

Moral-Dilemma Judgments by Individuals and Groups: Are Many Heads Really More Utilitarian than One?

Marta Rokosz¹, Michał Białek¹, Michał M. Stefańczyk¹, Bertram Gawronski²

¹ University of Wrocław, Poland

² University of Texas at Austin, USA

Moral dilemmas often involve a conflict between action-options that maximize outcomes for the greater good (*utilitarianism*) and inaction-options that conform to moral norms (*deontology*). Previous research suggests that, compared to individuals, groups show stronger support for outcome-maximizing actions that violate moral norms. The current study used a computational modeling approach to investigate whether this difference is driven by (1) stronger sensitivity to consequences, (2) weaker sensitivity to moral norms, or (3) weaker action aversion in moral-dilemma judgments made by groups. The results suggest that groups show a stronger sensitivity to consequences than individuals. Groups and individuals did not differ in terms of their sensitivity to moral norms and their general action aversion. The findings challenge the idea that groups are less action averse and less concerned about violating moral norms than individuals and instead suggest that group decisions are more strongly guided by outcomes for the greater good.

Keywords: CNI model, deontology, group decision-making, moral judgment, utilitarianism

Imagine a police officer interrogating a man who kidnapped several people. The man is hiding the hostages at an unknown location, and they are at risk of dying from dehydration if not found soon, but the kidnapper refuses to reveal where they are. Would it be morally acceptable for the police officer to use aggressive interrogation techniques deemed torture to obtain information about the hostages' whereabouts?

If your answer is *yes*, your response aligns with a utilitarian approach to resolving moral dilemmas, which prioritizes the greater good. According to a utilitarian view, actions are moral to the extent that they maximize overall well-being. Alternatively, if your answer is *no*, your response aligns with a deontological approach to resolving moral dilemmas, which prioritizes moral rules. According to a deontological view, the morality of an action depends on its consistency with moral norms, such as the norm that one should not cause harm. Expanding on prior work suggesting that, compared to individuals, groups show stronger support for outcome-maximizing actions that violate moral norms (Curşeu et al., 2020; Keshmirian et al., 2022), the current work investigated whether this difference is driven by (1) differential concerns about outcomes in a utilitarian sense, (2) differential concerns about moral norms in a deontological sense, or (3) differential levels of action aversion irrespective of outcomes and moral norms.

Moral Decision-Making in Groups

Although a considerable portion of morally relevant decisions is made in groups, few studies have investigated how moral decisions made by groups differ from those made by individuals. Regarding decisions in sacrificial moral dilemmas like the kidnapping scenario in the introductory paragraph, two studies found that, compared to individuals, groups show stronger support for outcome-maximizing actions that violate moral

norms (Curşeu et al., 2020; Keshmirian et al., 2022). However, like many other findings in research on moral-dilemma judgment, the meaning of the obtained difference remains unclear because judgments in the traditional moral-dilemma paradigm are shaped by multiple distinct factors (Gawronski & Ng, 2024).

First, sacrificial moral dilemmas typically present a forced choice between maximizing outcomes and adhering to moral norms. This approach confounds the measurement of utilitarian and deontological inclinations (Conway & Gawronski, 2013). In the kidnapping scenario, for example, supporting torture could indicate either concern for the greater good or disregard for moral norms. The forced-choice nature of these dilemmas makes it impossible to distinguish between the two possibilities, because endorsing one option necessarily means rejecting the other.

Second, the traditional moral-dilemma paradigm confounds conformity to moral codes with general action preferences (Crone & Laham, 2017), in that the utilitarian option typically involves action (e.g., using torture) and the deontological option involves inaction (e.g., not using torture). Hence, in addition to strong utilitarian concerns about outcomes and weak deontological concerns about moral norms, endorsement of the utilitarian option could also be driven by a strong general preference for action. Moreover, in addition to weak utilitarian concerns about outcomes and strong deontological concerns about moral norms, endorsement of the deontological option could also be driven by a strong general preference for inaction. Together, these considerations suggest that responses in sacrificial moral dilemmas can be shaped by at least three distinct factors: (1) concern about outcomes in a utilitarian sense, (2) concern about moral norms in a deontological sense, and (3) general

action aversion irrespective of outcomes and moral norms (Gawronski & Ng, 2024).

These concerns also apply to prior evidence suggesting that, compared to individuals, groups show stronger support for outcome-maximizing actions that violate moral norms. Specifically, the observed difference may be driven by three distinct mechanisms mediating the effect: (1) stronger concerns about outcomes, (2) weaker concerns about moral norms, or (3) weaker general action aversion in groups compared to individuals (see Figure 1).

The Current Research

The main goal of the current research was to investigate whether groups show stronger support for norm-violating outcome-maximizing actions because (1) group judgments are more strongly affected by outcomes, (2) group judgments are less strongly affected by moral norms, or (3) groups are less action averse. To this end, we used the CNI model of moral-dilemma responses, a mathematical model that disentangles sensitivity to consequences (C), sensitivity to moral norms (N), and general preference for inaction versus action (I) in responses to moral dilemmas (Gawronski et al., 2017). The CNI model's C parameter quantifies sensitivity consequences, defined as the tendency to support action when the benefits of the focal action are greater than the costs and to support inaction when the benefits of the focal action are smaller than the costs. The CNI model's N parameter quantifies sensitivity to moral norms, defined as the tendency to support inaction when the focal action is prohibited by a proscriptive norm and to support action when the focal action is proscribed by a prescriptive norm. Finally, the CNI model's I parameter quantifies general preference for inaction versus action, defined as the tendency to support inaction (vs. action) regardless of consequences and moral norms.

To quantify the three factors, the CNI model compares responses across four types of moral dilemmas: (1) dilemmas where a proscriptive norm prohibits action and the benefits of action for overall well-being are greater than the costs; (2) dilemmas where a proscriptive norm prohibits action and the benefits of action for overall well-being are smaller than the costs; (3) dilemmas where a prescriptive norm prescribes action and the benefits of action for overall well-being, are greater than the costs; and (4) dilemmas where a prescriptive norm prescribes action and the benefits of action for overall well-being are smaller than the costs (for an example, see Table 1). Based on the processing tree depicted in Figure 2, the CNI model provides four mathematical equations to estimate numerical values for the three model parameters (C , N , I) based on the empirically observed probabilities of action versus inaction responses on the four types of dilemmas (Gawronski et al., 2017). These equations

include the three model parameters as unknowns and the empirically observed probabilities of action versus inaction responses on the four types of dilemmas as known numerical values. Using maximum-likelihood statistics, the CNI model generates parameter estimates for the three unknowns that minimize the difference between the empirically observed probabilities of action versus inaction responses on the four types of dilemmas and the probabilities of action versus inaction responses predicted by the model equations using the identified parameter estimates. The adequacy of the model in describing the data can be evaluated by means of goodness-of-fit statistics, such that poor model fit would be reflected in a statistically significant deviation between the empirically observed probabilities of action (vs. inaction) responses on the four types of dilemmas and the corresponding probabilities predicted by the model. Differences in parameter estimates across experimental conditions can be tested by enforcing equal estimates for a given parameter across conditions. If setting a given parameter equal across conditions leads to a significant reduction in model fit, it can be inferred that the parameter estimates are significantly different across conditions.

To gain deeper insights into why groups show stronger support for outcome-maximizing actions that violate moral norms, participants in the current study responded to a series of moral dilemmas for research using the CNI model. Participants responded to the dilemmas either individually or collectively in small groups of three. We first tested the preregistered hypothesis that on dilemmas with a proscriptive norm where the benefits of action are greater than the costs (i.e., traditional dilemmas), preference for action will be greater when people make moral decisions in groups versus individually (Hypothesis 1). Expanding on this analysis, we tested the three preregistered hypotheses that sensitivity to consequences on the CNI model's C parameter will be greater when people make moral decisions in groups versus individually (Hypothesis 2a); sensitivity to moral norms on the CNI model's N parameter will be smaller when people make moral decisions in groups versus individually (Hypothesis 2b); and general preference for inaction versus action on the CNI model's I parameter will be smaller when people make moral decisions in groups versus individually (Hypothesis 2c).

Open Practices

We report how we determined our sample size, all data exclusions, all manipulations, and all measures. All materials, data, and analysis codes are publicly available at <https://osf.io/7yhuq/>. The study was formally preregistered at <https://osf.io/gvkec/>.

Methods

Participants

Based on the available resources (Lakens, 2022), we aimed for a sample of 400 participants. We preregistered that 100 participants would complete the study individually and 300 would complete the study in groups of three, providing 100 observations for each condition. A sensitivity analysis conducted using G*Power (Faul et al., 2007) indicated that 100 observations per condition provide 80% power for detecting a small-to-medium difference of $d = .40$ between two independent means with an α -level of .05 (two-tailed). This effect size is smaller than the effect size reported by Keshmirian et al. (2022) for differences in moral-dilemma judgments by individuals versus groups. For a χ^2 test, 100 observations per condition provide 80% power to detect a 0.2 success-rate difference across conditions (e.g., in 10,000 simulations with two groups of $N = 100$ each, success rate in Group1 = 50% and Group2 = 70%, the estimated power was 79.8%).

Participants ($N = 421$; 55% women; $M_{\text{age}} = 23.4$, $SD_{\text{age}} = 7.63$; age range: 16-58) were recruited via snowball sampling.¹ Trained research assistants ensured group integrity by screening for pre-existing relationships and verifying participant unfamiliarity. Of the total, 100 made decisions individually (60% women; $M_{\text{age}} = 26.6$, $SD_{\text{age}} = 10.3$; age range: 18-58) and 321 in 107 groups of three (54% women; $M_{\text{age}} = 22.5$; $SD_{\text{age}} = 6.28$; age range 16-51). Individual participants received ~US-\$2.50 in local currency, while group participants received triple this amount to compensate for longer study duration.

Procedure and Materials

The study was conducted online, with individuals completing a Qualtrics survey and groups participating via Zoom video calls. Participants in both conditions were presented with 24 moral dilemmas from Körner et al. (2020). The 24 dilemmas included six basic scenarios, each of which was presented in four variants, reflecting the manipulations of cost-benefit relations (i.e., benefits of action greater than costs vs. benefits of action smaller than costs) and moral norms (i.e., action prohibited by proscriptive norm vs. action prescribed by prescriptive norm). Following prior research (e.g., Białek et al., 2019; Körner et al., 2020; Paruzel-Czachura et al., 2023), the dilemmas were presented in a fixed semi-random order to mitigate potential proximity effects between similar dilemmas, ensuring a controlled presentation sequence. For each dilemma, participants were asked if they would perform the described action, using the response options *yes* or *no*. Participants were randomly assigned to either the group

or solo condition. In the group condition, discussions were limited to three minutes per dilemma. Group members were encouraged to discuss each scenario thoroughly before collectively submitting a single joint response for their group to the moderator, who registered the response in Qualtrics.

Preregistered Analysis Plan

To test Hypothesis 1, we calculated the sum of *yes* responses on traditional dilemmas (i.e., dilemmas with a proscriptive norm and benefits of action that are greater than the costs), which we interpreted as an index of support for outcome-maximizing actions that violate moral norms (Gawronski et al., 2017). Using Jamovi v.2.3.21.0 (The Jamovi Project, 2021) with the significance level set to $p < .05$, we tested whether scores on this index differed across the two experimental conditions (group vs. solo).

Hypotheses 2a-c were tested using the CNI model (Gawronski et al., 2017). Following our preregistered analysis plan, we fitted the CNI model to the aggregated moral-judgment data of each experimental condition to obtain estimates of sensitivity to consequences (*C*), sensitivity to moral norms (*N*), and general preference for inaction versus action (*I*) in moral-dilemma judgments provided by groups and individuals, respectively.

With two conditions, the model had a total of eight free categories (i.e., four types of dilemmas for each of the two conditions) and six parameters (i.e., three parameters estimated for each of the two conditions), resulting in two degrees of freedom. This model was used as a baseline to test differences in the three parameters across conditions. To this end, a new model was fitted to the data after constraining estimates for a given parameter to be equal across conditions (see Gawronski et al., 2017). The fit of this new model was then compared to the fit of the baseline model. This procedure was followed for all three model parameters. All analyses using the CNI model were conducted using the freeware multiTree (Moshagen, 2010). Following Gawronski et al. (2017), we used a fixed estimation algorithm with random start values, two replications, and a maximum of 90,000 iterations. Effect sizes for between-group differences on the three CNI model parameters were calculated based on means, standard errors, and sample sizes using Wilson's (2023) practical meta-analysis effect size calculator (see Gawronski et al., 2017).

Results

Traditional Analysis

Replicating earlier findings, groups showed stronger support for norm-violating outcome-

¹ When recruiting participants, we created more than 100 slots to accommodate potential no-shows. However, nearly all participants

who signed up for our study attended, leaving us with seven groups more than we had originally planned.

maximizing actions ($M = 3.48$, $SD = 1.31$) than participants who responded to the dilemmas individually ($M = 2.48$, $SD = 1.31$), $t(205) = 5.47$, $p < .001$, $d = 0.76$ (see Figure 2).

CNI Model Analysis

Although the baseline model showed suboptimal fit, $G^2(2) = 7.65$, $p = .022$, $\omega = 0.039$, the effect size for the deviation between predicted and observed data was far below the conventional benchmark for a small effect.² We therefore proceeded with our preregistered analyses to test whether the obtained difference in moral-dilemma judgments is driven by differences in (1) sensitivity to consequences, (2) sensitivity to moral norms, or (3) general preference for inaction versus action. The estimated parameter scores in the two conditions are depicted in Figure 3.

Constraining the C parameter to be equal across conditions led to a significant decrease in model fit, $\Delta G^2(1) = 55.55$, $p < .001$, $d = 1.04$, indicating that groups were more sensitive to consequences in their moral judgments than individuals deciding solo. Using the same analytic approach, there was no significant difference across conditions for the N parameter, $\Delta G^2(1) = 2.18$, $p = .139$, $d = 0.20$, and the I parameter, $\Delta G^2(1) = 0.07$, $p = .785$, $d = 0.04$, indicating that groups and individuals did not significantly differ in terms of their sensitivity to moral norms and in terms of their general preference for inaction versus action.

Discussion

The current study replicated earlier findings indicating that, compared to individuals, groups show stronger support for outcome-maximizing actions that violate moral norms (Curşeu et al., 2020; Keshmirian et al., 2022). Expanding on this finding, the current work suggests that this difference is driven by greater concerns about outcomes in groups compared to individuals. We did not find any evidence for the alternative hypotheses that groups would be less concerned about violating moral norms or that groups would be less action-averse than individuals. These conclusions are based on analyses using the CNI model of moral-dilemma responses, which revealed higher scores on the model's C parameter among groups compared to individuals, but no significant differences on the model's N and I parameters.

Initially, we suspected that groups' stronger support for norm-violating outcome-maximizing actions might be driven by increased pressure for action in group settings. Within the traditional dilemma approach, such pressure always supports the outcome-maximizing choice, but this does not necessarily mean

that groups are more concerned about outcomes. Using the CNI model to disentangle the two possibilities, the current results disproved this speculation. Groups and individuals showed no difference in general preference for inaction versus action. Instead, the key difference emerged in sensitivity to consequences, indicating that groups are genuinely more concerned about outcomes than individuals, rather than simply being more prone to action.

Our results also challenge the idea that groups' stronger support for norm-violating outcome-maximizing actions stems from reduced concern for moral norms. Despite suggestions that group settings might lead to moral disengagement and weaker negative emotions, potentially increasing willingness to violate moral norms (e.g., Lantos & Molenberghs, 2021), our findings show no difference between groups and individuals in sensitivity to moral norms.

Overall, our findings suggest that groups are more focused on the outcomes of their decisions than individuals. This conclusion is consistent with related findings in research on non-moral decisions, suggesting that groups are more concerned about the outcomes of risky choices than individuals (Hart et al., 2017). The current findings suggest that this differential concern about outcomes generalizes to decisions in moral dilemmas.

Expanding on our finding that, compared to individuals, groups show stronger concerns about maximizing outcomes for the greater good, an interesting question for future research is why groups and individuals differ in their concerns about outcomes. Several possible mechanisms warrant consideration. Based on prior work suggesting that deliberation and analytical thinking support utilitarian choices in the traditional dilemma paradigm (Greene et al., 2008; Patil et al., 2021), Keshmirian et al. (2022) suggested that, compared to individuals, groups may show stronger support for norm-violating outcome-maximizing actions because group interactions support deliberation and analytical thinking (e.g., Mercier & Sperber, 2011; Smith & Collins, 2009). Indeed, group decision-making often involves considering more alternatives compared to solo decision-making (Laughlin et al., 2006), and groups must justify preferred options to others, promoting more careful thinking (Lerner & Tetlock, 1999). Under this view, there may be nothing unique to group decision-making, as prompts to deliberate may similarly increase support for norm-violating outcome-maximizing actions among individuals deciding solo.

However, while these ideas align with widely held theoretical assumptions in the moral-dilemma

² According to Cohen (1988), an ω of 0.10 represents a small effect, an ω of 0.30 represents a medium effect, and an ω of 0.50 represents a large effect.

literature, the available evidence suggests otherwise. Meta-analytic evidence indicates that greater deliberation increases both utilitarian concerns about outcomes and deontological concerns about moral norms (Gawronski & Ng, 2024), which does not align with the current finding that groups and individuals differed only in terms of their sensitivity to consequences, but not in terms of their sensitivity to moral norms. Moreover, justifying decisions has been found to increase sensitivity to moral norms without affecting sensitivity to consequences (Ng et al., 2024), further challenging the deliberation explanation.

An alternative interpretation worth considering is that fatigue from longer group sessions might affect results. However, research suggests that fatigue either reduces support for norm-violating outcome-maximizing actions (Timmons & Byrne, 2019) or does not affect moral judgments (Stefańczyk et al., 2024). Moreover, in research using the CNI model, fatigue-induced noise would reduce scores on both the *C* and *N* parameters and push scores on the *I* parameter toward 0.5—patterns we did not observe in the current study. In fact, our results for the *C* parameter show the opposite pattern.

Next, discussions focusing on consequences may converge more easily given their objective nature, while discussions focusing on norms may diverge due to conflicting moral rules. This explanation would predict increased prominence of consequences in group decisions due to their measurability, while differing cultural or religious backgrounds might complicate norm-focused discussions. However, this explanation also falls short, because it predicts not only increased scores on the *C* parameter but also decreased scores on the *N* parameter—a pattern our data does not support.

The social context of group decision-making also deserves consideration. Participants might adjust their expressed opinions to make a favorable impression on other group members. Yet, given that people prefer collaborating with deontological (Everett et al., 2016), like-minded individuals (Bostyn et al., 2023), self-presentation would not sway people to present themselves as more utilitarian. If anything, prior research using the CNI model suggests that such a mechanism should increase sensitivity to moral norms in groups instead of increasing sensitivity to consequences (Gawronski, 2022).

Finally, another potential explanation involves the accessibility of utilitarian versus deontological arguments. While no direct evidence exists for differential articulation of outcome-based versus norms-based arguments, prior research has found varying accessibility of harm justifications (Cushman et al., 2006). These findings suggest that utilitarian arguments might be more easily expressed and

discussed in groups. Future research would be helpful to test this potential explanation.

Related to this point, studies using a qualitative approach may provide deeper insights into why groups show stronger concerns about outcomes than individuals. For example, one could include think-aloud protocols for solo decision-makers and recordings of group discussions and apply language-analysis tools to the transcripts to better understand the reasoning processes in the two conditions. Future studies could also incorporate theoretical developments suggesting that utilitarianism is a two-dimensional construct, comprising impartial beneficence and instrumental harm as two conceptually distinct components (Kahane et al., 2018). Impartial beneficence involves support for actions that benefit everyone equally, emphasizing the greatest good for the greatest number without favoritism. Instrumental harm, on the other hand, involves support for sacrificial harm if it increases overall well-being or reduces total harm. Future research testing the effects of group vs. solo settings on impartial beneficence and instrumental harm may provide further insights into how moral decisions made by groups differ from those made by individuals.

While the current findings provide more nuanced insights into why groups and individuals show different preferences in moral dilemmas, our study also has some limitations. First, our participants were from a WEIRD (western, educated, industrialized, rich and democratic) culture (Henrich et al., 2010), potentially limiting generalizability. Some non-WEIRD cultures (e.g., traditional Yali people of Papua) value deontological decisions more than WEIRD cultures (Sorokowski et al., 2020), suggesting potential differences in individual-group dynamics in non-WEIRD settings.

Second, our study's virtual format, necessitated by the COVID-19 pandemic, may have influenced results. While we replicated earlier in-person findings, face-to-face communication or unlimited time might alter the impact of consequences, moral norms, and action tendencies on group decisions. For example, online discussions, as opposed to in-person ones, have been found to involve more nuanced discussions of sensitive topics and allow for deeper sharing of personal stories (Woodyatt et al., 2016), which might influence the relative impact of moral norms.

Our study also did not consider the pre-existing moral inclinations of the participants who made moral judgments in groups. Thus, it is possible that participants with deontological inclinations were simply outnumbered by a utilitarian majority, which may promote utilitarian responses in groups. We deliberately chose the current between-subject design to prevent carryover effects or tendencies to adhere to initial choices made individually. While this concern

does not question our conclusions about a greater sensitivity to consequences in groups, future research may help to rule out the effects of numeric majorities.

In sum, the current study suggests that groups' (vs. individuals') stronger support for norm-violating outcome-maximizing actions is a robust and large effect (Funder & Ozer, 2019). More importantly, our findings suggest that this difference is driven by a stronger focus on consequences among groups. We did not find any evidence for the alternative ideas that groups are less concerned about moral norms or less action averse than individuals.

Author Note

This research was supported by grant no. 2020/38/E/HS6/00282 from the National Science Centre (NCN, Poland) to Michał Białek. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. The authors declare no conflicting interests.

Author Contributions

MB - Conceptualization; Data curation; Formal analysis; Funding acquisition; Methodology; Supervision; Validation; Roles/Writing - original draft; MR - Conceptualization; Data curation; Formal analysis; Investigation; Methodology; Project administration; Validation; Visualization; Roles/Writing - original draft; MMS - Conceptualization; Validation; Roles/Writing - original draft; BG - Conceptualization; Methodology; Writing – original draft

References

- Białek, M., Paruzel-Czachura, M., & Gawronski, B. (2019). Foreign language effects on moral dilemma judgments: An analysis using the CNI model. *Journal of Experimental Social Psychology, 85*, Article 103855.
- Bostyn, D.H., Chandrashekar, S.P. & Roets, A. (2023). Deontologists are not always trusted over utilitarians: revisiting inferences of trustworthiness from moral judgments. *Scientific Reports, 13*, Article 1665.
- Conway, P., & Gawronski, B. (2013). Deontological and utilitarian inclinations in moral decision making: A process dissociation approach. *Journal of Personality and Social Psychology, 104*(2), 216–235.
- Crone, D. L., & Laham, S. M. (2017). Utilitarian preferences or action preferences? De-confounding action and moral code in sacrificial dilemmas. *Personality and Individual Differences, 104*, 476–481.
- Curşeu, P. L., Fodor, O. C., A. Pavelea, A., & Meslec, N. (2020). “Me” versus “We” in moral dilemmas: Group composition and social influence effects on group utilitarianism. *Business Ethics: A European Review, 29*(4), 810–823.
- Everett, J. A., Pizarro, D. A., & Crockett, M. J. (2016). Inference of trustworthiness from intuitive moral judgments. *Journal of Experimental Psychology: General, 145*(6), 772–787.
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods, 39*, 175–191.
- Funder, D. C., & Ozer, D. J. (2019). Evaluating effect size in psychological research: Sense and nonsense. *Advances in Methods and Practices in Psychological Science, 2*(2), 156–168.
- Gawronski, B. (2022). Moral impressions and presumed moral choices: Perceptions of how moral exemplars resolve moral dilemmas. *Journal of Experimental Social Psychology, 99*, Article 104265.
- Gawronski, B., Armstrong, J., Conway, P., Friesdorf, R., & Hütter, M. (2017). Consequences, norms, and generalized inaction in moral dilemmas: The CNI model of moral decision-making. *Journal of Personality and Social Psychology, 113*(3), 343–376.
- Gawronski, B., & Ng, N. L. (2024). Beyond trolleyology: The CNI model of moral-dilemma responses. *Personality and Social Psychology Review*. Advance online publication.
- 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*(3), 1144–1154.
- Hart, E., Marciano, D., & Winter, E. (2017). Groups, group members and individuals: Choices and impulses in repeated risky decisions. *Maurice Falk Institute for Economic Research in Israel. Discussion Paper Series, 7*, 1–43.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). Most people are not WEIRD. *Nature, 466*(7302), 29.
- Kahane, G., Everett, J. A. C., Earp, B. D., Caviola, L., Faber, N. S., Crockett, M. J., & Savulescu, J. (2018). Beyond sacrificial harm: A two-dimensional model of utilitarian psychology. *Psychological Review, 125*(2), 131–164.
- Keshmirian, A., Deroy, O., & Bahrami, B. (2022). Many heads are more utilitarian than one. *Cognition, 220*, Article 104965.
- Körner, A., Deutsch, R., & Gawronski, B. (2020). Using the CNI model to investigate individual differences in moral dilemma judgments. *Personality and Social Psychology Bulletin, 46*(9), 1392–1407.

- Lakens, D. (2022). Sample size justification. *Collabra: Psychology*, 8(1), Article 33267.
- Lantos, D., & Molenberghs, P. (2021). The neuroscience of intergroup threat and violence. *Neuroscience and Biobehavioral Reviews*, 131, 77–87.
- Laughlin, P. R., Hatch, E. C., Silver, J. S., & Boh, L. (2006). Groups perform better than the best individuals on letters-to-numbers problems: Effects of group size. *Journal of Personality and Social Psychology*, 90(4), 644–651.
- Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, 125(2), 255–275.
- Mercier, H., & Sperber, D. (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences*, 34(2), 57-74.
- Moshagen, M. (2010). multiTree: A computer program for the analysis of multinomial processing tree models. *Behavior Research Methods*, 42(1), 42–54.
- Ng, N. L., Luke, D. M., & Gawronski, B. (2024). Thinking about reasons for one's choices increases sensitivity to moral norms in moral-dilemma judgments. *Personality and Social Psychology Bulletin*. Advance Online Publication.
- Paruzel-Czachura, M., Pypno, K., Everett, J. A. C., Białek, M., & Gawronski, B. (2023). The drunk utilitarian revisited: Does alcohol really increase utilitarianism in moral judgment? *Personality and Social Psychology Bulletin*, 49(1), 20-31.
- Patil, I., Zucchelli, M. M., Kool, W., Campbell, S., Fornasier, F., Calò, M., Silani, G., Cikara, M., & Cushman, F. (2021). Reasoning supports utilitarian resolutions to moral dilemmas across diverse measures. *Journal of Personality and Social Psychology*, 120(2), 443–460.
- Smith, E. R., & Collins, E. C. (2009). Contextualizing person perception: Distributed social cognition. *Psychological Review*, 116(2), 343–364.
- Stefańczyk, M., Zurek, G., Macyszyn, A., Ochman, A., Czajka, K., & Białek, M. (2024). Moral decision making is robust to physical fatigue: A Process Dissociation analysis. *PsyArXiv*. <https://doi.org/10.31234/osf.io/3g5ca>
- Sorokowski, P., Marczak, M., Misiak, M., & Białek, M. (2020). Trolley dilemma in Papua. Yali horticulturalists refuse to pull the lever. *Psychonomic Bulletin & Review*, 27(2), 398–403.
- The Jamovi Project. (2021). *Jamovi 2.2.5* [Computer software]. <https://www.jamovi.org/>
- Timmons, S., & Byrne, R. M. (2019). Moral fatigue: The effects of cognitive fatigue on moral reasoning. *Quarterly Journal of Experimental Psychology*, 72(4), 943-954.
- Wilson, D. B. (2023). Practical meta-analysis effect size calculator (Version 2023.11.27). <https://www.campbellcollaboration.org/escalc/html/EffectSizeCalculator-Home.php>
- Woodyatt, C. R., Finneran, C. A., & Stephenson, R. (2016). In-person versus online focus group discussions: A comparative analysis of data quality. *Qualitative Health Research*, 26(6), 741–749.

Table 1. Example of a moral dilemma involving either a proscriptive or a prescriptive norm where the benefits of action are either greater or smaller than the costs of action. Dilemmas adapted from Gawronski, et al., (2017). Reprinted with permission.

Proscriptive Norm Prohibits Action Benefits of Action Greater than Costs	Proscriptive Norm Prohibits Action Benefits of Action Smaller than Costs	Prescriptive Norm Prescribes Action Benefits of Action Greater than Costs	Prescriptive Norm Prescribes Action Benefits of Action Smaller than Costs
<p>You are the director of a hospital in a developing country. A foreign student who is volunteering in the country got infected with a rare virus.</p> <p>The virus is highly contagious and deadly to seniors and children. The only medication that can effectively stop the virus from spreading has severe side-effects. Although the virus will not kill her, the student suffers from a chronic immune deficiency that will make her die from these side-effects.</p> <p>Would you give the student the medication?</p>	<p>You are the director of a hospital in a developing country. A foreign student who is volunteering in the country got infected with a rare virus.</p> <p>The virus is highly contagious and can cause severe stomach cramps. The only medication that can effectively stop the virus from spreading has severe side-effects. Although the virus will not kill her, the student suffers from a chronic immune deficiency that will make her die from these side-effects.</p> <p>Would you give the student the medication?</p>	<p>You are the director of a hospital in a developing country. A foreign student who is volunteering in the country got infected with a rare virus.</p> <p>The virus is highly contagious and can cause severe stomach cramps. The student suffers from a chronic immune deficiency that will make her die from the virus if she is not returned to her home country for special treatment. However, taking her out of quarantine involves a considerable risk that the virus will spread.</p> <p>Would you take the student out of quarantine to return her to her home country for treatment?</p>	<p>You are the director of a hospital in a developing country. A foreign student who is volunteering in the country got infected with a rare virus.</p> <p>The virus is highly contagious and deadly to seniors and children. The student suffers from a chronic immune deficiency that will make her die from the virus if she is not returned to her home country for special treatment. However, taking her out of quarantine involves a considerable risk that the virus will spread.</p> <p>Would you take the student out of quarantine to return her to her home country for treatment?</p>

Figure 1. Three potential mechanisms underlying differential support for outcome-maximizing actions that violate moral norms. The dashed line depicts the prior finding that groups (vs. individuals) show stronger support for outcome-maximizing actions that violate moral norms. Solid lines depict the three potential mechanisms underlying this effect, in that groups (vs. individuals) may show (1) stronger concerns about outcomes, (2) weaker concerns about moral norms, or (3) weaker action aversion.

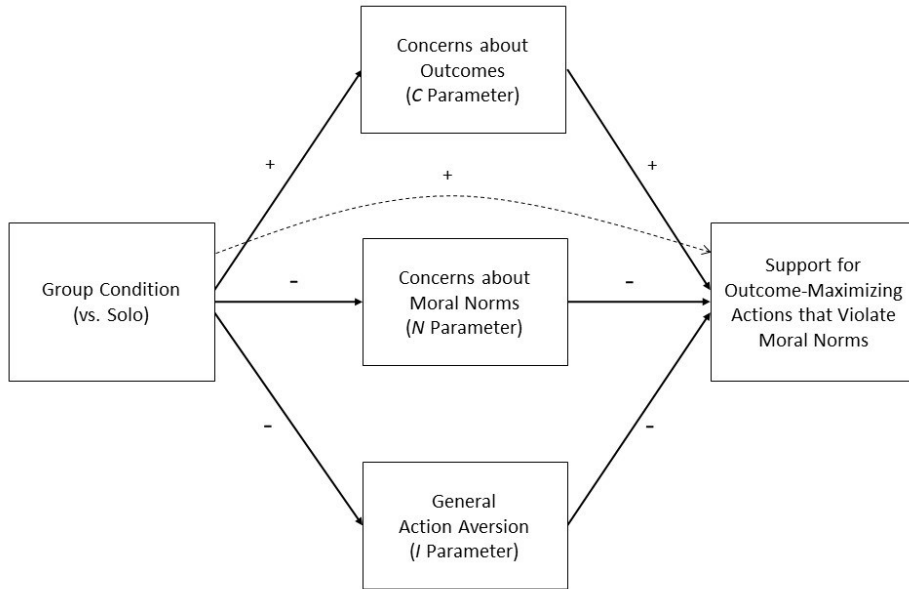


Figure 2. CNI model of moral decision-making predicting action versus inaction responses in moral dilemmas with proscriptive and prescriptive norms and consequences involving benefits of action that are either greater or smaller than the costs of action. Reproduced from Gawronski, Armstrong, Conway, Friesdorf, and Hütter (2017). Reprinted with permission from the American Psychological Association.

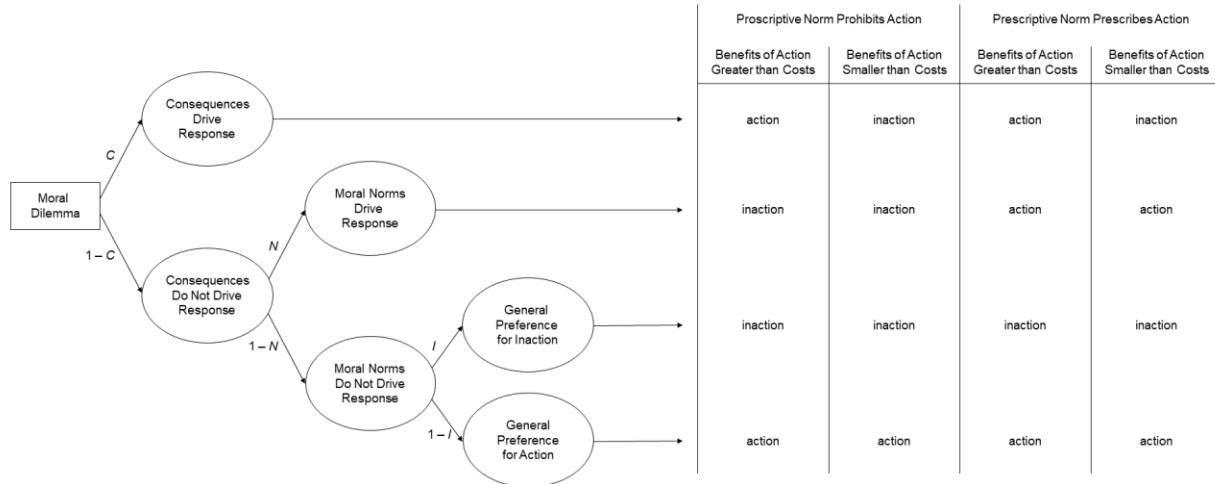
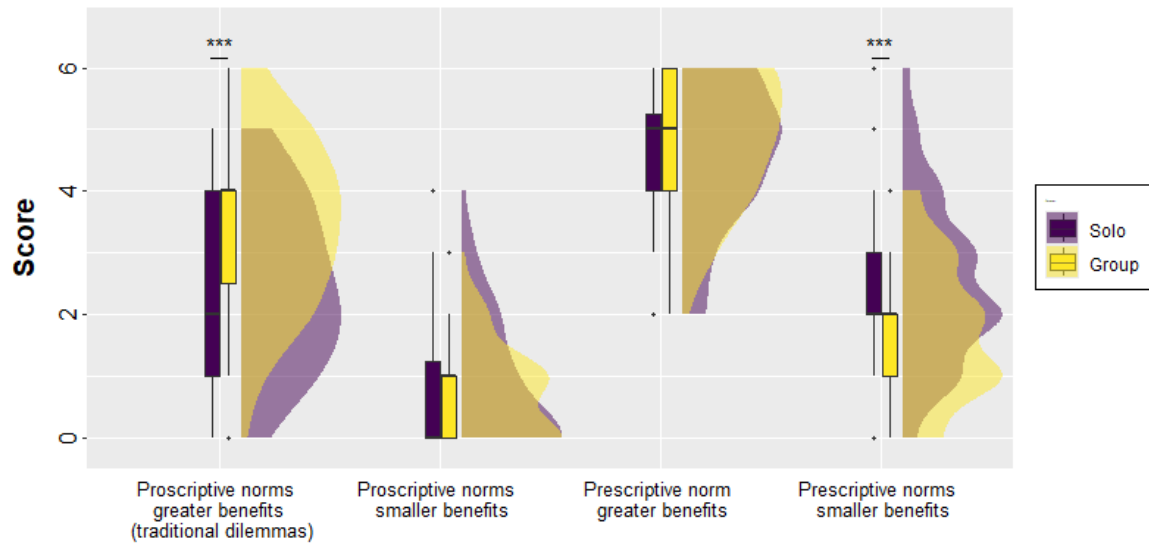



Figure 2. Preference for action (vs. inaction) as a function of group condition (solo vs. group), moral norms (proscriptive vs. prescriptive), and cost-benefit ratios (benefits of action greater than costs vs. benefits of action smaller than costs).



Note. *** $p < .001$. Minimum score is 0. Maximum score is 6.

Figure 3. Mean parameter estimates for sensitivity to consequences (*C Parameter*), sensitivity to moral norms (*N Parameter*), and general preference for inaction versus action (*I Parameter*) as a function of group condition (solo vs. group).



Note. *** $p < .001$. Error bars represent 95% confidence intervals.