

Choosing victims: Human fungibility in moral decision-making

Michał Białek*

Jonathan Fugelsang[†]

Ori Friedman[‡]

Abstract

In considering moral dilemmas, people often judge the acceptability of exchanging individuals' interests, rights, and even lives. Here we investigate the related, but often overlooked, question of how people decide who to sacrifice in a moral dilemma. In three experiments (total $N = 558$), we provide evidence that these decisions often depend on the feeling that certain people are fungible and interchangeable with one another, and that one factor that leads people to be viewed this way is shared nationality. In Experiments 1 and 2, participants read vignettes in which three individuals' lives could be saved by sacrificing another person. When the individuals were characterized by their nationalities, participants chose to save the three endangered people by sacrificing someone who shared their nationality, rather than sacrificing someone from a different nationality. Participants do not show similar preferences, though, when individuals were characterized by their age or month of birth. In Experiment 3, we replicated the effect of nationality on participant's decisions about who to sacrifice, and also found that they did not show a comparable preference in a closely matched vignette in which lives were not exchanged. This suggests that the effect of nationality on decisions of who to sacrifice may be specific to judgments about exchanges of lives.

Keywords: moral decision-making, fungibility, mental accounting

1 Introduction

Moral dilemmas often require people to consider the acceptability of exchanging individuals' interests, rights, and even lives. For example, in "trolley" dilemmas, people consider whether it would be acceptable to exchange one person's life in order to save the lives of five others (e.g., Foot, 1967; Thomson, 1985). The moral acceptability of such exchanges is often questionable, and much research has investigated factors that influence people's willingness to endorse exchanges of this kind (see Greene, 2015; Gürçay & Baron, 2017). However, moral dilemmas involving such exchanges may also involve a further decision – *who* should be sacrificed?

The question of how people decide who to sacrifice is often overlooked in research on moral-decision making. Some research has examined whether people's judgments about sacrificing one individual to save others is influenced by the individuals' race, political beliefs, and eminence (e.g., Petrinovich, O'Neill & Jorgensen, 1993; Uhlmann, Pizarro, Tannenbaum & Ditto, 2009) and on whether the individual to be sacrificed is related to those who will be saved

(Kurzban, DeScioli & Fein, 2012). However, this research mostly concerns people's judgments about the acceptability of the sacrifices, and does not touch on how people choose who to sacrifice. These are distinct questions and may be subject to different influences given that people reason differently when evaluating a single option than when considering multiple options (for reviews see Hsee, Zhang & Chen, 2004 and Gilbert, 2006, pp. 155-158).

To our knowledge, the closest relevant research has investigated people's choices about which of two individuals should be saved in situations where it is possible to save only one of them (Goodwin & Landy, 2014; Keren & Teigen, 2010). This research finds that these decisions are influenced by the ages of the individuals. For example, one study found that when considering a vignette in which only one of two people critically injured in a car accident could be saved, participants favored saving a young person who caused the accident over saving an older person who was its victim (Keren & Teigen, 2010).¹

Here, we suggest that people's decisions about who to sacrifice may also depend on the feeling that certain people are fungible, and interchangeable with one another. The notion of fungibility usually arises in the context of discussions about exchanges of money and other resources that are considered fungible (e.g., Thaler, 1990). For example, if you want to repay a loan of ten dollars to a friend, you need not return the particular bills she gave you; any bills summing

This research was supported by separate grants from the Natural Sciences and Engineering Research Council of Canada awarded to OF and to JF.

Copyright: © 2018. The authors license this article under the terms of the Creative Commons Attribution 3.0 License.

*Department of Psychology, University of Waterloo, Waterloo, ON, Canada; Department of Economic Psychology, Kozminski University, Warsaw, Poland.

[†]Department of Psychology, University of Waterloo, Waterloo, ON, Canada.

[‡]Corresponding author: Department of Psychology, University of Waterloo, 200 University Avenue W, Waterloo, Ontario, Canada N2L 3G1. E-mail: friedman@uwaterloo.ca.

¹As we review in the General Discussion, the findings of Kurzban et al. (2012) are relevant to our hypothesis. However, participants in their study did not choose who to sacrifice or save—instead, they make judgments about the acceptability of sacrificing different characters.

to ten dollars should be acceptable.² Similar considerations may apply in moral decision-making. If a sacrifice is necessary to save several people, we may prefer to sacrifice a potential victim who we view as interchangeable with them over a potential victim who is not interchangeable.

We investigated this possibility in three experiments. In our experiments, we characterized the individuals in moral dilemmas in terms of their nationalities. We expected this would make people view the individuals as fungible. Groups created by nationalities are entitative — they are viewed as an entity, and not just as an aggregate of individuals (Campbell, 1958; Lickel et al., 2000). Previous research shows that group entitativity affects intergroup relations, including intergroup aggression (Lickel, Miller, Stenstrom, Denson & Schmader, 2006), retaliation (Gaertner, Iuzzini & O'Mara, 2008), revenge (Sjöström & Gollwitzer, 2015), and punishment (Newheiser, Sawaoka & Dovidio, 2012). The main conclusion from this research is that members of a group might be seen as interchangeable, and suffer from aggressive actions that are intended to harm the group as a whole by targeting random members of this group. However, this research has not fully addressed the effects of group entitativity on decisions made by agents who are not a part of any of the groups involved in a moral decision. If decisions about who to sacrifice depend on feelings that certain people are interchangeable with one another, this should even arise when people make decisions for groups to which they do not belong.

In Experiments 1 and 2, we show that when individuals are characterized by their nationalities, participants choose to save three endangered people by sacrificing a victim who shares their nationality over a potential victim with a different nationality. However, participants do not show similar preferences when agents are characterized by age (Experiment 1) or month of birth (Experiment 2). In Experiment 3, we then show that this effect of nationality on decisions about who to sacrifice is specific to decisions about exchanging people's lives, and does not occur in contexts that do not require viewing lives as interchangeable.

2 Experiment 1

2.1 Method

Participants. In all experiments, we used TurkPrime (Litman, Robinson & Abberbock, 2017) to recruit participants from the US. All participants had to have completed at least 100 experiments with a 95% acceptance rate. Participants

²There are cases in which money is not fully fungible. For example, people have an aversion to receiving money with a tainted moral history (Tasimi & Gelman, 2017). They likewise judge that it is better to return a lost coin to its owner, instead of giving the owner an identical coin (Uhlmann & Zhu, 2013).

were not allowed to participate in two or more experiments in this line of research.

We analyzed data from 150 participants ($M_{\text{age}} = 36$ years, $SD = 11$ years, range = 20–75, 65 female). We initially recruited 200 participants, but excluded 47 because they failed to correctly respond to two comprehension questions, and a further 3 whose first language was not English (see Costa et al., 2014; Muda, Niszczota, Bialek & Conway, 2017, for foreign language effects on moral judgments).

Procedure and materials. Participants read a single scenario, where they were a NATO general who had to sacrifice one soldier to save three others. Participants could either choose to sacrifice a soldier who shared a feature with the three endangered soldiers, or to sacrifice a soldier who did not share this feature. Participants were randomly assigned to one of two versions of the task. In one version, the feature concerned the soldiers' nationalities. For example, participants could save three Estonians by sacrificing an Estonian or a Latvian.³ In the other version, the feature concerned the soldiers' ages (e.g., save three 23-year-olds by sacrificing a 23-year-old or a 21-year-old). We counterbalanced which feature (Estonian or Latvian; 21- or 23-years-old) was shared by four of the soldiers, and which characterized the dissimilar soldier. Below is a sample story and test question:

You are a NATO general overseeing an international military operation. Three soldiers, who all come from Latvia, are trapped in a building that an enemy tank is approaching. If nothing is done, the three soldiers from Latvia will be killed. The soldiers have recovered valuable information, and it is crucial that they survive. You decide to send another soldier on a suicide mission to slow down the tank, and allow the three soldiers from Latvia to escape. You can either send another soldier who is from Latvia or a soldier who is from Estonia.

Which soldier should you send on the suicide mission that will save the three soldiers from Latvia from being killed?

Participants responded to the test question using a 6 point Likert scale, where 1 indicated they would "definitely" sacrifice the soldier who was similar to the three endangered soldiers (e.g., "1 – Definitely the Latvian soldier" in the example above), and 6 indicated they would definitely sacrifice the soldier who was dissimilar (e.g., 6 – Definitely the Estonian soldier" in the example above); only the terminal points

³In choosing pairs of nationalities (Estonian, Latvian) and ages (21, 23) we had two aims. First, we hoped participants would view the members of each pair as distinct from one another (e.g., view Estonia and Latvia as different countries). Second, we hoped participants would not know much about these countries, and would not favor one member of the pair over the other (e.g., Estonia and Latvia are liked about equally), as such differences might mask the effects of nationality under investigation.

of the scale were labelled. After this, participants responded to two multiple choice comprehension questions, and some basic demographic questions.

2.2 Results and Discussion

If members of national groups are viewed as interchangeable and fungible, participants should choose to sacrifice a person from the same nationality as the people to be saved over someone with a different nationality. We expected participants would be less likely to show an equivalent pattern, though, when the people were characterized by age. Because our predictions are directional, we examined them using one-tailed tests.

As predicted, participants were more likely to sacrifice the soldier who shared a feature with the endangered soldiers when the soldiers were characterized by nationality, rather than by age, $t(148) = 1.98, p = .025, d = 0.32, B_{10} = 2.05$.⁴ To follow up on this difference, we used single-sample tests to examine whether ratings in each condition differed from the midpoint value of 3.5. When soldiers were characterized by nationality, participants were more inclined to sacrifice a soldier from the same country as the endangered soldiers to be saved, $t(82) = 3.65, p < .001, d = .56, B_{10} = 95.72$. However, when soldiers were characterized by age, participants did not show this inclination, and their responses did not differ from the midpoint value, $t(66) = 0.445, p = .813, d = 0.04, B_{01} = 5.08$.

Finally, we examined whether participants were biased to sacrifice members of a particular group (e.g., biased to sacrifice Latvians regardless of whether Latvians or Estonians would be saved). To test this, we compared scores across the counterbalancing groups in each condition. When the soldiers were characterized by nationality, there was a marginal effect of counterbalancing group, $t(81) = 1.93, p = .057, d = .42, B_{10} = 1.15$. However, when soldiers were characterized by their ages, responses did differ across counterbalancing groups, $t(65) = 3.03, p = .004, d = .74, B_{10} = 10.79$. This difference resulted because participants preferred sacrificing the 23-year-old soldier over the 21-year-old, regardless of the ages of the three soldiers who would be saved. This is consistent with previous findings showing an overall preference to save younger rather than older individuals (e.g., Goodwin & Landy, 2014).

Together, these findings suggest that participants view members of a national group as more interchangeable than other people. These findings also suggest this tendency

⁴In all Bayesian between group t-tests we used a default prior of Cauchy = .707, which is rather conservative and predicts relatively big effect sizes of the tested effects. B_{10} quantifies support for the tested hypothesis over the null; B_{01} quantifies support for the null over the tested hypothesis. In general, $B > 3$ is interpreted as moderate evidence, > 10 as strong evidence, > 30 as very strong evidence, and > 100 as extreme evidence. Results where $B < 3$ are interpreted as inconclusive (Jeffreys, 1961, Lee & Wagenmakers, 2013).

does not reflect a low-level bias to match descriptions of the soldier to be sacrificed and those who were endangered, as participants' did not respond similarly when the soldiers were characterized by their ages. However, we cannot reject this "matching" strategy conclusively, as participants might have shown a corresponding pattern for age groups, if they had not been biased to sacrifice the older soldier. To more conclusively rule out responses reflecting a "matching" bias, we conducted a follow-up experiment. In this experiment, the soldiers in the control condition were characterized by month of birth.

3 Experiment 2

3.1 Method

Participants. We analyzed data from 206 participants ($M_{\text{age}} = 36$ years, $SD = 10$ years, range = 20–74, 81 female). In this experiment we initially recruited 253 participants—we wanted to obtain a sample of 100 participants per condition, and based on Experiment 1, we assumed we would have to exclude 25% of recruited participants. Of the original 253 participants, 44 were excluded for answering at least one of the comprehension questions incorrectly, and 3 were excluded because English was not their first language.

Procedure and materials. The procedure was almost identical to that in Experiment 1, except soldiers in the control condition were characterized by the month in which they were born. For example, in one counterbalancing group, participants read about three endangered soldiers who were all born in March, and then could choose whether to sacrifice another soldier born in March or a soldier born in April. As can be seen in the Supplementary materials, the stories were also reworded slightly to make them read more naturally.

3.2 Results and Discussion

Consistent with Experiment 1, participants were more likely to sacrifice the soldier who shared a feature with the endangered soldiers when the soldiers were characterized by nationality, rather than by month of birth, $t(204) = 5.36, p < .001, d = 0.75, B_{10} > 100,000$. Single-sample tests against the midpoint value found that participants mostly sacrificed the soldier who was similar to the endangered soldiers when the soldiers were characterized by nationality, $t(95) = 6.512, p < .001, d = .66, B_{10} > 100,000$, but not by month of birth, $t(109) = -0.583, p = .561, d = 0.05, B_{01} = 13.91$. Additionally, we tested for a bias to sacrifice members of a particular group. We found none: nationalities, $t(94) = 0.807, p = .386, d = .17, B_{01} = 3.32$; month of birth, $t(108) = 1.09, p = .277, d = .21, B_{01} = 2.89$.

These findings again suggest that participants view members of a national group as more interchangeable than other people. Moreover, the findings rule out use of a low-level bias to match descriptions of the soldier to be sacrificed and those who were endangered.

Nonetheless, our findings do not *specifically* show that members of a national group are viewed as more interchangeable than other people — perhaps participants would respond similarly in scenarios where individuals are not exchanged for one another. We addressed this concern in the next experiment by comparing two kinds of scenario: an “exchange” version (similar to those used so far) where one soldier’s life is exchanged for the lives of three others, and a “non-exchange” version where the sacrifice does not save the lives of the three soldiers.

4 Experiment 3

4.1 Method

Participants. We analyzed data from 202 participants ($M_{\text{age}} = 29$ years, $SD = 11$ years, range = 20–73, 103 female). We originally recruited 250 participants, but excluded 43 participants because they failed at least one of two comprehension questions, and a further 5 participants because their first language was not English.

Procedure and materials. Participants again read a scenario in which they were a NATO general. Participants were randomly assigned to either read an exchange version of the scenario or a no-exchange version. In the exchange version, three soldiers from the same country were endangered, but could be saved by either sending another soldier from the same country on a dangerous mission, or by sending a soldier from a different country. In the no-exchange version, three soldiers from the same country had been killed while undertaking an important operation, and this operation could either be seen to succeed by either sending another soldier from the same country on a very dangerous mission, or by sending a soldier from a different country. As in the previous experiments, the soldiers were described as coming from Latvia and Estonia, and we counterbalanced the nationalities of the four similar soldiers and the other dissimilar soldier. An example scenario is presented below:

You are a NATO general overseeing an international military operation involving soldiers from Estonia and Latvia. Three soldiers, who all come from Estonia have obtained valuable information that is essential for the success of the operation. Unfortunately, the three Estonian soldiers were killed behind enemy lines, before they could transmit the information. An enemy tank is on its way to a building where they hid the information. If

nothing is done, the tank will demolish the building, and the information will be lost. You decide to send another soldier on a very dangerous mission to transmit the information. If he succeeds, the information will be obtained. But the soldier you send is unlikely to make it back alive. You can either send another soldier from Estonia on the mission, or you can send a soldier from Latvia.

4.2 Results and Discussion

If members of a national group are seen as more interchangeable than other people, participants in the exchange condition should be more likely to save the three endangered soldiers by endangering the soldier from their country than by endangering the soldier from another country. However, a similar pattern should not occur in the no-exchange condition. In that condition, the three soldiers are dead, and so the soldier sent on the mission is not exchanged for them.

Consistent with these predictions, participants were more likely to choose the similar soldier when his life could be exchanged for the three soldiers, compared with when no exchange was possible, $t(200) = 3.84$, $p < .001$, $d = 0.54$, $B_{10} = 129.3$. Testing the ratings against the midpoint value, we found that participants mostly chose the similar soldier in the exchange condition, $t(95) = 19.22$, $p < .001$, $d = 1.96$, $B_{10} = 313.9$, but not in the no-exchange condition, $t(105) = -1.44$, $p = .138$, $d = 0.14$, $B_{01} = 21.57$.

These findings again show that participants’ responses do not simply reflect a low-level matching strategy. More importantly, though, the findings also suggest that nationality-based choices are specific to contexts where people can be exchanged for one another, and not to similar situations not involving exchange.

5 General Discussion

In three experiments, we examined how participants decide who to sacrifice to secure a positive outcome. In all experiments, participants chose to save three endangered people by sacrificing a victim who shared their nationality, instead of choosing to sacrifice a potential victim with a different nationality. However, participants did not show similar preferences when the individuals were characterized by age or month of birth (Experiments 1 and 2). Also, participants did not base decisions on nationality when considering scenarios that did not require exchanging lives (Experiment 3).

Together, these findings suggest that decisions about who to sacrifice reflect beliefs that people are sometimes fungible, and hence interchangeable with one another. Further, the findings suggest that characterizing individuals in terms of their national groups contributes to the sense that people are mutually interchangeable. This conclusion is broadly con-

sistent with previous work showing that people often treat members of entitative groups as interchangeable — they advocate punishing members of a group for the wrongdoings of another member (e.g., Cushman, Durwin & Lively, 2012; Pereira & van Prooijen, 2018). This conclusion is also generally consistent with the fact that people are sometimes willing to sacrifice themselves for the benefit of their group, as occurs when a person gives their lives to save their group members, or to punish other groups in suicide terrorist attacks (for a review and theoretical discussion, see Whitehouse, in press).

A related explanation for our findings comes from mental accounting theory, which claims that people's perceptions of gains and losses depend on whether these events are viewed as affecting a single pool of resources or separate pools (Linville & Fischer, 1991; Thaler, 1985). For example, if someone gains \$30 but then has to pay \$10, they should be happier if the \$10 cost is viewed as a deduction from the \$30 (same pool of resources) than if it is viewed as a separate transaction (different pool of resources) (Thaler, 1999). When the gain and loss are viewed as relating to a single pool, they are perceived as a single event in which \$20 is gained; but when they are viewed as relating to different pools, they are perceived separately as a gain and a loss, and losses loom larger than gains (Novemsky & Kahneman, 2005; Tversky & Kahneman, 1991). From this perspective, participants may have preferred sacrificing a victim from the same group as the endangered people because this increased the sense of a single event in which two lives were gained. Sacrificing someone from a different nationality would have increased the sense of two events, including the loss of life.

However, our findings also give reason to doubt this explanation. In the no-exchange condition of Experiment 3, participants read a vignette where three soldiers had been killed on a mission, and one more soldier was needed to be sacrificed to fulfill the mission. Here participants showed no preference between choosing a soldier from the same nationality as the three dead soldiers and one from a different nationality. However, if participants applied mental accounting to this scenario, they should have again preferred sacrificing the soldier from the same nationality — a loss of resources from one pool is viewed less negatively than an equivalent loss tapping two pools (Thaler, 1985; Thaler, 1999). As such, our third experiment differentiates our fungibility account from mental accounting, and also favors our proposed account.

In our experiments, we examined how people choose victims in situations where sacrificing one person ensures a gain. However, most studies of moral decision-making instead focus on the different question of whether such sacrifices are morally permissible (e.g., Białek & De Neys, 2016, 2017; Gürçay & Baron, 2017; Kahane et al., 2018; Millar, Turri & Friedman, 2014; Millar, Starmans, Fugelsang & Friedman, 2016). As noted in the Introduction, these are different questions (i.e., acceptability of sacrifice vs choosing who to

sacrifice). As such, our findings do not address whether our manipulations of group membership would affect people's judgments of acceptability.⁵

One noteworthy aspect of our experiments is that participants made judgments about groups to which they did not belong. The participants were Americans, but they made judgments about story characters who were Estonian and Latvian. As such, our findings do not simply reflect differences between how people view ingroup and outgroup members. Moreover, the findings show that participants treat members from different outgroups as distinct from one another. Although they treated the members of each national group as interchangeable with other members of that group (e.g., Estonians as interchangeable with other Estonians), they did not treat them as interchangeable with members of other groups. It is as if the members of each nationality are viewed as separate fungible resources (see Bernard, Ottenberg & Redl, 1965, for how this leads to dehumanization).

However, the military setting of our story likely contributed to the findings, as soldiers who share the same nationality might be perceived as more coherent than civilians from that country. Likewise, soldiers might be viewed as resources *belonging* to their armies or countries, and so the military setting might have prompted fairness-based reasoning. Participants may have felt it would be unfair to deprive one "owner" of resources to benefit a different one (e.g., Millar et al., 2014). On both accounts, civilians might be less likely to be viewed as interchangeable. It is also possible that characterizing soldiers by their age (or month of birth) increased perceptions of them as a collection of unique individuals, as when people are the same age they are born separately rather than jointly. This claim is premised on evidence that joint vs independent activity affects the extent to which people are viewed as part of a group or as distinct individuals (Bartels & Burnett, 2011).

Finally, our findings may also help make sense of some potentially puzzling findings reported by Kurzban et al. (2012). In their study, participants read vignettes in which they imagined themselves in situations where they could sacrifice one person to save five. Across conditions, the vignettes specified different relations between participants, the victim, and the endangered people. For example, some vignettes specified that all of these parties were family members, while other vignettes specified they were all strangers. Interestingly, participants were more likely to accept the sacrifice when the victim was from the same family as themselves and the five endangered people, compared with when all the characters

⁵In two further studies, not reported here, we found no effect of nationality on judgments of acceptability — people found it equally acceptable (or unacceptable) to save several endangered people by sacrificing a victim regardless of whether the victim was from the same nationality or from a different nationality. This suggests that these two kinds of decisions (i.e., acceptability of sacrifice versus who to sacrifice) might be affected by different factors.

were strangers. Further, they were even more willing to sacrifice one family member to save others than to sacrifice a stranger for this purpose. Kurzban et al. noted that these findings are puzzling as it would be evolutionarily adaptive to maximize inclusive fitness, and to prefer preserving one's own genes by prioritizing one's family members (Axelrod & Hamilton, 1981). However, from the perspective of the current experiments, their findings may indicate that people view members of a single family as more interchangeable with one another than with unrelated people.⁶ As with nationalities, characterizing people as members of a family may increase perceptions that they are fungible. Further research is needed to test the boundary conditions of fungibility.

6 References

- Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211(4489), 1390–1396.
- Bartels, D. M., & Burnett, R. C. (2011). A group construal account of drop-in-the-bucket thinking in policy preference and moral judgment. *Journal of Experimental Social Psychology*, 47(1), 50–57.
- Bernard, V., Ottenberg, P., & Redl, F. (1965). Dehumanization: A composite psychological defense in relation to modern war. In M. Schwebel (Ed.), *Behavioral science and human survival* (pp. 64–82). Palo Alto, CA: Science and Behavior Books.
- Białek, M., & De Neys, W. (2016). Conflict detection during moral decision-making: Evidence for deontic reasoners' utilitarian sensitivity. *Journal of Cognitive Psychology*, 28(5), 631–639.
- Białek, M., & De Neys, W. (2017). Dual processes and moral conflict: Evidence for deontological reasoners' intuitive utilitarian sensitivity. *Judgment and Decision Making*, 12(2), 148–167.
- Campbell, D. T. (1958). Common fate, similarity, and other indices of the status of aggregates of persons as social entities. *Systems Research and Behavioral Science*, 3(1), 14–25.
- Costa, A., Foucart, A., Hayakawa, S., Aparici, M., Apesteguía, J., Heafner, J., & Keysar, B. (2014). Your morals depend on language. *PloS One*, 9(4), e94842.
- Cushman, F., Durwin, A. J., & Lively, C. (2012). Revenge without responsibility? Judgments about collective punishment in baseball. *Journal of Experimental Social Psychology*, 48(5), 1106–1110.
- Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review*, 5, 5–15.
- Gaertner, L., Iuzzini, J., & O'Mara, E. M. (2008). When rejection by one fosters aggression against many: Multiple-victim aggression as a consequence of social rejection and perceived groupness. *Journal of Experimental Social Psychology*, 44(4), 958–970.
- Gilbert, D. (2006). *Stumbling on happiness*. New York: Knopf.
- Goodwin, G. P., & Landy, J. F. (2014). Valuing different human lives. *Journal of Experimental Psychology: General*, 143(2), 778–803.
- Greene, J. D. (2015). The rise of moral cognition. *Cognition*, 135, 39–42.
- Gürçay, B., & Baron, J. (2017). Challenges for the sequential two-system model of moral judgement. *Thinking and Reasoning*, 23(1), 49–80.
- Hsee, C. K., Zhang, J., & Chen, J. (2004). Internal and substantive inconsistencies in decision-making. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 360–378). Oxford, England: Blackwell.
- Jeffreys, H. (1961). *Theory of probability* (3rd ed.). New York: Oxford University Press
- Kahane, G., Everett, J. A. C., Earp, B. D., Caviola, L., Faber, N. S., Crockett, M. J., & Savulescu, J. (2018). Beyond sacrificial harm: A two-dimensional model of utilitarian psychology. *Psychological Review*, 125(2), 131–164.
- Keren, G., & Teigen, K. H. (2010). Decisions by coin toss: Inappropriate but fair. *Judgment and Decision Making*, 5(2), 83–101.
- Kurzban, R., DeScioli, P., & Fein, D. (2012). Hamilton vs. Kant: Pitting adaptations for altruism against adaptations for moral judgment. *Evolution and Human Behavior* 33(4), 323–333.
- Lee, M. D., & Wagenmakers, E.-J. (2013). *Bayesian modeling for cognitive science: A practical course*. Cambridge, England: Cambridge University Press.
- Lickel, B., Hamilton, D. L., Wierzchowska, G., Lewis, A., Sherman, S. J., & Uhles, A. N. (2000). Varieties of groups and the perception of group entitativity. *Journal of Personality and Social Psychology*, 78(2), 223–246.
- Lickel, B., Miller, N., Stenstrom, D. M., Denson, T. F., & Schmader, T. (2006). Vicarious retribution: The role of collective blame in intergroup aggression. *Personality and Social Psychology Review*, 10(4), 372–390.
- Linville, P. W., & Fischer, G. W. (1991). Preferences for separating or combining events. *Journal of Personality and Social Psychology*, 60(1), 5–23.
- Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49(2), 433–442.

⁶This suggests similar (counter-intuitive) findings that could arise in our vignettes. For example, Estonian participants might prefer saving three Estonians by sacrificing another Estonian over accomplishing this by sacrificing a Latvian, even though this would go against group interests. However, even if true, group interests likely trump the tendency to view members of a nationality as fungible in some instances (e.g., if the soldiers from the two armies were in conflict with one another).

- Millar, C., Starmans, C., Fugelsang, J., & Friedman, O. (2016). It's personal: The effect of personal value on utilitarian moral judgments. *Judgment and Decision Making, 11*(4), 326–331.
- Millar, C., Turri, J., & Friedman, O. (2014). For the greater goods? Ownership rights and utilitarian moral judgment. *Cognition, 133*(1), 79–84.
- Muda, R., Niszczoła, P., Bialek, M., & Conway, P. (2017). Reading dilemmas in a foreign language reduces both deontological and utilitarian response tendencies. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 44*(2), 321–326.
- Newheiser, A., Sawaoka, T., & Dovidio, J. F. (2012). Why do we punish groups? High entitativity promotes moral suspicion. *Journal of Experimental Social Psychology, 48*(4), 931–936.
- Novemsky, N., & Kahneman, D. (2005). The boundaries of loss aversion. *Journal of Marketing Research, 42*(2), 119–128.
- Pereira, A., & van Prooijen, J. (2018). Why we sometimes punish the innocent: The role of group entitativity in collective punishment. *PloS One, 13*(5), e0196852.
- Petrinovich, L., O'Neill, P., & Jorgensen, M. (1993). An empirical study of moral intuitions: Toward an evolutionary ethics. *Journal of Personality and Social Psychology, 64*(3), 467–478.
- Sjöström, A., & Gollwitzer, M. (2015). Displaced revenge: Can revenge taste “sweet” if it aims at a different target? *Journal of Experimental Social Psychology, 56*, 191–202.
- Tasimi, A., & Gelman, S. A. (2017). Dirty money: The role of moral history in economic judgments. *Cognitive Science, 41*(S3), 523–544.
- Thaler, R. H. (1985). Mental accounting and consumer choice. *Marketing Science, 4*(3), 199–214.
- Thaler, R. H. (1990). Anomalies: Saving, fungibility, and mental accounts. *Journal of Economic Perspectives, 4*(1), 193–205.
- Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making, 12*(3), 183–206.
- Thomson, J. J. (1985). Double effect, triple effect and the trolley problem: Squaring the circle in looping cases. *Yale Law Journal, 94*(6), 1395–1415.
- Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *Quarterly Journal of Economics, 106*(4), 1039–1061.
- Uhlmann, E. L., Pizarro, D. A., Tannenbaum, D., & Ditto, P. H. (2009). The motivated use of moral principles. *Judgment and Decision Making, 4*(6), 476–491.
- Uhlmann, E. L., & Zhu, L. L. (2013). Money is essential: Ownership intuitions are linked to physical currency. *Cognition, 127*(2), 220–229.
- Whitehouse, H. (in press). Dying for the group: Towards a general theory of extreme self-sacrifice. *Behavioral and Brain Sciences*.