

New Evidence for Utilitarianism in Everyday Moral Judgment

Emily Rosenzweig

Tulane University

Erik G. Helzer

The Johns Hopkins Carey Business School

Abstract

When, if ever, do people apply utilitarian logic to their moral judgments? In response to recent critiques of the methods traditionally used to answer this question, we offer a new approach by investigating whether people show evidence of engaging in utilitarian moral calculus vis-à-vis the weighting of harm and benefit in assessments of wrongness. In three studies, we show that in judging everyday immoral behavior, people weight an action's benefits against its harms, and use the former as a mitigating factor in formulating moral judgments. Strikingly, we demonstrate that benefits reaped by the perpetrator of the immoral act lessen perceived wrongness: the more benefit a perpetrator accrues from their own bad act, the less wrong people perceive that act to be. Our work both advances the field's understanding of people's commitment to utilitarianism, and paves the way for additional research on the way benefits are factored into everyday moral judgment.

New Evidence for Utilitarianism in Everyday Moral Judgment

As a normative ethical theory, utilitarianism stipulates that the moral permissibility of an act depends upon its overall utility; one must consider the amount of benefit afforded to all involved, subtract the amount of harm that results, and from this generate a normative judgment about the (im)morality of the act. There is much to like about the utilitarian framework, particularly as an arbiter of public policy, wherein a concern for the “greater good” can bridge divides caused by ideology, bias, and competing value claims (Sidgwick, 1874). Perhaps as a result, many psychologists have attempted to determine when, if ever, people will endorse utilitarianism over other moral concerns (Greene, Morelli, Lowenberg, Nystrom, & Cohen, 2008; Greene, Cushman, Stewart, Lowenberg, Nystrom, & Cohen, 2009; Uhlmann, Pizarro, Tannenbaum, & Ditto, 2009, Bartels, 2008; Cushman, Young, & Hauser, 2006; Nichols & Mallon, 2006).

This research has yielded great insight into the variety of social, cognitive, and emotional factors that push people towards or away from acting in accordance with the greater good. However it has relied almost exclusively on sacrificial dilemmas, such as the trolley problem, as a proxy for utilitarian choice. Recent critiques have raised serious questions about whether responses to these sacrificial scenarios are reflective of genuine utilitarianism (Kahane, 2015; Kahane et al., 2015; Bartels & Pizarro, 2011; Djeriouat & Tremoliere, 2014; for review, see Helzer et al., 2016). Theoretical work has highlighted the fact that rejecting a deontological position is not the same as embracing a utilitarian one, and argues that sacrificial dilemmas fail

to test for the true markers of utilitarianism (Kahane, 2015); empirical work has found that ostensibly utilitarian responses to the trolley problem are more common among individuals who are lacking in moral character (e.g., those higher in psychopathy), suggesting these responses may be motivated by something far less high-minded than true commitment to utilitarianism. (Djeriouat & Tremoliere, 2014; Kahane et al., 2015).

What might be a better test of people's commitment to utilitarianism, particularly in their judgment of moral wrongdoing? One way forward might be to directly examine whether people engage in utilitarian calculus: do they counterweight harm and benefit when making moral judgments? There is no question that people attend to harm when judging an action's wrongness (e.g., Gray, Schein, & Ward, 2014): the more harm a particular action brings, all things being equal, the more wrong it seems. But what about benefit? Do people assess the benefits that result from moral wrongdoing and treat them as a factor that mitigates the wrongness of a harmful act, as utilitarianism would recommend?

Surprisingly, only a small handful of existing papers have explored how benefits affect moral judgments of a morally-troubling act. We know from this work that people are more willing to sacrifice an innocent person when 100 people are saved than when only five people are saved, and that more sophisticated moral thinkers will allow Heinz to steal pharmaceuticals in order to save his wife (Bonnenfon, Shariff, & Rahwan, 2016; Kohlberg, 1981). But these papers only include scenarios in which an agent is acting in order to bring benefits to an innocent or needy third party. In reality, the most common reason people commit moral violations is not to help innocent strangers tied to a train track or ailing family members, but to help *themselves*.

People typically lie, cheat, and steal in order to reap *personal* gain, and for this reason we suspect that adherence to utilitarian logic in everyday moral judgment will prove more challenging than these demonstrations would suggest.

Utilitarian logic requires that any benefits, including those reaped by the perpetrator of the act, should be tallied “toward the good.” Thus the prediction implied by a strong utilitarian account is that bad acts that benefit their perpetrator greatly should be perceived as *less wrong* than bad acts that confer little benefit. Do people actually treat the benefit a moral violation yields to its perpetrator as a factor that *mitigates* the wrongness of that action? Or are they repelled by the idea of ‘rewarding’ a wrongdoer for reaping greater benefit from someone else’s harm (cf., Inbar, Pizarro, & Cushman, 2012), resulting in harsher evaluation of those acts? This tension suggests that looking at cases where benefit accrues to a perpetrator is a particularly strong test of their implicit commitment to utilitarian logic in moral judgment; our paper is the first to look for evidence of whether they pass it.

Determining whether and how people are responsive to benefit is complicated by the fact that benefit is often naturally confounded with the harm caused by the moral violation. If John steals \$10 from Mary, his gaining \$10 is incumbent upon Mary’s losing it. Even in the case of less quantifiable moral wrongs, this is often true: if I gossip about my department chair to endear myself to my colleagues, the more salacious the information I share, the more social capital I gain, and the more the chair’s reputation will be harmed by my action.

Notably, this confound exists even in the stripped-down sacrificial dilemmas that have dominated the study utilitarian judgment. Imagine someone who makes the ‘utilitarian’ choice to kill one person in order to save five in a trolley problem. Although utilitarianism requires a calculation of benefit weighed against harm, this person may have been swayed only by potential benefit (saving one versus five), only by potential harm (killing one versus five), or by some mix of the two. Because harm and benefit are confounded, trolley problems as they are typically administered can offer little insight into whether people are actually engaging in the implicit calculation utilitarianism requires. A new approach is needed in order to address this question.

We introduce several methodological advancements in response. First, in order to test for the unique causal effect of perpetrator benefit on judgments of moral wrongness, we found novel ways to experimentally manipulate perceptions of benefit—both increasing and decreasing it—that were independent of the stimuli themselves. This approach has the advantage of leaving the actual moral violations participants were exposed to, as well as the amount of harm incurred by the victim, identical across conditions. Second, in order to promote external validity, we used realistic moral violations as our stimuli, rather than baroque scenarios or thought experiments. All of our studies used stimuli drawn directly from a comprehensive list of everyday immoral acts generated by participants in a large-scale momentary assessment study (Hoffman et al., 2014).

In sum, our broad purpose was to test whether people show an implicit commitment to utilitarian logic by, first, independently weighting benefit and harm in their computations of

immoral actions, and second, treating benefit as a mitigating factor of wrongness. In all studies, we focus on benefits accrued only by the perpetrator of the immoral act, which provides a particularly strong test of these questions. It is altogether possible, though inconsistent with utilitarian logic, that people may attend only to harm in their moral judgments, or may attend to benefit, but see acts that confer great benefit to the perpetrator as *more wrong* than those that confer little benefit. To preview our findings, we find the opposite. Benefits to the perpetrator independently feed into moral judgments, and (perhaps surprisingly) mitigate perceived wrongness of their acts, thereby offering the first evidence that people engage in sophisticated utilitarian calculus in their everyday moral judgments.

Study 1

Method

Participants. Eighty one Americans participated in Study 1 in exchange for a small payment on Amazon's Mechanical Turk. Given Study 1's fully within-subjects correlational design, this sample size offered 80% power to detect an effect size of .4 or larger (Westfall, Kenny, & Judd, 2014).

Materials. A primary goal of this project was to test our research questions using a wide range of naturalistic moral violations as stimuli. Recently, Hoffman, Wisnecki, Brandt, & Skitka (2014) conducted an ecological momentary assessment study in which they asked 1252 people to report (at multiple time points) any moral or immoral act that they had "committed, were the target of, witnessed, or learned about within the past hour" (p. 1341). We requested and received access to the resulting dataset of 3,828 actions that participants recorded and coded

as either moral or immoral. We filtered out all actions that participants had coded as moral, leaving a sample of 1,799 immoral actions. We also filtered out all acts that people reported having learned about on television or in other media like radio or books, out of a desire to exclude the type of extreme moral violations (e.g. serial killers) commonly found there. This left us with a final sample of 1268 everyday immoral actions.

In order to select a range of stimuli that would reflect the breadth of moral violations represented in the filtered dataset, we first grouped the remaining violations based on Hoffman et al.'s coding of the moral foundation to which they belonged (using Graham et al.'s [2013] five foundation model) and then looked for common themes within each foundation. For example, within the domain of authority, we identified four main themes: breaking the law, breaking policy or disrespecting one's boss, slacking on the job, and disrespecting one's parents or family. Ultimately we identified 29 themes that captured the large majority of the 1268 violations, and selected one moral violation from Hoffman et al.'s dataset that was representative of each theme. For example, for the theme of 'disrespecting one's family', we selected the entry "my nephew spoke disrespectfully to his mother;" for the theme 'breaking the law' we chose the entry "car failed to yield and then ran through a red light." We modified the text of each moral violation as little as possible in the process of turning them into stimuli for our participants to evaluate. Thus the stimuli generated from these two items were "J.R. spoke disrespectfully to his mother" and "J.R. failed to yield while driving and ran a red light." The end result is a stimuli set of 29 moral violations that reflect the wide range of everyday immoral acts that people report being exposed to (see Appendix A for list of stimuli).

Procedure. Participants first saw all 29 violations, in random order, one at a time. For each violation they were asked “how morally wrong is this,” and responded on a scale ranging from 1 (*not at all*) to 5 (*extremely*). Next, participants were shown the same 29 moral violations and were asked to tell us how much each act benefitted the person who committed it.

Participants read:

Acts can benefit the person who commits them in many ways, from financially, emotionally, sexually, to simply by making life easier for themselves. Thus there are no right or wrong answers to this question; please just use your own judgment to tell us how much you feel each of these behaviors benefits the person who commits them, on a scale ranging from No Benefit At All to Extremely Beneficial

Participants then rated those same stimuli based on how harmful each act was to others. As with benefit, we emphasized that “an act can harm others in many ways, from direct physical injury, to emotional harm, to damaging society as a whole,” and asked participants to report how harmful each act was on a scale ranging from 1 (*Not Harmful At All*) to 5 (*Extremely Harmful*). Finally, participants reported how frequent or common those 29 violations were, on a scale ranging from 1 (*Not common at all*) to 5 (*Extremely common*).

Results

Descriptive statistics for each of the 29 moral violations are included in Table 1.

Violations spanned the full range of the scale in wrongness, from the least wrong (*J.M. lusted*

after women) to the most wrong (*J.M. made sexual advances toward a minor*). On average, the set of stimuli was seen as moderately harmful ($M_{\text{harm}} = 3.06$, $SD_{\text{harm}} = 1.03$) and as conferring modest benefit to its perpetrator ($M_{\text{benefit}} = 1.99$, $SD_{\text{benefit}} = 0.98$).

Because Study 1 employed a fully-crossed within-subjects design, in which each participant evaluated all 29 stimuli, we fit multi-level models treating participant and stimuli as crossed random factors (Judd, Westfall, & Kenny, 2012). Our primary question was whether the amount of benefit people believed an act offered to its perpetrator was related to their moral judgment of those acts, and if so, what direction that relationship was in. Was it positive, such that acts seen as conferring greater benefits were also seen as more wrong? Or was it negative, suggesting people were viewing the acts as utilitarians would, reducing the wrongness of the act in proportion to the benefit it conferred?

Our first model included only perceptions of benefit as a predictor of judgments of moral wrongness. Supporting the utilitarian account, the more benefit participants believed a moral agent received from their action, the less wrong they saw that action as being ($B_{\text{benefit}} = -.164$, 95% CI = $[-.209, .119]$, $t(2329) = -7.24$, $p < .000$). Next we tested whether the weight given to benefit was independent of the weight participants gave to other constructs, such as harm or the commonness associated of the act. As expected based upon past work and theory (Gray et al., 2014), actions that were seen as more harmful were also seen as more wrong ($B_{\text{harm}} = .480$, 95% CI = $[.443, .517]$, $t(2276) = 25.57$, $p < .000$). In addition, actions that were seen as more commonplace were seen as less wrong ($B_{\text{common}} = -.129$, 95% CI = $[-.172, -.086]$, $t(2324) = -5.85$, $p < .001$). However even when both harm and commonness were entered into the model as covariates, benefit to the perpetrator continued to negatively predict judgments of how wrong

their act was ($B_{\text{benefit}} = -.091$, 95% CI = $[-.130, -.052]$, $t(2197) = -4.47$, $p < .000$). Results from this full model reveal remarkable sophistication in participants' moral judgment: these independent effects suggest that participants parsed and weighted the contributions of three different features of each moral violation when making their moral judgments, and when doing so treated benefit in a way that is consistent with utilitarian logic.

Study 2a

A strict utilitarian account predicts that benefits should not merely be negatively correlated with moral wrongness, but that increasing or decreasing the perceived benefits of an act should change perceptions of wrongness. Thus, our goal in Study 2a was to manipulate perceptions of benefit while holding the specific moral violations participants were exposed to constant, so as to test whether perceptions of benefit are causally linked to perceptions of moral wrongness.

Method

Participants. Two-hundred twenty Americans (124 women, $M_{\text{age}} = 35$) participated in exchange for a small payment through Amazon's Mechanical Turk. Because the manipulation used in this study was novel and its effect size was unknown, and because this study was a two-condition between-subjects design, we opted for a sample with at least 100 participants per condition.

Materials: We hypothesized that by narrowing the scope of benefit that we allowed participants to consider—namely focusing them only on material gain—we might be able to reduce perceptions of benefit relative to the broader definition of benefit used in Study 1. We

conducted a pre-test, in which we asked 91 participants to rate all 29 violations from Study 1 using a narrower definition of benefit restricted to material gain. From these data we identified eight moral violations where these restricted ratings of benefit were substantially lower than participants' ratings for those items using the broader definition of benefit used in Study 1. These eight moral violations (items # 1, 3, 4, 9, 12, 15, 17, and 27 from Table 1) became our stimuli set for Study 2.

Procedure. In the main study, participants were randomly assigned to either the broad benefit or narrow benefit condition. All participants started by reporting how much benefit they felt each of 8 actions provided to the person who committed them. Participants in the broad benefit condition made this rating using the same broad definition of benefit as in Study 1. Participants in the narrow benefit condition read instructions including the following text: "People often think about benefits too broadly—for example, by thinking about an array of possible, abstract benefits associated with a behavior...Please consider only immediate, material benefits, such as money or material goods that [John] stands to gain as a result of his behavior." Participants rated harm and wrongness using the same instructions as in Study 1, and participants rated benefit, harm, and wrongness on the same scales as in Study 1.

Results

Because our between-subjects manipulation was designed to influence perceptions of benefit for all of the moral violations participants were exposed to, we created an average benefit, average wrongness, and average harm composite for each participant. Our manipulation was effective: focusing participants only on material benefits reduced the amount

of benefit they felt moral agents received across the eight bad acts ($M_{\text{broad}} = 1.67$, $M_{\text{narrow}} = 1.46$), $t(218) = 2.78$, $p = .006$, $d = .376$. Additionally, replicating our results from Study 1 at the between-subjects level, the more benefit participants saw our moral agents as receiving from their immoral acts, the less wrong they believed those acts were, $B_{\text{benefit}} = -.213$, 95% CI = [.061, .364] $t(218) = -2.78$, $p = .006$. And as in Study 1, this effect persisted when perceived harm was included in the model as a predictor of wrongness, $B_{\text{benefit}} = -.214$, 95% CI = [.102, .326], $t(217) = -3.77$, $p < .000$.

Most critically, as the utilitarian account would predict, perceived benefit mediated the relationship between our manipulation and perceptions of moral wrongness. We tested the significance of this indirect effect using a bootstrapping procedure with 10,000 bootstrapped samples. The estimate of the indirect effect of condition on moral wrongness through benefit was .046, 95% CI = [.013, .104]. Including harm as a covariate, the indirect effect remained essentially the same, with an estimate of .045, 95% CI = [.014, .093]. Thus a manipulation that focused people only on the material benefits of an action significantly reduced perceptions of benefit, which in turn significantly increased perceptions of how wrong those actions were. This demonstrates that perceptions of benefit are not just correlational, but causally linked to perceptions of wrongness.

Study 2b

To provide further support for the utilitarian account, and in particular the causal link between perceived benefit and moral wrongness, in Study 2b, we experimentally *increased*

how much benefit participants saw in a set of moral violations, to examine whether that increase caused a reduction in the perceived wrongness of the act

Method

Participants. Two-hundred two Americans (90 women, $M_{\text{age}} = 33$) participated in exchange for payment on Amazon's Mechanical Turk. We chose this sample size based on the sample size from Study 2a.

Materials. Because our instructions defining benefit for participants in Study 1 were already quite broad, there was little room to further expand them. Instead, we hoped to capitalize on the fact that while those instructions named a variety of types of benefits a moral agent might receive from their action, participants may not have thought through them all when evaluating our stimuli. By forcing participants to consider the full array of benefits a moral agent might receive from each action, we aimed to increase perceptions of benefit while again leaving the specific stimuli participants were evaluating identical across conditions. Thus in developing the materials for Study 2b, we first analyzed the existing set of 29 moral violations from Study 1, to identify those that could plausibly offer multiple types of benefit to their perpetrator. We then conducted a pre-test, in which we presented 41 participants with a list of different types of benefits that could accrue to a person who committed each act, and asked them to report what they felt the primary benefit was and how much benefit they felt was associated with each action. Based on this information we selected 12 moral violations for in Study 2b that had room for movement in the amount of perceived benefit they conferred.

These twelve moral violations (items # 1, 3, 5, 10, 12, 17, 18, 21, 24, 26, 27, and 28 from Table 1) became our stimuli set for Study 2.

Procedure. All participants evaluated 12 moral violations, each attributed to someone with the initials J.R. Participants were randomly assigned to either the boost-benefit condition or the control condition. In the boost-benefit condition, prior to reporting how much benefit each moral violation offered its perpetrator, participants first answered a question asking them to identify the *primary* benefit that moral violation offered. This offered us the opportunity to present participants with a list of five different benefits that they had to read, consider, and choose from. We expected that engaging in this task would heighten participants' awareness of the amount of potential benefit conferred to the perpetrator.

For each violation, participants chose from a list of five different benefits, including "J.R. stands to benefit financially from this action (i.e. to make money or avoid losing money)," "J.R. stands to benefit psychologically or emotionally from this action (i.e. to feel better in some way or to avoid feeling worse)," "J.R. stands to benefit reputationally from this action (i.e. to improve reputation with some people or to avoid losing face)," "J.R. stands to benefit socially from this action (i.e. to gain status or the approval of someone else, or to avoid losing status), and "J.R. stands to gain in other ways from this action (e.g. saving time)". Participants then reported how much benefit they thought J.R. gained from each action. In the control condition, participants simply answered the same benefit question from Study 1 without first explicitly considering the full array of benefits the action might have offered. After reporting how much benefit they felt each moral infraction conferred, all participants then reported how wrong

each action was and how much harm it caused, using the same instructions and scales as in Study 1 and 2a.

Results

As in Study 2a, given the between-subjects design, we averaged together each participant's 12 benefit ratings, 12 wrongness ratings, and 12 harm ratings, resulting in three composite variables. Supporting the effectiveness of our manipulation, requiring participants to take the time to think through the full range of benefits that could be associated with each moral violation led to an increase in the total amount of benefit they thought the moral agent received from their act ($M_{\text{benefit-boost}} = 1.81$, $M_{\text{control}} = 1.41$, $t(200) = -5.10$, $p < .000$, $d = .72$). And as predicted by the utilitarian account, the more benefit an immoral act was seen as conferring to the person who committed it, the less wrong that act was seen to be ($B_{\text{benefit}} = -.283$, 95% CI [.139, .427], $t(200) = -3.87$, $p < .000$). This relationship once again remained when perceived harm was included into the model as a covariate ($B_{\text{benefit}} = -.165$, 95% CI [.064, .266], $t(199) = -3.25$, $p = .001$).

And as in Study 2a, perceived benefit mediated the relationship between our manipulation and perceptions of moral wrongness. We tested the significance of this indirect effect using a bootstrapping procedure with 10,000 bootstrapped samples. The estimate of the indirect effect of condition on moral wrongness through benefit was .131, with a 95% confidence interval ranging from -.218 to -.072. This mediation remained significant when harm was included as a covariate in the model, where the estimate of the indirect effect was -.077 and the 95% confidence interval ranged from -.132 to -.213. Thus a manipulation that directed

people's attention to the range of benefits moral agents might derive from their bad acts significantly *increased* perceptions of benefit, which in turn significantly *reduced* perceptions of how wrong those acts were.

General Discussion

Across three studies, looking both between and within subjects, we find evidence that people independently weight both harm and benefit when making judgments of real-world examples of moral wrongdoing. While it is easy to imagine participants recoiling at the idea of rewarding a moral perpetrator for benefitting from their own bad act—and perhaps even punishing them more because of it—participants instead treated benefit as a mitigating factor that reduced the perceived wrongness of a violation.

This research thus offers initial evidence suggesting people apply genuine utilitarian logic in their everyday moral judgments. Our approach to studying these questions responds to some of the most pointed recent critiques of the existing literature, and offers many improvements over traditional sacrificial dilemmas: a method of independently measuring harm and benefit; a procedure for manipulating these factors without altering the objective features of the scenario; and increased external validity.

Although we did not set out to test it directly, our work also addresses another distinguishing feature of utilitarianism—its requirement that the wellbeing of all people be valued equally (i.e., *impartiality*; Mill, 1863/1998). It is striking that participants treated a moral wrongdoer—arguably the person to whom moral leniency might be least palatable—as someone nonetheless worthy of having his benefit factored in to the greater good. This

suggests (albeit indirectly) at least some degree of impartiality in participants' deployment of the moral calculus.

These advancements notwithstanding, one question we cannot address is whether participants weighted harm and benefit equivalently in their computation of moral wrongness. While it may be tempting to compare the beta weights of harm and benefit, we caution against this, since there is no reason to believe participants applied the same psychological "ruler" to their assessments of each (and loss aversion suggests this is quite unlikely). One direction for future research would be to devise measures that would allow for direct comparisons between harm and benefit, although within normative ethics there is much debate about how such comparisons ought to be made (Smart & Williams, 1973).

One of the most remarkable features of our findings is just how sophisticated people's everyday moral judgments appear. Participants in Study 1 made rapid successive moral judgments, and rated the benefit and harm associated with each act only after having decided on its moral severity. Yet despite this, and despite the natural tendency for harm and benefit to be confounded, participants nonetheless parsed and counterweighted the unique effects of each factor. This suggests that when people encounter immoral acts in their everyday lives, the speed with which these judgements unfold is not necessarily indicative of a lack of nuance in them.

Beyond utilitarianism, it is notable that this work is among the first to empirically investigate the weight given to perpetrator benefit in moral judgment. Potential benefits are arguably the primary motivator of immoral behavior for most perpetrators; thus, it is surprising

that there is so little research on the topic. Going forward, there are many unanswered questions that may serve as the basis for future research. One set of questions pertains to the type of benefit reaped; are all benefits weighted equally when evaluating the wrongness of an act or are some benefits more palatable than others? There are also applied implications of our work that can now be tested: although people give implicit weight to benefit in their moral judgments, would they explicitly endorse it as a mitigating factor? For example, would individuals at a sentencing hearing respond positively to arguments that a perpetrator's actions brought them benefit, and thus were less wrong?

Ultimately, the role that benefit plays in moral judgment is central to an understanding of people's commitment to utilitarianism, but its value is not limited to one ethical tradition or set of questions about moral psychology. We believe the approach introduced here offers a new lens and method to explore important questions related to the role of self-interested motives in people's moral judgments of others acts.

Author's Note

Both authors contributed to the research concept and study designs. Data collection and analyses were performed by E. Rosenzweig. E. Rosenzweig drafted the bulk of the manuscript and E. Helzer provided critical revisions. Both authors approved the final version of the manuscript for submission.

We would like to thank David Pizarro for his helpful insights as we developed the manuscript.

References

- Bartels, D. (2008). Principled moral sentiment and the flexibility of moral judgment and decision making. *Cognition, 108*, 381–417.
- Bartels, D., & Pizarro, D. A. (2011). The mismeasure of morals: Antisocial personality traits predict utilitarian responses to moral dilemmas. *Cognition, 121*, 154–161.
- Bonnefon, J.F., Shariff, A., & Rahwan, I. (2016). The social dilemma of autonomous vehicles. *Science, 352*, 1573-1576.
- Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. *Psychological Science, 17*, 1082–1089.
- Djeriouat, H., & Tremolière, B. (2014). The Dark Triad of personality and utilitarian moral judgment: The mediating role of honesty/ humility and harm/care. *Personality and Individual Differences, 67*, 11–16.
- Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S., & Ditto, P.H. (2013). Moral foundations theory: the pragmatic validity of moral pluralism. *Advances in Experimental Social Psychology, 47*, 55-130.

Gray, K., Schein, C., & Ward, A. F. (2014). The myth of harmless wrongs in moral cognition: Automatic dyadic completion from sin to suffering. *Journal of Experimental Psychology: General, 143*, 1600-1615.

Greene, J.D., Cushman, F.A., Stewart, L.E., Lowenberg, K., Nystrom, L.E., & Cohen, J.D. (2009). Pushing moral buttons: The interaction between personal force and intention in moral judgment. *Cognition, 111*, 364-371.

Greene, J.D., Morelli, S.A., Lowenberg, K., Nystrom, L.E., & Cohen, J.D. (2008). Cognitive load selectively interferes with utilitarian moral judgment. *Cognition, 107*, 1144-1154.

Helzer, E.G., Fleeson, W., Furr, R.M., Meindl, P., & Barranti, M. (2016). Once a utilitarian, consistently a utilitarian? Examining principledness in moral judgment via the robustness of individual differences. *Journal of Personality*, in press

Hofmann, W., Wisneski, D.C., Brandt, M.J., & Skitka, L.J. (2014). Morality in everyday life. *Science, 345*, 1340-1343.

Inbar, Y., Pizarro, D.A., & Cushman, F. (2012). Benefiting from misfortune: When harmless actions are judged to be morally blameworthy. *Personality and Social Psychology Bulletin, 38*, 52-62.

Kahane, G. (2015). Sidetracked by trolleys: Why sacrificial moral dilemmas tell us little (or nothing) about utilitarian judgment. *Social Neuroscience*, *10*, 551-560.

Kahane, G., Everett, J. A. C., Earp, B. D., Farias, M., & Savulescu, J. (2015). "Utilitarian" judgments in sacrificial moral dilemmas do not reflect impartial concern for the greater good. *Cognition*, *134*, 193–209.

Kohlberg, Lawrence (1981). *Essays on Moral Development, Vol. I: The Philosophy of Moral Development*. San Francisco, CA: Harper & Row.

Mill, J. S. (1863/1998). *Utilitarianism*. New York: Oxford University Press.

Nichols, S., & Mallon, R. (2006). Moral dilemmas and moral rules. *Cognition*, *100*, 530–542.

Sidgwick, H. (1874). *The Methods of Ethics*. Cambridge: Cambridge University Press.

Smart, J. J. C., & Williams, B. (1973). *Utilitarianism for and against*. New York: Cambridge University Press.

Uhlmann, E. L., Pizarro, D. A., Tannenbaum, D., & Ditto, P. H. (2009). The motivated use of moral principles. *Judgment and Decision Making, 4*, 476–491.

Westfall, J., Kenny, D. A., & Judd, C. M. (2014). Statistical power and optimal design in experiments in which samples of participants respond to samples of stimuli. *Journal of Experimental Psychology: General, 143*, 2020-2045.

Table 1:

Descriptive statistics for Study 1 (N=81)

| Item # | Stimuli | Wrong | | Benefit | | Harm | | Common | |
|--------|--|-------|------|---------|------|------|------|--------|------|
| | | M | SD | M | SD | M | SD | M | SD |
| 1 | J.M. got his/her co-worker in trouble (N=80 for benefit) | 3.16 | 1.08 | 1.89 | 1.04 | 3.60 | 0.89 | 2.81 | 0.85 |
| 2 | J.M. aborted a baby | 2.72 | 1.56 | 2.94 | 1.37 | 3.41 | 1.60 | 2.85 | 1.18 |
| 3 | J.M. intentionally got crumbs on his/her roommate's blanket | 2.15 | 0.92 | 1.23 | 0.66 | 2.01 | 0.86 | 1.91 | 0.91 |
| 4 | J.M. lusted after women (N=80 for harm) | 1.64 | 1.05 | 2.17 | 1.18 | 1.78 | 0.98 | 4.23 | 1.02 |
| 5 | J.M. told a friend's secret | 3.26 | 0.96 | 1.48 | 0.76 | 3.31 | 0.92 | 3.23 | 1.05 |
| 6 | J.M. made inappropriate comments to women | 3.42 | 0.97 | 1.43 | 0.82 | 3.23 | 1.04 | 3.58 | 1.01 |
| 7 | J.M. threw garbage on the street | 2.88 | 0.90 | 1.47 | 0.73 | 2.85 | 0.85 | 3.42 | 1.19 |
| 8 | J.M. drove recklessly and rear-ended someone's car | 3.70 | 0.94 | 1.21 | 0.67 | 4.30 | 0.87 | 2.77 | 0.98 |
| 9 | J.M. lied about having to work to avoid in order to avoid spending time with family | 2.68 | 1.02 | 2.65 | 1.28 | 2.64 | 0.91 | 2.93 | 1.09 |
| 10 | J.M. badmouthed someone behind their back | 2.70 | 0.90 | 1.48 | 0.88 | 2.90 | 0.96 | 3.74 | 1.05 |
| 11 | J.M. didn't get up to give his/her seat to an elderly person because he/she was tired from working out | 2.77 | 1.24 | 2.40 | 1.25 | 2.60 | 0.92 | 2.67 | 1.00 |
| 12 | J.M. spoke disrespectfully to his/her mother | 2.99 | 1.10 | 1.47 | 0.91 | 2.79 | 1.05 | 3.54 | 1.08 |
| 13 | J.M. looked at hardcore porn | 2.01 | 1.23 | 2.59 | 1.29 | 1.95 | 1.19 | 4.00 | 1.06 |
| 14 | J.M. made sexual advances towards a minor | 4.68 | 0.74 | 1.69 | 1.02 | 4.38 | 1.02 | 2.21 | 1.19 |
| 15 | J.M. used offensive slang in conversation | 2.38 | 1.06 | 1.30 | 0.64 | 2.37 | 1.05 | 4.00 | 1.00 |
| 16 | J.M. looked up a woman's skirt on the subway | 3.84 | 1.09 | 2.20 | 1.14 | 3.30 | 1.11 | 2.28 | 1.12 |
| 17 | J.M. posted on Facebook that burger flippers don't deserve a living wage | 2.56 | 1.22 | 1.26 | 0.63 | 2.60 | 1.23 | 2.73 | 1.16 |
| 18 | J.M. talked loudly while in a quiet study room at the library | 2.31 | 0.98 | 1.31 | 0.74 | 2.37 | 0.97 | 2.78 | 1.05 |
| 19 | J.M. placed blame on a co-worker for his/her own mistake | 3.68 | 0.97 | 2.84 | 1.26 | 3.64 | 0.99 | 2.95 | 0.99 |
| 20 | J.M. stripped for money | 2.10 | 1.33 | 3.63 | 1.20 | 2.02 | 1.37 | 2.93 | 1.22 |

| | | | | | | | | | |
|----|--|------|------|------|------|------|------|------|------|
| 21 | J.M. threw rocks at a dog | 4.17 | 0.93 | 1.22 | 0.65 | 4.21 | 0.83 | 1.91 | 0.96 |
| 22 | J.M. passed off his/her own work and responsibilities to a subordinate | 3.04 | 1.10 | 3.07 | 1.22 | 3.17 | 1.00 | 3.28 | 0.94 |
| 23 | J.M. conducted personal business on company time | 2.23 | 0.91 | 3.14 | 1.17 | 2.59 | 1.00 | 3.72 | 0.99 |
| 24 | J.M. made a homophobic comment in conversation | 3.23 | 1.16 | 1.22 | 0.61 | 2.98 | 1.16 | 3.38 | 1.06 |
| 25 | J.M. had sex with a married man / woman | 3.72 | 1.16 | 2.58 | 1.32 | 3.79 | 1.17 | 3.17 | 1.15 |
| 26 | J.M. failed to yield while driving and ran a red light | 2.89 | 1.05 | 1.64 | 0.90 | 3.58 | 1.18 | 3.00 | 1.21 |
| 27 | J.M. hit someone in a bar fight | 3.23 | 1.06 | 1.89 | 0.97 | 3.93 | 0.82 | 2.80 | 1.07 |
| 28 | J.M. yelled at his/her spouse using crass language | 2.77 | 1.03 | 1.44 | 0.76 | 3.31 | 1.04 | 3.23 | 1.05 |
| 29 | J.M. stole from a grocery store (<i>N</i> =80 for benefit) | 3.78 | 0.95 | 2.86 | 1.25 | 3.28 | 1.06 | 2.90 | 1.11 |
