



Moral impressions and presumed moral choices: Perceptions of how moral exemplars resolve moral dilemmas[☆]

Bertram Gawronski

Department of Psychology, University of Texas at Austin, 108 E Dean Keeton A8000, Austin, TX 78712, USA

ARTICLE INFO

Keywords:

CNI model
Deontology
Moral exemplar
Moral impressions
Utilitarianism

ABSTRACT

Previous research suggests that individuals who prefer deontological over utilitarian choices in moral dilemmas are perceived to have stronger moral character than individuals who show the reverse preference. To gain deeper insights into the link between moral choices and moral impressions, the current research used a formal modeling approach to examine whether morally exceptional figures are perceived to differ from others in their sensitivity to consequences, sensitivity to moral norms, or general action tendencies when resolving moral dilemmas. Findings from four studies ($N = 980$) suggest that perceived morality is associated with greater presumed adherence to moral norms in the resolution of moral dilemmas. For sensitivity to consequences and general action tendencies, findings were mixed and attributable to characteristics confounded with perceived morality. The findings suggest a hitherto unexplored mechanism underlying moral-dilemma judgments by which moral judgments are based on mental simulations of decisions by morally exceptional figures.

Normative theories of morality in philosophy have often been used as guides for empirical research on moral judgment in psychology. One prominent line of work in this area has examined how people resolve moral dilemmas in which choices maximizing overall consequences for the greater good (*utilitarianism*) conflict with choices conforming to moral norms (*deontology*). A dominant question in this work concerns the processes underlying outcome-maximizing and norm-conforming judgments (Crockett, 2013; Cushman, 2013; Greene, 2008; Holyoak & Powell, 2016). Expanding on this work, a growing body of research has investigated moral impressions of individuals who make either outcome-maximizing or norm-conforming judgments in moral dilemmas. A central finding of this research is that individuals who make norm-conforming, deontological judgments are perceived as having a stronger moral character than those who make outcome-maximizing, utilitarian judgments (for a review, see Crockett, Everett, Gill, & Siegel, 2021). The goal of the current research was to provide deeper insights into the link between moral choices and moral impressions by examining whether morally exceptional figures are perceived to differ from others in their sensitivity to consequences, sensitivity to moral norms, or general action tendencies when resolving moral dilemmas.

1. Moral choices and moral impressions

A prominent line of research in moral psychology has drawn on the philosophical traditions of utilitarianism and deontology to examine how people resolve moral dilemmas that pit overall consequences for the greater good against adherence to relevant moral norms (Greene, Nystrom, Engell, Darley, & Cohen, 2004; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001). The most well-known dilemma from this line of research is the trolley problem, a scenario in which a runaway trolley is set on a collision course with five railroad workers (Foot, 1967). In a variant called the *footbridge dilemma*, it is possible to push a large man in front of the runaway trolley, killing the man but obstructing the trolley from killing the five workers (Thomson, 1976). Judgments in favor of pushing the man have been described as *characteristically utilitarian* in the sense that they maximize overall outcomes for the greater good, while judgments in opposition of pushing the man have been described as *characteristically deontological* in the sense that they conform to moral rules or duties relevant to the situation (Conway, Goldstein-Greenwood, Polacek, & Greene, 2018). Research using this paradigm revealed that people differ in their relative preference for utilitarian over deontological judgments (e.g., Gleichgerrcht & Young, 2013; Moore, Stevens, & Conway, 2011; Patil, 2015), that these individual differences are relatively stable over time (e.g., Hannikainen,

[☆] This paper has been recommended for acceptance by Dr. Joris Lammers.
E-mail address: gawronski@utexas.edu.

Machery, & Cushman, 2018; Helzer, Fleenon, Furr, Meindl, & Barranti, 2017), and that dilemma-unrelated contextual factors (e.g., time pressure, mood states) can influence relative preferences for one over the other kind of judgment (e.g., Suter & Hertwig, 2011; Valdesolo & DeSteno, 2006). Expanding on this work, a growing body of research has examined moral impressions of others based on how they resolve moral dilemmas. A central finding of this work is that people who prefer norm-adhering, deontological choices in moral dilemmas are perceived as being more moral and trustworthy than those who prefer outcome-maximizing, utilitarian choices (e.g., Bostyn & Roets, 2017; Critcher, Helzer, & Tannenbaum, 2020; Everett, Faber, Savulescu, & Crockett, 2018; Everett, Pizarro, & Crockett, 2016; Rom, Weiss, & Conway, 2017; Sacco, Brown, Lustgraaf, & Hugenberg, 2017; Turpin et al., 2021; Uhlmann, Zhu, & Tannenbaum, 2013).

2. Conceptual ambiguities

While the available evidence suggests a systematic relation between moral impressions and preference for deontological over utilitarian judgments, the meaning of this relation remains conceptually ambiguous. Across the majority of past research, deontological and utilitarian judgments have been measured using moral dilemmas similar in structure to the trolley problem, which carry with them two notable confounds. First, these dilemmas pit maximization of outcomes and adherence to moral norms against one another, such that endorsement of one necessarily implies rejection of the other. As a result, maximization of outcomes and adherence to moral norms are confounded with one another, even though they have been claimed to be the product of distinct psychological mechanisms (Conway & Gawronski, 2013). Second, maximization of outcomes is typically conflated with action (e.g., pushing the man), while adherence to moral norms is conflated with inaction (e.g., not pushing the man). As a result, maximization of outcomes and adherence to moral norms are further confounded with general preferences for acting and not acting, respectively (Crone & Laham, 2017). Together, the two confounds render relations between moral impressions and preference for deontological over utilitarian judgments conceptually ambiguous, because they could be driven by (1) a negative relation between perceived morality and maximization of outcomes, (2) a positive relation between perceived morality and adherence to moral norms, or (3) a positive relation between perceived morality and general preference for inaction over action (or a complex combination of relations).

All three possibilities seem plausible in light of prior work. First, it is possible that perceived morality is negatively related to maximization of outcomes. In line with this idea, the cost-benefit analyses associated with maximizing overall welfare may be perceived as producing erratic behavior (Sacco et al., 2017) and as being divorced from moral or empathic concern (Kreps & Monin, 2014; Uhlmann et al., 2013). Second, it is possible that perceived morality is positively related to adherence to moral norms. In line with this idea, people who adhere to moral norms may be perceived as reliable and predictable (Everett et al., 2016; Turpin et al., 2021) and as having strong empathic concern for others (Everett et al., 2016; Rom et al., 2017). Finally, it is possible that perceived morality is positively related to general preference for inaction over action. In line with this idea, a general preference for inaction may be perceived as reflecting a person's concern about potentially harmful effects of their actions, which should lead to a general bias against action regardless of the situation (Baron & Goodwin, 2020; Cushman, Young, & Hauser, 2006). The goal of the current research was to investigate these different possibilities by using the CNI model of moral decision-making (Gawronski, Armstrong, Conway, Friesdorf, & Hütter, 2017) to overcome the conceptual ambiguities of the traditional approach to moral-dilemma judgments.

3. The CNI model

The CNI model is a multinomial model that quantifies (1) sensitivity to consequences, (2) sensitivity to moral norms, and (3) general preference for inaction versus action in responses to moral dilemmas (Gawronski et al., 2017). To quantify these three distinct factors, the CNI model relies on responses to four types of moral dilemmas that differ in terms of cost-benefit ratios (the benefits of the focal action are either greater or smaller than the costs) and salient moral norms (the focal action is either prohibited or prescribed by a moral norm). Sensitivity to consequences is captured by the model's *C* parameter, which quantifies the extent to which actions are favored when they produce greater benefits than costs and opposed when they produce smaller benefits than costs (see first row in Fig. 1). Sensitivity to moral norms is captured by the model's *N* parameter, which quantifies the extent to which actions are favored when they are prescribed by a moral norm and opposed when they are prohibited by a moral norm (see second row in Fig. 1). General preference for inaction versus action is captured by the model's *I* parameter, which quantifies the extent to which actions are generally opposed (see third row in Fig. 1) or generally favored (see fourth row in Fig. 1). Previous research using the CNI model has provided valuable insights into the effects of cognitive resources (Gawronski et al., 2017), personal involvement (Gawronski et al., 2017), incidental emotions (Gawronski, Conway, Armstrong, Friesdorf, & Hütter, 2018), testosterone (Brannon, Carr, Jin, Josephs, & Gawronski, 2019), language use (Bialek, Paruzel-Czachura, & Gawronski, 2019), social power (Gawronski & Brannon, 2020), political orientation (Luke & Gawronski, 2021a), basic personality traits (Kroneisen & Heck, 2020; Luke & Gawronski, in press), psychopathy (Luke & Gawronski, 2021b; Luke, Neumann, & Gawronski, in press), and alcohol (Paruzel-Czachura, Pypno, Everett, Bialek, & Gawronski, in press) on moral-dilemma judgments. In the current research, we used the CNI model to investigate whether differences in perceived morality are linked to (1) differences in the presumed sensitivity to consequences, (2) differences in the presumed sensitivity to moral norms, or (3) differences in the presumed general preference for inaction versus action (or a complex combination of the three).

4. The current research

To address this question, the current research adopted an experimental procedure by Furr, Prentice, Hawkins, and Fleenon (2020) to investigate the presumed choices of morally exceptional nominees in moral dilemmas. This approach is based on a growing body of research examining how morally exceptional figures are perceived in terms of various personality characteristics (e.g., Hardy, Walker, Olsen, Skalski, & Basinger, 2011; Lapsley & Lasky, 2001; Walker, 1999; Walker & Hennig, 2004; Walker & Pitts, 1998). In the current studies, we used this approach to compare the presumed choices of morally exceptional nominees to those of the average person (Studies 1–2), morally average nominees (Studies 2–4), and socially influential nominees (Studies 3–4).¹ To investigate the replicability of our findings in a formal manner,

¹ An alternative approach to investigate the link between moral-dilemma judgments and moral impressions is to present participants with moral-dilemma judgments of hypothetical figures and ask participants about their moral impressions of these figures (e.g., Everett et al., 2016; Rom et al., 2017). Although this approach works well for the traditional dilemma paradigm, it is not possible to combine it with the CNI model, because the CNI model requires observations of multiple responses on different kinds of dilemmas for the separation of the three judgmental determinants.

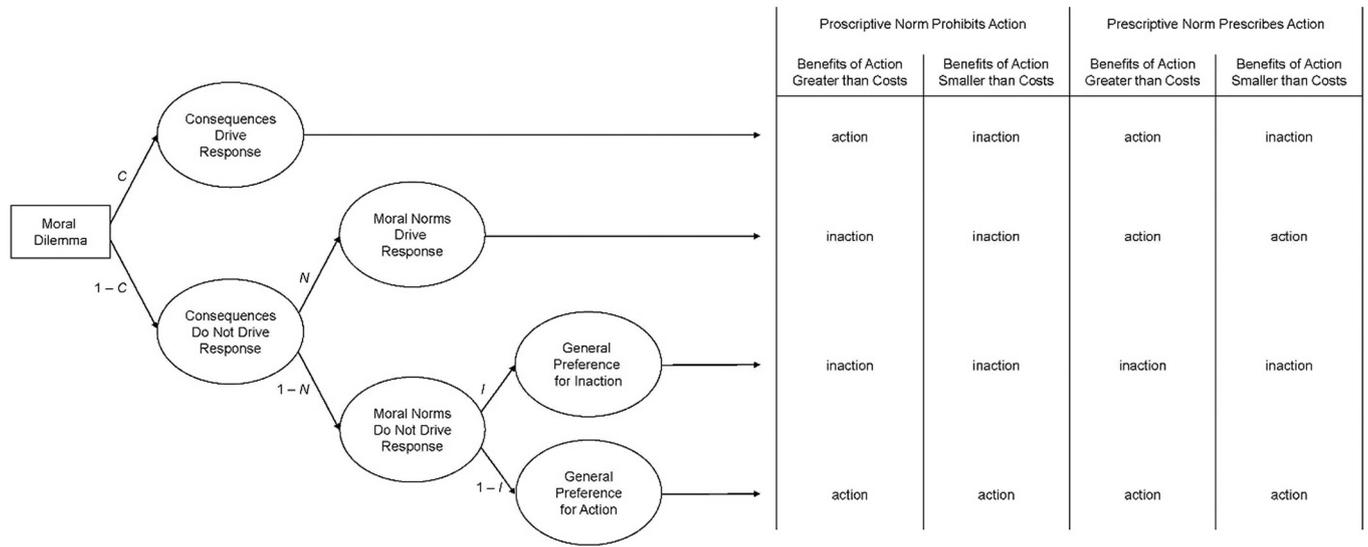


Fig. 1. CNI model of moral decision-making predicting action versus inaction responses in moral dilemmas with proscriptive and prescriptive norms and consequences involving benefits of action that are either greater or smaller than costs of action. Reproduced from Gawronski et al. (2017). Reprinted with permission from the American Psychological Association.

the final study was preregistered prior to data collection (Study 4). For all studies, we aimed to recruit 100 participants per condition which provides a statistical power of 80% in detecting a small-to-medium sized difference of $d = 0.40$ between two independent means (two-tailed).² By default, we excluded participants who failed to pass an instructional attention check (see Oppenheimer, Meyvis, & Davidenko, 2009). We report all data, all measures, and all experimental conditions. The data, analysis codes, and materials for the current studies are available at <https://osf.io/k3f9u/>.

5. Study 1

To provide more nuanced insights into the link between perceived morality and moral choices, participants in Study 1 were asked to nominate a public figure they considered highly moral, and to indicate whether their nominated exemplar would perform the actions described in a series of moral dilemmas. Participants in a control group were asked to indicate whether the average person would perform the described actions. Responses were analyzed using the CNI model to investigate whether the presumed choices of morally exceptional nominees differ from those of the average person in terms of (1) sensitivity to consequences, (2) sensitivity to moral norms, or (3) general preference for inaction over action.

5.1. Method

5.1.1. Participants

Participants were recruited in September 2018 using Amazon's Mechanical Turk (MTurk). Eligibility for participation was restricted to MTurk workers from the United States who were at least 18 years of age, successfully completed at least one prior assignment, had an approval rate of at least 95% across prior assignments, and had not participated in a prior study from the Principal Investigator's lab using the same moral dilemmas. To increase data quality, the assessment included a bot

² Because power analyses within multinomial modeling require simulations with expected population values for the three parameters and any specific expectations in this regard would be arbitrary, we made our a priori sample-size decision in a heuristic fashion based on simple comparisons of mean values using *t*-tests.

prevention question at the beginning of the study, which asked participants to solve a simple addition problem (e.g., “8 + 7”). In addition, the assessment included a reading intensive attention check at the end, which required participants to ignore the question. Participants who failed to solve the addition problem correctly were not allowed to complete the study. Data from participants who failed the attention check were excluded from analyses. Of the 206 participants who completed the survey in its entirety,³ 22 participants failed the attention check, leaving a final sample of 184 participants (47.28% female, 52.17% male, 0.54% other; $M_{age} = 34.27$, $SD_{age} = 10.62$). Of these participants, 85.87% identified as Caucasian, 9.78% as African American, 4.35% as Asian, 2.17% as Native American, and 2.17% as other ethnicities. The final sample of 184 participants provided 80% power in detecting a between-group difference of $d = 0.42$. Participants were compensated \$2.00 for their time.

5.1.2. Procedure and materials

After providing informed consent and completing the bot prevention question, participants were randomly assigned to one of two experimental conditions. Participants in the *moral-exemplar* condition were asked to nominate a public figure they considered highly moral (see Appendix). To avoid inducing artificial relations between perceptions of morality and presumed moral-dilemma judgments, we deliberately did not define *morality* or *moral character* in the nomination prompts, permitting a high degree of subjectivity about what makes a person moral. After the nomination, participants were asked to indicate for a series of moral dilemmas whether their nominated exemplar would perform the action described in the dilemma. The procedure was identical for participants in the *average-person* condition, the main difference being that participants were asked to indicate whether the average person would perform the action described in the scenario (without nominating a moral exemplar). The moral dilemmas were adapted from Körner, Deutsch, and Gawronski (2020), comprising 48 scenarios that varied in terms of cost-benefit ratios (i.e., the benefits of the described action are either greater or smaller than its costs) and salient moral norms (i.e., the described action is either prohibited or prescribed by a

³ Six participants completed the assessment but either did not submit a request for compensation or were denied compensation because they submitted an incorrect completion code.

moral norm). The 48 dilemmas were presented in the same random order for all participants. Each dilemma ended with a question, asking participants whether their nominated moral exemplar or the average person would perform the described action. Responses were measured with dichotomous *yes vs. no* forced-choice options. After responding to all dilemmas, participants completed a set of demographic questions regarding their age, gender, ethnicity, race, and education. Lastly, participants completed a reading intensive attention check, were thanked for their participation, and given a code for compensation.

5.1.3. Attention check

The attention check required participants to read a set of instructions, which ask participants not to answer a question (see Oppenheimer et al., 2009). Participants passed the attention check by *not* answering the question and moving on to the next question without selecting any answer choices. The attention check included the following question:

Many variables can greatly impact decision-making. In order to facilitate our research on decision-making we are interested in knowing certain factors about you, the decision maker. Specifically, we are interested in whether you actually take the time to read the directions; if not, then some of our manipulations that rely on changes in the instructions will be ineffective. So, in order to demonstrate that you have read the instructions, please ignore the sports items below. Instead, simply continue on to the next page after the options. Thank you very much. Which of these activities do you engage in regularly? (check all that apply):

The response options were: *Football, Soccer, Dancing, Watersports, Triathlon, Running, Volleyball, and I don't play sports*. If a participant answered the question by selecting any of the eight response options, the participant failed to correctly follow instructions and was therefore excluded from analyses.

5.1.4. Analyses

To provide more nuanced insights into the link between perceptions of morality and presumed moral-dilemma judgments, we analyzed participants' responses using the CNI model. Because the statistical details of CNI model analyses are described in detail elsewhere (Gawronski et al., 2017), we only summarize the basic steps of the modeling analysis. Based on the processing tree depicted in Fig. 1, the CNI model provides four non-redundant equations that include the three model parameters as unknowns and the empirically observed probabilities of action versus inaction responses on the four dilemma variants as known numerical values. Using maximum likelihood statistics, it is possible to estimate numerical values for the three unknowns, such that the discrepancy between the observed probabilities of action versus inaction responses across the four dilemma variants and the predicted probabilities of action versus inaction responses based on the model equations is minimized. The adequacy of the model in describing the data can be evaluated by means of goodness-of-fit statistics, in that poor model fit would be reflected in a statistically significant discrepancy between the empirically observed probabilities in a given data set and the probabilities predicted by the model for this data set. Differences in parameter estimates across groups can be tested by enforcing equal estimates for a given parameter across groups and comparing the fit of the constrained model to the fit of the baseline model. If setting a given parameter equal across groups leads to a significant reduction in model fit, it can be inferred that the parameter estimates for the two groups are significantly different. If setting a given parameter equal across groups does not lead to a significant reduction in model fit, the parameters for the two groups are not significantly different from each other. Following Gawronski et al. (2017), the analyses used a fixed estimation algorithm with random start values, two replications, and a maximum of 90,000 iterations. All analyses were conducted using the freeware multiTree (Moshagen, 2010) and the template files for CNI model analyses

provided by Gawronski et al. (2017) at <https://osf.io/xt66w/>. Following Gawronski et al. (2017), effect sizes of between-group differences were calculated with Lipsey and Wilson's (2001) online companion to their practical introduction to meta-analysis at <https://www.campbellcollaboration.org/escalc/html/EffectSizeCalculator-SMD8.php> using means, standard errors, and sample sizes.

5.2. Results

Responses were aggregated by summing the number of *action* and *inaction* responses for each of the four dilemma variants for each participant. With 12 scenarios for each of the four dilemma variants, aggregate scores could range from 0 to 12. Means and 95% confidence intervals of aggregate scores in the two conditions can be seen in Table 1. CNI parameter estimates and 95% confidence intervals by condition can be seen in Fig. 2.

The CNI model fit the data well, $G^2(2) = 0.07, p = .964$. There was a significant effect of figure type on the *C* parameter, $\Delta G^2(1) = 16.72, p < .001, d = 0.61$, indicating that sensitivity to consequences was lower in the moral exemplar condition than in the average person condition. There was also a significant effect of figure type on the *N* parameter, $\Delta G^2(1) = 10.16, p = .001, d = 0.47$, indicating that sensitivity to moral norms was higher in the moral exemplar condition than in the average person condition. There was no significant effect of figure type on the *I* parameter, $\Delta G^2(1) = 1.56, p = .21, d = 0.19$. Together, these results suggest that moral exemplars were perceived to be less sensitive to consequences and more sensitive to moral norms than the average person.

5.3. Discussion

The results of Study 1 provide more nuanced insights into the link between perceived morality and moral choices, suggesting that moral exemplars are perceived to be less sensitive to consequences and more sensitive to moral norms than the average person. Nevertheless, there are two notable limitations. The first limitation is that the abstraction of target figures differed across conditions. Whereas participants in the moral-exemplar condition were asked to make judgments about a specific person, participants in the average-person condition were asked to indicate how an abstract prototype (i.e., the average person) would respond to the moral dilemmas. Prior research suggests that person perception can differ in notable ways when considering an exemplar of a category (e.g., Albert Einstein) or an abstracted prototype of a category (e.g., an intelligent person) (e.g., Paulhus, 2000). Therefore, it remains unclear whether the obtained differences are due to different levels of perceived morality or different levels of abstraction. The second limitation is that the target figures might differ in their prominence in society. While the target figures in the moral-exemplar condition were constrained to be widely known in society, the target figure in the average-person condition was not constrained in this way. Therefore, it remains unclear whether the obtained differences are due to perceived morality or social prominence. Study 2 aimed to address these questions.

6. Study 2

The purpose of Study 2 was twofold. The first purpose was to replicate the obtained difference in perceptions of moral exemplars and the average person. The second purpose was to resolve the confounds in Study 1. Toward this end, Study 2 included an *average-exemplar* condition in addition to the *moral-exemplar* and *average-person* conditions of Study 1. In the average-exemplar condition, participants were asked to nominate a public figure who they perceived as "morally average" and to indicate how this person would respond to the same moral dilemmas. Because the target figure in the average exemplar condition is an actual public figure, this condition controls for both target abstraction and social prominence.

Table 1

Means and 95% confidence intervals of action (vs. inaction) responses on moral dilemmas with proscriptive and prescriptive norms and consequences involving benefits of action that are either greater or smaller than costs of action by condition. Scores can range from 0 to 12. The neutral reference value of equal numbers of action and inaction responses is 6.

	Proscriptive norm prohibits action				Prescriptive norm prescribes action			
	Benefits of action greater than costs		Benefits of action smaller than costs		Benefits of action greater than costs		Benefits of action smaller than costs	
	M	95% CI	M	95% CI	M	95% CI	M	95% CI
Study 1								
Moral Exemplar	4.44	[3.87, 5.02]	3.20	[2.55, 3.86]	9.05	[8.59, 9.50]	7.90	[7.40, 8.40]
Average Person	5.64	[5.10, 6.17]	3.48	[2.80, 4.15]	9.13	[8.74, 9.51]	6.96	[6.45, 7.47]
Study 2								
Moral Exemplar	4.74	[4.06, 5.43]	3.30	[2.55, 4.05]	9.81	[9.44, 10.19]	8.38	[7.89, 8.88]
Average Person	5.41	[4.80, 6.02]	3.21	[2.51, 3.91]	9.60	[9.17, 10.03]	7.13	[6.60, 7.65]
Average Exemplar	5.62	[4.98, 6.25]	3.42	[2.72, 4.11]	8.90	[8.43, 9.36]	7.01	[6.47, 7.56]
Study 3								
Moral Exemplar	4.62	[4.07, 5.18]	2.93	[2.27, 3.59]	9.36	[8.92, 9.81]	7.67	[7.11, 8.23]
Average Exemplar	4.94	[4.42, 5.46]	2.73	[2.09, 3.38]	9.34	[8.89, 9.79]	6.65	[6.15, 7.15]
Influential Exemplar	4.80	[4.19, 5.41]	2.95	[2.29, 3.62]	9.24	[8.67, 9.81]	7.44	[6.87, 8.00]
Study 4								
Moral Exemplar	4.83	[4.14, 5.53]	3.39	[2.62, 4.16]	9.92	[9.54, 10.31]	8.30	[7.83, 8.77]
Average Exemplar	5.84	[5.19, 6.48]	3.71	[2.91, 4.51]	9.58	[9.19, 9.98]	7.64	[7.08, 8.19]
Influential Exemplar	5.12	[4.42, 5.82]	3.55	[2.82, 4.29]	9.27	[8.69, 9.84]	7.51	[6.91, 8.10]

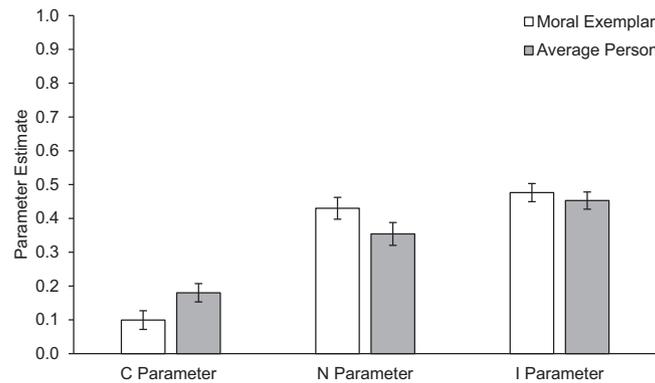


Fig. 2. Parameter estimates of sensitivity to consequences (C), sensitivity to moral norms (N), and general preference for inaction over action (I) as a function of figure type (moral exemplar vs. average person), Study 1. Error bars reflect 95% confidence intervals.

6.1. Method

6.1.1. Participants

Participants were recruited in September 2018 using Amazon's MTurk. Participation eligibility criteria and data quality precautions were the same as in Study 1. Of the 304 participants who completed the survey in its entirety,⁴ 38 participants failed the attention check, leaving a final sample of 266 participants (53.76% female, 45.49% male, 0.75% prefer not to answer; $M_{age} = 36.21$, $SD_{age} = 11.43$). Of these participants, 79.32% identified as Caucasian, 13.16% as African American, 6.77% as Asian, 1.50% as Native American, and 1.50% as other ethnicities. The final sample of 266 participants provided 80% power in detecting a difference of $f = 0.19$ across the three experimental groups. Participants were compensated \$2.00 for their time.

6.1.2. Procedure and materials

The procedure and materials were identical to those used in Study 1

⁴ Four participants completed the assessment but either did not submit a request for compensation or were denied compensation because they submitted an incorrect completion code.

with two exceptions. First, Study 2 included an average-exemplar condition in addition to the moral-exemplar and average-person conditions of Study 1. In the average-exemplar condition, participants were asked to nominate a public figure they considered morally average (see Appendix) and then complete the same moral dilemma battery by making judgments about whether their nominated exemplar would perform the action described in the scenario. Second, to ensure that the figures in the three conditions differed in terms of perceived morality, a one-item manipulation check was included (*How moral do you consider [figure]?*). Participants in the moral-exemplar and average-exemplar conditions were asked to rate the morality of their nominated exemplar; participants in the average-person condition were asked to rate the morality of the average person. Perceived morality was assessed with a 7-point rating scale ranging from 1 (*not at all moral*) to 7 (*extremely moral*).

6.1.3. Analyses

The analyses for Study 2 were identical to Study 1 with two exceptions. First, as a manipulation check, analyses were conducted to determine whether the perceived morality of figures was successfully manipulated across conditions. Second, differences in parameter estimates were tested by enforcing equal estimates for a given parameter

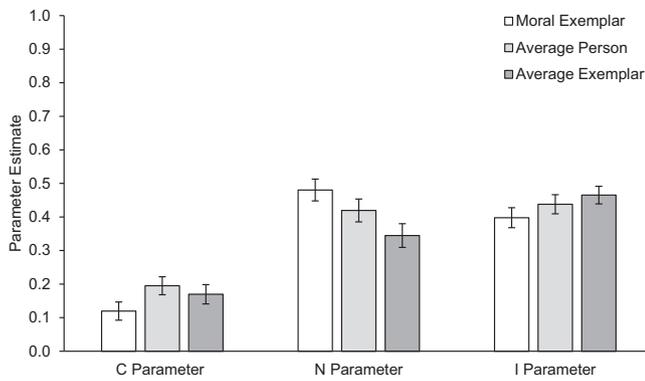


Fig. 3. Parameter estimates of sensitivity to consequences (C), sensitivity to moral norms (N), and general preference for inaction over action (I) as a function of figure type (moral exemplar vs. average person vs. average exemplar), Study 2. Error bars reflect 95% confidence intervals.

across all three groups and comparing the fit of the constrained model to the fit of the baseline model. If setting a given parameter equal across the three groups leads to a significant reduction in model fit, it can be inferred that the parameter estimates are significantly different across groups. In such cases, we conducted follow-up analyses in which we constrained the focal parameter to be equal across different pairs of conditions. If the fit of any of the constrained models was significantly worse compared to the baseline model, a significant difference was inferred to exist between the pair of conditions for the focal parameter.

6.2. Results

Moral judgment data were aggregated in line with the procedures in Study 1. Means and 95% confidence intervals of the aggregated moral judgment data by condition can be seen in Table 1. CNI parameter estimates and 95% confidence intervals by condition can be seen in Fig. 3.

6.2.1. Manipulation check

To test whether the manipulation of perceived morality was successful, morality ratings were submitted to a one-way ANOVA with three levels (average person, average exemplar, moral exemplar). Perceived morality significantly differed across the three conditions, $F(2, 262) = 56.37, p < .001, \eta^2 = 0.30$.⁵ Tukey's HSD post-hoc tests revealed that moral exemplars ($M = 6.50, SD = 0.65$) were perceived to be significantly more moral than average exemplars ($M = 5.36, SD = 1.25$), Tukey's HSD = 1.14, $p < .001, d = 1.14$, as well as the average person ($M = 4.91, SD = 1.07$), Tukey's HSD = 1.59, $p < .001, d = 1.78$. In addition, average exemplars were perceived to be more moral than the average person, Tukey's HSD = 0.45, $p = .011, d = 0.38$. These results indicate that the manipulation of perceived morality was successful in that moral exemplars were perceived as significantly more moral than either average exemplars or the average person.

6.2.2. CNI analysis

The CNI model fit the data well, $G^2(3) = 1.45, p = .693$. There was a significant difference across conditions on the C parameter, $\Delta G^2(2) = 15.52, p < .001$, the N parameter, $\Delta G^2(2) = 30.68, p < .001$, and the I parameter, $\Delta G^2(2) = 11.11, p = .004$.

Further analyses with the C parameter revealed that sensitivity to consequences was significantly lower in the moral-exemplar condition compared to the average-person condition, $\Delta G^2(1) = 15.03, p < .001, d$

$= 0.58$, and the average-exemplar condition, $\Delta G^2(1) = 6.17, p = .013, d = 0.38$. Sensitivity to consequences did not significantly differ between the average-person condition and the average-exemplar condition, $\Delta G^2(1) = 1.60, p = .206, d = 0.19$. These findings replicate and extend the findings of Study 1, suggesting that moral exemplars are perceived to be less sensitive to consequences than the average person and average exemplars.

Analyses with the N parameter revealed that sensitivity to moral norms was significantly higher in the moral-exemplar condition compared to the average-person condition, $\Delta G^2(1) = 6.45, p = .011, d = 0.38$, and the average-exemplar condition, $\Delta G^2(1) = 30.67, p < .001, d = 0.85$. In addition, there was a significant difference between the average-person and average-exemplar conditions, $\Delta G^2(1) = 8.92, p = .003, d = 0.45$, indicating that sensitivity to moral norms was higher in the average-person condition than in the average-exemplar condition.⁶ Replicating and extending the findings of Study 1, these results suggest that moral exemplars were perceived to be more sensitive to moral norms than the average person and average exemplars.

Finally, analyses with the I parameter revealed that general preference for inaction in the moral-exemplar condition was significantly lower compared to the average-exemplar condition, $\Delta G^2(1) = 11.10, p < .001, d = 0.51$, and marginally lower compared to the average-person condition, $\Delta G^2(1) = 3.64, p = .057, d = 0.29$. General preference for inaction did not significantly differ between the average-person and average-exemplar conditions, $\Delta G^2(1) = 1.92, p = .166, d = 0.21$. Different from Study 1, these results suggest that moral exemplars are perceived to be less action averse than average figures.

6.3. Discussion

Using the CNI model to disentangle three distinct factors contributing to moral-dilemma judgments, Study 2 revealed differences between moral exemplars and average figures with respect to all three factors. Replicating the findings of Study 1, moral exemplars were perceived to be less sensitive to consequences and more sensitive to moral norms than the average person. Extending the findings of Study 1, moral exemplars were also perceived to be less sensitive to consequences and more sensitive to moral norms than average exemplars, providing evidence for the robustness of these perceived differences after controlling for target abstraction and social prominence.

Unexpectedly, moral exemplars were also perceived to show a weaker general preference for inaction in comparison to average exemplars and a marginally weaker preference for inaction in comparison to the average person. Counter to the idea that moral exemplars might be perceived as more action averse due to concerns about harmful effects of their actions (see Baron & Goodwin, 2020; Cushman et al., 2006), this result suggests that moral exemplars are perceived as less (not more) action averse than others when resolving moral dilemmas. However, because the difference between the moral-exemplar and average-person conditions was only marginal and not obtained in Study 1, we exercise some caution in interpreting this finding.

Though Study 2 addressed several limitations of Study 1, one remaining ambiguity is the possibility that moral and average exemplars

⁵ Levene's test indicated a violation of homogeneity of error variances across conditions, $F(2, 262) = 15.94, p < .001$. However, analyses were consistent when using Welch's ANOVA and Games-Howell post-hoc tests.

⁶ It is worth noting that average exemplars were perceived to be more moral than the average person in the manipulation check, but less sensitive to moral norms in the main analyses, which seems inconsistent with a potential link between perceived morality and sensitivity to moral norms. Although comparing the presumed choices of figures other than moral exemplars was not the purpose of the current research, we further investigated the link between perceived morality (across figures) and sensitivity to moral norms by examining the correlation between the manipulation check assessing perceived morality and the N parameter across studies (see Integrative Data Analysis). In contrast to the discrepant finding noted here, results of this integrated analysis support a positive association between perceived morality and sensitivity to moral norms.

differ with respect to their perceived influence in society. That is, although the nominated moral and average exemplars were constrained to be prominent and well-known in society, moral exemplars may be perceived as having a greater societal impact than average exemplars. In this case, differences between moral and average exemplars might be driven by differences in perceived influence rather than differences in perceived morality. To address this concern, Study 3 substituted the average-person condition with an influential-exemplar condition.

7. Study 3

The purpose of Study 3 was twofold. The first purpose was to replicate the perceived differences between moral exemplars and average exemplars found in Study 2. The second purpose was to address the social-influence confound in Studies 1 and 2. Toward this end, Study 3 used the basic design of Study 2, the only difference being that the average-person condition was substituted with an influential-exemplar condition. Participants in the influential-exemplar condition were asked to nominate a public figure who they consider to be extremely influential in society and then indicate how this person would respond in the same moral dilemmas.

7.1. Method

7.1.1. Participants

Participants were recruited in October 2018 using Amazon's MTurk. Participation eligibility criteria and data quality precautions were the same as in Studies 1 and 2. Of the 302 participants who completed the survey in its entirety,⁷ 36 participants failed the attention check, leaving a final sample of 266 participants (53.38% female, 45.86% male, 0.38% prefer not to answer; $M_{age} = 36.22$, $SD_{age} = 11.12$). Of these participants, 82.33% identified as Caucasian, 10.90% as African American, 4.51% as Asian, 1.88% as Native American, and 3.01% as other ethnicities. The final sample of 266 participants provided 80% power in detecting a difference of $f = 0.19$ across the three experimental groups. Participants were compensated \$3.00 for their time.

7.1.2. Procedure and materials

The procedure and materials were identical to those used in Study 2 with two exceptions. First, the average-person condition was substituted with an influential-exemplar condition. Similar to the moral-exemplar and average-exemplar conditions, participants were asked to nominate a public figure and then complete the same moral dilemma battery by making judgments about whether their nominated exemplar would perform the action described in the scenario. Yet, different from the moral-exemplar and average-exemplar conditions, participants in the influential-exemplar condition were asked to nominate a figure who is highly influential and prominent in society (see Appendix). Second, to investigate differences in perceived influence across the three conditions, an additional one-item manipulation check was included (*How influential do you consider [figure]?*). Perceived influence was assessed with a 7-point rating scale ranging from 1 (*not at all influential*) to 7 (*extremely influential*).

7.2. Results

Moral judgment data were aggregated in line with the procedures in Studies 1 and 2. Means and 95% confidence intervals of the aggregated moral judgment data by condition can be seen in Table 1. CNI parameter estimates and 95% confidence intervals by condition can be seen in Fig. 4.

⁷ Two participants completed the assessment but either did not submit a request for compensation or were denied compensation because they submitted an incorrect completion code.

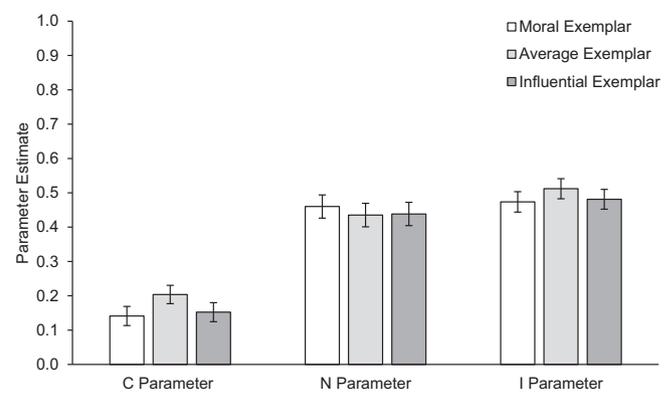


Fig. 4. Parameter estimates of sensitivity to consequences (C), sensitivity to moral norms (N), and general preference for inaction over action (I) as a function of figure type (moral exemplar vs. average exemplar vs. influential exemplar), Study 3. Error bars reflect 95% confidence intervals.

7.2.1. Manipulation check

To determine whether the manipulation of perceived morality was successful, morality ratings were submitted to a one-way ANOVA with three levels (average exemplar, moral exemplar, influential exemplar). Perceived morality significantly differed across the three conditions, $F(2, 263) = 21.05$, $p < .001$, $\eta^2 = 0.14$. Moral exemplars ($M = 6.27$, $SD = 0.90$) were perceived as significantly more moral compared to influential exemplars ($M = 5.80$, $SD = 1.28$), Tukey's HSD = 0.47, $p = .017$, $d = 0.42$, and average exemplars ($M = 5.20$, $SD = 1.09$), Tukey's HSD = 1.07, $p < .001$, $d = 1.06$. In addition, influential exemplars were perceived as significantly more moral than average exemplars, Tukey's HSD = 0.60, $p < .001$, $d = 0.51$. These results indicate that the manipulation of perceived morality was successful, in that moral exemplars were perceived as significantly more moral than either average exemplars or influential exemplars.

Submitted to the same ANOVA, ratings of perceived influence significantly differed across the three conditions, $F(2, 263) = 28.84$, $p < .001$, $\eta^2 = 0.18$.⁸ Average exemplars ($M = 5.50$, $SD = 1.16$) were perceived as significantly less influential than moral exemplars ($M = 6.32$, $SD = 0.90$), Tukey's HSD = -0.82 , $p < .001$, $d = 0.78$, and influential exemplars ($M = 6.52$, $SD = 0.74$), Tukey's HSD = -1.02 , $p < .001$, $d = 1.03$. Moral and influential exemplars were not perceived to differ in influence, Tukey's HSD = -0.20 , $p = .361$, $d = 0.24$. These results support concerns that perceived morality might have been confounded with perceived influence in Studies 1 and 2, raising questions about whether the obtained results are driven by differences in perceived morality or differences in perceived influence.

7.2.2. CNI analysis

The CNI model fit the data well, $G^2(3) = 2.26$, $p = .521$. There was a significant difference across conditions on the C parameter, $\Delta G^2(2) = 11.60$, $p = .003$, but no significant differences across conditions on either the N parameter, $\Delta G^2(2) = 1.22$, $p = .544$, or the I parameter, $\Delta G^2(2) = 3.67$, $p = .159$.

Further analyses with the C parameter revealed that sensitivity to consequences was lower in the moral-exemplar condition compared to the average-exemplar condition, $\Delta G^2(1) = 10.09$, $p = .001$, $d = 0.48$, but not the influential-exemplar conditions, $\Delta G^2(1) = 0.31$, $p = .580$, $d = 0.08$. Moreover, sensitivity to consequences was significantly lower in the influential-exemplar condition compared to the average-exemplar condition, $\Delta G^2(1) = 6.85$, $p = .009$, $d = 0.39$. Together, these results

⁸ Levene's test indicated a violation of homogeneity of error variances across conditions, $F(2, 263) = 13.43$, $p < .001$. However, analyses were consistent when using Welch's ANOVA and Games-Howell post-hoc tests.

suggest that both moral and influential exemplars are perceived to be less sensitive to consequences than average exemplars.

7.3. Discussion

Study 3 revealed differences between figures with respect to sensitivity to consequences, but not with respect to sensitivity to moral norms and general preference for inaction versus action. Replicating the findings of Study 2, moral exemplars were perceived to be less sensitive to consequences than average exemplars. However, this difference seems to be driven by differences in perceived influence rather than perceived morality, given that (1) moral exemplars were perceived to be more moral compared to both average and influential exemplars, (2) both moral and influential exemplars were perceived to be more influential compared to average exemplars, (3) both moral and influential exemplars were perceived to be less sensitive to consequences than average exemplars, and (4) moral and influential exemplars were not perceived to differ in their sensitivity to consequences. In contrast to the findings of Study 2, moral exemplars were not perceived to be more sensitive to moral norms or to be less action averse. Together, these findings cast doubt on the relations between perceived morality and the three factors of moral dilemma judgments. As a final test of these relations, we conducted Study 4 as a preregistered direct replication of Study 3.

8. Study 4

The purpose of Study 4 was twofold. The first purpose was to clarify the mixed findings of Studies 2 and 3 regarding differences in perceptions of moral exemplars and average exemplars. Based on the findings from Study 2, we expected that moral exemplars will be perceived as less sensitive to consequences (preregistered Hypothesis 1) and more sensitive to moral norms (preregistered Hypothesis 2) than average exemplars (confirmatory analyses). The second purpose was to investigate whether the predicted differences are driven by perceived morality or perceived influence (exploratory analyses). To the extent that (1) moral exemplars are perceived to be more moral than both average and influential exemplars, (2) both moral and influential exemplars are perceived to be more influential than average exemplars, (3) both moral and influential exemplars are perceived to differ in their moral judgments compared to average exemplars, and (4) moral and influential exemplars are not perceived to differ from each other in their moral judgments, perceived differences in the moral judgments of moral and average exemplars in the confirmatory tests would have to be attributed to perceived influence rather than perceived morality.

8.1. Method

Participants were recruited in November 2018 using Amazon's MTurk. Participation eligibility criteria and data quality precautions were the same as in Studies 1–3. Of the 304 participants who completed the survey in its entirety,⁹ 40 participants failed the attention check, leaving a final sample of 264 participants (44.70% female, 54.92% male, 0.38% other; $M_{\text{age}} = 36.00$, $SD_{\text{age}} = 11.28$). Of these participants, 78.79% identified as Caucasian, 12.50% as African American, 6.82% as Asian, 2.27% as Native American, and 1.14% as other ethnicities. The final sample of 264 participants provided 80% power in detecting a difference of $f = 0.19$ across the three experimental groups. Participants were compensated \$3.00 for their time. The procedure, materials, and analysis plans were identical to Study 3 and preregistered at <https://osf.io/pb724/>.

⁹ Four participants completed the assessment but either did not submit a request for compensation or were denied compensation because they submitted an incorrect completion code.

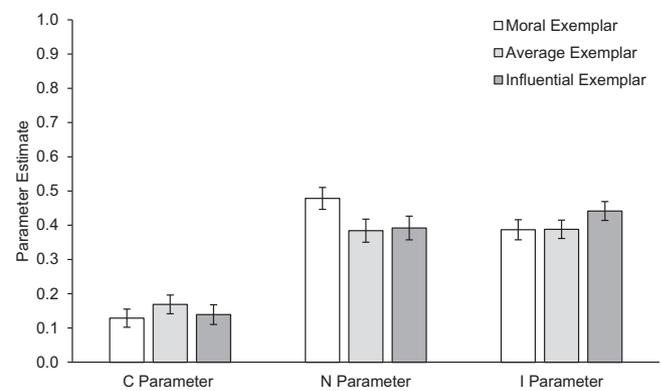


Fig. 5. Parameter estimates of sensitivity to consequences (C), sensitivity to moral norms (N), and general preference for inaction over action (I) as a function of figure type (moral exemplar vs. average exemplar vs. influential exemplar), Study 4. Error bars reflect 95% confidence intervals.

8.2. Results

Moral judgment data were aggregated in line with the procedures in Studies 1–3. Means and 95% confidence intervals of the aggregated moral judgment data by condition can be seen in Table 1. CNI parameter estimates and 95% confidence intervals by condition can be seen in Fig. 5.

8.2.1. Manipulation check

Submitted to a one-way ANOVA with three levels (average exemplar, moral exemplar, influential exemplar), perceived morality significantly differed across the three conditions, $F(2, 261) = 14.59$, $p < .001$, $\eta^2 = 0.10$.¹⁰ Supporting the effectiveness of the morality manipulation, moral exemplars ($M = 6.41$, $SD = 0.82$) were perceived as significantly more moral compared to average exemplars ($M = 5.55$, $SD = 1.08$), Tukey's HSD = 0.86, $p < .001$, $d = 0.90$, and influential exemplars ($M = 5.81$, $SD = 1.36$), Tukey's HSD = 0.60, $p = .001$, $d = 0.54$. Influential and average exemplars were not perceived to differ in morality, Tukey's HSD = 0.26, $p = .272$, $d = 0.21$. These results indicate that the manipulation of perceived morality was successful, in that moral exemplars were perceived as significantly more moral than either average exemplars or influential exemplars.

Perceived influence also significantly differed across conditions, $F(2, 261) = 14.61$, $p < .001$, $\eta^2 = 0.10$.¹¹ Average exemplars ($M = 5.76$, $SD = 1.03$) were perceived as significantly less influential compared to moral exemplars ($M = 6.34$, $SD = 0.80$), Tukey's HSD = -0.59 , $p < .001$, $d = 0.64$, and influential exemplars ($M = 6.42$, $SD = 0.84$), Tukey's HSD = -0.66 , $p < .001$, $d = 0.70$. Moral and influential exemplars were not perceived to differ in influence, Tukey's HSD = -0.08 , $p = .838$, $d = 0.09$. These results corroborate concerns that perceived morality is confounded with perceived influence, raising questions about whether differences in the presumed judgments of moral and average exemplars are driven by differences in perceived morality or differences in perceived influence.

8.2.2. CNI analysis

The CNI model fit the data well, $G^2(3) = 0.87$, $p = .833$. The C parameter did not significantly differ across conditions, $\Delta G^2(2) = 4.54$,

¹⁰ Levene's test indicated a violation of homogeneity of error variances across conditions, $F(2, 261) = 10.70$, $p < .001$. However, analyses were consistent when using Welch's ANOVA and Games-Howell post-hoc tests.

¹¹ Levene's test indicated a violation of homogeneity of error variances across conditions, $F(2, 261) = 5.99$, $p = .003$. However, analyses were consistent when using Welch's ANOVA and Games-Howell post-hoc tests.

$p = .103$. In contrast, both the N parameter, $\Delta G^2(2) = 19.51, p < .001$, and the I parameter, $\Delta G^2(2) = 9.73, p = .008$, significantly differed across conditions.

Further analyses with the N parameter revealed that sensitivity to moral norms was higher in the moral-exemplar condition compared to the average-exemplar condition, $\Delta G^2(1) = 15.80, p < .001, d = 0.59$, and the influential-exemplar condition, $\Delta G^2(1) = 12.91, p < .001, d = 0.55$. There was no significant difference between the average-exemplar and influential-exemplar conditions, $\Delta G^2(1) = 0.10, p = .750, d = 0.05$. These results suggest that moral exemplars were perceived to be more sensitive to moral norms than either average or influential exemplars.

Further analyses with the I parameter revealed that general preference for inaction versus action was higher in the influential-exemplar condition compared to the average-exemplar condition, $\Delta G^2(1) = 7.42, p = .006, d = 0.42$, and the moral-exemplar condition, $\Delta G^2(1) = 7.12, p = .008, d = 0.41$. There was no significant difference between the average-exemplar and moral-exemplar conditions, $\Delta G^2(1) = 0.004, p = .947, d = 0.01$. These results suggest that influential exemplars were perceived to be more action averse than both average exemplars and moral exemplars.

8.3. Discussion

The results of Study 4 suggest that moral exemplars are perceived to differ from others in their sensitivity to moral norms. Confirming Hypothesis 2, moral exemplars were perceived to be more sensitive to moral norms than average exemplars. In conjunction with the findings of Studies 1 and 2, these results suggest that the null effect in Study 3 might have been a false negative. Exploratory analyses further suggest that this difference is indeed driven by differences in perceived morality rather than perceived influence. Although both moral and influential exemplars were perceived to be more influential than average exemplars, moral exemplars were perceived to be more sensitive to moral norms than either average or influential exemplars.

Counter to Hypothesis 1, there was no significant difference between figures with respect to sensitivity to consequences. These results stand in contrast to the findings of Studies 2 and 3, suggesting that moral exemplars are perceived to be less sensitive to consequences than average exemplars. Yet, even if the null effect in Study 4 is a false negative, results from Study 3 suggest that the obtained difference in our previous studies is driven by perceived influence rather than perceived morality. From this perspective, the primary difference between moral exemplars and other individuals seems to be that moral exemplars are perceived to be more sensitive to moral norms than others.

Finally, both moral exemplars and average exemplars were perceived to be less action averse than influential exemplars. However, because moral exemplars were not perceived to differ in their levels of action aversion in comparison to average exemplars, this finding is irrelevant for the main question of the current research and therefore not discussed any further.

9. Integrative data analysis

To gain further confidence in our conclusion that perceived morality is linked to sensitivity to moral norms, but not to sensitivity to consequences and general action tendencies, we conducted four sets of integrative data analyses (IDA) using the data from all four studies (see Curran & Hussong, 2009). In a first set of analyses, we followed the analytic approach used in the four individual studies by aggregating moral-dilemma responses across participants within each figure condition and fitting the CNI model to moral judgment responses at the group-level (see Gawronski et al., 2017). The results of these analyses are summarized in Table 2. With respect to the C parameter, moral exemplars were perceived to be less sensitive to consequences than the average person and average exemplars, but not in comparison to influential exemplars. This pattern of results is consistent with our conclusion

that the weaker sensitivity to consequences in the presumed choices of moral exemplars is driven by perceived influence rather than perceived morality. With respect to the N parameter, moral exemplars were perceived to be more sensitive to moral norms than the average person, average exemplars, and influential exemplars, supporting the proposed link between perceived morality and sensitivity to moral norms. With respect to the I parameter, moral exemplars were perceived to be less action averse than average exemplars and influential exemplars, but not in comparison to the average person. These results might suggest a negative association between perceived morality and action aversion. However, given the non-significant difference between the moral-exemplar condition and the average-person condition, we exercise caution in drawing strong conclusions from these results.

In a second set of analyses, we fit the CNI model to responses at the individual-level (rather than the group-level) by aggregating moral-dilemma responses for each individual participant (rather than across participants within each condition), resulting in unique CNI parameter estimates for each individual participant (see Körner et al., 2020).¹² Using this approach, we conducted independent samples t -tests to examine differences in individual-level parameter estimates between conditions. The results of these analyses are summarized in Table 3. With respect to the C parameter, moral exemplars were perceived to be less sensitive to consequences than the average person and average exemplars, but not in comparison to influential exemplars. These findings corroborate our conclusion that the weaker sensitivity to consequences in the presumed responses of moral exemplars is driven by perceived influence rather than perceived morality. With respect to the N parameter, moral exemplars were perceived to be more sensitive to moral norms than average exemplars. However, in contrast to the group-level IDA, moral exemplars were not perceived to differ in their sensitivity to moral norms in comparison to either the average person or influential exemplars. Therefore, while the group-level IDA provided clear support for a link between perceived morality and sensitivity to moral norms, the individual-level IDA raises questions about the robustness and the conceptual meaning of this link. Regarding perceived action aversion, presumed responses of moral exemplars were not significantly different from any comparison figure.

To provide clarity regarding the link between perceived morality and sensitivity to moral norms, we conducted a third set of analyses directly examining the association between perceived morality and presumed choices in moral dilemmas. To this end, we analyzed correlations between responses to our manipulation check assessing perceived figure morality and CNI parameters estimated at the individual-level. To identify unique links with perceived morality independent of perceived influence, we first analyzed zero-order correlations between perceived morality and CNI parameters (Studies 2–4), followed by partial correlations between perceived morality and CNI parameters controlling for perceived influence (Studies 3–4). With respect to the C parameter, perceived morality showed a significant negative zero-order correlation with sensitivity to consequences, $r(793) = -0.14, p < .001$, but this association was not statistically significant when controlling for perceived influence, $r(527) = -0.08, p = .080$, again suggesting that the relation between perceived morality and sensitivity to consequences is driven by perceived influence rather than perceived morality. With respect to the N parameter, perceived morality showed a significant positive zero-order correlation with sensitivity to moral norms, $r(793) = 0.22, p < .001$, and this association remained statistically significant when controlling for perceived influence, $r(527) = 0.24, p < .001$, providing further support for the postulated link between perceived

¹² An advantage of analyses using individual-level estimates is that they can account for heterogeneity in responses across participants. However, a notable downside of this approach is that moral-judgment parameters are estimated using a smaller number of responses, which leads to greater measurement error in the resulting estimates (see Luke & Gawronski, in press).

Table 2

Integrative data analysis across Studies 1–4 examining differences in CNI model parameters as a function of figure type using group-level approach.

Comparison	C Parameter	Difference	N Parameter	Difference	I Parameter	Difference
Average Person vs. Moral Exemplar ^a Average Person (n = 190)	0.19 [0.17, 0.21]	$\Delta G^2(1) = 31.71$, $p < .001$, $d = 0.59$	0.39 [0.36, 0.41]	$\Delta G^2(1) = 16.36$, $p < .001$, $d = 0.42$	0.45 [0.43, 0.47]	$\Delta G^2(1) = 0.19$, $p = .661$, $d = 0.05$
Moral Exemplar (n = 174)	0.11 [0.09, 0.13]		0.45 [0.43, 0.48]		0.44 [0.42, 0.46]	
Average Exemplar vs. Moral Exemplar ^b Average Exemplar (n = 271)	0.18 [0.17, 0.20]	$\Delta G^2(1) = 20.66$, $p < .001$, $d = 0.40$	0.39 [0.37, 0.41]	$\Delta G^2(1) = 35.85$, $p < .001$, $d = 0.52$	0.45 [0.44, 0.47]	$\Delta G^2(1) = 8.46$, $p = .004$, $d = 0.25$
Moral Exemplar (n = 261)	0.13 [0.11, 0.15]		0.47 [0.45, 0.49]		0.42 [0.40, 0.44]	
Influential Exemplar vs. Moral Exemplar ^c Influential Exemplar (n = 170)	0.15 [0.13, 0.17]	$\Delta G^2(1) = 0.62$, $p = .431$, $d = 0.09$	0.42 [0.39, 0.44]	$\Delta G^2(1) = 9.90$, $p = .002$, $d = 0.34$	0.46 [0.44, 0.48]	$\Delta G^2(1) = 4.55$, $p = .033$, $d = 0.23$
Moral Exemplar (n = 175)	0.13 [0.12, 0.15]		0.47 [0.45, 0.49]		0.43 [0.41, 0.45]	

C = sensitivity to consequences; N = sensitivity to moral norms; I = general preference for inaction over action. Numbers in brackets depict 95% confidence intervals. ^a Studies 1–2. ^b Studies 2–4. ^c Studies 3–4.

Table 3

Integrative data analysis across Studies 1–4 examining differences in CNI model parameters as a function of figure type using individual-level approach.

Comparison	C Parameter	Difference	N Parameter	Difference	I Parameter	Difference
Average Person vs. Moral Exemplar ^a Average Person (n = 190)	0.20 [0.17, 0.23]	$t(351.05) = 4.51$ $p < .001$, $d = 0.47$	0.44 [0.39, 0.48]	$t(362) = -1.53$ $p = .127$, $d = 0.16$	0.54 [0.49, 0.58]	$t(362) = -0.01$ $p = .993$, $d = 0.00$
Moral Exemplar (n = 174)	0.12 [0.10, 0.14]		0.49 [0.44, 0.54]		0.54 [0.49, 0.58]	
Average Exemplar vs. Moral Exemplar ^b Average Exemplar (n = 271)	0.19 [0.17, 0.22]	$t(516.09) = 3.57$ $p < .001$, $d = 0.31$	0.44 [0.40, 0.48]	$t(530) = -2.07$ $p = .039$, $d = 0.18$	0.55 [0.51, 0.58]	$t(530) = 1.85$ $p = .065$, $d = 0.16$
Moral Exemplar (n = 261)	0.14 [0.12, 0.16]		0.50 [0.46, 0.54]		0.50 [0.46, 0.54]	
Influential Exemplar vs. Moral Exemplar ^c Influential Exemplar (n = 170)	0.15 [0.13, 0.18]	$t(343) = 0.46$ $p = .645$, $d = 0.05$	0.47 [0.42, 0.53]	$t(343) = -0.67$ $p = .503$, $d = 0.07$	0.54 [0.49, 0.58]	$t(343) = 0.94$ $p = .348$, $d = 0.10$
Moral Exemplar (n = 175)	0.15 [0.12, 0.17]		0.50 [0.45, 0.55]		0.50 [0.46, 0.55]	

C = sensitivity to consequences; N = sensitivity to moral norms; I = general preference for inaction over action. Numbers in brackets depict 95% confidence intervals. ^a Studies 1–2. ^b Studies 2–4. ^c Studies 3–4.

morality and sensitivity to moral norms. Finally, with respect to the I parameter, perceived morality showed a significant negative zero-order correlation with action aversion, $r(793) = -0.07, p = .049$, but this association was not statistically significant when controlling for perceived influence, $r(527) = -0.01, p = .746$, casting further doubts about a potential link between perceived morality and general action tendencies.

Finally, in a fourth set of analyses, we aimed to integrate the links between exemplar status in our experimental manipulation, measured perceptions of morality and influence, and the three CNI parameters. Toward this end, we conducted two sets of multiple regression analyses with the combined data of the studies that included measures of both perceived morality and perceived influence (Studies 3 and 4). In a first set of multiple regression analyses, we regressed perceived morality onto dummy-coded exemplar status of being morally exceptional vs. “other” (i.e., morally average or socially influential) and dummy-coded exemplar status of being socially influential vs. “other” (i.e., morally exceptional or morally average). Correspondingly, we regressed perceived influence onto dummy-coded exemplar status of being morally exceptional vs. “other” (i.e., morally average or socially influential) and dummy-coded exemplar status of being socially influential vs. “other” (i.e., morally exceptional or morally average). In a second set of multiple regression analyses, we regressed each of the three CNI parameters onto perceived morality and perceived influence. The combined results of the multiple regression analyses are presented in Fig. 6. The results corroborate concerns about a potential confound between perceived morality and perceived influence in the exemplar-nomination manipulation, in that (1) the nominated moral exemplars were perceived to be more moral and more influential compared to the

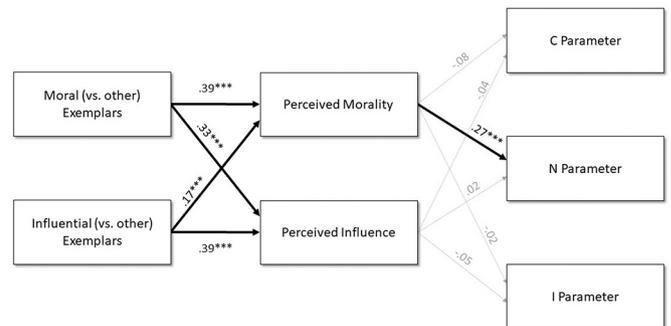


Fig. 6. Results of multiple-mediator analyses on the relations between exemplar type, perceived exemplar characteristics, and moral-judgment parameters. Combined data from Studies 3 and 4.

Note. * $p < .05$, ** $p < .01$, *** $p < .001$; C parameter = sensitivity to consequences; N parameter = sensitivity to moral norms; I parameter = general preference for inaction over action. Bold black arrows depict statistically significant relations; dotted gray arrows depict statistically non-significant relations.

nominated other exemplars and (2) the nominated influential exemplars were perceived to be more influential and more moral compared to the nominated other exemplars. Yet, despite the lack of unique links between exemplar-status in the nomination manipulation and subjectively perceived exemplar-characteristics, perceived morality was uniquely related to the N parameter, in that exemplars who were rated higher in morality were perceived to be more sensitive to moral norms than exemplars who were rated lower in morality. Perceived morality was not

significantly related to the *C* or *I* parameters. Perceived influence was not significantly related to any of the three CNI parameters. Together, these results suggest that, although our exemplar-nomination approach is prone to confounds in the characteristics of the nominated exemplars, perceived morality is uniquely linked to sensitivity to moral norms, in that individuals perceived to be higher in morality are presumed to adhere more to moral norms than those perceived to be lower in morality.

Expanding on the results of the multiple regression analyses, we also explored the particular exemplars that were nominated in the different figure conditions. A closer inspection of the nominated exemplars further clarifies why the links between exemplar-status and perceived exemplar-characteristics show overlap between non-matching dimensions despite the unique link between perceived morality and sensitivity to moral norms. Table 4 provides an overview of the 10 most frequently nominated exemplars in each category across studies. Although some of the nominated moral exemplars were unique in the sense that they were not nominated in any of the other conditions (e.g., Gandhi, Jesus, Mother Teresa, Billy Graham), a notable aspect of the three lists is that there is considerable overlap between the exemplars that have been nominated as morally exceptional, morally average, and socially influential. Indeed, Barack Obama turned out to be the most frequently nominated exemplar in all three categories. For the comparison between morally exceptional and socially influential exemplars, the overlap between the nominated exemplars explains the absence of unique links between exemplar status and perceived exemplar characteristics, in that moral exemplars were perceived to be more influential than non-moral exemplars (in addition to being perceived as more moral) and influential exemplars were perceived to be more moral than non-influential exemplars (in addition to being perceived as more influential). Moreover, for the comparison between morally exceptional and morally average exemplars, the overlap between the nominated exemplars illustrates the inherent subjectivity of moral impressions, in that the same person may be perceived as morally exceptional by some participants and morally average by others. Yet, regardless of such disagreements about the morality of particular exemplars, there seems to be considerable agreement about the link between a person's morality and their presumed choices in moral dilemmas, in that those who are perceived to be more moral are presumed to be more sensitive to moral norms.

10. General discussion

Past research on moral-dilemma judgment has focused predominantly on the processes underlying outcome-maximizing and norm-conforming judgments (Crockett, 2013; Cushman, 2013; Greene, 2008; Holyoak & Powell, 2016). A newly emerging question in moral-dilemma research is how people form moral impressions of individuals who make either outcome-maximizing or norm-conforming judgments

Table 4
Most frequently nominated exemplars across studies as a function of nomination prompt (morally exceptional vs. morally average vs. socially influential). Exemplars are listed in rank-order of most frequent nominations within each category.

	Morally Exceptional (Studies 1–4)	Morally Average (Studies 2–4)	Socially Influential (Studies 3–4)
#1	Barack Obama	Barack Obama	Barack Obama
#2	Abraham Lincoln	Donald Trump	Donald Trump
#3	Donald Trump	George Bush	Abraham Lincoln
#4	Gandhi	Joe Biden	Martin Luther King
#5	Martin Luther King	Oprah Winfrey	Albert Einstein
#6	Oprah Winfrey	Abraham Lincoln	Elon Musk
#7	Billy Graham	George Washington	George Washington
#8	Bernie Sanders	Chris Pratt	Hillary Clinton
#9	Jesus	Elon Musk	Michelle Obama
#10	Mother Teresa	Hillary Clinton	Bill Gates

in moral dilemmas. A central finding of this work is that individuals who make norm-conforming, deontological judgments are perceived as having a stronger moral character than those who make outcome-maximizing, utilitarian judgments (for a review, see Crockett et al., 2021). The current research aimed to provide deeper insights into the link between moral impressions and moral choices by examining whether morally exceptional figures are perceived to differ from others in their (1) sensitivity to consequences, (2) sensitivity to moral norms, or (3) general action tendencies when resolving moral dilemmas.

Collectively, our findings suggest a positive association between perceived morality and sensitivity to moral norms. Although the relevant effects did not reach statistical significance in one of the four studies (Study 3), moral exemplars were perceived to be more sensitive to moral norms compared to the average person (Studies 1 and 2), average exemplars (Studies 2 and 4), and influential exemplars (Study 4). This conclusion was further supported by the results of our IDA using the same group-level approach, with moral exemplars being perceived to have a stronger sensitivity to moral norms than all three comparison groups. Moreover, while these differences only partially replicated when using an individual-level approach, perceptions of morality as assessed by our manipulation check showed a significant positive association with sensitivity to moral norms, and this association remained statistically significant after controlling for perceived influence. Taken together, these results suggest a link between perceptions of morality and adherence to moral norms and duties.

Although moral exemplars were perceived to be less sensitive to consequences than average figures in three of the four studies, this difference seemed to be driven by the perceived influence of moral exemplars rather than their perceived morality. This conclusion is supported by the findings that (1) moral exemplars were perceived to be more moral compared to both average and influential exemplars, (2) both moral and influential exemplars were perceived to be more influential compared to average exemplars, (3) both moral and influential exemplars were perceived to be less sensitive to consequences than average exemplars, and (4) moral and influential exemplars were not perceived to differ in their sensitivity to consequences. This conclusion is further bolstered by the results of our IDA, which found no differences in sensitivity to consequences between moral and influential exemplars using either the group-level or individual-level approach and no relation between perceived morality and sensitivity to consequences after controlling for perceived influence. Thus, although moral exemplars were perceived to be more sensitive to consequences compared to average figures in three of the four studies, these differences do not seem to be driven by perceptions of morality.

The current studies did not obtain any evidence for the idea that moral exemplars would be perceived as more action averse compared to others. If anything, our findings suggest that moral exemplars are perceived as less action averse compared to average exemplars (Study 2) and influential exemplars (Study 4). With that said, moral exemplars were not perceived to differ from others in terms of their action aversion in two of the four studies (Studies 1 and 3). Moreover, results from our IDA revealed either mixed or no evidence for perceived differences in action aversion between moral exemplars and other figures. Taken together, the current findings provide no reliable evidence for a potential link between perceptions of morality and presumed action aversion.

10.1. Implications for social perception

Together, the current findings provide valuable insights into the relation between perceived morality and the central aspects of utilitarianism and deontology. Across several different lines of research, there is growing evidence for a systematic relation between perceived moral character and preference for deontological over utilitarian judgments (for a review, see Crockett et al., 2021). However, given the methodological limitations inherent in the traditional dilemma paradigm (Conway & Gawronski, 2013; Crone & Laham, 2017), the conceptual

meaning of this relation is ambiguous. One possibility is that the observed relation is rooted in a negative association between perceived morality and sensitivity to consequences (see [Kreps & Monin, 2014](#); [Sacco et al., 2017](#); [Uhlmann et al., 2013](#)). An alternative possibility is that the observed relation is rooted in a positive association between perceived morality and sensitivity to moral norms (see [Everett et al., 2016](#); [Rom et al., 2017](#)). Finally, a third possibility is that the observed relation is rooted in a positive association between perceived morality and general preference for inaction over action (see [Baron & Goodwin, 2020](#); [Cushman et al., 2006](#)). The findings of the current research support the second possibility, suggesting that perceived morality is positively associated with sensitivity to moral norms.

Based on the findings of prior research, there are two potential reasons for the link between perceived morality and sensitivity to moral norms. One possibility is that those adhering to structured moral norms and rules are viewed as predictable in their future behavior, fostering a greater sense of trust. Consistent with this view, those making deontological as opposed to utilitarian judgments on moral dilemmas are perceived as more predictable ([Turpin et al., 2021](#)) and elicit higher degrees of cooperation in economic games ([Bostyn & Roets, 2017](#); [Everett et al., 2016](#)). Another possibility is that adherence to moral norms signals strong empathic concern for the welfare of others, supporting perceptions of compassion and kindness. Consistent with this view, individuals high in psychopathy—who are known to lack empathic concern—show rather low sensitivity to moral norms in the resolution of moral dilemmas (e.g., [Luke et al., in press](#); [Luke & Gawronski, 2021b](#)), and decision-makers who prefer deontological over utilitarian judgments on moral dilemmas are perceived as more strongly engaging in affective processing and empathy ([Rom et al., 2017](#); [Uhlmann et al., 2013](#)). It is worth noting that these two possibilities are not mutually exclusive, in that perceived predictability may promote perceptions of morality via trait inferences of trustworthiness and perceived empathic concern may promote perceptions of morality via trait inferences of compassion. Consistent with this idea, trustworthiness and compassion are among the traits most closely linked to perceptions of morality ([Brambilla, Sacchi, Rusconi, & Goodwin, 2021](#); [Landy & Uhlmann, 2018](#)). Thus, adherence to moral norms may be related to perceived morality, because norm-congruent behavior drives inferences about multiple traits considered to be at the heart of morality.

10.2. Implications for moral judgment

Although the current research focused primarily on the link between moral-dilemma judgments and perceptions of morality, the obtained results also raise interesting new questions for research on the mechanisms underlying moral-dilemma judgments. A common assumption in the moral-dilemma literature is that utilitarian judgments are the product of controlled cognitive analyses of costs and benefits, whereas deontological judgments are rooted in automatic emotional responses to the idea of causing harm ([Greene et al., 2001](#); [Greene et al., 2004](#)). Although the validity of this account is the subject of ongoing debates, it deserves credit for advancing the idea that norm-congruent judgments may result from processes that do not involve a conscious consideration of moral norms (see also [Haidt, 2001](#)). The current findings suggest an alternative mechanism that may lead to norm-congruent judgments without conscious consideration of moral norms. When faced with a moral dilemma, people may think about individuals they deem morally exceptional, simulate how these individuals might respond, and use their presumed response to guide their own decision (see [Fleeson, 2019](#)). To the extent that the presumed responses of moral exemplars conform to moral norms, using these responses as guides can lead to norm-congruent decisions without conscious consideration of moral norms.

To the extent that the role of moral exemplars in moral-dilemma judgments can be empirically confirmed, the presumed mechanism raises a number of interesting follow-up questions. One important question is whether mental simulation of choices by moral exemplars is

a universal mechanism, or whether its use is limited to individuals with certain characteristics (e.g., individuals with strong religious beliefs). Relatedly, another important question concerns the factors that determine the choice of moral exemplars and the extent to which their central characteristics are similar or different across individuals. Expanding on these ideas, we have conducted a series of follow-up studies (one pre-registered) examining whether perceptions of how moral exemplars resolve moral dilemmas in comparison to others depend on person-related characteristics of the participants. A preliminary finding of these studies is that there seems to be high consensus about the link between perceived morality and sensitivity to moral norms, in that moral exemplars were perceived to be more sensitive to moral norms compared to average exemplars regardless of potentially relevant characteristics of the participants (e.g., religiosity, political attitudes). Consistent with the current findings, consensus regarding potential differences in sensitivity to consequences and general action tendencies was relatively low.

Directly related to this point, it is worth noting that the current research was concerned with central tendencies at the sample level, which does not conflict with the idea of individual differences in perceptions of morality (see [Funder, 2006](#)). For example, although our findings suggest that, on average, perceptions of morality are linked to stronger sensitivity to moral norms, it is possible that a subset of “utilitarian” participants perceives a link between morality and stronger sensitivity to consequences (see [Conway et al., 2018](#)). To the extent that this subset is relatively small, it may not produce a central-tendency effect at the sample level, but this does not mean that sensitivity to consequences is unrelated to perceptions of morality for everyone in the sample. That being said, it is worth noting that our findings at the sample level would suggest the opposite, in that moral exemplars were perceived to show weaker (rather than stronger) sensitivity to consequences in three of the four studies (although this link was driven by perceived influence rather than perceived morality). Moreover, our unpublished follow-up studies on person-related characteristics suggest that, while there seems to be considerable agreement about the link between perceived morality and sensitivity to moral norms, agreement regarding sensitivity to consequences and general action tendencies seems to be much lower, leaving room for systematic individual differences in perceptions of moral character. Future research may go beyond central-tendency effects at the sample level by investigating individual differences in perceptions of morality.

10.3. Potential objections

While the current research did not put forward a specific conception of morality to avoid inducing artificial relations, it is possible that people conceive of a variety of different types of moral exemplars. For example, [Walker and Hennig \(2004\)](#) examined different conceptions of moral exceptionalism in caring, just, and brave moral exemplars. While these different exemplars were perceived to share a common core of personality traits, they were perceived to be unique in terms of their personality profiles. Maintaining a broad conception of moral exceptionalism (as in the current studies) has the advantage of not constraining the definition of morality to emphasize any single trait, which could pose a problem if people do not view the specified trait as morally relevant. Nevertheless, there seems to be consensus about the moral status of some traits (e.g., trustworthiness, compassion; see [Brambilla et al., 2021](#); [Landy & Uhlmann, 2018](#)), and moral exemplars characterized by these specific traits may be perceived to differ in their responses to moral dilemmas. Future research might investigate this question further by examining whether moral exemplars with different moral traits (e.g., trustworthy, compassionate) are perceived to differ in their responses to moral dilemmas.

Another issue concerns recent criticisms regarding conceptual and methodological aspects of the CNI model ([Baron & Goodwin, 2020](#)). Although the majority of these criticisms have been refuted as being

based on mischaracterizations of the model and flawed statistical analyses (Gawronski et al., 2020), it seems appropriate to address three valid criticisms that are relevant to the findings of the current research. First, Baron & Goodwin, 2020 pointed out that a general response bias favoring inaction on the *I* parameter could be interpreted as an instance of deontological responding in the sense that it is consistent with the broad norm *first, do no harm*. Given that moral exemplars were not perceived to reliably differ from others with respect to the *I* parameter, it may therefore be argued that perceived morality is not related to every potential instance of norm-congruent responding. Based on these considerations, it seems important to clarify that the link between perceived morality and norm adherence obtained in the current research reflects an unconditional adherence to both proscriptive and prescriptive norms surrounding harm and care (see Gawronski et al., 2020). It does not reflect differences in the adherence to the broad norm *first, do no harm*, which is reflected in greater levels of action aversion on the *I* parameter.¹³

Second, Baron & Goodwin, 2020 expressed concerns that relations between parameter estimates and external variables might depend on the hierarchical position of the *C* and the *N* parameter in the processing tree (see Fig. 1). Although Baron and Goodwin's reanalyses of existing data to demonstrate this possibility included major flaws and a correct reanalyses of the same data did not reveal any meaningful difference as a function of model specifications (Gawronski et al., 2020), we reconducted the group-level IDA using a modified model in which the hierarchical position of the *C* and the *N* parameters was reversed (see Table S1 in Supplemental Online Materials). Consistent with the findings by Gawronski et al. (2020), the results obtained with the modified model were entirely consistent with those obtained using the original model.

A final criticism by Baron & Goodwin, 2020 is that the moral dilemmas for research using the CNI model may differ in terms of their validity in capturing the central manipulations and consequences and moral norms. In response to this concern, Gawronski et al. (2020) tested whether consequences and moral norms were validly manipulated across the variants of each basic dilemma in the original battery. Overall, the manipulations of consequences and moral norms were well captured across moral dilemmas with the exception of one dilemma (*abduction dilemma*). Because this dilemma was included in all of the current analyses in accordance with the analytic plan we generated before data collection, we reconducted the group-level IDA excluding responses to the one dilemma with questionable validity (see Table S2 in the Supplemental Online Materials). Consistent with the findings of corresponding analyses by Gawronski et al. (2020), our results remained unchanged in terms of statistical significance and the interpretation of our main findings were unaffected after excluding the problematic dilemma.

11. Conclusion

Drawing on the prominent philosophical traditions of utilitarianism

¹³ Within the CNI model, the difference between proscriptive and prescriptive norms is conceptualized via pairs of morally relevant actions and inactions that have the same overall outcome (e.g., killing Person A and letting Person A die both result in the loss of Person A's life). Whereas the identified action within each action-inaction pair is conceptually linked to a proscriptive norm (i.e., killing someone is morally prohibited), the opposite of the identified inaction is conceptually linked to a prescriptive norm (i.e., saving someone's life is morally prescribed). Scores on the CNI model's *N* parameter reflect the extent to which participants adhere to the proscriptive and prescriptive norms identified in this manner. Scores on the *I* parameter reflect the extent to which participants show a general preference for inaction versus action irrespective of these pairs of norms and the described consequence. For a more detailed discussion of the operationalization of moral norms within the CNI model, see Gawronski et al. (2020).

and deontology, moral psychology has provided valuable insights into how people make judgments about right and wrong. The primary goal of the current research was to examine how perceived morality is related to the central aspects of utilitarianism and deontology by examining whether morally exceptional figures are perceived to differ from others in their sensitivity to consequences, sensitivity to moral norms, or general action tendencies when resolving moral dilemmas. Using the CNI model to disentangle the three factors underlying moral-dilemma judgments, we found evidence for a positive association between perceived morality and presumed sensitivity to moral norms. For sensitivity to consequences and general action tendencies, findings were mixed and attributable to characteristics confounded with perceived morality.

Open practices

The data, analysis codes, and materials for the current studies are available at <https://osf.io/k3f9u/>. The preregistration for Study 4 is available at <https://osf.io/pb724/>.

Author note

A very special thanks to Dillon Luke for his invaluable assistance and feedback regarding the current research.

Appendix A. Appendix: Nomination prompts

A.1. Moral exemplar prompt

We now ask that you please think of a public figure who you consider a highly moral person. This person should be a current or historical figure who is well known in society. The person should have a strong moral character and have shown an extraordinary commitment to morality in their life. The person should be one of the most moral people you can think of and should be far more moral than others. Please write this person's name below.

A.2. Average exemplar prompt

We now ask that you please think of a public figure who you consider a morally average person. This person should be a current or historical figure who is well known in society. The person should have average moral character and have shown an ordinary level of morality in their life. The person should be one of the most ordinary public figures you can think of and should be about as moral as anyone else. Please write this person's name below.

A.3. Influential exemplar prompt

We now ask that you please think of a public figure who you consider a highly influential person. This person should be a current or historical figure who is well known in society. The person should have a strong assertive character and have shown an extraordinary impact on world events in their life. The person should be one of the most influential people you can think of and should be far more influential than others. Please write this person's name below.

Appendix B. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jesp.2021.104265>.

References

- Baron, J., & Goodwin, G. P. (2020). Consequences, norms, and inaction: A critical analysis. *Judgment and Decision Making*, 15, 421–442.

- Bialek, 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., & Roets, A. (2017). Trust, trolleys, and social dilemmas: A replication study. *Journal of Experimental Psychology: General, 146*, e1–e7.
- Brambilla, M., Sacchi, S., Rusconi, P., & Goodwin, G. P. (2021). The primacy of morality in impression development: Theory, research, and future directions. *Advances in Experimental Social Psychology, 64*, 187–262.
- Brannon, S. M., Carr, S., Jin, E. S., Josephs, R. A., & Gawronski, B. (2019). Exogenous testosterone increases sensitivity to moral norms in moral dilemma judgments. *Nature Human Behavior, 3*, 856–866.
- Conway, P., & Gawronski, B. (2013). Deontological and utilitarian inclinations in moral decision making: A process dissociation approach. *Journal of Personality and Social Psychology, 104*, 216–235.
- Conway, P., Goldstein-Greenwood, J., Polacek, D., & Greene, J. D. (2018). Sacrificial utilitarian judgments do reflect concern for the greater good: Clarification via process dissociation and the judgments of philosophers. *Cognition, 179*, 241–265.
- Critcher, C. R., Helzer, E. G., & Tannenbaum, D. (2020). Moral character evaluation: Testing another's moral-cognitive machinery. *Journal of Experimental Social Psychology, 87*, Article 103906.
- Crockett, M. (2013). Models or morality. *Trends in Cognitive Sciences, 17*, 363–366.
- Crockett, M., Everett, J. A. C., Gill, M., & Siegel, J. Z. (2021). The relational logic of moral inference. *Advances in Experimental Social Psychology, 64*, 1–64.
- Crone, D. L., & Laham, S. M. (2017). Utilitarian preferences or action preferences? Deconfounding action and moral code in sacrificial dilemmas. *Personality and Individual Differences, 104*, 476–481.
- Curran, P. J., & Hussong, A. M. (2009). Integrative data analysis: The simultaneous analysis of multiple data sets. *Psychological Methods, 14*, 81–100.
- Cushman, F. (2013). Action, outcome, and value: A dual-system framework for morality. *Personality and Social Psychology Review, 17*, 273–292.
- Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. *Psychological Science, 17*, 1082–1089.
- Everett, J. A. C., Faber, N. S., Savulescu, J., & Crockett, M. J. (2018). The costs of being consequentialist: Social inference from instrumental harm and impartial beneficence. *Journal of Experimental Social Psychology, 79*, 200–216.
- Everett, J. A. C., Pizarro, D. A., & Crockett, M. J. (2016). Inference of trustworthiness from intuitive moral judgments. *Journal of Experimental Psychology: General, 145*, 772–787.
- Fleeson, W. (2019). The suitability of the morally exceptional as guides to moral development. *Keynote address at the Exemplar Interventions to Develop Character conference, Winston-Salem, NC.*
- Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review, 5*, 5–15.
- Funder, D. C. (2006). Towards a resolution of the personality triad: Persons, situations, and behaviors. *Journal of Research in Personality, 40*, 21–34.
- Furr, R. M., Prentice, M., Hawkins, A., & Fleeson, W. (2020). *What does it mean to be seen as moral? Lay perceptions of the personality characteristics of immoral, moral, and exceptionally moral people* [Manuscript in preparation].
- 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*, 343–376.
- Gawronski, B., & Brannon, S. M. (2020). Power and moral dilemma judgments: Distinct effects of memory recall versus social roles. *Journal of Experimental Social Psychology, 86*, Article 103908.
- Gawronski, B., Conway, P., Armstrong, J., Friesdorf, R., & Hütter, M. (2018). Effects of incidental emotions on moral dilemma judgments: An analysis using the CNI model. *Emotion, 18*, 989–1008.
- Gawronski, B., Conway, P., Hütter, M., Luke, D. M., Armstrong, J., & Friesdorf, R. (2020). On the validity of the CNI model of moral decision-making: Reply to Baron and Goodwin (2020). *Judgment and Decision Making, 15*, 1054–1072.
- Gleichgerrcht, E., & Young, L. (2013). Low levels of empathic concern predict utilitarian moral judgment. *PLoS One, 8*, Article e60418.
- Greene, J. (2008). The secret joke of Kant's soul. In W. Sinnott-Armstrong (Ed.), *Moral psychology: The neuroscience of morality* (pp. 35–79). Cambridge: MIT Press.
- Greene, J. D., Nystrom, L. E., Engell, A. D., Darley, J. M., & Cohen, J. D. (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron, 44*, 389–400.
- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science, 293*, 2105–2108.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review, 108*, 814–834.
- Hannikainen, I. R., Machery, E., & Cushman, F. A. (2018). Is utilitarian sacrifice becoming more morally permissible? *Cognition, 170*, 95–101.
- Hardy, S. A., Walker, L. J., Olsen, J. A., Skalski, J. E., & Basinger, J. C. (2011). Adolescent naturalistic conceptions of moral maturity. *Social Development, 20*, 562–586.
- Helzer, E. G., Fleeson, W., Furr, R. M., Meindl, P., & Barranti, M. (2017). Once a utilitarian, consistently a utilitarian? Examining principledness in moral judgment via the robustness of individual differences. *Journal of Personality, 85*, 505–517.
- Holyoak, K. J., & Powell, D. (2016). Deontological coherence: A framework for commonsense moral reasoning. *Psychological Bulletin, 142*, 1179–1203.
- 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*, 1392–1407.
- Kreps, T. A., & Monin, B. (2014). Core values versus common sense: Consequentialist views appear less rooted in morality. *Personality and Social Psychology Bulletin, 40*, 1529–1542.
- Kronenisen, M., & Heck, D. W. (2020). Interindividual differences in the sensitivity to consequences, moral norms, and preferences for inaction: Relating basic personality traits to the CNI model. *Personality and Social Psychology Bulletin, 46*, 1013–1026.
- Landy, J. F., & Uhlmann, E. L. (2018). Morality is personal. In K. Gray, & J. Graham (Eds.), *The atlas of moral psychology: Mapping good and evil in the mind* (pp. 121–132). New York: Guilford Press.
- Lapsley, D. K., & Lasky, B. (2001). Prototypic moral character. *Identity: An International Journal of Theory and Research, 1*, 345–363.
- Lipsey, M. W., & Wilson, D. S. (2001). *Practical meta-analysis*. Thousand Oaks, CA: Sage.
- Luke, D. M., & Gawronski, B. (2021a). Political ideology and moral dilemma judgments: An analysis using the CNI model. *Personality and Social Psychology Bulletin, 47*, 1520–1531.
- Luke, D. M., & Gawronski, B. (2021b). Psychopathy and moral dilemma judgments: A CNI model analysis of personal and perceived societal standards. *Social Cognition, 39*, 41–58.
- Luke, D. M., & Gawronski, B. (in press). Temporal stability of moral dilemma judgments: A longitudinal analysis using the CNI model. *Personality and Social Psychology Bulletin*.
- Luke, D. M., Neumann, C. S., & Gawronski, B. (in press). Psychopathy and moral dilemma judgment: An analysis using the four-factor model of psychopathy and the CNI model of moral decision-making. *Clinical Psychological Science*.
- Moore, A. B., Stevens, J., & Conway, A. R. (2011). Individual differences in sensitivity to reward and punishment predict moral judgment. *Personality and Individual Differences, 50*, 621–625.
- Moshagen, M. (2010). MultiTree: A computer program for the analysis of multinomial processing tree models. *Behavior Research Methods, 42*, 42–54.
- Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology, 45*, 867–872.
- Paruzel-Czachura, M., Pypno, K., Everett, J. A. C., Bialek, M., & Gawronski, B. (in press). The drunk utilitarian revisited: Does alcohol really increase utilitarianism in moral judgment? *Personality and Social Psychology Bulletin*.
- Patil, I. (2015). Trait psychopathy and utilitarian moral judgment: The mediating role of action aversion. *Journal of Cognitive Psychology, 27*, 349–366.
- Paulhus, D. L. (2000). Nominations of intelligent exemplars: Contributions of target achievement and prototype fit. *Social Cognition, 18*, 319–328.
- Rom, S. C., Weiss, A., & Conway, P. (2017). Judging those who judge: Perceivers infer the roles of affect and cognition underpinning others' moral dilemma judgments. *Journal of Experimental Social Psychology, 69*, 44–58.
- Sacco, D. F., Brown, M., Lustgraaf, C. J. N., & Hugenberg, K. (2017). The adaptive utility of deontology: Deontological moral decision-making fosters perceptions of trust and likeability. *Evolutionary Psychological Science, 3*, 125–132.
- Suter, R. S., & Hertwig, R. (2011). Time and moral judgment. *Cognition, 119*, 454–458.
- Thomson, J. J. (1976). Killing, letting die, and the trolley problem. *The Monist, 59*, 204–217.
- Turpin, M. H., Walker, A. C., Fugelsang, J. A., Sorokowski, P., Grossman, I., & Bialek, M. (2021). The search for predictable moral partners: Predictability and preferences in moral character. *Journal of Experimental Social Psychology, 97*, Article 104196.
- Uhlmann, E. L., Zhu, L., & Tannenbaum, D. (2013). When it takes a bad person to do the right thing. *Cognition, 126*, 326–334.
- Valdesolo, P., & DeSteno, D. (2006). Manipulations of emotional context shape moral judgment. *Psychological Science, 17*, 476–477.
- Walker, L. J. (1999). The perceived personality of moral exemplars. *Journal of Moral Education, 28*, 145–162.
- Walker, L. J., & Hennig, K. H. (2004). Differing conceptions of moral exemplarity: Just, brave, caring. *Journal of Personality and Social Psychology, 86*, 629–647.
- Walker, L. J., & Pitts, R. C. (1998). Naturalistic conceptions of moral maturity. *Developmental Psychology, 34*, 403–419.