obligations regardless of whether policymakers

themselves are responsible. However, the above experiments have implications for what individual policymakers ought to do, and so implications for the ethics of individual actions.

3 Bias

To understand the second value of experiments we describe, it is necessary to first defend a more basic claim: sometimes trying to follow a rule is wrong, even if the rule isn't. This claim has been most famously defended in the context of consequentialism; if we ought to maximize good consequences, trying to maximize good consequences can be wrong when this will not maximize good consequences (Sidgwick 1907; Parfit 2011). There is debate over whether this undermines certain types of consequentialism, and whether a principle is self-defeating when holding that individuals ought to bring about outcomes which can only be brought about if they do not try to bring about these outcomes. We needn't resolve this debate to accept a more basic claim: if we take it as a brute fact that there are a certain wrongful actions we ought to not engage in, then we have a pro tanto reason to not try to follow a principle when doing so will make it very difficult or impossible to avoid engaging in these wrongful actions. This is true even if the principle does not itself imply that we ought to engage in these wrongful actions, and merely trying to follow the principle has this effect.

For a concrete example, consider the principle of liability. There is debate over what this principle entails, but one prominent view holds that it is wrong to harm those not liable to harm, and one is not liable to harm if either posing no threat or not responsible for this threat (McMahan 2011). In the context of immigration, some claim migrants can pose various threats: they can threaten the democratic functioning of a state when entering (Miller 2016), or threaten citizens' freedom of association (Wellman 2008), or even threaten citizens' ability to refrain from taking on costly duties, as when migrants enter and must be provided healthcare (Blake 2019). Less controversially, migrants pose a threat if carrying a contagious disease, as during a global pandemic. When migrants pose a threat via migration, they are not liable to harm from enforcement if not responsible for their migration. They are not responsible if forced to enter a country because they are trafficked or fleeing life-threatening conditions.

Now, imagine a border official tried to follow a principle of liability - only targeting those posing a threat and responsible for this threat - but she overwhelmingly targeted migrants from the Middle East based stereotypes that they were more likely to be voluntarily migrating to pose a threat to citizens. Similarly, imagine voters trying to follow a principle of liability ended up endorsing candidates who targeted migrants from the Middle East, exaggerating the odds that such migrants posed a responsible threat, or targeting Middle Eastern migrants liable to harm more than non-Middle Eastern migrants liable to harm. Assuming that those trying to follow the principle of liability ended up drawing upon biases they had little control or awareness of, and this lead them to engage in wrongful actions – targeting migrants or voting in a wrongful manner – they would have moral reason to not try to follow this principle. This reason might not be decisive - trying to follow the principle might be worth it if yielding some morally valuable result - but the effects of trying to follow the principle matter in the overall assessments of what they ought to try to do.

To be clear: we are not claiming that one has reason to avoid following a principle if this will lead to wrongful actions; one could often just follow the principle and not engage in the wrongful actions. For example, one could often exclude migrants posing a responsible threat without using their ethnicity as a criteria for whether they are excluded. Our claim is more specific: when trying to follow a principle will make it nearly impossible to avoid acting based on implicit biases one has little awareness or control over, such that one cannot simply try following the principle and avoid the biases, this creates one reason to not try to follow the principle.¹ Of course, if all plausible principles one attempts to follow lead to the application of biases, it may be that one ought to just pick the principle which is better in other regards; but assuming that trying to follow other plausible principles do not have this effect, one might have decisive reason to try following one or more of these other principles instead.

A close variant of this claim has been made by Matthew Lindauer. Lindauer defends the claim that, when a principle has the effect of increasing bias, it is less “fruitful” in one way

¹There is a rich debate over whether individuals are responsible and/or blameworthy for their implicit biases, and the discriminatory actions resulting from these biases (Saul 2013; Holroyd, Scaife, and Stafford 2017; Washington and Kelly 2016). Our general claim is not dependent on the outcome of this debate: regardless of whether one is responsible and blameworthy for implicit biases and actions arises from these biases, one has moral reason to avoid trying to follow principles that result in biases that are truly difficult or impossible to resist.

compared to a principle which does not increase bias or reduces bias. A principle which reduces bias is “fighting back” against injustice (Lindauer 2019). Our claim is inspired by his, but more explicitly highlights the distinction between cases where the letter of the principle contributes to biases and cases trying to follow the principle contributes to biases. Even if faithfully following a principle does not lead to bias, because the letter of the principle clearly stipulates that bias is prohibited, it remains the case that establishing what principle someone should try to follow requires establishing whether trying to follow that principle will lead to bias. The principle of liability might not itself lead to bias, in that it specifies individuals’ liability based on whether they pose a responsible threat rather than whether they are of a given ethnicity or country, but trying to follow the principle may still lead to bias. When it does, one has reason to not try to follow the principle, at least when one has no control or awareness of this bias arising.

If the above is true, third parties also have reason to avoid persuading individuals to try to follow a principle that leads these individuals to hold or apply biases. For example, human rights organisations have reason to avoid persuading individuals to try to follow a principle of liability in immigration if this increases unjust biases against certain migrants.

If the above claims stand, experiments can serve the valuable function of providing evidence of what individuals and organisations have reason to try to do, and reason to not try to do. This is because experiments can provide evidence of whether bias arises when individuals try to follow a given principle.

To see how, it is worth noting that there is more than one way a person might try to follow a principle. One way is by being presented a principle, and instructed to try to follow this principle. An experiment can establish if individuals really are trying to follow the principle, and whether they demonstrate bias in doing so. For example, an experiment could present a treatment group with a principle of liability, and a control group with no such principle, and then present both groups with vignettes of fictional migrants who are attempting to cross the border. The vignettes would vary in the degree that migrants are responsible for their migration choices, with some forced to migrate due to war, and others making a purely voluntary choice. The vignettes would further vary in the degree that migrants pose threats, with some

being on a terrorist watch list, others carrying a contagious disease, and still others posing no specified threat. Importantly, migrants would vary in their ethnicity and/or region-of-origin. Subjects would then be asked whether they support deporting different migrants in the vignettes presented. If subjects instructed to follow the principle of liability were less likely compared to the control group to support deporting migrants posing no threat or forcibly migrating, this would suggest that subjects were following the principle of liability set out. If subjects who followed this principle were also more likely to discriminate against migrants from a given ethnicity or region as compared to subjects in the control group - more likely to think deporting these migrants justified even when no more likely to be liable - this would provide evidence that trying to follow the principle contributes to bias.

There is another way a person might try to follow a principle. A person might be trying to follow a principle without quite being aware they are. They might be responding to certain attributes when presented with specific cases, aware that these attributes are important, but never spelling out the precise principle explaining why these attributes are important. We take such a phenomenon to be common. For example, border officials and voters may avoid harming those who seem to pose no threat, or who seem forced to migrate, roughly following a principle of liability without being aware they are. An experiment can establish both if this is the case, and if bias arises in their application of the principle. This is because an experiment can assess how principles might be applied "organically" in the absence of any explicit instruction on the application of principles by the experimenter(s). The experiment would be designed in a similar manner as that above, using vignettes that vary descriptions of migrants, but without a treatment group providing instructions regarding a principle of liability². We now describe in more detail one such experiment we conducted.

²The distinction made in the previous two paragraphs between experiments that do and do not explicitly describe a principle is similar to what Lindauer refers to as "intervention-based" vs "non-intervention based" methods. These are apt phrases to describe the distinction we have in mind, but we avoid using them here, as an experiment could include an intervention other than that of a philosophical principle. For example, some subjects may be presented with cases of migrants who are forced to migrate while others are not. This is a type of intervention, but would be categorized under Lindauer's "non-intervention based" method because there is no presentation of a principle. We therefore thought it may be confusing to refer to such experiments as using a non-intervention based method. For more on this topic, see Lindauer 2020 *ibid* at 2134-2135

4 Immigration and liability: an experiment

In this section, we first present the principle of liability as applied to immigration in more detail. We then outline our precise hypotheses about whether UK and US citizens are sensitive to this principle and biased in its application, followed by our experimental methods to test these hypotheses, and our results. In the penultimate section 5 we present implications of these findings for what various agents ought to do.

4.1 Liability: some further details

For the purposes of this chapter, we adopt the principle of liability articulated in the last section: an individual is liable if she (a) poses a threat and (b) is responsible for their threat (McMahan 2011). In some cases, this view provides a fairly straightforward evaluation of whether a migrant is liable to harm: if a migrant is forced at gunpoint to enter a country, and poses a threat because she is potentially carrying a contagious disease, she is not liable to harm. It might still be all-things-considered permissible to require that she remain in quarantine for an extended period of time if necessary, but this would be permissible for reasons unrelated to liability.

Some cases are more complex, because they involve an individual who was responsible for one action, and they pose a risk in virtue of both this action and another action they are not responsible for. Consider a migrant who committed a crime, or joined a terrorist organisation, and then is forced by war to migrate to the US. These sorts of cases are difficult, because the choice to join the terrorist organisation or commit the crime may be voluntary, and the individual is therefore responsible for this choice, but not the choice to migrate.

We lack the room to completely address whether such migrants are liable to the harms of immigration enforcement, including deportation and detention. However, we will presume at least this: the fact that a migrant is forced to migrate is relevant for determining the harm she is liable to experience, assuming that her intention in migrating is to avoid life-threatening conditions, and the risks she poses are below a given threshold. While a person who will definitely commit a terrorist act may be liable to detention and deportation even if she was forced to leave her home country, a person who previously joined a terrorist organisation, but

shows no indication of still being a member of the organisation, is not liable to life-threatening or injurious deportation if she was forced to migrate.

4.2 Hypotheses

We hypothesize that citizens will care about these considerations, and while they won't perfectly align their opinions with the demands of justice - people rarely do (Hidalgo 2018; Vries et al. 2019) - they will care about whether a migrant poses a threat as well as whether they migrated because they had little other choice. More specifically, we hypothesize that:

Hypothesis 1 *The less a threat a migrant appears to pose, the more unreasonable citizens find the use of immigration enforcement.*

Given that liability also concerns justified harm against those who do pose a threat, but are not responsible for this threat, we hypothesize that respondents will be less supportive of enforcement against migrants who pose a threat in migrating but are not responsible for migrating. In particular, they are less supportive of enforcement against those who pose a threat but are forced to migrate due to conditions in their home countries:

Hypothesis 2 *There is less acceptance among citizens for the use of immigration enforcement against migrants posing a threat who are migrating to avoid harm in home countries, as compared to those who are posing a threat and not migrating to avoid harm in home countries.*

Though we hypothesize that the principle of liability is at least loosely followed, it is likely applied in a biased manner. Earlier experiments on immigration attitudes have found bias in general, with subjects often more supportive of migrants not from the Middle East (Aalberg, Iyengar, and Messing 2012; Hainmueller and Hopkins 2015a; Dustmann and Preston 2007; Lee 2008). We therefore hypothesize:

Hypothesis 3 *Citizens are more likely to support immigration control against migrants from the Middle East as compared to other migrants.*

4.3 Methodology

Sample Characteristics. To test these hypotheses, we conducted an online survey with national samples in both the United Kingdom and the United States. The British sample was collected via Prolific Academic (N=1,745) and was 67.0% female, 57.1% under the age of 35, and 52.4% of respondents had university degrees. The American sample was collected via Qualtrics (N=1,804) and was 52.5% female, 32.6% under the age of 35, and 44.6% had university degrees. In terms of political ideology, the British sample identified as more leftist ($M=5.9$, $sd=2.1$) on a 10pt self-reported ideology scale while the American sample leaned right of center ($M=4.1$, $sd=2.6$).

Factorial vignette design. Our design used experimental vignettes (Atzmüller and Steiner 2010; Hainmueller and Hopkins 2015b; Turper 2017) to evaluate the extent that respondents support a given instance of enforcement. The advantage of such designs is that it allows researchers to evaluate the impact of many elements in a complex scenario in parallel. Respondents were presented with five vignettes involving an unnamed fictional migrant who seeks to enter or remain in the United Kingdom or United States, depending on the sample, with varying types of enforcement utilised to compel them to leave or prevent their entrance, the harm arising from enforcement, and migrants varying in terms of their reasons for migrating, the threats they pose, and their region of origin. Respondents were asked to decide the extent that denying this migrant the ability to enter or remain in the US/UK was unreasonable on a scale of 1 to 7 (extremely reasonable to extremely unreasonable).

We randomly vary attributes of the immigration case in a single-profile factorial vignette. Below are more details of these varying attributes:³

- **The migrant's country of origin:** collapsed by region, for the UK: Africa, Eastern Europe, Middle East, or Southeast Asia and for the US: Africa, Middle East, Eastern Europe, Central and South America, East Asia;⁴
- **The migrant's reasons for attempting to enter the UK/US, indicating whether**

³Figures A.1 and A.2 in the Appendix show one example permutation of the factorial vignettes as shown to respondents in the UK and US respectively.

⁴Africa: Nigeria, Ethiopia, South Africa, Kenya; Eastern Europe: Poland, Romania, Russia, Ukraine; Middle East: Iran, Syria, UAE, Israel; Southeast Asia: India, Bangladesh, Pakistan, Myanmar; South and Central America: Venezuela, Nicaragua, Brazil, Mexico; East Asia: China, Philippines, Vietnam, Myanmar

the migrant was forced to migrate: voluntary migrants are entering/wishing to remain in the UK/US to seek non-necessitous *Economic Opportunity*. Forced migrants are entering/wishing to remain in the UK because they are fleeing *Extreme Poverty*, *Ethnic Persecution*, or to receive *Medical Treatment*;

- **The reasons for being denied a visa to live in the UK/US, indicating whether the migrant poses a threat:** we presume that individuals pose a threat if they are denied a visa because they have a *Criminal Record*, their name is on a *Terror Watchlist*, or they obtained a *COVID-19 Positive* test result. We presume that an individual does not clearly pose a threat if they are denied a visa because they are *Unemployed* or *Immigration quotas* do not permit the provision of further visas.
- **The potential consequences of removing the migrant:** migrants could face non-lethal bodily *Harm* due to *Persecution* in their home countries, due to *Malnutrition* in their home countries, due to *Illness* in their home countries, due to being in *Custody* during enforcement, or due to *Deportation* itself. They could, alternatively, face *Death* due to *Persecution* in their home countries, due to *Malnutrition* in their home countries, due to *Illness* in their home countries, due to being in *Custody* during enforcement, or due to *Deportation* itself.⁵

The attributes presented above are the ones relevant for the analyses in this chapter. However, in the full design, we also varied the gender the migrant and whether the migrant had already crossed the border. The data presented here is part of an ongoing project and the full analyses are included in other forthcoming work (Gerver, Lown, and Duell, forthcoming).⁶

Beyond the factorial vignette described above, we also implemented a between-respondent information treatment, which only serves as robustness check of the results presented here and is described in the Appendix.

4.4 Results

Throughout the discussion of our empirical results, we present marginal means, which are the means of our outcome measure for the group of respondents that saw a vignette including a specific attribute level (e.g. a female migrant) averaging over all other attributes. (Leeper, Hobolt, and Tilley 2020). In the survey, participants were asked the degree to which "Denying

⁵These are broadly divided into harm due to the actions of border agents and harms due to the migrant returning to their country of origin. In the latter case, it is therefore necessary that these covary with the migrant's initial reason for migrating, rather than being fully randomized. For example, a migrant who left their home country due to persecution and faced harm after being deported would face harm from persecution, rather than illness.

⁶Replication data are available from the project Dataverse:
<https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/ZFOTRO>.

migrant entry is unreasonable" in response to each vignette. To ease interpretation when presenting responses on this outcome measure, we reversed the survey coding throughout so that the high end is finding exclusion to be reasonable, which we interpret as support for excluding the migrant in question.⁷ We obtained the following key findings:

The less a threat a migrant appears to pose, the less reasonable respondents found the use of any enforcement. As noted above, we considered a migrant to pose a threat if, in the vignette, they were denied a visa because they were on a terrorist watch list, had a criminal history, or had contracted Covid-19. Migrants who did not pose these threats were depicted as being denied a visa because they had not secured employment or would exceed immigration quotas. Respondents were significantly less supportive of all enforcement for those posing no threat as compared to those who did.⁸ Figure 1 illustrates this pattern for both the US and the UK, though it is stronger in the latter.

UK respondents found enforcement against migrants posing a threat as less reasonable if migrants were forced to leave their home countries. No such findings arose in the US. As noted, we presumed that a migrant who engaged in a voluntary act which poses a threat (i.e. voluntarily committed a crime) is liable to less harm if they were forced to migrate. This claim is perhaps controversial, but we argued that subjects would generally view enforcement as less reasonable if migrants posing a threat were forced to migrate. This hypothesis was confirmed in the UK, but not in the US. In the UK, the marginal mean of reasonableness of exclusion of a voluntary migrant is 5.23 (SE=0.046) while the one of the forced migrant is only 4.75 (SE=0.052). The quantities in the US are 5.02 (SE=0.049) and 5.08 (SE=.046), respectively. The former difference is significant, while the latter is not.⁹

⁷Figure A.3 in the appendix shows the marginal mean of the outcome measure on all attribute levels included in the factorial vignette excluding separating the country attribute in all its levels but including a regional indicator for ease of display.

⁸We find a positive and significant coefficient at $\alpha = .05$ on the attribute levels terrorist watch list, criminal history, and had contracted Covid-19 in an OLS regression of the outcome measure on the fully factorized attribute reason for which entry was denied (taking either unemployed or immigration quotas as reference category); standard errors are clustered at the respondent level (see Table A.1 in the appendix).

⁹In the UK, we find a positive and significant coefficient at $\alpha = .05$ on the variable *forced* in an OLS regression of the outcome measure on the variable *forced* run on observations of migrants who are posing a threat (those on a terrorist watch list, with a criminal record, or who had contracted Covid-19). The variable *forced* takes on value *forced* for the attribute levels fleeing extreme poverty and fleeing ethnic persecution and *voluntary* otherwise;

Figure 2 below shows that British respondents gave responses consistent with the second hypothesis, in that they indicated enforcement against a threatening forced migrant (i.e. terror watch-list, criminal history, or Covid-19 positive) was less reasonable than enforcement against a threatening voluntary migrant. No such distinction was found amongst US respondents, who treat voluntary and forced migrants identically when both pose a threat.

Enforcement against Middle Eastern migrants is not seen as more reasonable than enforcement against non-Middle Eastern migrants. However, in the UK there was bias against Middle Eastern migrants who are posing a threat and facing no harm of injury or death from enforcement. Figure 3 shows that in both the US and UK, on the aggregate, there was no significant difference between subjects' views of whether enforcement against Middle Eastern migrants was reasonable, as compared to enforcement against non-Middle Eastern migrants.

However, there was one type of bias. As noted in the overview of our methods, the vignettes varied in the degree that migrants would face harm from both enforcement itself and conditions in home countries. Some vignettes included migrants dying or becoming injured in immigration detention, during deportation, and after returning to their home countries. Some vignettes included migrants who faced no such harms at all. When it comes to vignettes where harm against migrants took place (either during enforcement or in home countries), there was no region-based bias in either the US or UK, but where no harm took place, British respondents viewed it as more reasonable to deny entry or remove Middle-Eastern migrants posing a threat as compared to other migrants posing a threat.

In the UK, for migrants posing a threat who would not be harmed upon being denied entry, the marginal mean is 6.03 (SE=0.09) for Middle-Eastern migrants but only 5.69 (SE=0.08) for migrants from other regions. The difference between those two marginal means is significant while the difference between those quantities in the US is not; further, the difference of the marginal means for Middle-Eastern and other migrants when being harmed is also not significant.¹⁰

standard errors are clustered at the respondent level. The coefficient on *forced* is not significant in the US (see Table A.2 in the appendix).

¹⁰In the UK, we find a positive and significant coefficient at $\alpha = .05$ on the variable *Middle Eastern* in an

5 Implications

Before addressing potential implications of findings for what various agents ought to do, it is worth spelling out some limitations.

First, we only asked individuals to state their judgments about individual migrants. It may be that voters would not be sensitive to whether migrants posed a threat or were forced to migrate - i.e. would not be sensitive to attributes important for liability - if they were asked to judge cases involving far more migrants. This could be because of the "psychic-numbing" effect, where individuals are less likely to feel the need to stop harms against large numbers of individuals as compared to specific individuals (Slovic et al. 2013). Because there are a large number of non-labile migrants, a policymaker attempting to reduce harm for all such migrants may be voted out of office before able to do so. Or, alternatively, a policymaker may need to focus on presenting voters with cases of individual non-labile migrants to persuade them to support policies which avoid deporting all non-labile migrants; in other words, perhaps the psychic-numbing effect can be counteracted precisely by presenting all voters with the sorts of individual cases that appear in our experiment. Further experiments could assess if this is the case, and to more thoroughly understand what messaging persuades voters to support enforcement consistent with a principle of liability.

Second, this is but one experiment; further experiments are necessary to establish if the results are replicable, and so our findings here are not decisive. Relatedly, there is a possibility that subjects were sensitive to liability-related attributes in the vignettes, despite secretly or sub-consciously thinking such attributes irrelevant. We therefore cannot fully establish how people will act at the polls based on the result of this experiment.

Finally, it is important not to draw strong conclusions from differences between the British and American findings, given the different ways in which the samples were gathered. In particular, the Prolific sample in the UK seems to be more politically sophisticated (based in part on their open-ended responses not discussed here), which may account for some of the dis-

OLS regression of the outcome measure on the variable *region* run on observations of migrants who are posing a threat (those on a terrorist watch list, with a criminal record, or who had contracted Covid-19) separately for observations where the migrant was not harmed vs when the migrant was harmed upon being denied entry; standard errors are clustered at the respondent level (see Table A.3 in the appendix).

crepancies in the results.¹¹ Moreover, in both the UK and US subjects did not express a perfect commitment to a principle of liability. They simply were sensitive to this principle, in that they viewed harm inflicted on individuals not liable to harm as far less reasonable than those liable to harm.

With the above caveats in mind, below are some tentative implications from our findings for how three groups of actors ought to act.

5.1 Elected policymakers

If citizens are sensitive to liability in some regards - they distinguish between those who do and do not pose a threat via their migration, and those who are and are not forced to migrate - this suggests policymakers are not limited by lack of popular support to pass legislation sensitive to liability. Of course, not being constrained by popular support does not mean passing legislation is possible. Vocal minorities - such as far right voters not sensitive to liability - could sway the majority, and party heads might prevent party members from attempting to pass legislation (Hampshire 2013). Nonetheless, our findings weaken the claim that policymakers are unable to pass liability-sensitive legislation specifically because of general opposition.

This claim - that there is general support for liability, even if other barriers persist - is important. It suggests that at least some liability-based changes do not face a barrier faced by other principles endorsed by philosophers. For example, many philosophers claim states have a duty to accept far more refugees (Gibney 2004; Carens 2013; Hidalgo 2018; Miller 2016), but the majority of citizens in a range of countries oppose a major increase in refugees accepted (Cope and Crabtree 2020; Ferwerda, Flynn, and Horiuchi 2017; Kalogeraki 2019; Horiuchi and Ono 2018). Even in the United States, where a majority support accepting at least 95,000 refugees for resettlement each year (Da Silva 2021), the majority do not support accepting the millions that the US has a duty to accept according to a range of plausible theories of global justice. Many philosophers additionally support increases in immigration more generally (Gib-

¹¹For example, in both surveys we included post-experimental questions asking respondents their opinions about whether they supported granting medical care to unauthorised migrants, and permanent residency to migrants who had served as frontline workers during the Covid-19 pandemic. Both surveys included closed-ended questions concerning these policies, and open-ended questions which respondents could answer, asking them why they gave the answers they gave. As compared to the US Qualtrics sample, the UK prolific sample was far more likely to answer these open-ended questions, and far more likely to give long, multi-sentence responses.

ney 2004; Carens 2013; Brock 2020), despite opposition amongst the general public (Hidalgo 2018). Were a policymaker to attempt to begin the legislative process of significantly increasing the number of refugees or migrants provided visas, she might be quickly voted out of office before the process was complete. In contrast, reducing the harm instigated against non-labile migrants during enforcement may be popular enough to render legislation politically feasible.

If so, this could have implications for a range of policies. Consider a UK policy where those who have a criminal history can be subject to deportation, even if arriving in the UK to avoid life-threatening conditions, and even if they will face a high risk of death upon return. The only exception to this policy is if an individual is returning to a country where they will face specific human rights violations, including torture. Those who would face a life-threatening famine or lack of medical care can still be deported (Anderson 2017; *AS and DD v Secretary of State for the Home Department* 2008; *BB, PP, W.U, Y and Z v Secretary of State for the Home Department* 2016).

Respondents in our experiment, in contrast to the above policies, seem sensitive to the reasons an individual who committed a crime arrived in the UK, viewing enforcement against those forced to arrive as less reasonable than against those voluntarily arriving. Were policymakers to attempt to pass legislation which would provide additional protections to those forced to arrive in the UK, including those with a criminal history, they may be able to pass such legislation without being voted out of office.

This assumes, however, that such legislation could be passed and applied in a non-biased manner. As noted in our findings, there is some evidence of bias amongst UK subjects, who were more likely to support enforcement against Middle Eastern migrants posing a threat, so long as no injury or death arose in enforcement. When such bias arises, policymakers could face a dilemma: they could either support legislation more in line with liability, but only for some migrants, and thus resulting in discrimination; or they could attempt to pass such legislation without bias against Middle Eastern migrants, but likely unable to actually pass such legislation. We do not resolve this dilemma here, and it may not actually arise when it comes to liability itself. Liability concerns justified harm; bias did not arise in the UK when it came to migrants facing harm, in that subjects made no distinction between Middle Eastern

and non-Middle Eastern migrants facing harm. It was simply that UK subjects were biased between groups experiencing no harm. If so, then legislatures may be able to change policies specifically concerning migrants who would face harm from enforcement, without instituting any bias. For example, the UK could refrain from deporting those fleeing violence back to life-threatening conditions, even if they have committed a crime in the past.

The above analysis mostly assumes that the public is the only barrier to reforming policies. Even when there are other barriers - such as political parties and a vocal far right minority - the public's general attitude may be a good proxy for deciding which other barriers to try to overcome. For example, if there is especially strong opposition to deporting refugees to life-threatening violence after they commit a crime, but a political party head insists on such deportations in light of far-right protests, individual policymakers may have strong reason to try and persuade the political party head and/or vocal minority to change their minds, or refuse to acquiesce to the far-right demands. In contrast, when all oppose a change - when the general public, a vocal minority, and political party heads all insist that current policy is reasonable - overcoming such opposition may be beyond an individual policy-maker's capabilities.

5.2 Voters

As noted, we found that US respondents were both sensitive to one aspect of liability - they were less supportive of enforcement against those not posing a threat - and did not demonstrate bias against migrants from the Middle East. If one ought to generally follow a principle of liability, but one has reason to avoid doing so when this contributes to bias against Middle Easterners, this particular reason does not seem to apply in the US. If so, US voters should not avoid trying to follow a principle of liability based on fear that this will contribute to bias against Middle Easterners. This is true regardless of one's general views of immigration.

For example, imagine a given US voter thinks it is justified to limit the number of migrants arriving or remaining, because she thinks a major increase in immigration can have negative economic repercussions for disadvantages members of society in the short-term, and because a major increase in immigration can strain the state's ability to prevent crime. She doesn't think migrants are more likely to commit crimes, only that some limits in inward migration

can prevent the straining of law enforcement in the short-term. Put aside whether she is right about these facts; she thinks these facts are true, but is also against using harm against migrants not liable to harm. In other words, she thinks it is justified to deport or detain individuals to limit immigration, because of the effects of immigration on law enforcement, but still thinks that immigrants should not be deported or detained if they would face harm they are not liable to face. Scrolling on her news feed she learns about a migrant who faces deportation after committing a serious crime, and she is uncertain whether she thinks it is wrong to deport this migrant. She thinks the migrant is liable to the harm of imprisonment, but not the harm of deportation, because the migrant will return to life-threatening conditions if deported. If she is choosing between a candidate who supports the deportation of this migrant or migrants like him, and one who does not, she should account for these candidates' positions in her overall assessment of how she should vote. She should not avoid appealing to considerations of liability out of fear that accounting for this consideration will cause her to hold biased beliefs. She should not think, "in forming my opinion about whether migrants should be deported, and how I should vote, I should ignore whether they are liable to harm, because this will cause me to have implicit biases (which I have no awareness or control over) against migrants from the Middle East. I will end up judging them as more liable than migrants who are not from the Middle East." There is no evidence that accounting for liability has this effect.

In the UK, findings were more ambiguous. As noted, while there was no clear bias in general, there was a statistically significant difference between attitudes about non-harmful enforcement against Middle Eastern migrants posing a threat, and non-harmful enforcement against non-Middle Eastern migrants posing a threat. As a reminder, "non-harmful" refers to enforcement causing no death or injury, or no known death or injury.

It is not clear why respondents were biased when migrants posing a threat faced no risk of death or injury from enforcement, but they may have thought that such bias (if they were aware of it at all) was not wrong. We presume it is, because we presume that bias can be wrong even if there is no obvious harm (Slavny and Parr 2015). Moreover, enforcement against individuals from the Middle East likely always involves at least some harm (even if not death or injury), such as the harm of being denied the legal right to remain in a country in virtue of

one's region-of-origin (Mendoza 2014; Reed-Sandoval 2019).

If so, this raises the question of whether UK voters should try being sensitive to principles of liability. One good reason to be sensitive to liability is that such sensitivity might not itself contribute to bias; it may be that individuals are already biased against Middle Eastern migrants who would face no death or injury from immigration enforcement, and put aside such bias when exposed to cases of migrants who would face such harms. If this is the case, voters choosing to be sensitive to liability at least avoid bias for migrants who would face such harms. Put another way: ignoring liability may lead to bias against both Middle Eastern migrants who face harms of death or injury and those who would not face such harms, but being sensitive to liability at least avoids bias against Middle Eastern migrants who would face harms of death or injury. Alternatively, perhaps voters choosing to be sensitive to liability does lead to greater bias. For example, perhaps subjects exposed to cases with injury or death show no bias against Middle Eastern migrants in such cases, but as a result of being exposed to such cases show greater bias against Middle Eastern migrants who face no death or injury. This could serve as evidence that making these liability-related features salient, and being sensitive to these features, contributes to bias against Middle Eastern migrants who would face no harm, as compared to non-Middle Eastern migrants who would face no harm. Further experiments can establish if this is the case.

5.3 Organisations and Activists

The above analysis has related implications for organisations. Imagine a refugee rights organisation emphasizes that refugees who committed crimes are being deported to life-threatening conditions, and voters agree that such refugees should not be deported, but this causes voters to begin supporting the deportation of Middle Eastern non-refugee migrants who commit crimes and would face no risks in home countries, as compared to non-Middle Eastern migrants who commit crimes and would face no risks in home countries. Given that refugee rights organisations have a responsibility to not contribute to injustice, and given that discrimination is unjust even if not resulting in injury or death, such organisations would have reason to not emphasize life-threatening conditions refugees would face if deported. While

this reason might not be decisive - the benefits of preventing the deportation of refugees may outweigh the harms of discrimination - this reason is one of many in organisation's overall assessment of how to act.

Our findings suggest that the above is not relevant for US organizations. US subjects showed no bias in evaluating whether migrants who would face no injury or death from enforcement were wronged. If so, organizations engaging with US citizens act responsibly - in the sense of not contributing to bias - when emphasizing facts relevant for liability.

For a concrete example, consider a 2015 Human Rights Watch campaign aimed at stopping the deportation of individuals with a criminal history, but who did not seem liable to the harm of deportation. The campaign included profiles of these two migrants:

Abdulkhikim Haji-Eda, a refugee from Ethiopia who came to the US at the age of 13, was ordered removed as a drug trafficker for a single conviction for selling a small quantity of cocaine at the age of 18. Now 26 years old, he has no other convictions and is married to a US citizen and has two US citizen children.

"Mr. V.," a refugee and permanent resident from Vietnam, was ordered deported in 2008 for a 1999 conviction for possession of crack cocaine. Although he has since been granted a full and unconditional pardon from the state of South Carolina, Mr. V. remains under a deportation order and remains in the US only because of restrictions on the repatriation of certain Vietnamese nationals (Human Rights Watch 2015).

Some might fear that, by emphasizing that the above migrants are not liable to harm because they are refugees, the campaigners risked increasing support for the deportation of African and Southeast Asians who had a criminal history and were not refugees, as compared to non-African and non-Southeast Asian migrants who had a criminal history and were not refugees. Our findings found no such bias in the US. This suggests that the above campaigns were not objectionable for reasons relating to bias.

Because of potential bias against Middle Eastern migrants arising in the UK, the implications there are less clear. Moreover, in both the US and UK we did not include migrants from every possible region, and bias may indeed arise against migrants from regions not included.

For example, consider a particularly effective UK campaign to halt the deportation of fifty migrants back to Jamaica in 2020. The campaign did not shy away from discussing the crimes some individuals had committed, but emphasized that many left Jamaica because of risks, and would face risks if deported, including risks of death. This campaign, including fundraising from the campaign to cover legal costs, meant only thirteen of the original fifty immigrants were deported (Sky News 2020). It may be that such campaigns - emphasizing harm from deportation - contribute to bias against migrants from Jamaica who both committed crimes and would not face risks of injury or death from deportation, as compared to non-Jamaican migrants who committed crimes and would not face risks of death or injury from deportation. While our experiment could not establish if being sensitive liability *contributes* to bias, and we did not include migrants from Jamaica, establishing broader biases and effects would be valuable for future research, and for helping organizations determine future campaign messaging.

6 Conclusion

We set out to describe two ways that experiments are relevant for moral and political philosophy. First, they provide evidence of whether a policymaker has public support to pass a given piece of legislation without being voted out of office before the legislative process is complete; if she cannot pass the legislation, and ought implies can, she cannot have a duty to pass the legislation. Second, experiments can establish whether individuals trying to follow a given principle end up applying this principle in a biased manner. When the bias is difficult to counter, this can create one reason to avoid trying to follow the principle, and/or avoid encouraging others to do so.

We presented an example of an experiment aiming to fulfil the above functions. The experiment sought to understand whether applying a principle of liability to immigration has public support, and whether those sensitive to this principle demonstrate implicit bias in its application. We demonstrated that the general public is sensitive to considerations of liability. In both the US and UK, subjects viewed harm against migrants not posing a threat as less reasonable than harm against other groups of migrants, and in the UK they viewed harm against migrants posing a threat but forced to migrate as less reasonable than harm against migrants

posing a threat but voluntarily migrating. Moreover, subjects in the both the US and UK did not generally demonstrate any bias against migrants from a given region, though we found evidence of specific bias in the UK against Middle Eastern migrants posing a threat who would face no injury or death from enforcement.

These findings indicate that policymakers may be able to pass legislation which, at least in some ways, is sensitive to whether migrants are liable to harm. Moreover, US activists aiming to implement more ethical immigration policies needn't worry that appeals to liability will increase bias. At least, there is no evidence from the experiment to indicate that they need fear this effect, though this is less clear in the UK.

The experiment we presented was limited: it only pertained to attitudes about individual migrants and may not indicate how US and UK citizens will actually vote in the polls. However, the experiment provides one example of how empirical findings can impact what individuals ought to do, and what they ought to persuade others to do.

Figure 1: Marginal mean of how reasonable it is to deny a given migrant entry to the country by the reason for which entry was denied. We show 95% confidence bounds computed from standard errors clustered at the respondent-level.




Figure 2: Marginal mean of how reasonable it is to deny a given migrant entry based on whether they are a threat and whether they were forced to migrate. We show 95% confidence bounds computed from standard errors clustered at the respondent-level.





Figure 3: Marginal mean of how reasonable it is to deny a given migrant entry to the country by whether they would not be harmed as a consequence of returning home or whether harm they could be exposed to is known or unknown. We show estimates for threatening migrants only. We show 95% confidence bounds computed from standard errors clustered at the respondent-level.



References

- Aalberg, Toril, Shanto Iyengar, and Solomon Messing. 2012. "Who is a 'deserving' immigrant? An experimental study of Norwegian attitudes". *Scandinavian Political Studies* 35 (2): 97–116.
- Anderson, David. 2017. *Deportation with Assurances*.
- AS and DD v Secretary of State for the Home Department. 2008. EWCA Civ 289, [2008] HRLR 28.
- Atzmuller, Christiane, and Peter M. Steiner. 2010. "Experimental Vignette Studies in Survey Research". *Methodology* 6 (2): 128–138.
- BB, PP, W.U, Y and Z v Secretary of State for the Home Department. 2016. Appeal Nos. SC/39/2005 &c.
- Betts, Katharine. 2001. "Boat people and public opinion in Australia". *People and place* 9 (4): 34–48.
- Blake, Michael. 2019. *Justice, Migration, and Mercy*. Oxford University Press.
- Brock, Gillian. 2020. *Justice for People on the Move: Migration in Challenging Times*. Cambridge University Press.
- Bruneau, Emile, Nour Kteily, and Lasse Laustsen. 2018. "The unique effects of blatant dehumanization on attitudes and behavior towards Muslim refugees during the European 'refugee crisis' across four countries". *European Journal of Social Psychology* 48 (5): 645–662.
- Carens, Joseph. 2013. *The ethics of immigration*. Oxford University Press.
- Cope, Kevin L, and Charles Crabtree. 2020. "A nationalist backlash to international refugee law: Evidence from a survey experiment in Turkey". *Journal of Empirical Legal Studies* 17 (4): 752–788.
- Dustmann, Christian, and Ian P Preston. 2007. "Racial and economic factors in attitudes to immigration". *The BE Journal of Economic Analysis & Policy* 7 (1).
- Enoch, David. 2011. *Taking morality seriously: A defense of robust realism*. Oxford University Press.
- Ferwerda, Jeremy, DJ Flynn, and Yusaku Horiuchi. 2017. "Explaining opposition to refugee resettlement: The role of NIMBYism and perceived threats". *Science advances* 3 (9): e1700812.
- Gerver, Mollie, Patrick Lown, and Dominik Duell. "Proportional Immigration Control". *Journal of Politics*: forthcoming.
- Gibney, Matthew J. 2004. *The ethics and politics of asylum: Liberal democracy and the response to refugees*. Cambridge University Press.
- Hainmueller, Jens, and Daniel J Hopkins. 2015a. "The hidden American immigration consensus: A conjoint analysis of attitudes toward immigrants". *American Journal of Political Science* 59 (3): 529–548.
- . 2015b. "The Hidden American Immigration Consensus: A Conjoint Analysis of Attitudes Towards Immigrants". *American Journal of Political Science* 59 (3): 529–548.

- Hampshire, James. 2013. *The Politics of Immigration: Contradictions of the liberal state*. Cambridge, UK and Malden, MA: Polity.
- Hidalgo, Javier S. 2018. *Unjust borders: Individuals and the ethics of immigration*. Routledge.
- Holroyd, Jules, Robin Scaife, and Tom Stafford. 2017. "Responsibility for implicit bias". *Philosophy Compass* 12 (3): e12410.
- Horiuchi, Yūsaku, and Yoshikuni Ono. 2018. *Public opposition to refugee resettlement: the case of Japan*. RIETI.
- Hosein, Adam. 2019. *The ethics of migration: An introduction*. Routledge.
- Human Rights Watch. 2015. *A Price Too High: Detention and Deportation of Immigrants in the US for Minor Drug Offenses*. Accessed on 6 July 2021 at <https://www.hrw.org/report/2015/06/16/price-too-high/us-families-torn-apart-deportations-drug-offenses>.
- Kalogeraki, Stefania. 2019. "Opposition to Syrian refugees and immigrants during the refugee crisis in Greece". *Journal of Modern Greek Studies* 37 (2): 361–395.
- Koos, Sebastian, and Verena Seibel. 2019. "Solidarity with refugees across Europe. A comparative analysis of public support for helping forced migrants". *European Societies* 21 (5): 704–728.
- Lee, Taeku. 2008. "Race, immigration, and the identity-to-politics link". *Annu. Rev. Polit. Sci.* 11:457–478.
- Leeper, Thomas J, Sara B Hobolt, and James Tilley. 2020. "Measuring subgroup preferences in conjoint experiments". *Political Analysis* 28 (2): 207–221.
- Lindauer, Matthew. 2019. "Experimental philosophy and the fruitfulness of normative concepts". *Philosophical Studies*: 1–24.
- Lister, Matthew. 2020. "Enforcing immigration law". *Philosophy Compass* 15 (3): e12653.
- McMahan, Jeff. 2011. "Who is morally liable to be killed in war". *Analysis* 71 (3): 544–559.
- Mendoza, José Jorge. 2014. "Discrimination and the presumptive rights of immigrants". *Critical Philosophy of Race* 2 (1): 68–83.
- . 2015. "Enforcement matters: Reframing the philosophical debate over immigration". *The Journal of Speculative Philosophy* 29 (1): 73–90.
- Miller, David. 2005. "Immigration: The Case for Its Limits". In *Contemporary Debates in Applied Ethics*, ed. by A. Cohen and C. Wellman. Malden, MA: Blackwell Publishing.
- . 2016. *Strangers in Our Midst: The political philosophy of immigration*. Cambridge: Harvard University Press.
- Parfit, Derek. 2011. *On what matters*. Vol. 2. Oxford University Press.
- Reed-Sandoval, Amy. 2019. *Socially undocumented: Identity and immigration justice*. Oxford University Press.
- Saul, Jennifer. 2013. "Unconscious influences and women in philosophy". *Women in philosophy: What needs to change*: 39–60.
- Sidgwick, Henry. 1907. *The Methods of Ethics*. MacMillan.
- Silverman, Stephanie J. 2014. "Detaining immigrants and asylum seekers: a normative introduction". *Critical review of international social and political philosophy* 17 (5): 600–617.

- Slavny, Adam, and Tom Parr. 2015. "Harmless discrimination". *Legal Theory* 21 (2): 100–114.
- Slovic, Paul, et al. 2013. "Psychic numbing and mass atrocity". In *The behavioral foundations of public policy*, 126–142. Princeton University Press.
- Turper, Sedef. 2017. "Fearing what? Vignette experiments on anti-immigrant sentiments". *Journal of Ethnic and Migration Studies* 43 (11): 1792–1812.
- Vries, Ieke de, et al. 2019. "Anti-immigration sentiment and public opinion on human trafficking". *Crime, law and social change* 72 (1): 125–143.
- Washington, Natalia, and Daniel Kelly. 2016. "Who's responsible for this? Implicit bias and the knowledge condition". *Implicit bias and philosophy: Moral responsibility, structural injustice, and ethics* 2:11–36.
- Wellman, Christopher. 2008. "Immigration and Freedom of Association". *Ethics* 119:109–141.

Appendix

Figure A.1: Screen shot of one realization of the factorial vignette as shown to respondents in the UK sample

A male irregular migrant from South Africa is seeking entrance into the United Kingdom to seek economic opportunity. They were denied a visa which would allow them to enter the UK because immigration quotas do not permit it. In order to prevent them from entering the UK, it is necessary to forcibly bar them from a flight entering the UK. There is a small chance that they will die in custody.

Under the circumstances described in the scenario above, how reasonable or unreasonable is it to prevent the migrant in the scenario from entering the UK?

- Extremely reasonable
- Moderately reasonable
- Slightly reasonable
- Neither reasonable nor unreasonable
- Slightly unreasonable
- Moderately unreasonable
- Extremely unreasonable



Figure A.2: Screen shot of one realization of the factorial vignette as shown to respondents in the US sample

A female irregular migrant from Syria is seeking to remain in the United States to seek economic opportunity. They were denied a visa which would allow them to remain in the U.S. because their name is on a terrorism watch-list. In order to prevent them from remaining in the U.S., it is necessary to detain them in the U.S. until they agree to return home. The consequences of their returning home are unknown.

Under the circumstances described in the scenario above, how reasonable or unreasonable is it to remove the migrant in the scenario from the U.S.?

- Extremely reasonable
- Moderately reasonable
- Slightly reasonable
- Neither reasonable nor unreasonable
- Slightly unreasonable
- Moderately unreasonable
- Extremely unreasonable



Between Subjects Treatment: Certainty of Information Because a principle of liability relates to harm, we also sought to establish if respondents thought a given level of harm justified based on the threat and responsibility of a given migrant (the latter defined as whether they were forced to migrate). We have three experimental conditions relating to harm in a between-subject design. In Condition 1, respondents are provided with certain information about the harm arising from enforcement (*Certain*). In Condition 2, respondents are provided with information that is probabilistic (*Small chance, High chance, Near certain*). Finally, in Condition 3, respondents are provided with no information about what happens to the migrant in the vignette (*No outcome information*). For the purpose of the results presented here, the between-respondent information treatment serves as robustness check; we present results below pooling observations across treatments (See Figure A.3 below).

Figure A.3: Marginal mean of how reasonable it is to deny a given migrant entry to the country by immigration case attribute and the between-respondent probability of harm treatment. We show 95% confidence bounds computed from standard errors clustered at the respondent-level. The figure omits the country of origin attribute for ease of display but categorizes the country of origin into a region of origin indicator.




Table A.1: OLS Regression of the outcome measure on the fully factorized attribute reason for which entry was denied (taking either unemployed or immigration quotas as reference category); standard errors are clustered at the respondent level

<i>Reference category</i> <i>Sample</i>	<i>Unemployed</i>		<i>Immigration quotas</i>	
	<i>United States</i>	<i>United Kingdom</i>	<i>United States</i>	<i>United Kingdom</i>
(Intercept)	4.325*** (0.059)	3.855*** (0.059)	4.267*** (0.060)	3.682*** (0.058)
Immigration Quotas	-0.058 (0.069)	-0.173* (0.072)		
Unemployed			0.058 (0.069)	0.173* (0.072)
COVID-19 Positive	0.473*** (0.073)	0.406*** (0.080)	0.530*** (0.073)	0.580*** (0.079)
Criminal History	0.769*** (0.070)	1.058*** (0.072)	0.826*** (0.070)	1.232*** (0.073)
Terror Watchlist	0.923*** (0.071)	1.939*** (0.069)	0.981*** (0.071)	2.113*** (0.070)
R ²	0.038	0.132	0.038	0.132
Adj. R ²	0.037	0.132	0.037	0.132
Num. obs.	9019	8627	9019	8627
RMSE	2.009	1.988	2.009	1.988
N Clusters	1839	1728	1839	1728

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table A.2: OLS Regression of the outcome measure on the variable *forced* run on observations of migrants who are posing a threat (those on a terrorist watch list, with a criminal record, or who had contracted Covid-19). The variable *forced* takes on value *forced* for the attribute levels fleeing extreme poverty and fleeing ethnic persecution and *voluntary* otherwise; standard errors are clustered at the respondent level.

<i>Sample</i>	<i>United States</i>	<i>United Kingdom</i>
(Intercept)	5.077*** (0.046)	5.234*** (0.046)
Forced	-0.056 (0.054)	-0.481*** (0.057)
R ²	0.000	0.014
Adj. R ²	0.000	0.013
Num. obs.	5428	5178
RMSE	1.962	2.041
N Clusters	1805	1704

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table A.3: OLS Regression of the outcome measure on the variable *forced* run on observations of migrants who are posing a threat (those on a terrorist watch list, with a criminal record, or who had contracted Covid-19). The variable *forced* takes on value *forced* for the attribute levels fleeing extreme poverty and fleeing ethnic persecution and *voluntary* otherwise; standard errors are clustered at the respondent level.

<i>Sample</i>	<i>United States</i>		<i>United Kingdom</i>	
	Known or <i>Consequence</i>	No harm	Known or <i>Consequence</i>	No harm
(Intercept)	4.988*** (0.044)	5.403*** (0.079)	4.825*** (0.047)	5.697*** (0.080)
Middle East	-0.032 (0.073)	0.062 (0.141)	0.103 (0.071)	0.335** (0.115)
R ²	0.000	0.000	0.000	0.008
Adj. R ²	-0.000	-0.001	0.000	0.007
Num. obs.	4581	847	4374	804
RMSE	1.988	1.776	2.083	1.681
N Clusters	1753	635	1656	607

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$