Actions Speak Louder Than Outcomes in Judgments of Prosocial Behavior

Daniel A. Yudkin, Annayah M. B. Prosser, and Molly J. Crockett
Yale University

Recently proposed models of moral cognition suggest that people's judgments of harmful acts are influenced by their consideration both of those acts' consequences ("outcome value"), and of the feeling associated with their enactment ("action value"). Here we apply this framework to judgments of prosocial behavior, suggesting that people's judgments of the praiseworthiness of good deeds are determined both by the benefit those deeds confer to others and by how good they feel to perform. Three experiments confirm this prediction. After developing a new measure to assess the extent to which praiseworthiness is influenced by action and outcome values, we show how these factors make significant and independent contributions to praiseworthiness. We also find that people are consistently more sensitive to action than to outcome value in judging the praiseworthiness of good deeds, but not harmful deeds. This observation echoes the finding that people are often insensitive to outcomes in their giving behavior. Overall, this research tests and validates a novel framework for understanding moral judgment, with implications for the motivations that underlie human altruism.

Keywords: moral judgment, prosociality, altruism

Supplemental materials: http://dx.doi.org/10.1037/emo0000514.supp

A young professional takes a moment from her hurried commute home from work to help an elderly man cross the street. What might drive such seemingly selfless behavior? Recent theoretical advances have proposed two distinct possibilities. On one hand, the woman may be motivated by the desire to see the man arrive safely on the other side (Batson, 1994). Under this view, it is anticipated *consequences* that drive her behavior. On the other hand, the woman may be motivated by a desire to attain the internal satisfaction of having carried out a good deed (Cialdini, 1991). From this perspective, it is not so much the social consequences of performing the deed driving the behavior as it is the "warm glow" the act produces (Andreoni, 1990; Crumpler & Grossman, 2008; Nunes & Schokkaert, 2003).

Recent models of moral cognition have attempted to describe the unique psychological properties of each of these motivations. According to these models, any good deed has both an "outcome value," which corresponds to the benefit it produces for others, and an "action value," which corresponds to the feeling associated with its enactment (Crockett, 2013; Cushman, 2013; Gesiarz & Crockett, 2015). This echoes work in economic theory that differentiates between generosity motivated by concern for the welfare of the recipient ("pure altruism") and generosity motivated by a sense of satisfaction from doing good ("warm glow"; Andreoni, 1990, 1995). In the current work, we build off this conceptualization, defining "action value" as any positive feeling that arises in an actor as the result of performing a prosocial deed, and outcome value as any positive consequences conferred to the recipient.

Past research suggests that both action and outcome value may be used as a basis for moral decision-making. Demonstrating the importance of outcome value, research has found that people are willing to help others even in the absence of personal incentives (Batson et al., 1999; Franzen & Pointner, 2012); and that they tend to invest more resources in charities believed to be effective (Null, 2011). Other work suggests that people will attempt to help others even if they are led to believe their own moods are unalterable, suggesting that warm glow is not the only factor compelling good behavior (Schroeder, Dovidio, Sibicky, Matthews, & Allen, 1988; see also Dovidio, Allen, & Schroeder, 1990).

On the other hand, there is also evidence that action value motivates prosocial behavior. For instance, research shows that people enjoy the act of giving, even if they know their actions will result in no real benefit (Crumpler & Grossman, 2008). Other work shows that people typically feel more positive affect and higher self-esteem after they have performed an act of giving (Aknin et al., 2013; Hitlin, 2007; Weinstein & Ryan, 2010). And more generally, research finds that a high degree of charitable giving is motivated by the pursuit of personal satisfaction, regardless of outcomes for others (Ribar & Wilhelm, 2002; see also Andreoni, 1990; Konow, 2010).

From this research, it is clear that both action and outcome value play important roles in moral decision-making. However, the manner in which these factors influence moral judgment is less understood (Carlson & Zaki, 2018). Judgment and decision-making are often treated as interrelated concepts in psychology

(e.g., Bazerman & Moore, 2008), yet they are not strictly identical. For instance, while some research suggests that people are more likely to perform acts they evaluate positively (e.g., Crockett, Siegel, Kurth-Nelson, Dayan, & Dolan, 2017; Herr, 1986), other researchers argue that substantive differences exist between the psychological processes underlying decision-making for the self versus judgment of others (Monin, Pizarro, & Beer, 2007). Thus, the degree to which such findings from the moral decision-making literature can be applied to moral judgment is unclear. Moral judgment is an important psychological concept in its own right, insofar as it has important practical consequences. For instance, political disagreements often hinge on people's evaluations of others' actions, or policies that pertain thereto. And cultural mores are often determined by the collective approval or disapproval of certain actions. Thus, the factors that contribute to moral judgment are an important area of investigation distinct from those relating to moral decision-making. Accordingly, we set out to examine how outcome and action values influence people's judgments of prosocial behavior.

Action and Outcome Values Affect Judgment

The current work draws inspiration from recent research highlighting the importance of action and outcome value in people's moral evaluations of harmful acts. Miller, Hannikainen, and Cushman (2014; Study 5) examined the extent to which action and outcome values contributed to people's moral condemnation of harmful behaviors. Participants are told a story of a man named John who has a terminal illness and sincerely wants to die. John asks Carl to carry out a mercy killing on him, and Carl takes pity on John and agrees to carry out his wish. Participants then read 23 items depicting different ways that Carl might kill John, including "Holding a shotgun to John's head and pulling the trigger" and "Stabbing John in the throat." According to each condition, participants were asked to rate the *action value*, *outcome value*, or *moral wrongness* of each of the 23 acts. To assess action value, the researchers asked participants to rate how *upsetting* it would be to "act out" performing each behavior as though it were part of a movie script (this measure is based on the notion of "evaluative simulation," in which one adopts the point of view of the person performing the behavior to assess its moral content, Lieberman & Lobel, 2012). To assess outcome value, they asked them to rate how much *suffering* each act would impose. Finally, to assess moral judgment, they asked people how *wrong* each act was. Then, the researchers examined the zero-order correlations between these items' mean action and outcome ratings and the mean ratings of wrongness. Results indicated that both action and outcome values significantly predicted participants' wrongness judgments, and the size of these effects were equivalent. This finding suggests that people are equally sensitive to both action and outcome values when making moral evaluations of harmful behaviors (see also Cushman, Gray, Gaffey, & Mendes, 2012; Hannikainen, Miller, & Cushman, 2014; Miller & Cushman, 2013).

We sought to extend these findings from the harm domain to judgments of prosocial behavior. Past work suggests that positive versus negative acts are processed differently in the human mind (e.g., Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Carver, 2006; Higgins, 1998; Janoff-Bulman, Sheikh, & Hepp, 2009). Thus, it is of interest whether people's judgments of prosocial

behaviors are similar or different to those pertaining to harmful acts. We sought to observe whether people's evaluations of the moral goodness of prosocial deeds would be sensitive both to how beneficial those deeds are (outcome value) and to how good they would feel to perform (action value). We also tested the relative strengths of these two evaluative systems. Studies of prosocial decision-making suggests that people may be more swayed by action value than by outcome value. For example, research on the "identifiable victim effect," which shows that people tend to donate more to identifiable victims than to large crowds of sufferers, shows that people may be more tempted to seek the warm glow of giving to a single victim than of dispersing resources to the impersonal crowd (Small, Loewenstein, & Slovic, 2007). Furthermore, recent work shows that people's choice of charitable donations is influenced more by subjective preference (that presumably elicits more warm glow), than by known effectiveness (Berman, Barasch, Levine, & Small, 2018). On the basis of these findings, we expect that while both action and outcome value will influence people' judgments of prosocial behavior, they will be more influenced by the former than by the latter.

We conducted three studies to test this hypothesis. Our general method was to have people rate either the praiseworthiness, the action value, or the outcome value of a set of good deeds that were selected to vary in both action and outcome value. For example, giving a flower to a child in a park would have a relatively high action value (because it feels good to perform), but a relatively low outcome value (as the child will likely abandon the flower a few minutes later). Meanwhile, writing a check to the Against Malaria foundation has a relatively low action value (writing a check is a relatively cold, emotionless action) but a high outcome value (the money can be used to purchase bed nets that can save hundreds of lives). We then examined the extent to which people's praiseworthiness ratings of these good deeds were sensitive to their action and outcome values. Praiseworthiness is an established metric in the moral psychology literature for assessing people's evaluations of moral deeds (e.g., Nelkin, 2016; Nelson-le Gall, 1985). We also examined, in the final study, the associations between action and outcome sensitivities across the prosocial and harm domains. All collected data is available online at the Open Science Framework (osf.io/k3pca). A schematic of the experimental designs can be found in Figure 1. All studies were approved by the University of Oxford Central Ethics Committee (Ref: MS-IDREC-C1-2015–098).

Study 1

To test our hypotheses, we needed to create a set of prosocial behaviors that would vary widely in action and outcome values, and could, therefore, be used to efficiently gauge the extent to which each of these factors predicted praiseworthiness. Consequently, our first step was to recruit a sample of participants and ask them to provide either an action value or an outcome value rating for a large set of prosocial behaviors. From this set, we were able to select a subset of deeds whose action and outcome values were largely uncorrelated and well distributed. Next, to determine how much action and outcome values predicted praise judgments, we recruited a second sample of participants and asked them to rate the moral praiseworthiness of each item (Figure 1a). We expected that, if both action and outcome values predict praise-
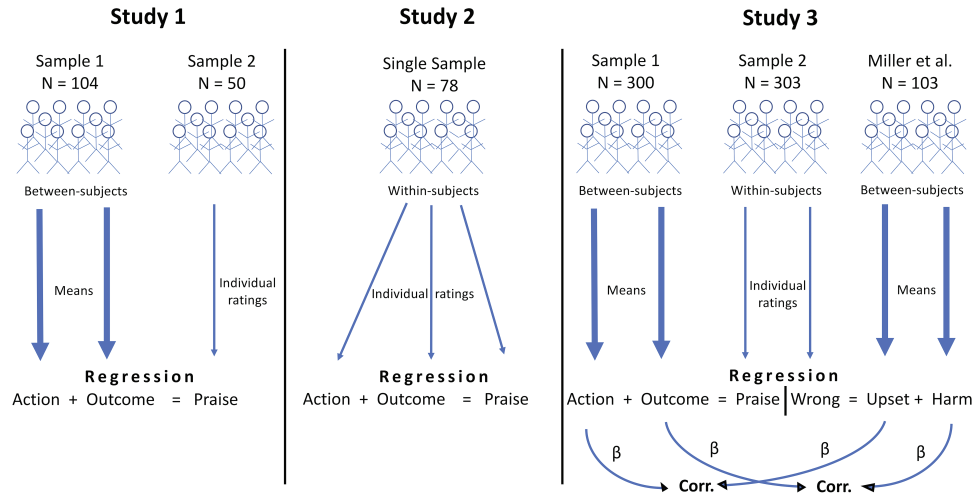
Figure 1. Schematic of experimental design for Studies 1–3. The graphic depicts the sample sizes used for each study, the experimental procedure used (within- or between-subjects design), and the type of rating (individual or mean) used in the participant-level regressions. See the online article for the color version of this figure.

worthiness, then variation in these two factors would be capable of explaining a significant proportion in the variance in people's praiseworthiness ratings.

Method

Participants. We collected two samples of participants. As we did not know what effect size to expect given lack of previous investigation on this exact topic, we followed the recommendations of Ledgerwood and Ratliff (2015) and attempted to collect approximately 50 participants per condition. Because participants in Sample 1 were randomly assigned to rate either action or outcome value, we recruited approximately 100 participants. We planned to exclude participants who demonstrated no variation in their ratings of action and outcome values (i.e., providing the same rating across the entire block), because our instructions specifically stated that the deeds varied in terms of how good they feel to perform and how much benefit they conferred, and that they should rate the deeds accordingly.

Our instructions stated that the deeds varied according to how good they feel to perform and how much benefit they confer, and so any participant who demonstrated no variation at all in their responses (i.e., providing the same rating for a whole block) was excluded for not following instructions correctly. This resulted in the exclusion of one participant from further analysis leading to a total of 104 participants (M_{age} 38.8, SD = 10.9, 51 men, 54 women). Sample 2 (praiseworthiness) consisted of 50 participants (M_{age} 35.6, SD = 11.2, 27 men, 23 women); no participants were excluded from this sample. Both samples were recruited from the TurkPrime platform (Litman, Robinson, & Abberbock, 2017) to participate in a study entitled "Easy and Interesting Survey, rate items on a characteristic!" in exchange for $2.

Procedure. To obtain a diverse set of prosocial behaviors, we first generated a list of 95 good deeds in a variety of domains through a review of literature, news reports, Web searches, and anecdotal experiences (for a complete list, see online supplementary materials). Then, according to randomized participant condi-

tion, we asked participants in Sample 1, after providing basic demographic information, to rate either the action or outcome value of each deed.

Instructions for both action and outcome measures were adapted from Miller et al. (2014). For action values, participants were told:

> In the following section, we will ask you to imagine that you are an actor playing a character in a movie. As part of the movie script, your character must perform various good deeds. Your task in this section is to rate **how good it would make you feel** to act out each good deed in the context of the movie set.

The scale on which participants responded went from 1 (*the LEAST good I can imagine from acting out a good deed*) to 10 (*the MOST good I can imagine from acting out a good deed*). For outcome values, participants were told:

> In the following section, you will be asked to evaluate a series of good deeds. For each good deed, please say how much of a positive effect you think doing that deed would have for any potential recipients. Remember: your ratings are not about how much the action would benefit *you*. The goal is to identify how much of a positive effect each action would have for the **any potential recipients of the act.**

The scale on which participants responded went from 1 (*the LEAST beneficial a good deed could be*) to 10 (*the MOST beneficial a good deed could be*). For praiseworthiness, participants were told:

> In the following section, you will be asked to evaluate a series of good deeds. Your task is to rate how morally praiseworthy it would be to perform each deed. While every good deed may be a least somewhat praiseworthy, some might be more praiseworthy than others.

The scale on which participants responded went from 1 (*the LEAST praiseworthy a deed could be*) to 10 (*the MOST praiseworthy a good deed could be*). All participants were reminded that there were no right or wrong answers to each of these items, and that we were only interested in how they would honestly feel about

each good deed described. Full scripts of participant instructions are to be found in the online supplementary materials.

Results

First we created indices of action value and outcome value by averaging the ratings provided by Sample 1. (A deed's mean action value was obtained by averaging across participants in the *action* condition; a deed's mean outcome value was obtained by averaging across participants in the *outcome* condition; see Figure 1). The overall average correlation in the full set of 95 items was $r = .739$, indicating a strong degree of association between the action and outcome value of prosocial deeds. Because we wanted to obtain a sample that would be best capable of differentiating the respective contributions of action and outcome value to moral judgment, we sought to obtain a subset of these behaviors that was less strongly associated. To do this, we created an algorithm in MATLAB that (a) selected a random subset of 23 items out of the 95 total; (b) computed the correlations between action and outcome values for the 23 items; (c) repeated this process 10,000 times; and (d) returned the subset with the lowest correlation. Twenty-three items were used because it was the same number of items used in previous research in the harm domain (Miller et al., 2014). The resulting set consisted of 23 items with a correlation between action and outcome values of $r = .23$, which is sufficiently low to be able to assess their independent contributions to. The action scores had a mean of 6.06, $SD = .65$; examples of high action-value items included "Comforting an old lady who has tripped and fallen" and "Helping a group of ducklings safely cross the street." The outcome scores had a mean of 5.42, $SD = 1.07$; examples of high outcome-value items included, "Convincing a friend to donate bone marrow to a stranger in the hospital" and "Asking a passerby to call 911 upon witnessing a car crash." This list of 23 items was used as our measure of Prosocial Action and Outcome (PAO) for the remainder of the research (see Table S1 and Figure S1 in the online supplementary material).

To test the hypothesis that action and outcome values each make independent contributions to moral praiseworthiness, we sought to examine whether mean action and outcome judgments on the PAO predicted judgments of praiseworthiness. To accomplish this, we conducted a series of linear regressions, one for each participant, that allowed mean-centered action and outcome values extracted from the PAO to predict that person's praiseworthiness judgments (Praise$_{individual}$ = β*Action$_{mean}$ + β*Outcome$_{mean}$ + c; see Figure 1). The analysis is akin to a multilevel approach with participant intercept and slope set as random factors. The resulting output produced two β weights, which we labeled "action sensitivity" and "outcome sensitivity," respectively. The individual constant in this equation denoted each person's predicted praiseworthiness at mean levels of action and outcome value.

To test our hypothesis that both action and outcome values contribute to praiseworthiness, we conducted one-sample *t* tests measuring the mean action β and outcome β against the test value of 0. Results indicated that both action value ($M = .59$, $SD = .79$, $t(49) = 5.62$, $p < .001$, $d = .80$, 95% confidence interval, CI [.38, .81]) and outcome value $M = .20$, $SD = .40$, $t(49) = 3.55$, $p < .001$, $d = .5$, 95% CI [.08, .31]) made significant and independent contributions to judgments of moral praise. Another way to test this was to run a linear regression allowing mean action and outcome values to predict mean praiseworthiness calculated for each item (hereafter dubbed "simple linear regression.") The analysis corroborated the result both for action (β = .60, $SE = .15$, $t(20) = 3.87$, $p < .001$) and outcome values (β = .20, $SE = .09$, $t(20) = 2.17$, $p = .043$).

We also tested the relationship between the βs obtained from the multilevel approach. Neither the action nor outcome sensitivity was correlated with the constant, $ps > .40$; however, action and outcome value were negatively correlated, $r = -.35$, $p = .011$. Finally, and most importantly, a paired-samples *t* test revealed that people were significantly more sensitive to action than to outcome values, $M_{diff} = .40$, $SD = .98$, $t(49) = 2.91$, $p = .005$, $d = .41$, 95% CI [.12, .67] (see Figure 2).

Discussion

The purpose of this experiment was to test whether action and outcome values predict praiseworthiness. To answer this question, we identified a set of 23 deeds that, because they demonstrated low correlation between action and outcome values, would prove an effective measure of the independent contributions of these evaluations to moral praise. Next, we used this measure to assess how strongly each of these factors affected praise. Results indicated that both factors made significant contributions. This is consistent with previous work demonstrating that both action and outcome values contribute to judgments of moral wrongness (Miller et al., 2014). In addition, we found that action and outcome βs were negatively correlated. This implies that people who were sensitive to action values were less sensitive to outcome values and vice versa. Finally, we found that people were more sensitive to action than to outcome values in their evaluations of praise. This suggests that the latter may be more influential than the former in people's determination of the moral worth of prosocial behavior.

Study 2

Study 1 established that the praiseworthiness of prosocial behaviors is predicted by group-average evaluations of the action and outcome values of those behaviors. In this second experiment we sought to replicate this result, as well as to examine whether *personal* ratings of action and outcome values for prosocial behaviors similarly predicted praiseworthiness of those behaviors. From a procedural perspective, the primary difference between this and the previous experiment was that it used a within-rather than a between-subjects design (see Figure 1). This allowed us to test, for each person, the extent to which the praiseworthiness ratings of various acts were determined by the action and the outcome value that *they themselves* assigned to those acts. Overall, then, the purpose of this study was to conceptually replicate the results of Study 1, and to determine if the same results would hold when considering individuals' own action and outcome ratings.

Method

Participants. A power analysis based on the smallest effect size in Study 1 ($d = .41$) indicated a required sample size of 80 (two-tailed paired samples *t* test, 95% power). We sought a sample of this size and, as in Study 1, excluded four participants who showed no variation in their ratings of action or outcome value.
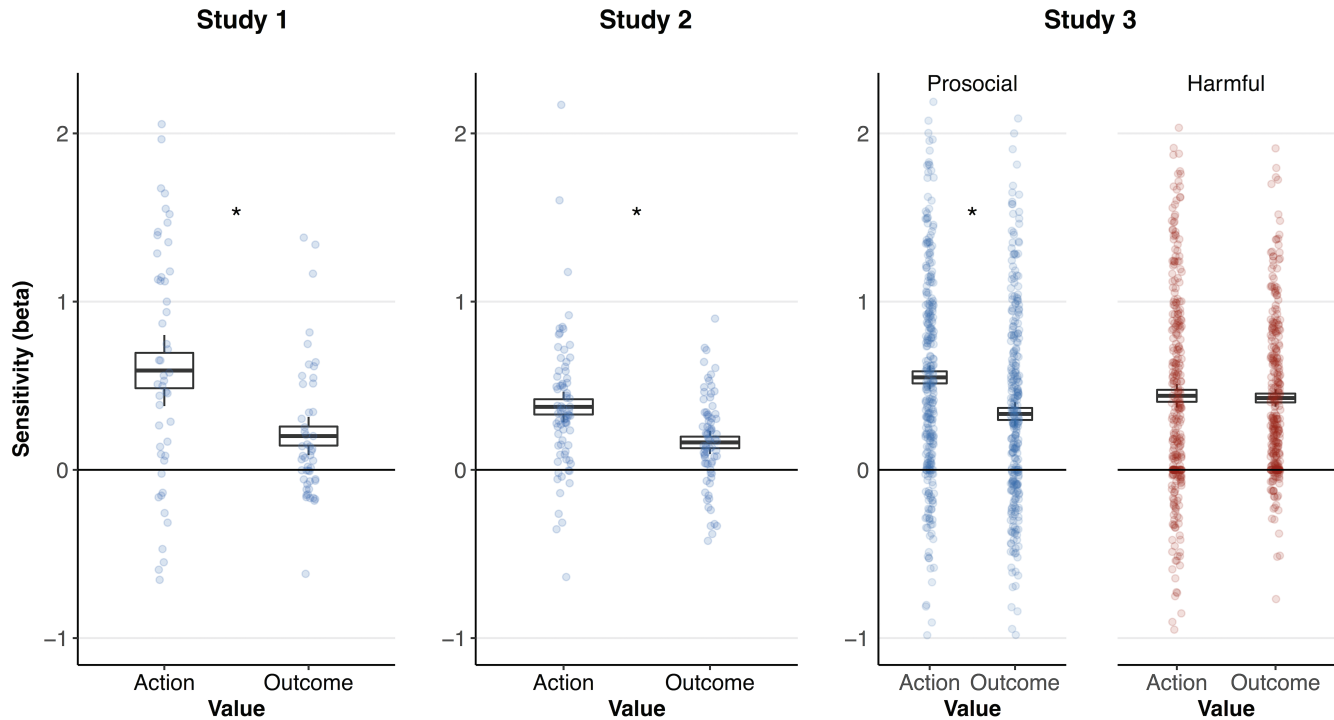
Figure 2. Individual action- and outcome-sensitivity, as measured by β weights extracted from within-subjects regressions, with praiseworthiness (and wrongness; see Study 3) as the dependent variable. Center boxplot lines represent means; outer lines *SEM*; whiskers 95% confidence interval (CI). Dots reflect individual data points. * $p \leq .005$. See the online article for the color version of this figure.

The final sample consisted of 78 people (41 men, 37 women, $M_{age} = 33.16$, $SD = 9.15$), who participated in a study entitled "Rate items, short and interesting study" through MTurk (Litman et al., 2017).

Procedure. Participants were asked to make judgments about a series of good deeds. To avoid influencing praiseworthiness judgments with action or outcome ratings, participants always rated praiseworthiness first, then rated action and outcome value in random order (specific instructions were identical to those in Study 1).

Analytic plan. To conceptually replicate Study 1, we ran a series of regressions at the individual level and extracted the resulting action and outcome β weights. The difference between this analysis and that used in Study 1 was that, instead of using *mean* action and outcome ratings, we allowed *each person's* individual ratings on both these measures to predict their praiseworthiness ratings (see Figure 1). Thus, the regression we ran for each participant was $Praise_{individual} = \beta^*Action_{individual} + \beta^*Outcome_{individual} + c$. The resulting β weights were extracted and used to indicate action and outcome sensitivity, respectively.

Results

Results indicated that both action ($M = .31$, $SD = .34$, $t(77) = 8.03$, $p < .001$, $d = .91$, 95% CI [.23, .39]) and outcome ($M = .12$, $SD = .30$, $t(77) = 3.53$, $p < .001$, $d = .40$, 95% CI [.05, .18]) values differed significantly from 0, suggesting that people's action and outcome evaluations significantly predicted their praise-

worthiness ratings. A simple linear regression allowing mean action ($\beta = .61$, $SE = .15$, $t(20) = 4.23$, $p < .001$) and outcome values ($\beta = .40$, $SE = .13$, $t(20) = 3.15$, $p = .005$) of the 23 items to predict mean praiseworthiness corroborated these results.

As in the previous experiment, action βs and outcome βs were negatively correlated, $r = -.37$, $p < .001$. Neither action nor outcome sensitivity was correlated with the constant, $ps > .05$. In addition, people were once again more sensitive to action than to outcome values, $M_{diff} = .19$, $SD = .52$, $t(77) = 3.27$, $p = .002$, $d = .37$, 95% CI [.07, .31]—replicating the findings of Experiment 1 (see Figure 2).

Discussion

Here we tested whether an individual's action and outcome evaluations were associated with their ratings of praise. In a conceptual replication of Study 1, we found that people's individual ratings of action and outcome value indeed predicted praiseworthiness. This provides further confirmation that both these factors are significantly associated with people's moral judgments. Action and outcome sensitivity were again negatively correlated suggesting that the greater people's sensitivity to one the less to the other. Finally, we found that action sensitivity was stronger than outcome sensitivity. This corroborates the observations in Study 1: that prosocial judgments are more influenced by action than by outcome value. Overall, then, we have provided additional evidence for the power of both action and outcome value to influence people's judgments of praiseworthiness.

Study 3

In this experiment we sought to confirm that action and outcome values both significantly predict praiseworthiness, and that action value has greater influence. We further wished to compare these results to those in the domain of harmful acts. As noted above, Miller et al. (2014) find that both action and outcome values predict wrongness, but it is unclear which is more influential. Finally, we were interested in seeing whether people exhibited *cross-domain consistency* in their sensitivity to action and outcome values. On one hand, it is possible that a person's sensitivity to action and outcome values for prosocial acts will be unrelated to their sensitivity to action and outcome values for harmful acts. On the other hand, it is possible that a person's sensitivity in one domain would predict his or her sensitivity in the other. This would provide evidence for a set of domain-general valuation mechanisms that use both action value and outcome value to calculate the moral worth of different behaviors, regardless of their content or valence. The current research set out to see if this was the case.

Our approach was initially similar to that performed in Study 1. We recruited one sample of participants to rate the action or outcome value of 23 prosocial items. Then we recruited a second sample to rate the praiseworthiness of those same items, and used the mean action and outcome ratings obtained from the first sample to predict each person's praiseworthiness ratings in the second sample.

We used the same procedure for a set of harmful behaviors. We obtained from Miller et al. (2014) mean action and outcome values for the 23 harmful behaviors used in their Study 5. Then we asked the same sample of participants who had rated the praiseworthiness of the prosocial items (Sample 2) to rate the wrongness of the harmful items. Based on the means obtained from the other authors, we could then compute people's sensitivities to action and outcome in the harm domain, and, finally, compare these to action and outcome sensitivities in the prosocial domain.

Method

Participants.   The action value and outcome value means provided for the Harm domain by Miller et al. (2014) had been obtained from 103 subjects (53 men, 50 women, $M_{age} = 34.3$). We obtained these means courtesy of the authors. The action or outcome means for prosocial acts were obtained from a new sample we recruited on TurkPrime (Sample 1). To increase our statistical precision regarding action and outcome value means, we sought about 300 participants to provide either action or outcome ratings for each of the 23 prosocial items (maximum anticipated standard error ~0.1). As in Studies 1 and 2, we slightly overcollected, then excluded four participants on the basis of the fact that they showed no variance in their action or outcome ratings and ended with a final sample of 300 participants, 169 men, 131 women, $M_{age} = 35.0$, $SD = 10.2$. Participants were randomly assigned to rate either action or outcome values.

Finally, the primary sample (Sample 2) consisted of those who rated, in random order, the praiseworthiness of the prosocial items and the wrongness of the harmful items. Because our main question was how action and outcome sensitivity differed across domain (prosocial vs. harmful), we planned to conduct a repeated-measures analysis of variance (ANOVA) with domain (prosocial vs. harmful) and value (action vs. outcome) as within-subjects factors. We anticipated a minimum correlation among the factors of $r = .1$ and a small effect size ($f = .12$); a power analysis in G*Power gave a suggested sample size of 270 to obtain our effect (power = 95%); we slightly overcollected to be conservative and ultimately succeeded in recruiting 303 participants (170 men, 131 women, 2 not reporting, $M_{age} = 35.6$, $SD = 11.0$) through MTurk (Litman et al., 2017) in exchange for $1.20. No participants were eliminated from this sample.

Procedure.   We first sought to obtain mean action and outcome ratings for both prosocial and harmful behaviors. For prosocial behaviors, as in previous studies, we asked Sample 1 to rate, for each of 23 items, how beneficial each action would be, as well as how good it would feel to perform. In our analysis, to be conservative, we planned to use only the first block that people responded to, in case that their responses to the first set of questions influenced their responses to the second set (it is plausible, for instance, that judging good deeds' outcome value first might create demand effects that influence subsequently judged action value, and vice versa). In other words, action means were obtained by averaging the ratings of those participants who had responded first to the action block in the survey; outcome means were obtained by averaging outcome ratings among those who had responded first to the outcome block. The full dataset was collected for exploratory analyses and is available online (osf.io/k3pca).

For the *harmful action* and *outcome* means, we obtained data previously collected by Miller et al. (2014; Study 5). As we have described, in that work, the researchers had examined how action- and outcome-values influenced people's perception of the wrongness of harmful acts. Specifically, they exposed people to a set of 23 items in which a harmful act was being committed. They then assigned them to either the action, outcome, or judgment condition. People in the *action* condition were asked how *upset* it would make them to act out each of the 23 actions. People in the *outcome* condition were asked to determine how much *suffering* would be caused by each of these actions. Finally, people in the *judgment* condition were asked to indicate how *wrong* each of these actions was. We requested only the mean *upset* and the *suffering* ratings obtained from the sample, as these items were functionally parallel to the *prosocial action* and *outcome* ratings obtained from Sample 1 described above.

For the primary sample (Sample 2), after completing consent forms and demographic questions, participants were informed that they would be making some judgments about a series of actions. They were then exposed in randomized order to the "praiseworthiness" block and the "wrongness" block, each with 23 items. After completion of both the prosocial and the harm block of items, participants were debriefed and dismissed.

Results

To assess people's level of sensitivity to action and outcome values in the prosocial domain, we ran a series of regressions that, as in Study 1, used action and outcome means (obtained from Sample 1) to predict the praiseworthiness rating of each participant in Sample 2 (see Figure 1). Results indicated that both action ($M = .55$, $SD = .62$, $t(302) = 15.34$, $p < .001$, $d = .88$, 95% CI [.48, .62]) and outcome ($M = .33$, $SD = .62$, $t(302) = 9.21$, $p < .001$, $d = .52$, 95% CI [.26, .40]) values differed significantly from 0,

suggesting that both valuations made significant and independent contributions to moral praise. A simple linear regression predicting average praiseworthiness ratings across the 23 items from mean action values ($\beta = .54$, $SE = .16$, $t(20) = 3.24$, $p = .004$) and mean outcome values ($\beta = .25$, $SE = .13$, $t(20) = 1.88$, $p = .074$) corroborated these results (though outcome values were only marginally significant). As before, action and outcome sensitivities were negatively correlated, $r = -.22$, $p < .001$. Both action values, $r = -.11$, $p = .067$ and outcome values, $r = -.13$, $p = .020$ showed a negative (or marginally negative) correlation with the constant.

We then performed the same analyses in the harm domain, allowing the mean action (upset) rating and harm (suffering) rating obtained from the other group to predict people's judgments of wrongness of each of the harmful acts. Corroborating the results obtained by Miller et al. (2014), both action ($M = .44$, $SD = .62$, $t(302) = 12.36$, $p < .001$, $d = .71$, 95% CI [.37, .51]) and outcome ($M = .43$, $SD = .46$, $t(302) = 16.33$, $p < .001$, $d = .94$, 95% CI [.38, .48]) values made significant and independent contributions to moral condemnation. A simple linear regression predicting average harm ratings across the 23 items from mean action values ($\beta = .44$, $SE = .09$, $t(20) = 4.77$, $p < .001$) and mean outcome values ($\beta = .42$, $SE = .07$, $t(20) = 6.12$, $p < .001$) corroborated these results. Action and outcome sensitivities were negatively correlated, $r = -.20$, $p < .001$, and a significant correlation emerged between the constant and the action β, $r = .15$, $p < .001$ and in the negative direction between the constant and the outcome β, $r = -.24$, $p < .001$, suggesting that people with higher average wrongness ratings were more sensitive to how upsetting an action would be to perform, while people with lower average wrongness ratings were more sensitive to suffering.

Next, we examined cross-domain associations between action and outcome values. Results indicated that both action sensitivity and outcome sensitivity were associated across domains: $r = .13$, $p = .015$ and $r = .18$, $p = .001$, respectively, suggesting that people who tended to be sensitive to prosocial action values were also sensitive to harmful action values, and people tended to be sensitive prosocial outcome values were also sensitive to harmful outcome values.

Finally, we examined whether the relative strength of action and outcome sensitivities differed according to domain. Confirming the results of Studies 1 and 2, results indicated that action sensitivity was stronger on average that outcome sensitivity for judgments of prosocial behavior, $M_{\text{diff}} = .21$, $SD = .98$, $t(302) = 3.90$, $p < .001$, $d = .22$, 95% CI [.11, .33]). By contrast, and confirming the results of Miller et al. (2014; Study 5), action sensitivity was not stronger than outcome sensitivity for judgments of harmful behavior, $M_{\text{diff}} = -.01$, $SD = .84$, $t(302) = .26$, $p = .78$, $d = .02$, 95% CI [−.08, .11]). A 2 (value: action vs. outcome) × 2 (domain: prosocial vs. harm) repeated measures ANOVA indicated the interaction between these factors was significant, $F(1, 302) = 8.45$, $p = .004$, $\eta_p^2 = .03$ (see Figure 2). Comparisons with value and across domain showed that while people were more sensitive to prosocial than harmful action values, $M_{\text{diff}} = .11$, $SD = .82$, $t(302) = 2.33$, $p = .020$, $d = .13$, 95% CI [.02, .20]), they were less sensitive to prosocial than to harmful outcome values, $M_{\text{diff}} = -.09$, $SD = .70$, $t(302) = -2.36$, $p = .019$, $d = -.13$, 95% CI [−.17, −.02]).

Discussion

In this experiment, we succeeded in replicating our finding from Studies 1 and 2 that action and outcome values both significantly predicted praise. This observation was echoed in the harm domain, replicating the findings of Miller et al. (2014). Next, we found evidence of a positive cross-domain relationship between action and outcome sensitivities: the extent to which people weighted action and outcome values when judging prosocial behaviors predicted the extent to which they weighted the same valuations when judging harmful behaviors. Finally, we found that, while action values were significantly more influential than outcome in judgments of praiseworthiness, there was no action-outcome difference in judgments of wrongness. This finding replicates and extends the observations from Studies 1 and 2, as well as those in Miller et al. (2014), which similarly found no action-outcome difference in the harm domain. Action values, apparently, outweigh outcome values in moral judgments of prosocial but not harmful acts.

General Discussion

The purpose of this research was to better understand people's judgments of prosocial behavior. Across three studies, mean action and outcome value made positive and independent contributions to the rated praiseworthiness of good deeds. In addition, people's sensitivity to action value was stronger than their sensitivity to outcome value—an effect that did not hold in the harm domain, where both types of values held equal weight. Finally, both action and outcome sensitivities were significantly correlated across the prosocial and harm domains, suggesting within-subject consistency in action and outcome sensitivity.

These findings are important in light of recent theoretical developments regarding the architecture underlying prosocial judgment. According to these accounts, good behavior is driven by at least two evaluative systems: one focusing on the consequences of good deeds; the other focusing on the feeling associated with their enactment (Crockett, 2013; Cushman, 2013; Gesiarz & Crockett, 2015). While recent work has demonstrated the importance of both these systems in the context of judgments of harms (Miller et al., 2014), the question of whether this framework applies to judgments of prosocial acts has not been tested until now. Here we show that both factors contribute to prosocial evaluation, providing this existing theoretical work with empirical validation.

The current findings also fit with research suggesting that people are not always sensitive to outcomes when making prosocial decisions (Berman et al., 2018; Cameron & Payne, 2011; Slovic, 2007). For instance, people will often perform the action of giving resources to charity, even if they know it will have no actual benefit (Crumpler & Grossman, 2008), and research finds that charitable giving is often motivated by the pursuit of warm glow (e.g., Ribar & Wilhelm, 2002; see also Andreoni, 1990; Konow, 2010). In line with these observations, researchers have found that people would be willing to pay approximately the same amount to save 2,000, 20,000, or 200,000 birds from oil-polluted ponds (Bloom, 2016; Desvousges et al., 1992). In addition, a survey of Canadians showed that people would be willing to pay approximately the same amount of money to clean up all the lakes in Ontario as to clean up the lakes in a narrower part of the province (Kahneman, 1986; see also Buechel & Morewedge, 2014). Finally, other research suggests that people do experience real and lasting

emotional benefits from performing altruistic behavior, further cementing the psychological importance of the warm glow of giving (Aknin et al., 2013; Weinstein & Ryan, 2010). The current findings corroborate these effects from the standpoint of the judgment of moral acts.

These observations are also theoretically relevant in light of a long-standing debate regarding the motivations underlying altruistic behavior. Some theorists (e.g., Batson et al., 1999; Schroeder et al., 1988) argue that a primary motivator for prosocial behavior is a desire to achieve a positive outcome or benefit for others. Others, by contrast, suggest that prosocial acts arise primarily from more self-directed motivations, such as the desire to achieve a warm glow from giving altruistically, or to alleviate personal distress (e.g., Andreoni, 1990; Cialdini, 1991). While this past work did not distinguish between action and outcome value, the notion of other-versus self-directed motivations for helping could be construed as loosely mapping onto the outcome-versus action-value distinction put forth in this article. Accordingly, the fact that people value action-over outcome-value might suggest that the former holds greater sway in motivating prosocial behavior.

We note, however, that the applicability of our findings to this debate is qualified by important differences in theoretical framing between our work and past research. First of all, the abovementioned research focuses predominantly on empathy-induced helping of others in distress. Such behavior focuses on the alleviation of negative outcomes, rather than the promotion of positive ones. As past research (e.g., Higgins, 1998) demonstrates, critical psychological differences exist between the prevention of negative outcomes and the promotion of positive ones. These differences suggest that our findings may not apply to negative contexts.

Furthermore, our results differ from these findings insofar as our work examines the evaluation rather than the enactment of prosocial behavior. These constructs are also psychologically distinct. For instance, Monin et al. (2007) suggest that enactment relies more heavily on sophisticated reasoning processes, while evaluation is more subject to quick, affect-laden processes. Applied to the current findings, this suggests that our observations about the primacy of action value should not be automatically assumed to apply in the domain of prosocial behavior. Rather, if prosocial behavior entails greater reliance on reasoning, this may in fact bolster the importance of outcome value; a claim which is supported by previous research (Crockett, 2013; Cushman, 2013; Gesiarz & Crockett, 2015). On the other hand, there is a clear association between action and judgment, in that people are more likely to perform behaviors they judge more positively (see Crockett et al., 2017). This association raises the possibility that an increased emphasis on action value may apply to the domain of behavior after all under particular circumstances.

In addition to findings about people's relative sensitivity to different valuations, we also found, in Study 3, evidence of consistency in sensitivity to both action and outcome values across both prosocial and harmful deeds. In other words, people swayed by action or outcome value when judging prosocial acts were similarly swayed when judging harmful acts. The observation further confirms the idea, put forward in recent theoretical work, that distinct cognitive mechanisms may drive moral evaluation regardless of valence (e.g., Crockett, 2013; Cushman, 2013), and raises the question of whether such tendencies may constitute stable individual differences.

An interesting question that arises in the context of these findings concerns the underlying psychological phenomena driving them. Specifically, why do people place more value on action than outcome when judging the praiseworthiness of prosocial behaviors? One possibility has to do with the "concreteness" of high action-value behaviors. Behaviors high in action value, by their very nature, are more "up close and personal" than those high in outcome value. Psychological proximity has, in turn, been tied closely to concrete construals; while psychological distance leads to more abstract construals (Liberman & Trope, 2014). The concrete quality of behaviors high in action value may make them more easy to mentally simulate; thus, increasing people's sensitivity to this particular dimension of behavior. However we caution that this conjecture is purely speculative and should be tested in further research.

On a broader level, past work has suggested that judging the wrongness of harmful actions involves a process of "evaluative simulation," whereby we evaluate the moral status of another's action by simulating the affective response that we would experience performing the action ourselves (Miller et al., 2014). Our results are consistent with the possibility that evaluative simulation also plays a role in judging the praiseworthiness of helpful actions. If people evaluate helpful actions by simulating what it feels like to perform the action, then we would expect to see similar biases in moral evaluation as those that exist for moral action. Previous work has shown that individuals often do not act to maximize the benefits that others receive, but instead to maximize the good feelings associated with performing good deeds (Berman et al., 2018; Gesiarz & Crockett, 2015; Ribar & Wilhelm, 2002). Thus, the asymmetry in moral evaluation seen in the present studies may reflect a correspondence between first-person moral decision-making and third-person moral evaluation.

Practical Implications

Our work can shed light on some of the challenges that might face the "effective altruism" (EA) movement, which emphasizes the importance of giving to charities that have the most concrete benefit for every dollar spent. In particular, our research suggests that this movement may be facing an uphill battle, because, while effective altruism encourages people to focus on consequences, we find that, in fact, deeds high in action value are considered most praiseworthy (Berman et al., 2018). As a result, it may be difficult for effective altruists to garner social praise for their actions. Indeed, the EA movement has suffered some criticism in precisely this vein; one piece, for instance, called the movement an instance of "elitist philanthropy . . . that can kill the very altruistic spirit it claims to foster" (Berger & Penna, 2013).

Of course, this does not mean that the movement should abandon its cause. It simply raises the notion that the movement may be best served by simultaneously playing up the appeal of action value in its effort to elicit charitable donations—for instance, by pairing instances of giving with pleasurable associations (such as increased social standing or praise), or by fostering a sense of identity or community around the act of giving (that is also known to lead to positive associations; Jetten, Haslam, & Alexander, 2012; Thoits & Hewitt, 2001). Research has found that people who perform charitable deeds on the basis of action-based considerations are rated as moral as those who act on the basis of other

considerations (Barasch, Levine, Berman, & Small, 2014; see also Everett, Pizarro, & Crockett, 2016). Acknowledging the value of warm glow in addition to outcome-based approaches may prove an effective means of encouraging charitable giving.

These results also raise the question of whether moral values can be changed. According to existing theory, action values may arise from a habitual association between carrying out an action and the positive feeling that this enactment elicits (Gesiarz & Crockett, 2015). It is possible, then, that people who establish positive habits of prosocial giving will exhibit increases in action value over time, thereby rating warm glow actions more praiseworthy (see Thoits & Hewitt, 2001). For example, Community Supported Agriculture (CSA) programs often require their members to participate in a few hours of garden work or volunteering per month. It would be interesting to see whether such programs increase people's sensitivity to the action value of prosocial behavior over time. Similarly, festivals such as Burning Man immerse their participants for a week in an environment where giving gifts is the only acceptable form of exchange (Chen, 2012). We are currently at work on a project that examines whether such experiences exert significant and lasting changes on people's evaluative systems.

Conclusion

Our data support the proposition that two distinct forms of evaluation contribute to people's overall assessments of moral praiseworthiness. People are more influenced by the action value of prosocial behaviors than by outcome value in determining moral praiseworthiness. Our observations support and extend the implications of a growing body of work aimed at understanding moral judgment, which ultimately may help shape policies aimed at increasing altruistic behavior.

References

Aknin, L. B., Barrington-Leigh, C. P., Dunn, E. W., Helliwell, J. F., Burns, J., Biswas-Diener, R., . . . Norton, M. I. (2013). Prosocial spending and well-being: Cross-cultural evidence for a psychological universal. *Journal of Personality and Social Psychology, 104,* 635–652. http://dx.doi.org/10.1037/a0031578

Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *Economic Journal, 100,* 464. http://dx.doi.org/10.2307/2234133

Andreoni, J. (1995). Warm-glow versus cold-prickle: The effects of positive and negative framing on cooperation in experiments. *The Quarterly Journal of Economics, 110,* 1–21.

Barasch, A., Levine, E. E., Berman, J. Z., & Small, D. A. (2014). Selfish or selfless? On the signal value of emotion in altruistic behavior. *Journal of Personality and Social Psychology, 107,* 393–413. http://dx.doi.org/10.1037/a0037207

Batson, C. D. (1994). Why act for the public good? Four answers. *Personality and Social Psychology Bulletin, 20,* 603–610. http://dx.doi.org/10.1177/0146167294205016

Batson, C. D., Ahmad, N., Yin, J., Bedell, S. J., Johnson, J. W., & Templin, C. M. (1999). Two threats to the common good: Self-interested egoism and empathy-induced altruism. *Personality and Social Psychology Bulletin, 25,* 3–16. http://dx.doi.org/10.1177/0146167299025001001

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology, 5,* 323–370. http://dx.doi.org/10.1037/1089-2680.5.4.323

Bazerman, M. H., & Moore, D. A. (2008). *Judgment in managerial decision making* (7th ed.). New York, NY: Wiley.

Berger, K., & Penna, R. M. (2013, November 25). The elitist philanthropy of so-called effective altruism. *Stanford Social Innovation Review*. Retrieved from https://ssir.org/articles/entry/the_elitist_philanthropy_of_so_called_effective_altruism

Berman, J. Z., Barasch, A., Levine, E. E., & Small, D. A. (2018). Impediments to effective altruism: The role of subjective preferences in charitable giving. *Psychological Science, 29,* 834–844. http://dx.doi.org/10.1177/0956797617747648

Bloom, P. (2016). *Against empathy*. New York, NY: Ecco.

Buechel, E. C., & Morewedge, C. K. (2014). The (relative and absolute) subjective value of money. In E. Bijleveld & H. Aarts (Eds.), *The psychological science of money* (pp. 93–120). New York, NY: Springer. http://dx.doi.org/10.1007/978-1-4939-0959-9_5

Cameron, C. D., & Payne, B. K. (2011). Escaping affect: How motivated emotion regulation creates insensitivity to mass suffering. *Journal of Personality and Social Psychology, 100,* 1–15. http://dx.doi.org/10.1037/a0021643

Carlson, R. W., & Zaki, J. (2018). Good deeds gone bad: Lay theories of altruism and selfishness. *Journal of Experimental Social Psychology, 75,* 36–40. http://dx.doi.org/10.1016/j.jesp.2017.11.005

Carver, C. S. (2006). Approach, avoidance, and the self-regulation of affect and action. *Motivation and Emotion, 30,* 105–110. http://dx.doi.org/10.1007/s11031-006-9044-7

Chen, K. K. (2012). Artistic prosumption: Cocreative destruction at Burning Man. *American Behavioral Scientist, 56,* 570–595. http://dx.doi.org/10.1177/0002764211429362

Cialdini, R. B. (1991). Altruism or egoism? That is (still) the question. *Psychological Inquiry, 2,* 124–126. http://dx.doi.org/10.1207/s15327965pli0202_3

Crockett, M. J. (2013). Models of morality. *Trends in Cognitive Sciences, 17,* 363–366. http://dx.doi.org/10.1016/j.tics.2013.06.005

Crockett, M. J., Siegel, J. Z., Kurth-Nelson, Z., Dayan, P., & Dolan, R. J. (2017). Moral transgressions corrupt neural representations of value. *Nature Neuroscience, 20,* 879–885. http://dx.doi.org/10.1038/nn.4557

Crumpler, H., & Grossman, P. J. (2008). An experimental test of warm glow giving. *Journal of Public Economics, 92,* 1011–1021. http://dx.doi.org/10.1016/j.jpubeco.2007.12.014

Cushman, F. (2013). Action, outcome, and value: A dual-system framework for morality. *Personality and Social Psychology Review, 17,* 273–292. http://dx.doi.org/10.1177/1088868313495594

Cushman, F., Gray, K., Gaffey, A., & Mendes, W. B. (2012). Simulating murder: The aversion to harmful action. *Emotion, 12,* 2–7. http://dx.doi.org/10.1037/a0025071

Desvousges, W. H., Johnson, F. R., Dunford, R. W., Boyle, K. J., Hudson, S. P., & Wilson, K. N. (1992). *Measuring nonuse damages using contingent valuation: An experimental evaluation of accuracy*. Research Triangle Park, NC: Research Triangle Institute.

Dovidio, J. F., Allen, J. L., & Schroeder, D. A. (1990). Specificity of empathy-induced helping: Evidence for altruistic motivation. *Journal of Personality and Social Psychology, 59,* 249–260. http://dx.doi.org/10.1037/0022-3514.59.2.249

Everett, J. A., Pizarro, D. A., & Crockett, M. J. (2016). Inference of trustworthiness from intuitive moral judgments. *Journal of Experimental Psychology: General, 145,* 772–787. http://dx.doi.org/10.1037/xge0000165

Franzen, A., & Pointner, S. (2012). Anonymity in the dictator game revisited. *Journal of Economic Behavior & Organization, 81,* 74–81. http://dx.doi.org/10.1016/j.jebo.2011.09.005

Gesiarz, F., & Crockett, M. J. (2015). Goal-directed, habitual and Pavlovian prosocial behavior. *Frontiers in Behavioral Neuroscience, 9,* 135.

Hannikainen, I. R., Miller, R. M., & Cushman, F. A. (2014). Act versus impact: Conservatives and liberals exhibit different structural emphases in moral judgment. *Ratio, 30,* 462–493.

Herr, P. M. (1986). Consequences of priming: Judgment and behavior. *Journal of Personality and Social Psychology, 51,* 1106–1115. http://dx.doi.org/10.1037/0022-3514.51.6.1106

Higgins, E. T. (1998). Promotion and prevention: Regulatory focus as a motivational principle. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 30, pp. 1–46). San Diego, CA: Academic Press. http://dx.doi.org/10.1016/S0065-2601(08)60381-0

Hitlin, S. (2007). Doing good, feeling good: Values and the self's moral center. *The Journal of Positive Psychology, 2,* 249–259. http://dx.doi.org/10.1080/17439760701552352

Janoff-Bulman, R., Sheikh, S., & Hepp, S. (2009). Proscriptive versus prescriptive morality: Two faces of moral regulation. *Journal of Personality and Social Psychology, 96,* 521–537. http://dx.doi.org/10.1037/a0013779

Jetten, J., Haslam, C., & Alexander, S. H. (Eds.). (2012). *The social cure: Identity, health and well-being.* Hove, England: Psychology Press.

Kahneman, D. (1986). Comments in R. G. Cummings, D. S. Brookshire, & W. D. Schulze (Eds.), *Valuing environmental goods: An assessment of the contingent valuation method.* Totowa, NJ: Rowman and Allanheld.

Konow, J. (2010). Mixed feelings: Theories of and evidence on giving. *Journal of Public Economics, 94,* 279–297. http://dx.doi.org/10.1016/j.jpubeco.2009.11.008

Ledgerwood, A., & Ratliff, K. A. (2015, February). *A practical guide to navigating our transitioning science.* Symposium conducted at the annual conference of the Society for Personality and Social Psychology, Long Beach, CA.

Liberman, N., & Trope, Y. (2014). Traversing psychological distance. *Trends in Cognitive Sciences, 18,* 364–369. http://dx.doi.org/10.1016/j.tics.2014.03.001

Lieberman, D., & Lobel, T. (2012). Kinship on the Kibbutz: Coresidence duration predicts altruism, personal sexual aversions and moral attitudes among communally reared peers. *Evolution and Human Behavior, 33,* 26–34. http://dx.doi.org/10.1016/j.evolhumbehav.2011.05.002

Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods, 49,* 433–442. http://dx.doi.org/10.3758/s13428-016-0727-z

Miller, R., & Cushman, F. (2013). Aversive for me, wrong for you: First-person behavioral aversions underlie the moral condemnation of harm. *Social and Personality Psychology Compass, 7,* 707–718. http://dx.doi.org/10.1111/spc3.12066

Miller, R. M., Hannikainen, I. A., & Cushman, F. A. (2014). Bad actions or bad outcomes? Differentiating affective contributions to the moral condemnation of harm. *Emotion, 14,* 573–587. http://dx.doi.org/10.1037/a0035361

Monin, B., Pizarro, D. A., & Beer, J. S. (2007). Deciding versus reacting: Conceptions of moral judgment and the reason-affect debate. *Review of General Psychology, 11,* 99–111. http://dx.doi.org/10.1037/1089-2680.11.2.99

Nelkin, D. K. (2016). Difficulty and degrees of moral praiseworthiness and blameworthiness. *Noûs, 50,* 356–378. http://dx.doi.org/10.1111/nous.12079

Nelson-le Gall, S. A. (1985). Motive–outcome matching and outcome foreseeability: Effects on attribution of intentionality and moral judgments. *Developmental Psychology, 21,* 332–337. http://dx.doi.org/10.1037/0012-1649.21.2.332

Null, C. (2011). Warmglow, information and inefficient charitable giving. *Journal of Public Economics, 95,* 455–465. http://dx.doi.org/10.1016/j.jpubeco.2010.06.018

Nunes, P. A., & Schokkaert, E. (2003). Identifying the warm glow effect in contingent valuation. *Journal of Environmental Economics and Management, 45,* 231–245. http://dx.doi.org/10.1016/S0095-0696(02)00051-7

Ribar, D. C., & Wilhelm, M. O. (2002). Altruistic and joy-of-giving motivations in charitable behavior. *Journal of Political Economy, 110,* 425–457. http://dx.doi.org/10.1086/338750

Schroeder, D. A., Dovidio, J. F., Sibicky, M. E., Matthews, L. L., & Allen, J. L. (1988). Empathic concern and helping behavior: Egoism or altruism? *Journal of Experimental Social Psychology, 24,* 333–353. http://dx.doi.org/10.1016/0022-1031(88)90024-8

Slovic, P. (2007). "If I look at the mass I will never act": Psychic numbing and genocide. *Judgment and Decision Making, 2,* 79–95.

Small, D. A., Loewenstein, G., & Slovic, P. (2007). Sympathy and callousness: The impact of deliberative thought on donations to identifiable and statistical victims. *Organizational Behavior and Human Decision Processes, 102,* 143–153.

Thoits, P. A., & Hewitt, L. N. (2001). Volunteer work and well-being. *Journal of Health and Social Behavior, 42,* 115–131. http://dx.doi.org/10.2307/3090173

Weinstein, N., & Ryan, R. M. (2010). When helping helps: Autonomous motivation for prosocial behavior and its influence on well-being for the helper and recipient. *Journal of Personality and Social Psychology, 98,* 222–244. http://dx.doi.org/10.1037/a0016984