

Consequences, Norms, and Generalized Inaction in Moral Dilemmas: The CNI Model of Moral Decision-Making

Bertram Gawronski
University of Texas at Austin

Joel Armstrong
University of Western Ontario

Paul Conway
Florida State University

Rebecca Friesdorf
Wilfrid Laurier University

Mandy Hütter
Eberhard Karls Universität Tübingen

Research on moral dilemma judgments has been fundamentally shaped by the distinction between utilitarianism and deontology. According to the principle of utilitarianism, the moral status of behavioral options depends on their consequences; the principle of deontology states that the moral status of behavioral options depends on their consistency with moral norms. To identify the processes underlying utilitarian and deontological judgments, researchers have investigated responses to moral dilemmas that pit one principle against the other (e.g., trolley problem). However, the conceptual meaning of responses in this paradigm is ambiguous, because the central aspects of utilitarianism and deontology—consequences and norms—are not manipulated. We illustrate how this shortcoming undermines theoretical interpretations of empirical findings and describe an alternative approach that resolves the ambiguities of the traditional paradigm. Expanding on this approach, we present a multinomial model that allows researchers to quantify sensitivity to consequences (*C*), sensitivity to moral norms (*N*), and general preference for inaction versus action irrespective of consequences and norms (*I*) in responses to moral dilemmas. We present 8 studies that used this model to investigate the effects of gender, cognitive load, question framing, and psychopathy on moral dilemma judgments. The findings obtained with the proposed CNI model offer more nuanced insights into the determinants of moral dilemma judgments, calling for a reassessment of dominant theoretical assumptions.

Keywords: deontology; moral judgment; multinomial modeling; omission bias; utilitarianism

In the summer of 2014, Dr. Kent Brantly, an American doctor working for a medical mission group in Liberia, got infected with Ebola. Dr. Brantly's chances of surviving his infection were extremely low unless he was returned to the United States for treatment. Yet, bringing him back involved a risk that the virus might spread and cause an Ebola outbreak in the United States. During the weeks before Dr. Brantly was returned and cured of the deadly disease, his case became the subject of heated debates in the popular media (Blinder & Grady, 2014). Whereas some claimed a moral duty to save Dr. Brantly's life by returning him to the United States for treatment, others argued that it would be better to let him die in Liberia in order to avoid the potential death of a larger number of people that might be caused by an Ebola outbreak in the United States.

The debate surrounding Dr. Brantly's case illustrates potential conflicts between two moral principles that play a central role in research on moral judgment. The first principle, often associated with the moral philosophy of Immanuel Kant, emphasizes the irrevocable universality of rights and duties. According to the principle of *deontology*, the moral status of a behavioral option is derived from its consistency with moral norms (often called *rule-based morality*). From a deontological view, a behavioral option is morally acceptable if it is consistent with moral norms, but it is morally unacceptable if it is inconsistent with moral norms (e.g., not returning Dr. Brantly to the United States is morally unacceptable from a deontological view, because it violates a moral duty to

save his life). The second principle, often associated with the moral philosophy of John Stuart Mill, emphasizes the greater good. According to the principle of *utilitarianism*, the moral status of a behavioral option depends on its consequences for overall well-being (also called *consequentialist morality*). To the extent that a behavioral option increases overall well-being, it is deemed morally acceptable from a utilitarian view. Yet, if the same behavioral option decreases overall well-being, it is deemed morally unacceptable from a utilitarian view (e.g., not returning Dr. Brantly to the United States is morally acceptable from a utilitarian view, because it prevents the potential death of a larger number of people).

For decades, moral psychology was dominated by rationalist theories assuming that moral judgments arise from deliberate thought processes involving the reasoned application of abstract moral principles (Kohlberg, 1969). More recently, this rationalist approach gave way to theories that attribute a central role to affective and intuitive processes (Greene & Haidt, 2002). Consistent with this development, several recent theories argue that moral judgments often stem from psychological processes that do not involve reasoned applications of abstract moral principles. For example, a widespread assumption in moral psychology is that utilitarian judgments result from a deliberate cognitive analysis of costs and benefits, whereas deontological judgments are the product of automatic emotional processes that do not involve a reasoned application of moral norms (Greene,

Morelli, Lowenberg, Nystrom, & Cohen, 2008; Greene, Nystrom, Engell, Darley, & Cohen, 2004; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001).

To test these hypotheses, numerous studies have investigated responses to moral dilemmas that pit one principle against the other, the most prominent example being the so-called trolley problem (e.g., Bartels, 2008; Greene et al., 2001, 2004; Koenigs, Young, Adolphs, Tranel, Cushman, Hauser, & Damasio, 2007; Moore, Clark, & Kane, 2008; Nichols & Mallon, 2006; Starcke, Ludwig, & Brand, 2012; Strohminger, Lewis, & Meyer, 2011; Suter & Hertwig, 2011; Valdesolo & DeSteno, 2006). In the current article, we argue that the conceptual meaning of responses in this paradigm is ambiguous, because the defining aspects of utilitarianism and deontology, consequences and norms, are not manipulated (see Gawronski & Beer, in press). Therefore, theoretical interpretations of the observed judgments are premature and prone to inaccurate conclusions about the psychological underpinnings of moral judgments. To address this limitation, we propose an alternative approach in which utilitarian responses are inferred from the sensitivity of judgments to morally relevant consequences, and deontological responses are inferred from the sensitivity of judgments to moral norms. Expanding on this approach, we present a mathematical model that allows researchers to quantify sensitivity to consequences (C), sensitivity to moral norms (N), and general preference for inaction versus action irrespective of consequences and norms (I) in responses to moral dilemmas.¹ To demonstrate the more nuanced insights that can be gained with this approach, we present 8 studies that have used this model to investigate the effects of gender, cognitive load, question framing, and psychopathy on moral dilemma judgments.

The Moral Dilemma Paradigm

The traditional moral dilemma paradigm is based on the idea that utilitarian and deontological responses can be measured with scenarios that pit one principle against the other. The most prominent example is the trolley problem in which a runaway trolley would kill a group of five workers unless participants engage in actions to redirect or stop the trolley. In the original switch dilemma, participants could pull a lever to redirect the trolley to another track, where it would kill only one person instead of five (Foot, 1967). Other variants of the trolley problem include the footbridge dilemma, in which the five workers could be saved by pushing a man from a bridge to stop the trolley (Thomson, 1976). According to the principle of utilitarianism, pulling the lever or

pushing the man would be morally acceptable, because either action maximizes overall well-being (i.e., it is morally acceptable to kill one person, if it helps to save the lives of five). According to the principle of deontology, both actions are morally unacceptable, because they are in conflict with the moral norm that one should not kill other people (i.e., it is morally unacceptable to kill another person, regardless of the consequences). Thus, participants who view these actions as acceptable are usually claimed to have made a utilitarian judgment, whereas participants who view them as unacceptable are claimed to have made a deontological judgment.

Despite their striking popularity, research using trolley problems has been criticized for many reasons, one prominent critique objecting to the unrealistic, comical nature of the employed scenarios. For example, Bauman, McGraw, Bartels, and Warren (2014) argued that trolley problems tend to be amusing rather than sobering, and therefore unlikely to elicit the same psychological processes as moral situations in the real-world. This criticism could be addressed by using more serious dilemmas of higher real-world relevance. For example, after 9/11 there were heated debates regarding whether it would be acceptable to shoot down abducted passenger planes to prevent terrorists from crashing them into densely populated areas (Whitlock, 2006). Carefully constructed scenarios of this kind would address many of the concerns against trolley problems while keeping their original dilemma structure. However, there is a much more fundamental limitation that still applies to any such dilemmas. The main problem is that the defining aspects of utilitarian and deontological responses, consequences and norms, are not manipulated in the traditional dilemma paradigm, which undermines theoretical interpretations of the observed responses.

Identifying Utilitarian and Deontological Responses

Conceptually, utilitarian judgments are defined by the property that they are sensitive to consequences. Thus, to categorize a given judgment as utilitarian, it is essential to confirm its property of being sensitive to consequences, which requires a comparison of judgments across scenarios with different consequences (Gawronski & Beer, in press). Yet, somewhat surprisingly, the specific consequences of a given behavioral option are rarely manipulated in moral dilemma research (for notable exceptions, see Cao, Zhang, Song, Wang, Miao, & Peng, in press; Conway & Gawronski, 2013; Nichols & Mallon, 2006; Trémolière & Bonnefon, 2014; Piazza, Sousa, & Holbrook, 2013). The failure to manipulate consequences renders theoretical interpretations of traditional dilemma responses ambiguous. On the one hand, it is possible that participants are willing to sacrifice the life of one person only when it saves the lives of several others. On the other hand, it is possible

¹ The qualifier *general* in *general preference for inaction versus action* is meant to specify that the described preference is independent of consequences and moral norms (see Figure 1). It is not meant to describe a psychological trait or a judgmental preference that necessarily generalizes to non-moral decision-making (although neither one is ruled out as an empirical possibility).

that participants are willing to sacrifice the life of one person even if no lives are saved.

To illustrate this ambiguity, consider a participant who views the action of redirecting the trolley in the switch dilemma as acceptable. In terms of the traditional approach, this participant would be claimed to have made a utilitarian judgment. Now imagine that the same participant views the action of redirecting the trolley as acceptable even when this action saves the life of only one person (or no lives at all). In this case, the observed pattern of judgments would not qualify as utilitarian in the sense of the above definition, because it is insensitive to morally relevant consequences. Consistent with this concern, several studies have shown that participants with sub-clinical levels of psychopathy show a greater willingness to accept harmful actions in the traditional trolley paradigm than non-psychopathic participants (Bartels & Pizarro, 2011; Kahane, Everett, Earp, Farias, & Savulescu, 2015; Patil, 2015; Pletti, Lotto, Buodo, & Sarlo, in press). Yet, it seems highly implausible that this difference reflects a greater sensitivity to morally relevant consequences among psychopaths. Instead, it seems more likely that psychopaths are willing to accept harmful actions regardless of their consequences. Thus, acceptance of harmful action in trolley dilemmas (and structurally similar dilemmas) may reflect either (1) a genuine sensitivity to consequences in the utilitarian sense, or (2) a general willingness to accept harmful actions irrespective of their consequences. A clear distinction between the two possibilities requires experimental manipulations of consequences, which tend to be absent in traditional moral dilemma research.

Interpretations of deontological judgments entail a similar ambiguity. Conceptually, deontological judgments are defined by the property that they are sensitive to moral norms. Thus, to categorize a given judgment as deontological, it is essential to confirm its property of being sensitive to moral norms, which requires a comparison of judgments across scenarios involving different moral norms (Gawronski & Beer, in press). Yet, somewhat surprisingly, the judgmental implications of moral norms have hardly ever been manipulated in moral dilemma research (for a notable exception, see Broeders, van den Bos, Müller, & Ham, 2011). The most significant limitation in this regard is the exclusive focus on proscriptive norms (i.e., norms that specify what people should not do) without any consideration of prescriptive norms (i.e., norms that specify what people should do) (Janoff-Bulman, Sheikh, & Hepp, 2009). Experimental manipulations of dilemmas involving proscriptive versus prescriptive norms are essential for theoretical interpretations of moral dilemma responses, because the exclusive use of dilemmas involving proscriptive norms conflates sensitivity to moral norms with general preference for inaction (see Crone & Laham, 2017; van den Bos, Müller, & Damen, 2011). On the one hand, it is possible that rejection of

morally proscribed actions reflects a sensitivity of judgments to moral norms. On the other hand, it is possible that rejection of morally proscribed actions reflects a general preference for inaction regardless of moral norms.

Although previous research equated deontological judgments with preference for inaction, this confound does not reflect the way sensitivity to moral norms has to play out in moral dilemma judgments (see Crone & Laham, 2017; Gawronski & Beer, in press). To illustrate this point, consider the debate surrounding Dr. Brantly's Ebola infection. As we noted at the beginning of this article, some people claimed a moral duty to return Dr. Brantly to the United States to save his life, whereas others pointed out that his return risked many lives if it caused an Ebola outbreak in the United States. Different from the structure of the trolley problem, a deontological judgment in the Ebola debate supports action (i.e., a moral duty to return Dr. Brantly to the United States to save his life), whereas a utilitarian judgment supports inaction (i.e., not returning Dr. Brantly to prevent potential harm to a larger number of people).

As the Ebola example illustrates, a general preference for inaction has to be distinguished from sensitivity to moral norms, because the latter would suggest action in cases involving a prescriptive norm. In moral psychology, the role of general action aversion has been studied under the label *omission bias*, which refers to the finding that harm caused by action is perceived as worse than equivalent harm caused by inaction (Cushman, Young, & Hauser, 2006; Spranca, Minsk, & Baron, 1991). In the broader field of decision-making, similar asymmetries have been found in research on the *status quo bias* and the *action effect*. The status quo bias describes a preference for inaction in decision contexts with a status quo alternative (e.g., option to maintain versus change a prior decision), which leads to a perpetuation of the current state of affairs (Samuelson & Zeckhauser, 1988). The action effect refers to the finding that negative outcomes are regretted more when they are a result of action compared to inaction (Kahneman & Tversky, 1982). Applied to research using trolley problems (and structurally similar dilemmas), the prevalence of action-inaction asymmetries in decision-making suggests that rejection of harmful action may reflect either (1) a genuine sensitivity of judgments to moral norms or (2) general preference for inaction regardless of moral norms. A clear distinction between the two possibilities requires experimental manipulations of dilemmas involving proscriptive versus prescriptive norms, which tend to be absent in traditional moral dilemma research.

Our discussion of consequences and norms as determinants of moral dilemma judgments further illustrates why it is imperative to consider general action tendencies as a potential third determinant. Although general preference for inaction versus action can lead to

judgments that are congruent with either the utilitarian or the deontological principle (see Baron, 1994), the specific patterns of congruence depend on features of the situation. For example, in the typical scenario of trolley dilemmas, a general preference for inaction would lead to a judgment that is congruent with the deontological principle and incongruent with the utilitarian principle. Yet, this pattern differs from the one in the Ebola dilemma, where a general preference for inaction would lead to a judgment that is congruent with the utilitarian principle and incongruent with the deontological principle. Moreover, a general preference for inaction would conflict with both principles when (1) the benefits of a given action outweigh its costs and (2) a prescriptive norm prescribes action. In such cases, both sensitivity to consequences and sensitivity to moral norms would suggest action, which conflicts with a general preference for inaction.

Resolving Interpretational Ambiguities

Our analysis suggests that unambiguous interpretations of moral dilemma judgments require confirmations of their defining properties. Categorizations of moral dilemma judgments as “utilitarian” presuppose that the observed judgment is sensitive to consequences, which requires experimental manipulations of consequences. Categorizations of moral dilemma judgments as “deontological” presuppose that the observed judgment is sensitive to moral norms, which requires experimental manipulations of moral norms. Together, these conclusions imply that conceptual interpretations of moral dilemma judgments require a comparison of responses to four types of dilemmas involving different consequences and norms: (1) dilemmas in which a proscriptive norm prohibits action and the benefits of action for overall well-being are greater than the costs of action; (2) dilemmas in which a proscriptive norm prohibits action and the benefits of action for overall well-being are smaller than the costs of action; (3) dilemmas in which a prescriptive norm prescribes action and the benefits of action for overall well-being are greater than the costs of action; (4) dilemmas in which a prescriptive norm prescribes action and the benefits of action for overall well-being are smaller than the costs of action (Gawronski & Beer, in press). Variants of the four types of dilemmas used in the current research are presented in Appendix A. To avoid the concern that moral dilemmas tend to be unrealistic and amusing rather than sobering (e.g., Bauman et al., 2014), all of these dilemmas were inspired by real-world cases that ignited moral debates about the most appropriate courses of action (e.g., Blinder & Grady, 2014). For example, the immune deficiency dilemma was inspired by Dr. Kent Brantly’s Ebola infection, using the severity of action outcomes to manipulate consequences (i.e., minor vs. severe consequences) and salient

behavioral options to manipulate moral norms (i.e., norm to not harm vs. norm to help).²

The table on the right side of Figure 1 depicts the patterns of responses across the four types of dilemmas that would qualify as utilitarian and deontological judgments in terms of their definitions as being sensitive to consequences and norms, respectively. A pattern of utilitarian responding is reflected in a preference for action when the benefits of action for overall well-being are greater than the costs of action and a preference for inaction when the benefits of action for overall well-being are smaller than the costs of action (see first row in table of Figure 1). A pattern of deontological responding is reflected in a preference for inaction when a proscriptive norm prohibits action and a preference for action when a prescriptive norm prescribes action (see second row in table of Figure 1). Either of these patterns has to be distinguished from a general preference for inaction versus action regardless of consequences and norms (see third and fourth rows in table of Figure 1). If the focus is limited to dilemmas in which a proscriptive norm prohibits action and the benefits of action are greater than the costs (e.g., trolley problem), sensitivity to consequences becomes indistinguishable from a general preference for action regardless of consequences and sensitivity to moral norms become indistinguishable from a general preference for inaction regardless of moral norms (see first column in table of Figure 1).³

² The *immune deficiency* dilemma was inspired by Dr. Kent Brantly’s Ebola’s infection in Liberia and the debate about his return to the United States for treatment. The *assisted suicide* dilemma was inspired by the case of Brittany Maynard and the resulting debate about assisted suicide. The *abduction dilemma* was inspired by the beheading of the American journalist James Foley by the terrorist group ISIS. The *transplant dilemma* was inspired by debates about a potential denial of life support to obtain organs for transplants. The *torture dilemma* was inspired by whistleblowers’ revelation of the use of illegal interrogation techniques by the CIA and a child abduction case in Canada that ignited a similar debate. The *vaccine dilemma* was inspired by the debate about the potential use of unapproved vaccines to fight the Ebola epidemic in various African countries.

³ A philosophical objection against our conceptualization of deontological judgments is that general preference for inaction is consistent with the deontological doctrine of doing and allowing (DDA), which states that actively causing harm is morally worse than passively allowing harm. From this perspective, both sensitivity to moral norms and general inaction represent patterns of deontological responding, suggesting that they should be combined under the unifying umbrella of *deontological judgments*. However, from a psychological view, the two cases are still fundamentally different, because they lead to distinct patterns of moral judgments. Whereas sensitivity to moral norms favors inaction in dilemmas involving a proscriptive norm and action in dilemmas involving a prescriptive norm (see second row in table of Figure 1), a general preference for inaction favors inaction regardless of proscriptive and prescriptive norms (see third row in table of Figure 1). Because the two patterns of responses are likely driven by different psychological processes, we deem it important to distinguish between sensitivity to moral norms and general preference for inaction in moral dilemma judgments, even though both patterns may be described as deontological. We address the DDA argument in more detail in the General Discussion when we discuss theoretical implications of our findings.

To avoid potential misunderstandings of our argument, it may help to clarify the primary source of disagreement with the traditional way of interpreting moral dilemma judgments. In the traditional approach, it would be *sufficient* for a judgment to qualify as utilitarian if it supports action in dilemmas where a proscriptive norm prohibits action and the benefits of action are greater than the costs (e.g., trolley problem). Conversely, it would be *sufficient* for a judgment to qualify as deontological if it supports inaction in dilemmas where a proscriptive norm prohibits action and the benefits of action are greater than the costs. What matters is the mere consistency of a given judgment with either the utilitarian or the deontological principle (see Greene, 2007). Our conceptualization differs from the traditional approach by treating these criteria as *necessary*, but *insufficient*, for the identification of utilitarian and deontological judgments. To qualify as utilitarian, a judgment needs to meet the criterion of being sensitive to consequences, which cannot be inferred from a single judgment but requires a comparison across dilemmas with different consequences (Gawronski & Beer, in press). That is, moral judgments should reflect a preference for action when the benefits of action are greater than the costs and a preference for inaction when the benefits of action are smaller than the costs. In other words, a preference for action on dilemmas in which a proscriptive norm prohibits action and the benefits of action are greater than the costs is *necessary* but *insufficient* to categorize moral judgments as utilitarian. Similarly, to qualify as deontological, a judgment needs to meet the criterion of being sensitive to moral norms, which cannot be inferred from a single judgment but requires a comparison across dilemmas with different moral norms (Gawronski & Beer, in press). That is, moral judgments should reflect a preference for inaction when a proscriptive norm prohibits action and a preference for action when a proscriptive norm prescribes action. In other words, a preference for inaction on dilemmas in which a proscriptive norm prohibits action and the benefits of action are greater than the costs is *necessary* but *insufficient* to categorize moral judgments as deontological.⁴

Another important aspect of our conceptualization is that it refers to properties of overt judgments rather than underlying mechanisms (see also Greene, 2007). Although utilitarian judgments are conceptually defined

by the property that they are sensitive to consequences, this definition does not presuppose that the effect of consequences on moral judgments is mediated by a deliberate analysis of costs and benefits. After all, it is possible that the effect of consequences on moral judgments is mediated by a process that does not involve a deliberate analysis of costs and benefits (see Greene & Haidt, 2002). Similarly, deontological judgments are conceptually defined by the property that they are sensitive to moral norms, but this definition does not presuppose that the effect of norms on moral judgments is mediated by a reasoned application of moral norms. After all, it is possible that the effect of moral norms is mediated by a process that does not involve a reasoned application of moral norms, such as automatic emotional reactions (see Greene & Haidt, 2002). The psychological mechanism underlying each type of effect is a theoretical issue that can only be addressed via empirical data (see De Houwer, 2011). They are not implied by the conceptual definition of utilitarian and deontological judgments in terms their sensitivity to consequences and moral norms.

Quantification of Response Tendencies with Multinomial Modeling

Expanding on the proposition that utilitarian and deontological judgments should be identified on the basis of their sensitivity to consequences and norms, we developed a multinomial processing tree model that allows researchers to quantify (1) sensitivity to consequences, (2) sensitivity to moral norms, and (3) general preference for inaction versus action in responses to moral dilemmas (see Figure 1). Multinomial modeling is a statistical technique designed to disentangle the simultaneous contribution of multiple factors to categorical responses (for reviews, see Batchelder & Riefer, 1999; Hütter & Klauer, 2016). In social psychology, multinomial modeling has been used to investigate the effect of stereotypes on memory (e.g., Gawronski, Ehrenberg, Barse, Zukova, & Klauer, 2003; Klauer & Wegener, 1998), processes underlying implicit measures (e.g., Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005; Meissner & Rothermund, 2013; Nadarevic & Erdfelder, 2011; Stahl & Degner, 2007), attitude formation via evaluative conditioning (e.g., Hütter & Sweldens, 2013; Hütter, Sweldens, Stahl, Unkelbach, & Klauer, 2012), the activation and application of stereotypes (e.g., Krieglmeier & Sherman, 2012), intentional and unintentional empathy for pain (e.g., Cameron, Spring, & Todd, in press), and intentional and unintentional moral evaluations (e.g., Cameron, Payne, Sinnott-Armstrong, Scheffer, & Inzlicht, 2017). In the current work, we adapted a multinomial modeling approach to measure the strength of specific response patterns in participants' judgments across moral dilemmas with different consequences and norms. By comparing participants' responses to the four types of

⁴ To further illustrate this argument, consider the case of a police officer shooting an unarmed Black man holding an ambiguous object. Although the officer's response may be categorized as an instance of racial bias, such a categorization presupposes that the officer's response would have been different if the target was White. Similar to our arguments about the categorization of moral dilemma judgments, a categorization of a response as racial bias requires that the response is sensitive to the race of the target. Such sensitivity cannot be inferred from a single response, but requires a comparison of responses across cases involving targets with different racial backgrounds (see Payne, 2001; Correll, Judd, Park, & Wittenbrink, 2002).

dilemmas across multiple scenarios (see Appendix A), our multinomial model allows researchers to quantify (1) the degree to which participants' judgments are sensitive to consequences (see first row of Figure 1), (2) the degree to which participants' judgments are sensitive to moral norms (see second row of Figure 1), and (3) the degree to which participants' judgments reflect a general preference for inaction versus action irrespective of consequences and norms (see third versus fourth rows of Figure 1).

The mathematical technique for these quantifications can be illustrated by mapping the response patterns in Figure 1 to a processing tree in which they are depicted as the outcomes of four processing paths (see Figure 1). The three parameters in this processing tree capture the likelihoods that the response to a given dilemma is driven by consequences (C), moral norms (N), and a general preference for inaction versus action irrespective of consequences and norms (I), hence termed the *CNI model of moral decision-making*. The four paths in the processing tree capture the four cases that lead to the depicted patterns of responses in Figure 1. The utilitarian response pattern in the first row of the table should emerge when responses to the four types of dilemmas are driven by consequences. In the processing tree, this case is captured by the path C , which depicts the case that consequences drive responses. The deontological response pattern in the second row of the table should emerge when responses to the four types of dilemmas are driven by moral norms. In the processing tree, this case is captured by the path $(1 - C) \times N$, which depicts the case that moral norms drive responses when consequences do not drive responses. The pattern of general inaction in the third row of the table should emerge when responses to the four types of dilemmas reflect a general preference for inaction. In the processing tree, this case is captured by the path $(1 - C) \times (1 - N) \times I$, which depicts the case of general inaction when neither moral norms nor consequences drive responses. Finally, the pattern of general action in the fourth row of the table should emerge when responses to the four types of dilemmas reflect a general preference for action. In the processing tree, this case is captured by the path $(1 - C) \times (1 - N) \times (1 - I)$, which depicts the case of general action when neither moral norms nor consequences drive responses.⁵

Using the four processing paths on the left side of Figure 1, it is possible to derive mathematical equations that capture the statistical likelihood of an action (inaction) response on each of the four types of dilemmas. These equations can be derived by (1)

identifying all cases within a given column that involve an action (inaction) response, (2) linking each case to its corresponding processing path on the left side of the figure, and (3) summing the mathematical representations of these paths in a single equation. For example, in dilemmas where a proscriptive norm prohibits action and the benefits of action for overall well-being are greater than the costs of action, an inaction response should occur when either (1) moral norms drive the response when consequences do not drive the response, $(1 - C) \times N$, or (2) general preference for inaction drives the response when neither moral norms nor consequences drive the response, $(1 - C) \times (1 - N) \times I$. Algebraically, this probability can be depicted as:

$$p(\text{inaction} \mid \text{proscriptive norm, benefits} > \text{costs}) = [(1 - C) \times N] + [(1 - C) \times (1 - N) \times I]$$

Conversely, in dilemmas where a proscriptive norm prohibits action and the benefits of action for overall well-being are greater than the costs of action, an action response should occur when either (1) consequences drive the response, C , or (2) general preference for action drives the response when neither moral norms nor consequences drive the response, $(1 - C) \times (1 - N) \times (1 - I)$. Algebraically, this probability can be depicted as:

$$p(\text{action} \mid \text{proscriptive norm, benefits} > \text{costs}) = C + [(1 - C) \times (1 - N) \times (1 - I)]$$

The same procedure can be used to derive the equations for action and inaction responses on the other three types of dilemmas. For the four types of dilemmas and the two kinds of potential responses (i.e., action vs. inaction) this procedure leads to a total of eight equations (see Appendix B). Because the likelihood of an action response on a given dilemma type is equal to 1 minus the likelihood of an inaction response on that dilemma type, there are a total of four non-redundant equations that include three unknowns (i.e., C , N , I). Using maximum likelihood statistics, multinomial modeling generates parameter estimates for the three unknowns that aim to minimize the difference between the empirically observed probabilities of action versus inaction in participants' responses on the four types of dilemmas and the probabilities of action versus inaction responses predicted by the model equations using the identified parameter estimates. The adequacy of the model in describing the data can be evaluated by means of goodness-of-fit statistics, such that poor model fit would be reflected in a statistically significant deviation between empirically observed probabilities and the probabilities predicted by the model.⁶ Differences in

⁵ In the processing tree, general preference for action is conceptualized as the opposite of general preference for inaction. Hence, general action and general inaction are captured by the same parameter (I). To keep the coding of general action tendencies consistent with the concept of *omission bias* (see Cushman et al., 2006), we chose a scoring procedure in which higher scores reflect a stronger preference for inaction and lower scores reflect a stronger preference for action.

⁶ Poor goodness-of-fit in multinomial modeling usually suggests that a basic assumption of the model is violated (see Klauer, 2015). Such violations of model assumptions may reflect the general inadequacy of the model in describing the data of a given paradigm, which should lead to poor model fit across all applications of the model to this particular paradigm. Alternatively, violations of model assumptions may reflect incidental effects of additional factors that are not captured by the model, which should lead to poor model fit only when these factors are present, but not when they are absent. Moreover, because large sample

parameter estimates across groups can be tested by enforcing equal estimates for a given parameter across groups. 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.

Similar tests can be conducted to investigate whether a given parameter estimate significantly differs from a reference value. For example, to test the impact of consequences on moral judgments, the *C* parameter is set equal to zero and the resulting model fit is compared to the fit of the model that does not include any restrictions for the *C* parameter. To the extent that enforcing a parameter estimate of zero leads to a significant reduction in model fit, it can be inferred that consequences significantly influenced participants' responses to the four types of moral dilemmas. The same approach can be used to test the influence of moral norms captured by the *N* parameter. For the *I* parameter, comparisons to reference values are equivalent, except that the reference value reflecting the absence of a general action tendency is 0.5. Whereas values higher than 0.5 on the *I* parameter reflect a general preference for inaction, values lower than 0.5 reflect a general preference for action.⁷

sizes lead to smaller confidence intervals for the estimated parameters, the likelihood of significant deviations between actual and predicted response probabilities increases as a function of sample size. Thus, poor model fit does not necessarily question the general validity of a given model in describing the data to the extent that (1) significant deviations between actual and predicted probabilities are limited to a small number of studies in a series of multiple studies with large samples and (2) the obtained deviations are small in terms of their effect size (i.e., when sample size is taken into account in the evaluation of significant deviations). The size of significant deviations can be evaluated with the effect size measure Cohen's *w*. According to Cohen (1988), a *w* of 0.10 represents a small effect, a *w* of 0.30 represents a medium effect, and a *w* of 0.50 represents a large effect.

⁷ It is worth noting that the particular position of the three parameters in the processing tree does not affect the goodness-of-fit of the model. All six combinatorically possible models have the same degrees of freedom and impose the same equality restrictions on the probabilities for showing a particular response on the four kinds of dilemmas. Hence, there is no empirical basis to compare the relative validity of the six models in describing the data by means of model fit (K. C. Klauer, personal communication, March 17, 2014). An important consideration in choosing the current model specification was that the *I* parameter should be set as the lowest one in the hierarchy, because the response patterns predicted for the absence of all processes are defined as the opposite pattern of the parameter that has the lowest position in the hierarchy (i.e., general action as the opposite of general inaction). This feature permits an estimation of general action preferences along a bipolar continuum of inaction versus action (rather than a unipolar dimension of general preference for inaction). To investigate potential differences in the observed results when the positions of the *C* and the *N* parameter are reversed, we reran all reported analyses with the reversed model. All of the reported effects replicated with the reversed model. If anything, the results were somewhat stronger in that two effects that were only marginally significant with the model in Figure 1 reached statistical significance with the reversed model (i.e., the marginal effect on the *N* parameter in Study 3a and the marginal effect on the *C*

Relation to Process Dissociation

An important question is how the CNI model relates to earlier work by Conway and Gawronski (2013) who used Jacoby's (1991) process dissociation (PD) procedure to disentangle the independent contributions of utilitarian and deontological inclinations to moral judgments. Mathematically, PD is very similar to multinomial modeling, in that both procedures allow researchers to quantify the contribution of multiple factors to behavioral outcomes (see Hütter & Klauer, 2016; Payne & Bishara, 2009). The primary difference is that PD is limited to two factors whose contributions are calculated directly via linear algebra using two equations with two unknowns. Multinomial modeling permits the quantification of more than two factors whose contributions are estimated through maximum likelihood statistics using more equations than unknowns.

Aside from these technical differences, a major content-related difference is that Conway and Gawronski's (2013) PD approach is limited to moral dilemmas involving proscriptive norms. It does not capture the difference between moral dilemmas involving proscriptive versus prescriptive norms. As such, the outcomes captured by the PD model are limited to the four cells in the top-left corner of Figure 1. This limitation leads to two major confounds in the parameters of the PD model (see Gawronski, Conway, Armstrong, Friesdorf, & Hütter, 2016; Hütter & Klauer, 2016): (1) a confound between sensitivity to moral norms and general preference for inaction in the PD parameter that is supposed to capture deontological inclinations (*D* parameter) and (2) a confound between sensitivity to consequences and general preference for action in the PD parameter that is supposed to capture utilitarian inclinations (*U* parameter). Our CNI model resolves the confounds in Conway and Gawronski's (2013) PD model, which can (1) lead to spurious effects on the two PD parameters when a given factor influences general preference for inaction and (2) conceal meaningful effects when a given factor influences general preference for inaction in a manner that counteracts simultaneous effects on sensitivity to consequences or sensitivity to moral norms.

The Current Research

To demonstrate the more nuanced insights that can be gained from our CNI model, we present the results of 8 studies that tested effects of gender (Studies 1a and 1b), cognitive load (Studies 2a and 2b), question framing (Studies 3a and 3b), and psychopathy (Studies 4a and 4b) on moral dilemma judgments. Our concern with gender differences was inspired by the question of whether women are more deontological or less utilitarian than men (or both), which remains ambiguous in the

parameter in Study 4a were statistically significant with the reversed model).

traditional dilemma approach (see Friesdorf, Conway, & Gawronski, 2015). Our studies on cognitive load were based on the hypothesis that utilitarian judgments are the product of cognitively effortful processes (see Greene et al., 2008). Our studies on framing effects were inspired by research that utilized manipulations of personal force to demonstrate enhanced deontological responding as a result of increased emotional engagement (e.g., Greene et al., 2001). Finally, our concern with psychopathy was based on research showing that participants with subclinical levels of psychopathy show a stronger tendency for utilitarian responding in trolley dilemmas, which has been cited as an argument against the validity of the trolley paradigm (e.g., Bartels & Pizzaro, 2011; Kahane et al., 2015).

Based on recent concerns about the reproducibility of psychological findings (Open Science Collaboration, 2015), we conducted one initial study and one replication for each of the four variables. To avoid conclusions from false positives, we limit interpretations to parameter effects that replicate across the two studies.⁸ To ensure that the scenarios in our dilemmas are perceived as morally relevant by our participants, we also conducted a pilot study in which participants were asked to rate the moral relevance of the behaviors in our dilemmas.

To assess the goodness-of-fit of the model in describing the data and to calculate estimates for the three parameters, we used the free software multiTree by Moshagen (2010).⁹ In addition to goodness-of-fit statistics and estimates for the three parameters, the software provides standard errors and 95% confidence intervals for the estimated parameter values.¹⁰ All of the reported studies used the same estimation algorithm with

⁸ In addition to the eight studies reported in the main text, the Supplementary Materials report the results of two studies that investigated effects of images of the focal targets in the moral dilemmas. The experimental manipulation in these studies was adapted from Conway and Gawronski (2013) who aimed to manipulate salience of harm through pictures of the focal targets. The basic idea underlying their manipulation was that increased salience of harm should enhance automatic emotional responses to the idea of causing harm, which should increase deontological judgments according to the dual-process model of moral dilemma judgment (see Amit & Greene, 2012). Although the CNI model fit the data well in both studies, the effects on the three parameters were inconsistent across the two studies. In the first study, presenting images of the focal targets failed to produce any significant effects at all. In the second study, the same manipulation significantly decreased scores on the *N* and the *C* parameters without affecting the *I* parameter. In line with our stated practice of limiting interpretations to parameter effects that replicate across studies, we refrain from drawing any conclusions from these findings. Because the two studies also failed to reproduce the original findings by Conway and Gawronski (2013), we attribute the inconsistent results to unreliable effects of Conway and Gawronski's (2013) picture manipulation.

⁹ The software can be downloaded for free at: <http://psycho3.uni-mannheim.de/Home/Research/Software/multiTree/>

¹⁰ Effect sizes of between-group differences in the current studies were calculated with David Wilson's online effect size calculator using means, standard errors, and sample sizes (<https://www.campbellcollaboration.org/escalc/html/EffectSizeCalculator-SMD8.php>). The online tool is an official companion to Lipsey and Wilson's (2001) practical guide to meta-analyses.

random start values, two replications, and a maximum of 90,000 iterations. With two groups in each study, our model has a total of 8 free categories (i.e., responses to four types of dilemmas for each of the two groups) and a total of 6 parameters (i.e., three parameters estimated for each of the two groups), which result in a difference of 2 for the degrees of freedom of the model. A zip-file with a multiTree template and a tutorial on how to analyze moral judgment data with our CNI model are available at: http://www.bertramgawronski.com/documents/CNI-Model_Materials.zip. The zip-file also includes copies of the dilemma materials (see Appendix A), a template file for lab studies with our moral dilemmas using the psychological lab software MediaLab by Empirisoft, and an SPSS syntax file for the aggregation of data obtained with our MediaLab template file.

For each study, we aimed to recruit 200 participants (i.e., 100 participants per condition), which provides a statistical power of .80 to detect a small between-group effect of $d = .40$ in the difference between two independent mean values. By default, we excluded participants who failed to pass an instructional attention check (4.6% across all studies) that was included in all of the reported studies (see Oppenheimer, Meyvis, & Davidenko, 2009). The data for each study (or study session in Studies 4a and 4b) were collected in one shot without intermittent statistical analyses. We report all measures, all conditions, and all data exclusions. The raw data and syntax files for all studies are publicly available at <https://osf.io/xt66w/>.

To compare the findings obtained with our CNI model to the results of the traditional approach, we first report participants' responses on moral dilemmas where a proscriptive norm prohibits action and the benefits of action for overall well-being outweigh its costs. In the traditional approach, action responses on this type of dilemma have been treated as a direct indicator of utilitarian responding and inaction responses have been treated as a direct indicator of deontological responding. Thus, larger proportions of action responses are typically interpreted as reflecting the degree to which participants show a preference for utilitarian over deontological responses. To illustrate the more nuanced insights that can be gained from our CNI model compared to Conway and Gawronski's (2013) PD approach, we also analyzed the data using PD. Expanding on the analysis using the traditional and the PD approach, each study presents an analysis of the obtained data with our CNI model, which offers more nuanced insights into whether the focal variable in a given study influenced moral dilemma judgments via (1) sensitivity to consequences, (2) sensitivity to moral norms, or (3) general preference for inaction versus action (or a combination of the three).

Pilot Study

To confirm that participants perceive the themes of our dilemmas to be morally relevant (cf. Baumann et al.,

2014), we conducted a pilot study in which participants were asked to rate the moral relevance of the behaviors in our scenarios. To obtain reference points for morally relevant and morally irrelevant behaviors, participants were additionally asked to rate the moral relevance of various behaviors mentioned by participants in an experience sampling study on morality in everyday life (Hofmann, Wisniewski, Brandt, & Skitka, 2014) and a set of behaviors in non-moral decision problems (Greene et al., 2001).

Participants. Participants were recruited for a study on “how people make moral judgments” via Amazon’s Mechanical Turk (MTurk) crowdsourcing platform (Buhrmester, Kwang, & Gosling, 2011). Eligibility for participation was limited to English native speakers with a HIT approval rate of at least 97%. Participants received compensation of \$1.00 for completing the study. Of the 273 MTurk workers who initially began the study, 200 completed the study until the end. Of these participants, 5 failed to pass an instructional attention check (see below). Data from these participants were excluded from the statistical analyses, leaving us with a final sample of 195 participants (110 women, 83 men, 2 neither male nor female; $M_{age} = 36.58$ years, $SD_{age} = 12.06$).¹¹

Procedure and materials. Participants were asked to read descriptions of 56 behaviors (see Appendix C) and rate them for their moral relevance on 5-point scales ranging from 1 (*not at all*) to 5 (*very much*). The behavioral statements included three sets of items: (1) brief descriptions of the 12 behaviors in our basic dilemmas involving either a proscriptive or a prescriptive norm, (2) the 32 sample behaviors listed in the Supplementary Online Materials of Hofmann et al.’s (2014) experience sampling study on morality in everyday life, and (3) a set of 12 behaviors in non-moral decision problems adapted from Greene et al. (2001). Participants were presented with the following introduction before they completed the rating task:

In this study, we are testing materials for a project on moral judgment. Toward this end, you will be asked to read 56 statements regarding different behaviors. We would like you to indicate the extent to which each of these behaviors represents an issue of moral relevance. A behavior would be morally relevant if you consider it as either moral or immoral. A behavior would be morally irrelevant if you consider it neither moral nor immoral. Please consider each of the statements independently. There are no right or wrong answers--just choose what seems best to you.

Upon completion of the rating task, participants were presented with a one-item instructional attention check to ensure that they thoroughly read the behavioral descriptions (see Oppenheimer et al., 2009). The attention check included the following instructions:

Most modern theories of decision making recognize the fact that decisions do not take place in a vacuum. Individual preferences and knowledge, along with situational variables can greatly impact the decision process. 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 questionnaires that require you to understand the instructions will be filled out inaccurately. So, in order to demonstrate you have read the instructions, please answer 'very bad' on the mood item below. Thank you very much.

Below the instructions, participants were presented with the question *What is your current mood?* and the response options: (1) *very bad*, (2) *bad*, (3) *poor*, (4) *neither good nor bad*, (5) *fair*, (6) *good*, (7) *very good*. By default, we excluded all participants from the analyses who did not follow the instruction to respond *very bad* on this item.

Results and Discussion

To compare the perceived moral relevance of the three sets of behaviors, we calculated three aggregate scores by averaging participants’ ratings of (1) the behaviors in our moral dilemmas (Cronbach’s $\alpha = .88$), (2) the behaviors of Hofmann et al.’s (2014) experience sampling study on morality in everyday life (Cronbach’s $\alpha = .85$), and (3) the behaviors in Greene et al.’s (2001) non-moral decision problems (Cronbach’s $\alpha = .93$). An analysis of variance (ANOVA) with Behavior Type as a within-subjects factor revealed that participants’ ratings of moral relevance significantly differed across the three sets of behaviors, $F(1, 194) = 103.23, p < .001, \eta_p^2 = .347$ (see Table 1). Post-hoc analyses revealed that the behaviors in our moral dilemmas were rated higher on moral relevance compared to both the moral behaviors in Hofmann et al.’s (2014) experience sampling study, $t(194) = 12.59, p < .001, d = 0.961$, and the behaviors in Greene et al.’s (2001) non-moral decision problems, $t(194) = 17.18, p < .001, d = 1.231$. Moreover, the moral behaviors in Hofmann et al.’s (2014) experience sampling study were rated higher on moral relevance compared to the behaviors in Greene et al.’s (2001) non-moral decision problems, $t(194) = 16.47, p < .001, d = 1.206$.

To explore potential differences between thematic subsets of items, we also calculated subscores for the central themes in the three sets of behaviors (see Table 1). Although there was considerable variation within each set of behavior, almost all behaviors in our moral dilemmas were rated higher on moral relevance compared to the moral behaviors in Hofmann et al.’s (2014) experience sampling study. The only comparisons that were not statistically significant involved the difference between the abduction dilemma and the liberty/oppression domain, $t(193) = 1.21, p = .227, d =$

¹¹ Because participants had the option to skip items, missing data for skipped items led to lower samples sizes for some of the statistical tests reported below.

0.088, the difference between the assisted suicide dilemma and the liberty/oppression domain, $t(194) = 0.82$, $p = .413$, $d = 0.054$, the difference between the vaccine dilemma and the liberty/oppression domain, $t(194) = -1.35$, $p = .179$, $d = 0.102$, and the difference between vaccine dilemma and the honesty/dishonesty domain, $t(189) = 1.89$, $p = .060$, $d = 0.141$. For all other comparisons, the behaviors in our moral dilemmas were rated significantly higher on moral relevance compared to the moral behaviors in Hofmann et al.'s (2014) experience sampling study (all t s > 2.54, all p s < .02). A similar pattern emerged for the comparisons of the behaviors in our moral dilemmas and the behaviors in Greene et al.'s (2001) non-moral decision problems, in that all behaviors in our moral dilemmas were rated higher on moral relevance (all t s > 12.52, all p s < .001). The same was true for the comparisons of the moral behaviors in Hofmann et al.'s (2014) experience sampling study and the behaviors in Greene et al.'s (2001) non-moral decision problems, in that all behaviors in Hofmann et al.'s (2014) experience sampling study were rated higher on moral relevance (all t s > 7.47, all p s < .001). Together, these results support the assumption that participants perceive the general themes of our dilemmas as morally relevant.

Study 1a

The main goal of Study 1a was to provide a first test of the validity of our CNI model in describing patterns of responses to the four types of moral dilemmas. To the extent that the CNI model provides an accurate description of the obtained response patterns, it should reveal appropriate estimates of goodness-of-fit (i.e., the probabilities predicted by the model on the basis of the estimated parameter values should not significantly deviate from the empirically observed probabilities of action versus inaction responses to the four types of dilemmas).

In addition to establishing the overall fit of the CNI model in describing the data, Study 1a aimed at exploring its usefulness in providing deeper insights into the nature of previously obtained findings, using gender differences in moral dilemma judgments as an example. Using the traditional trolley paradigm, several studies found that men, compared to women, show a stronger willingness to accept harmful actions when such actions lead to better consequences for a larger number of people (e.g., Arutyunova, Alexandrov, & Hauser, 2016; Fumagalli et al., 2010). Using Conway and Gawronski's (2013) PD approach, Friesdorf et al. (2015) further investigated whether this gender difference reflects (1) stronger deontological inclinations among women or (2) stronger utilitarian inclinations among men (or both). In their study, women showed higher scores on the D parameter than men, but there was no gender difference on the U parameter. Yet, as we explained earlier in this article, the D parameter confounds sensitivity to moral norms with

general preference for inaction and the U parameter confounds sensitivity to consequences with general preference for action (see Gawronski et al., 2016; Hütter & Klauer, 2016). These confounds can (1) lead to spurious effects on the two PD parameters when a given factor influences general preference for inaction versus action and (2) conceal meaningful effects when a given factor influences general preference for inaction in a manner that counteracts simultaneous effects on sensitivity to consequences or sensitivity to moral norms. Our CNI model resolves these confounds, thereby providing more nuanced insights into whether gender differences in moral dilemma judgments reflect differences in (1) sensitivity to consequences, (2) sensitivity to moral norms, or (3) general preference for inaction versus action (or any combination of the three).

Method

Participants. Participants were recruited for a study on "how people make moral judgments" via Amazon's MTurk. Eligibility for participation was limited to English native speakers with a HIT approval rate of at least 97%. Participants received compensation of \$1.00 for completing the study. Of the 228 MTurk workers who initially began the study, 206 completed all measures.¹² Of these participants, 5 failed to pass an instructional attention check. Data from these participants were excluded from the statistical analyses, leaving us with a final sample of 201 participants (95 women, 106 men; $M_{age} = 32.20$ years, $SD_{age} = 10.96$).

Procedure and materials. Participants were asked to read and respond to a total of 24 moral dilemmas (see Appendix A). Following Greene et al. (2001), the dilemmas were phrased in a second-person view, each depicting participants as actors who must choose whether or not to perform a described action (see also Conway & Gawronski, 2013). Dilemmas were presented in a fixed random order, with each dilemma being presented individually on a separate screen.¹³ Participants were asked to indicate for each dilemma whether or not it is acceptable to perform the described action (yes vs. no). The dilemmas included 4 parallel versions of 6 basic scenarios that varied in terms of whether (1) the dilemma involved a proscriptive norm that prohibits action or a prescriptive norm that prescribes action and (2) the benefits of the described action for overall well-being were either greater or smaller than its costs for overall well-being. Participants received the following

¹² Six participants completed all measures, but did not submit a request for compensation.

¹³ We used a fixed random order instead of a fully randomized order, because full randomization increases measurement error, which in turn reduces statistical power. In addition, we aimed to avoid incidental confounds that can lead to spurious effects when our dilemmas are presented in different orders to different groups or in different conditions.

instructions before they were presented with the dilemmas:

On the following screens you will see a series of scenarios that people may come across in life. Please read them carefully. Even though some scenarios may seem similar, each scenario is different in important ways. After each scenario, you will be asked to make a judgment about whether you find the described action appropriate or inappropriate. Please note that some scenarios refer to things that may seem unpleasant to think about. This is because we are interested in people's thoughts about difficult, real-life issues.

After completion of the moral dilemmas, participants were asked to respond to the same one-item attention check used in the Pilot Study.

Results

The data were aggregated by calculating the sum of *action* responses to the four types of moral dilemmas. With a total of 6 scenarios for each dilemma type, aggregate scores could range from 0 to 6. Means and 95% confidence intervals are presented in Table 2.

Traditional analysis. The traditional approach focuses exclusively on moral dilemmas involving a proscriptive norm that prohibits action in cases where the benefits of action outweigh its costs to well-being. A preference for action over inaction on this type of dilemma is typically interpreted as a preference for utilitarian over deontological responses (see Greene et al., 2001). Averaged across the six dilemmas of this type, participants in the current study showed a slight preference for action over inaction on this type of dilemma ($M = 3.06$, $SD = 1.53$). However, the overall preference score was not significantly different from the neutral reference point of 3, $t(200) = 0.55$, $p = .580$, $d = 0.039$, which reflects an equal preference for action and inaction across the six dilemmas. Replicating the well-established gender difference in responses to moral dilemmas, men showed a significantly stronger preference for action versus inaction in this type of dilemma than women, $t(199) = 2.89$, $p = .004$, $d = 0.408$ (see Table 2). In terms of the traditional approach, this difference would be interpreted as a stronger preference for utilitarian over deontological responses among men compared to women.

PD analysis. PD scores were calculated in line with the procedures described by Conway and Gawronski (2013), using the probabilities of *action* responses on moral dilemmas with proscriptive norms involving benefits of action that are either greater than the costs of action (i.e., incongruent dilemmas) or smaller than the costs of action (i.e., congruent dilemmas). Replicating the findings of Friesdorf et al. (2015), women showed significantly higher scores on the *D* parameter than men, $t(199) = 3.38$, $p = .001$, $d = 0.477$ (see Table 3), but there was no significant gender difference on the *U* parameter, $t(199) = 0.69$, $p = .492$, $d = 0.097$ (see Table 3).

CNI model. The CNI model fit the data well when the data were analyzed without considering participants' gender, $G^2(1) = 1.23$, $p = .268$. Both the *C* parameter ($M = 0.203$, 95% CI [0.176, 0.240]) and the *N* parameter ($M = 0.231$, 95% CI [0.196, 0.265]) were significantly greater than zero, indicating that participants were highly sensitive to both consequences and norms in responding to the moral dilemmas, $\Delta G^2(1) = 207.14$, $p < .001$ for the *C* parameter and $\Delta G^2(1) = 171.13$, $p < .001$ for the *N* parameter. The *I* parameter did not significantly deviate from its neutral reference point of 0.5, $\Delta G^2(1) = 1.30$, $p = .254$, indicating an equal distribution of action and inaction responses in the total sample ($M = 0.513$, 95% CI [0.491, 0.535]).

The CNI model also fit the data well when parameter scores were estimated separately for men and women, $G^2(2) = 1.32$, $p = .517$. Whereas women showed significantly higher scores on the *N* parameter than men, $\Delta G^2(1) = 26.00$, $p < .001$, $d = 0.726$ (see Figure 2), there were no significant gender differences on the *C* parameter, $\Delta G^2(1) = 1.34$, $p = .247$, $d = 0.164$ (see Figure 2). Moreover, there was a significant difference on the *I* parameter, in that women showed a stronger general preference for inaction than men, $\Delta G^2(1) = 12.34$, $p < .001$, $d = 0.504$ (see Figure 2). Together, these results suggest that gender differences in moral dilemma judgments are due to (1) a stronger sensitivity to norms among women and (2) a stronger general preference for inaction among women. There seem to be no gender differences in the sensitivity to consequences.

Discussion

The main goal of Study 1a was to provide a first test of the validity of our CNI model in describing patterns of responses to the four types of moral dilemmas. In addition, we aimed at exploring its usefulness in providing deeper insights into the nature of previously obtained findings, using gender differences in moral dilemma judgments as an example. Overall, the CNI model fit the data well. The probabilities predicted by the model on the basis of the estimated parameter values did not significantly deviate from the empirically observed probabilities of action versus inaction responses to the four types of dilemmas. Moreover, replicating the well-established gender difference in moral dilemma judgments (e.g., Arutyunova et al., 2016; Friesdorf et al., 2015; Fumagalli et al., 2010), men showed a stronger acceptance of action than women when a proscriptive norm prohibited action and action led to better consequences for a larger number of people. Further analyses with Conway and Gawronski's (2013) PD model suggested that this gender difference is due to stronger deontological inclinations among women, but there seem to be no gender differences in utilitarian inclinations (see also Friesdorf et al., 2015). Our CNI model provided deeper insights into the nature of gender differences in moral dilemma judgments, showing that

they are due to (1) a stronger sensitivity to moral norms among women, and (2) a stronger general preference for inaction among women. There was no evidence for gender differences in the sensitivity to consequences. Together, these results provide preliminary evidence for the validity of the CNI model in describing responses to the four types of moral dilemmas and for its ability to provide more nuanced insights into the nature of existing findings.

Study 1b

Based on recent concerns about the reproducibility of psychological findings (Open Science Collaboration, 2015), we aimed to replicate the findings of Study 1a in a follow-up study using the same materials. To avoid conclusions from false positives, we limit interpretations of gender effects on the three parameters of the CNI model to those that replicate across the two studies.

Method

Participants. Participants were recruited for a study on “how people make moral judgments” via Amazon’s MTurk. Participants received compensation of \$1.00 for completing the study. Eligibility for participation was limited to English native speakers who (1) had a HIT approval rate of at least 97% at the time of the study and (2) had not participated in prior studies from our lab using the same set of moral dilemmas. Of the 212 MTurk workers who initially began the study, 202 completed all measures.¹⁴ Of these participants, 5 failed to pass an instructional attention check. Data from these participants were excluded from the statistical analyses, leaving us with a final sample of 197 participants (102 women, 95 men; $M_{age} = 35.77$, $SD_{age} = 11.47$).

Procedure and materials. Participants were asked to read and respond to the 24 moral dilemmas of Study 1a, using the same fixed random order. After completion of the moral dilemmas, participants were asked to respond to the same one-item attention check used in the Pilot Study.

Results

The data were aggregated in line with the procedures of Study 1a. Means and 95% confidence intervals are presented in Table 2.

Traditional analysis. In line with the traditional approach, we first investigated participants’ responses on moral dilemmas involving a proscriptive norm that prohibits action in cases where the benefits of action outweigh its costs to well-being. In the traditional approach, a preference for action over inaction on this type of dilemma would be interpreted as a preference for utilitarian over deontological responses. Averaged across the six dilemmas of this type, participants in the current

study showed a slight preference for inaction over action ($M = 2.88$, $SD = 1.37$). However, the overall preference score did not significantly differ from the neutral reference point of 3, $t(196) = 1.25$, $p = .213$, $d = 0.089$, which reflects an equal preference for action and inaction across the six dilemmas. Different from the results in Study 1a, men and women did not significantly differ in their preference for action versus inaction in this type of dilemma, $t(195) = 1.31$, $p = .191$, $d = 0.187$ (see Table 2).

PD analysis. PD scores were calculated in line with the procedures of Study 1a. Replicating the findings of Study 1a, women showed significantly higher scores than men on the D parameter, $t(195) = 3.93$, $p < .001$, $d = 0.560$ (see Table 3). Different from Study 1a, there was also a significant effect on the U parameter, such that women showed significantly higher U scores than men, $t(195) = 2.59$, $p = .010$, $d = 0.369$ (see Table 3).

CNI model. The CNI model fit the data well when the data were analyzed without considering participants’ gender, $G^2(1) = 0.85$, $p = .355$. Both the C parameter ($M = 0.197$, 95% CI [0.170, 0.224]) and the N parameter ($M = 0.329$, 95% CI [0.295, 0.363]) were significantly larger than zero, indicating that participants were highly sensitive to both consequences and norms in responding to the moral dilemmas, $\Delta G^2(1) = 199.16$, $p < .001$ for the C parameter and $\Delta G^2(1) = 347.62$, $p < .001$ for the N parameter. The I parameter did not significantly deviate from its neutral reference point of 0.5, $\Delta G^2(1) = 0.78$, $p = .377$, indicating an equal distribution of action and inaction responses in the total sample ($M = 0.489$, 95% CI [0.464, 0.514]).

The CNI model also fit the data well when parameter scores were estimated separately for men and women, $G^2(2) = 1.51$, $p = .469$. Replicating the findings of Study 1a, women showed significantly higher scores than men on the N parameter, $\Delta G^2(1) = 17.43$, $p < .001$, $d = 0.599$ (see Figure 3) and the I parameter, $\Delta G^2(1) = 9.12$, $p = .003$, $d = 0.428$ (see Figure 3). Moreover, there was a significant difference on the C parameter, indicating that women showed a stronger sensitivity to consequences than men, $\Delta G^2(1) = 6.43$, $p = .011$, $d = 0.364$ (see Figure 3). Together, these results suggest that women show (1) a stronger sensitivity to norms, (2) a stronger sensitivity to consequences, and (3) a stronger general preference for inaction compared to men.

Discussion

Addressing potential concerns about the reproducibility of psychological findings (Open Science Collaboration, 2015), Study 1b aimed to replicate the key findings of Study 1a. As in Study 1a, the CNI model fit the data well. Although there were no gender differences in the traditional analysis of moral dilemma responses, the CNI model replicated the findings of Study 1a showing (1) a stronger sensitivity to norms among women and (2) a stronger general preference for inaction among women. Different from Study 1a, there was also a

¹⁴ Two participants completed all measures, but did not submit a request for compensation.

significant effect on the *C* parameter, indicating that women showed a stronger sensitivity to consequences than men. A similar effect emerged in the PD analysis, which showed higher scores on both the *D* parameter and the *U* parameter among women compared men. However, because there were no effects on the *U* parameter of the PD model and the *C* parameter of the CNI model in Study 1a, we refrain from drawing any conclusions from this particular finding. Nevertheless, the current results provide further evidence for the validity of the CNI model in describing responses to the four types of moral dilemmas and for its ability to provide more nuanced insights into the nature of existing findings.

Study 2a

Study 2a aimed at providing deeper insights into the processes underlying moral dilemma judgments. A widespread assumption in moral psychology is that deontological judgments are the product of automatic emotional responses to the idea of causing harm; utilitarian judgments are assumed to result from cognitively effortful processes involving a deliberate analysis of costs and benefits and a suppression of automatic emotional responses favoring a deontological judgment (e.g., Greene et al., 2001, 2004, 2008). Consistent with the overarching hypothesis that utilitarian judgments are the product of effortful processes, Suter and Hertwig (2011) found that time pressure reduced participants' willingness to accept harmful action in the traditional trolley paradigm (see also Greene et al., 2008). Using the PD approach, Conway and Gawronski (2013) further showed that cognitive load selectively reduced scores on the *U* parameter, without affecting scores on the *D* parameter. Yet, as we explained earlier in this article, the two PD parameters are confounded with general preference for inaction versus action, which undermines the possibility of drawing conclusions about whether cognitive load influences moral dilemma judgments via (1) sensitivity to consequences, (2) sensitivity to moral norms, or (3) general preference for inaction versus action (or any combination of the three). Thus, to provide more nuanced insights into the effect of limited cognitive resources on moral dilemma judgments, Study 2a tested the effect of cognitive load on the three parameters of the CNI model.

Method

Participants. Participants were recruited for a study on "how people make moral judgments" via Amazon's MTurk. Participants received compensation of \$1.00 for completing the study. Eligibility for participation was limited to English native speakers who (1) had a HIT approval rate of at least 97% at the time of the study and (2) had not participated in prior studies from our lab using the same set of moral dilemmas. Participants were randomly assigned to either a *low load* or *high load* condition. Of the 242 MTurk workers who initially began

the study, 203 completed all measures.¹⁵ Of these participants, 9 failed to pass an instructional attention check. Data from these participants were excluded from the statistical analyses, leaving us with a final sample of 194 participants (97 women, 96 men, 1 missing; $M_{age} = 34.26$, $SD_{age} = 11.90$).

Procedure and materials. Participants were asked to read and respond to the 24 moral dilemmas of Study 1, using the same fixed random order. To manipulate the amount of residual cognitive capacity for the moral judgment task, participants in the two experimental conditions were asked to perform a secondary task while reading and responding to the dilemmas. Toward this end, participants were presented with unique meaningless digit strings before each of the 24 moral dilemmas (see Conway & Gawronski, 2013). Participants were asked to concentrate on the digit string, commit it to memory, and report it after they read and provided their answer to the dilemma. Participants in the *high load* condition were presented with a unique seven-digit string before each moral dilemma; participants in the *low load* condition were presented with a unique two-digit string before each moral dilemma. Each of the digit strings in the *high load* condition contained at least one uppercase letter, one lowercase letter, one number, and one punctuation mark (e.g., *n63#m1Q*). The digit strings in the *low load* condition included the first two digits of the seven-digit string in the *high load* condition for the same moral dilemma (e.g., *n6*).¹⁶ Digit strings and moral dilemmas were matched through a random procedure that was kept constant for all participants in the two conditions. After completion of the moral dilemmas, participants were asked to respond to the same one-item attention check used in the Pilot Study.

Results

The data were aggregated in line with the procedures of Study 1a. Means and 95% confidence intervals are presented in Table 2.

Manipulation check. To test the effectiveness of our experimental manipulation in differentially taxing participants' cognitive resources, we coded for each digit string of the cognitive load task whether participants correctly reproduced it after they provided their dilemma judgment. Responses were coded as *correct* when

¹⁵ Three participants completed all measures, but did not submit a request for compensation.

¹⁶ Different from the widespread use of a control condition without a secondary task, the *low load* condition in Study 2a was used to avoid a commonly neglected confound between cognitive resources and processing goals (see Gast, Gawronski, & De Houwer, 2012). That is, *high load* and *no load* conditions differ not only in terms of the residual amount of cognitive resources left for the primary task, but also in terms of participants' processing goals (i.e., goal of completing the primary task versus goal of simultaneously completing two different tasks). A *low load* control condition avoids this confound by keeping participants' processing goals equal across the two conditions (e.g., Yzerbyt, Coull, & Roucher, 1999).

participants correctly reproduced all digits of a given digit string including the formatting of upper and lower cases. Responses that deviated from the original digit string were coded as *incorrect*. We then generated accuracy scores for each participant by calculating the proportion of correct responses among the 24 items of the cognitive load task. Consistent with the assumption that the memorization task was more difficult, and thus more cognitively taxing, in the high load compared to the low load condition, participants in the high load condition showed a lower accuracy score than participants in the low load condition ($M_s = 0.51$ vs. 0.86), $t(192) = 10.88$, $p < .001$, $d = 1.607$.

Traditional analysis. In line with the traditional approach, we first investigated participants' responses on moral dilemmas involving a proscriptive norm that prohibits action in cases where the benefits of action outweigh its costs to well-being. In the traditional approach, a preference for action over inaction on this type of dilemma would be interpreted as a preference for utilitarian over deontological responses. Consistent with the results of previous studies, participants in the high load condition tended to show a weaker preference for action versus inaction on this type of dilemma than participants in the low load condition, but this difference was only marginally significant, $t(192) = 1.67$, $p = .097$, $d = 0.239$ (see Table 2).

PD analysis. PD scores were calculated in line with the procedures of Study 1a. There were no significant effects of cognitive load on the U parameter, $t(192) = 0.97$, $p = .335$, $d = 0.139$ (see Table 3), and the D parameter, $t(192) = 0.52$, $p = .604$, $d = 0.075$ (see Table 3).

CNI model. In the current study, the CNI model showed suboptimal fit, in that the probabilities predicted by the model showed a marginally significant deviation from the empirically observed probabilities, $G^2(2) = 4.98$, $p = .083$. However, with an effect size of Cohen's $w = 0.033$, this deviation was negligible (Cohen, 1988; see Footnote 6). The only significant effect on the three parameters was obtained for the I parameter, which showed a stronger general preference for inaction in the high load condition compared to the low load condition, $\Delta G^2(1) = 5.19$, $p = .023$, $d = 0.328$ (see Figure 4). There were no statistically significant effects of cognitive load on the N parameter, $\Delta G^2(1) = 0.01$, $p = .927$, $d = 0.013$ (see Figure 4), and the C parameter, $\Delta G^2(1) = 1.35$, $p = .245$, $d = 0.168$ (see Figure 4). Together, these results suggest that cognitive load influenced moral dilemma judgments by increasing participants' general preference for inaction. There seem to be no effects of cognitive load on the sensitivity to consequences and norms.

Discussion

The main goal of Study 2a was to provide deeper insights into the effect of limited cognitive resources on moral dilemma judgments. A widespread assumption in

moral psychology is that utilitarian judgments are the product of cognitively effortful processes, which is consistent with the results of earlier studies (e.g., Conway & Gawronski, 2013; Greene et al., 2008; Suter & Hertwig, 2011). To provide more nuanced insights into the role of cognitive resources in moral dilemma judgments, Study 2a tested the effect of cognitive load on (1) sensitivity to consequences, (2) sensitivity to moral norms, and (3) general preference for inaction versus action. Consistent with previous findings, participants in the high load condition, compared to participants in the low load condition, tended to be less willing to accept action on moral dilemmas where a proscriptive norm prohibits action and action leads to better consequences for a larger number of people. There were no significant effects of cognitive load on either of the two PD parameters. Further analyses with the CNI model suggest that cognitive load influenced moral dilemma judgments by increasing participants' general preference for inaction versus action. There were no significant effects of cognitive load on participants' sensitivity to consequences and norms.

Study 2b

To investigate the reliability of the obtained effects of cognitive load, Study 2b aimed to replicate the findings of Study 2a using the same procedure and materials.

Method

Participants. Participants were recruited for a study on "how people make moral judgments" via Amazon's MTurk. Participants received compensation of \$1.00 for completing the study. Eligibility for participation was limited to English native speakers who (1) had a HIT approval rate of at least 97% at the time of the study and (2) had not participated in prior studies from our lab using the same set of moral dilemmas. Participants were randomly assigned to either a *low load* or *high load* condition. Of the 253 MTurk workers who initially began the study, 204 completed all measures.¹⁷ Of these participants, 10 failed to pass an instructional attention check. Data from these participants were excluded from the statistical analyses, leaving us with a final sample of 194 participants (91 women, 103 men; $M_{age} = 36.35$, $SD_{age} = 12.40$).

Procedure and materials. The procedure and materials were identical to the ones in Study 2a.

Results

The data were aggregated in line with the procedures of Study 1a. Means and 95% confidence intervals are presented in Table 2.

¹⁷ Four participants completed all measures, but did not submit a request for compensation.

Manipulation check. To test the effectiveness of our experimental manipulation in differentially taxing cognitive resources, we aggregated participants' performance on the cognitive load task in line with the procedures of Study 2a. Consistent with the assumption that the memorization task was more difficult, and thus more cognitively taxing, in the high load compared to the low load condition, participants in the high load condition showed a lower accuracy score than participants in the low load condition ($M_s = 0.57$ vs. 0.87), $t(192) = 7.90$, $p < .001$, $d = 1.172$.

Traditional analysis. In line with Study 2a, we first investigated participants' responses on moral dilemmas involving a proscriptive norm that prohibits action in cases where the benefits of action outweigh its costs to well-being. Replicating the pattern of Study 2a, participants in the high load condition showed a weaker preference for action in this type of dilemma than participants in the low load condition. Although this effect was only marginally significant in Study 2a, it did reach the conventional level of statistical significance in Study 2b, $t(192) = 2.18$, $p = .030$, $d = 0.314$ (see Table 2).

PD analysis. PD scores were calculated in line with the procedures of Study 1a. Replicating the findings of Study 2a, there were no significant effects of cognitive load on the U parameter, $t(192) = 1.29$, $p = .199$, $d = 0.185$ (see Table 3), and the D parameter, $t(192) = 1.37$, $p = .173$, $d = 0.199$ (see Table 3).

CNI model. The CNI model fit the data well in the current study, $G^2(2) = 1.29$, $p = .524$. Moreover, replicating the central finding of Study 2a, there was a significant effect of cognitive load on the I parameter, such that participants showed a stronger general preference for inaction in the high load condition compared to the low load condition, $\Delta G^2(1) = 13.77$, $p < .001$, $d = 0.535$ (see Figure 4). There were no significant effects of cognitive load on the N parameter, $\Delta G^2(1) = 0.05$, $p = .826$, $d = 0.032$ (see Figure 5) and the C parameter, $\Delta G^2(1) = 2.08$, $p = .149$, $d = 0.209$ (see Figure 4). These results confirm our conclusion that cognitive load influences moral dilemma judgments by increasing participants' general preference for inaction. Yet, there seem to be no effects of cognitive load on participants' sensitivity to consequences and norms.

Discussion

The current study replicated the effect of cognitive load obtained in Study 2a: participants showed a stronger general preference for inaction under conditions of high load compared to low load, but there were no significant effects of cognitive load on participants' sensitivity to consequences and norms. A potential interpretation of this finding is that participants under high load feel that they do not have the capacity to make a well-informed decision. Thus, they may prefer not to engage in any action regardless of consequences and norms, because harm caused by action is typically perceived as more

severe than the same amount of harm caused by inaction (Cushman et al., 2006). Applied to extant theories of moral judgment (Greene et al., 2001, 2004, 2008), these results suggest that cognitive load influences moral dilemma judgments by enhancing the omission bias, not by reducing sensitivity to consequences in a utilitarian sense. Moreover, the current findings suggest that both consequences and moral norms influence moral judgments via efficient processes, in that neither of them was affected by cognitive load.

Study 3a

The main goal of Study 3a was to demonstrate the usefulness of the CNI model in providing deeper insights into framing effects on moral dilemma judgments. Drawing on the hypothesis that use of personal force enhances automatic emotional responses to the idea of causing harm, Greene et al. (2001) found that participants were less willing to accept harmful action in the traditional trolley paradigm when the described action involved personal contact (i.e., pushing a person from a bridge to stop the trolley from killing five people) than when it did not involve personal contact (i.e., switching a lever to redirect the trolley to a different track where it would kill only one person instead of five people). Interestingly, this framing effect emerged although the moral nature of the described action (i.e., killing one person) as well as the consequences of the action (i.e., saving five people) were exactly the same in the two conditions. Drawing on the traditional approach, this framing effect has been interpreted as evidence for the hypothesis that deontological judgments are the product of automatic emotional responses to the idea of causing harm, which should increase as a result of personal contact with the target of one's actions.

In the current study, we aimed to demonstrate the value of the CNI model in providing more nuanced insights into the nature of framing effects on moral dilemma judgments. Similar to Greene et al.'s (2001) approach, we manipulated the framing of our dilemma questions while keeping the moral nature of the described actions as well as their consequences identical across conditions. Yet, different from Greene et al.'s (2001) focus on the use of personal force, participants in the current study were asked to indicate for each dilemma either (1) if it is morally acceptable to perform the described action (i.e., moral judgment) or (2) if they would perform the described action (i.e., moral action). To the extent that personal involvement is greater for concrete decisions about personal actions compared to abstract judgments of moral acceptability, the two kinds of framings may lead to different patterns of moral dilemma judgments (see Pletti et al., in press). Using our CNI model, we were interested in whether such framing effects involve differences in the sensitivity to consequences, the sensitivity to moral norms, or the

general preference for inaction versus action (or any combination of the three).

Method

Participants. Participants were recruited for a study on “how people make moral judgments” via Amazon’s MTurk. Participants received compensation of \$1.00 for completing the study. Eligibility for participation was limited to English native speakers who (1) had a HIT approval rate of at least 97% at the time of the study and (2) had not participated in prior studies from our lab using the same set of moral dilemmas. Participants were randomly assigned to either a *moral judgment* or *moral action* condition. Of the 221 MTurk workers who initially began the study, 202 completed all measures.¹⁸ Of these participants, 16 failed to pass an instructional attention check. Data from these participants were excluded from the statistical analyses, leaving us with a final sample of 186 participants (100 women, 86 men; $M_{age} = 35.77$, $SD_{age} = 12.79$).

Procedure and materials. Participants were asked to read and respond to the 24 moral dilemmas of Study 1a, using the same fixed random order. Participants in the *moral judgment* condition were asked to indicate for each dilemma if it is morally acceptable to perform the described action, using the same question wording as in Study 1a. Participants in the *moral action* condition were asked to indicate for each dilemma if they would perform the action described in the scenario. After completion of the moral dilemmas, participants were asked to respond to the same one-item attention check used in the Pilot Study.

Results

The data were aggregated in line with the procedures of Study 1a. Means and 95% confidence intervals are presented in Table 2.

Traditional analysis. We first investigated participants’ responses on moral dilemmas involving a proscriptive norm that prohibits action in cases where the benefits of action outweigh its costs to well-being. In the traditional approach, a preference for action over inaction on this type of dilemma would be interpreted as a preference for utilitarian over deontological responses. There was no statistically significant effect of question framing on judgments in this type of moral dilemma, $t(184) = 1.59$, $p = .114$, $d = 0.233$ (see Table 2).

PD analysis. PD scores were calculated in line with the procedures of Study 1a. There were no significant effects of question framing on the U parameter, $t(184) = 0.76$, $p = .448$, $d = 0.111$ (see Table 3), or the D parameter, $t(184) = 1.32$, $p = .190$, $d = 0.193$ (see Table 3).

CNI model. In the current study, the CNI model did not fit the data, in that the response patterns predicted by the model significantly deviated from the observed response patterns in the data, $G^2(2) = 11.93$, $p = .003$. However, the effect size of this deviation was rather small, Cohen’s $w = 0.052$, indicating negligible misfit when controlling for statistical power (Cohen, 1988; see Footnote 6). Further analyses revealed that there was no statistically significant effect of question framing on the C parameter, $\Delta G^2(1) = 2.44$, $p = .118$, $d = 0.230$ (see Figure 6). The N parameter revealed a marginally significant effect, such that participants tended to show a weaker sensitivity to moral norms in the moral action condition than the moral judgment condition, $\Delta G^2(1) = 3.31$, $p = .069$, $d = 0.268$ (see Figure 6). Moreover, there was a significant effect on the I parameter, showing a stronger general preference for inaction in the moral action condition compared to the moral judgment condition, $\Delta G^2(1) = 35.18$, $p < .001$, $d = 0.713$ (see Figure 6). These results suggest that a focus on moral action (as opposed to moral acceptability) increases general preference for inaction and reduces sensitivity to moral norms. There seems to be no framing effect on sensitivity to consequences.

Discussion

The main goal of Study 3a was to demonstrate the usefulness of the CNI model in providing deeper insights into framing effects on moral judgments. Toward this end, participants were asked to indicate for each dilemma either (1) if it is morally acceptable to perform the described action (i.e., moral judgment) or (2) if they would perform the described action (i.e., moral action). Interestingly, there was no significant effect of question framing on moral dilemma judgments when the data were analyzed using the traditional approach. There were also no significant effects of question framing on the two parameters of Conway and Gawronski’s (2013) PD approach. Yet, analyses with our CNI model revealed that a focus on moral action (as opposed to moral acceptability) had two simultaneous effects that cancelled each other out in the traditional approach as well as in the PD approach. First, participants in the moral action condition showed a stronger general preference for inaction compared to participants in the moral judgment condition. Second, participants in the moral action condition tended to show a weaker sensitivity to moral norms compared to participants in the moral judgment condition. These effects cannot be detected with the traditional and the PD approach, because both approaches focus exclusively on moral dilemmas involving proscriptive norms. In the traditional approach, the two effects cancel each other out, because enhanced general preference for inaction supports inaction and reduced sensitivity to moral norms supports action. Similarly, in the PD approach, the two effects cancel each other out, because enhanced general preference for inaction

¹⁸ Two participants completed all measures, but did not submit a request for compensation.

increases scores on the *D* parameter, whereas reduced sensitivity to moral norms decreases scores on the *D* parameter. These confounds are disentangled in the CNI model, which provides separate parameter estimates for sensitivity to consequences, sensitivity to moral norms, and general preference for inaction versus action. However, because the CNI model failed to show adequate fit in the current study and the obtained effect on the *N* parameter was only marginally significant, the obtained effects should be treated with caution in the absence of additional data. The main goal of Study 3b was to address this limitation.

Study 3b

To investigate the reliability of the obtained effects of question framing, Study 3b aimed to replicate the findings of Study 3a using the same procedure and materials.

Method

Participants. Participants were recruited for a study on “how people make moral judgments” via Amazon’s MTurk. Participants received compensation of \$1.00 for completing the study. Eligibility for participation was limited to English native speakers who (1) had a HIT approval rate of at least 97% at the time of the study and (2) had not participated in prior studies from our lab using the same set of moral dilemmas. Participants were randomly assigned to either a *moral judgment* or *moral action* condition. Of the 225 MTurk workers who initially began the study, 200 completed all measures. Of these participants, 11 failed to pass an instructional attention check. Data from these participants were excluded from the statistical analyses, leaving us with a final sample of 189 participants (98 women, 91 men; $M_{age} = 34.72$, $SD_{age} = 10.69$).

Procedure and materials. The procedure and materials were identical to the ones in Study 3a.

Results

The data were aggregated in line with the procedures of Study 1a. Means and 95% confidence intervals are presented in Table 2.

Traditional analysis. We first investigated participants’ responses on moral dilemmas involving a proscriptive norm that prohibits action in cases where the benefits of action outweigh its costs to well-being. Replicating the results of Study 3a, there was no statistically significant effect of question framing on this type of moral dilemma, $t(187) = 1.57$, $p = .118$, $d = 0.229$ (see Table 2).

PD analysis. PD scores were calculated in line with the procedures of Study 1a. Replicating the findings of Study 3a, there were no significant effects of question framing on the *U* parameter, $t(187) = 0.69$, $p = .493$, $d = 0.100$ (see Table 3), and the *D* parameter, $t(187) = 0.68$, $p = .497$, $d = 0.099$ (see Table 3).

CNI model. The CNI model fit the data well, $G^2(2) = 4.19$, $p = .123$. Replicating the main finding of Study 3a, there was a significant effect on the *I* parameter, which showed a stronger general preference for inaction in the moral action condition compared to the moral judgment condition, $\Delta G^2(1) = 29.50$, $p < .001$, $d = 0.799$ (see Figure 7). This time, the obtained difference on the *N* parameter did reach statistical significance, indicating a weaker sensitivity to moral norms in the moral action condition than the moral judgment condition, $\Delta G^2(1) = 6.15$, $p = .013$, $d = 0.363$ (see Figure 7). There was no significant effect on the *C* parameter, $\Delta G^2(1) = 0.09$, $p = .767$, $d = 0.043$ (see Figure 7). These results support our conclusion that a focus on moral action (as opposed to moral acceptability) increases general preference for inaction and reduces sensitivity to moral norms. There seems to be no framing effect on sensitivity to consequences.

Discussion

Study 3b replicated the obtained effects of question framing in Study 3a. First, participants in the moral action condition showed a stronger general preference for inaction compared to participants in the moral judgment condition. Second, participants in the moral action condition tended to show a weaker sensitivity to moral norms compared to participants in the moral judgment condition. These effects cannot be detected with the traditional and the PD approach, because both approaches exclusively focus on moral dilemmas involving proscriptive norms. In the traditional approach, the two effects cancel each other out, because enhanced general preference for inaction supports inaction and reduced sensitivity to moral norms supports action. Similarly, in the PD approach, the two effects cancel each other out, because enhanced general preference for inaction increases scores on the *D* parameter, whereas reduced sensitivity to moral norms decreases scores on the *D* parameter. Together, the results of Study 3a and 3b demonstrate the usefulness of the CNI model in providing deeper insights into framing effects on moral dilemma judgments by uncovering effects on moral judgments that cannot be detected with existing approaches.

Study 4a

In Study 4a, we aimed to demonstrate the usefulness of the CNI model in providing deeper insights into previous findings that may be deemed counterintuitive. Several studies using the traditional approach have shown that participants with sub-clinical levels of psychopathy are more willing to accept harmful action in the traditional dilemma paradigm compared to non-psychopathic participants (e.g., Bartels & Pizarro, 2011; Kahane et al., 2015; Patil, 2015; Pletti et al., in press). In terms of the traditional interpretation, this finding would suggest that participants with high levels of psychopathy show a stronger preference for utilitarian over

deontological judgments compared to participants with low levels of psychopathy. However, it seems highly implausible that these results reflect a greater sensitivity to morally relevant consequences among psychopaths. Instead, it seems more likely that psychopaths are willing to accept harmful actions regardless of their consequences. The traditional approach is unable to capture the difference between the two cases, because it does not include manipulations of consequences and norms as the defining features of utilitarian and deontological responding.

The main goal of Study 4a was to demonstrate the value of the CNI model in providing deeper insights into the effects of (sub-clinical) psychopathy on moral dilemma judgments. Toward this end, we asked a broad sample of participants to complete a measure of psychopathy in a first session. Based on their responses, we identified those participants whose psychopathy scores fell into either the highest or the lowest quartile of the sample, and invited them to complete our battery of moral dilemmas in a second session. Using the CNI model, we then tested whether participants with high versus low scores on the psychopathy measure differ in terms of their sensitivity to consequences, their sensitivity to moral norms, and their general preference for inaction versus action. Because men and women differ in terms of their responses to moral dilemmas (see Studies 1a and 1b) and because psychopathy is more prevalent among men than women (see Cale & Lilienfeld, 2002), we limited participation to men in order to avoid potential confounds between psychopathy and gender.

Method

Participants. Participants were recruited for a two-session study on “personality and moral judgments” via Amazon’s MTurk. The first session included the psychopathy measure; the second session included the moral dilemmas. Participants received compensation of \$0.30 for completing the first session and \$1.00 for completing the second session. Eligibility for participation was limited to English native speakers who (1) had a HIT approval rate of at least 97% at the time of the study and (2) had not participated in prior studies from our lab using the same set of moral dilemmas. As another eligibility criterion, the advertisement on the MTurk website noted that participation in the study is restricted to male adults. Of the 522 MTurk workers who initially began the first session, 503 completed all measures.¹⁹ Of these participants, 47 reported being female and 2 reported being neither male nor female, leaving us with a sample of 454 male participants. Based on their scores on the psychopathy measure, we identified 121 participants with scores in the lowest quartile and 122 participants with scores in the highest quartile. For

one of these participants, the MTurk ID provided in the first session did not match any MTurk IDs in Amazon’s data base, leaving us with a sample of 242 participants who were invited for the second session. Approximately two weeks after completion of the first session, these participants were sent a follow-up email through MTurk that included an invitation to complete the second part of the study. Participation in the second session was restricted to the 242 participants who were invited to complete the second part. Of these participants, 196 completed the moral dilemmas within our predetermined time window of two weeks. Four additional participants started the study, but did not complete it. Of the 196 participants who completed all measures, 12 failed to pass an instructional attention check, leaving us with a final sample of 184 participants. Eighty-nine of these participants had psychopathy scores in the highest quartile; 95 had psychopathy scores in the lowest quartile.

Procedure and materials. Psychopathy was measured with Paulhus, Neumann, and Hare’s (2009) 30-item SRP-III scale, which was administered in the first part of our two-session study. Responses were measured with 7-point scales ranging from 1 (*strongly disagree*) to 7 (*strongly agree*), which showed high internal consistency in the full sample at Time 1 (Cronbach’s $\alpha = .94$). In the second part, participants were asked to read and respond to the 24 moral dilemmas from Study 1a, using the same fixed random order. Based on previous research showing differences between psychopaths and non-psychopaths for action decisions, but not for judgments of morality (Pletti et al., in press), the current study used the moral action framing of Studies 3a and 3b. After completion of the moral dilemmas in the second session, participants were asked to complete the same instructional attention check as in the Pilot Study.

Results

Moral dilemma responses were aggregated in line with the procedures of Study 1a. The cut-off for the identification of participants with low levels of psychopathy on the SRP-III was 2.7; the cut-off for the identification of participants with high levels of psychopathy was 3.8. Means and 95% confidence intervals of moral dilemma responses for the two groups are presented in Table 2.

Traditional analysis. In line with the traditional approach, we first investigated participants’ responses on moral dilemmas involving a proscriptive norm that prohibits action in cases where the benefits of action outweigh its costs to well-being. In the traditional approach, a preference for action over inaction on this type of dilemma would be interpreted as a preference for utilitarian over deontological responses. Although participants with high levels of psychopathy showed a slightly stronger preference for action over inaction in this type of dilemma than participants low levels of

¹⁹ Three participants completed all items of the psychopathy measure, but did not submit a request for compensation.

psychopathy, this difference did not reach statistical significance, $t(182) = 1.28, p = .202, d = 0.191$ (see Table 2).

PD analysis. PD scores were calculated in line with the procedures of Study 1a. Analyses revealed a marginally significant effect of psychopathy on the *D* parameter, indicating that participants high in psychopathy tended to show lower *D* scores than participants low in psychopathy, $t(182) = 1.67, p = .096, d = 0.247$ (see Table 3). There was no significant effect of psychopathy on the *U* parameter, $t(182) = 1.29, p = .200, d = 0.190$ (see Table 3).

CNI model. The CNI model fit the data well, $G^2(2) = 0.29, p = .864$. Participants with high levels of psychopathy showed significantly lower scores on the *N* parameter than participants with low levels of psychopathy, $\Delta G^2(1) = 12.35, p < .001, d = 0.521$ (see Figure 8). Moreover, there were marginally significant effects on the *C* parameter, $\Delta G^2(1) = 2.77, p = .096, d = 0.247$, and the *I* parameter, $\Delta G^2(1) = 3.15, p = .076, d = 0.262$ (see Figure 8). Specifically, participants with high levels of psychopathy tended to show a lower sensitivity to consequences and a weaker general preference for inaction than participants with low levels of psychopathy. Together, these results suggest that psychopaths show (1) a weaker sensitivity to moral norms, (2) a weaker sensitivity to consequences, and (3) a weaker general preference for inaction compared to non-psychopaths.

Discussion

Study 4a demonstrates the value of the CNI model in providing deeper insights into counterintuitive findings obtained with the traditional approach. Previous research using the traditional approach has shown that participants with sub-clinical levels of psychopathy are more willing to accept harmful action in the trolley paradigm compared to non-psychopathic participants (e.g., Bartels & Pizarro, 2011; Kahane et al., 2015; Patil, 2015; Pletti et al., in press). Although we found a similar pattern in the current study, there was no significant effect of psychopathy in the traditional analysis. However, counter to the conclusion that psychopathy did not have any effect on moral dilemma judgments, analyses with the CNI model suggest that participants with high levels of psychopathy showed a weaker sensitivity to moral norms compared to participants with low levels of psychopathy. There were also marginally significant differences on the other two parameters, suggesting a weaker sensitivity to consequences and a weaker tendency for general inaction among participants with high levels of psychopathy. Together, these findings reconcile the counterintuitive nature of previous findings, which seem to suggest that psychopaths show a stronger preference for utilitarian over deontological judgments compared to non-psychopaths (e.g., Bartels & Pizarro, 2011; Kahane et al., 2015; Patil, 2015; Pletti et al., in press). Counter to this conclusion, the current findings suggest that psychopaths

show (1) a weaker sensitivity to consequences, (2) a weaker sensitivity to moral norms, and (3) a weaker general preference for inaction compared to non-psychopaths.

Study 4b

Study 4b aimed to replicate the main findings of Study 4a to test the reproducibility of the obtained effects (Open Science Collaboration, 2015). A successful replication seems particularly important for the current study, because two of the three effects in Study 4a were only marginally significant. In addition, we aimed to replicate the findings of Study 4a with a different measure of psychopathy in order to provide converging evidence for the generality of our findings.

Method

Participants. Participants were recruited for a two-session study on “personality and moral judgments” via Amazon’s MTurk. The first session included the psychopathy measure; the second session included the moral dilemmas. Participants received compensation of \$0.30 for completing the first session and \$1.00 for completing the second session. Eligibility for participation was limited to English native speakers who (1) had a HIT approval rate of at least 97% at the time of the study and (2) had not participated in prior studies from our lab using the same set of moral dilemmas. As another eligibility criterion, the advertisement on the MTurk website noted that participation in the study is restricted to male adults. Of the 555 MTurk workers who initially began the first session, 504 completed all measures.²⁰ All of these participants reported being male.²¹ Based on their scores on the psychopathy measure, we identified 138 participants with scores in the lowest quartile and 139 participants with scores in the highest quartile, leaving a sample of 277 participants for the second session. Approximately two weeks after completion of the first session, these participants were sent a follow-up email through MTurk that included an invitation to complete the second part of the study. Participation in the second session was restricted to the 277 participants who were invited to complete the second part. Of these participants, 209 completed the moral dilemmas within our predetermined time window of two weeks. Seven additional participants started the study, but did not complete it. Of the 209 participants who completed all measures, 11 failed to pass an instructional attention check, leaving us with a final sample of 198

²⁰ Four participants completed all items of the psychopathy measure, but did not submit a request for compensation.

²¹ Because a considerable proportion of participants in the first session of Study 4a reported being female, we added an additional restriction, such that (1) participants had to report their gender before completing the psychopathy measure, (2) participants were told that they are ineligible for participation in this study if they did not report being male, and (3) the study was terminated if they reported being female.

participants. Eighty-nine of these participants had psychopathy scores in the highest quartile; 109 had psychopathy scores in the lowest quartile.

Procedure and materials. Psychopathy was measured with Levenson, Kiehl, and Fitzpatrick's (1995) 16-item Primary Psychopathy Scale (PPS), which was administered in the first part of our two-session study. Responses were measured with 4-point scales using the response options (1) *disagree strongly* (2) *disagree somewhat*, (3) *agree somewhat*, and (4) *agree strongly*, which showed high internal consistency in the full sample at Time 1 (Cronbach's $\alpha = .90$). In the second part, participants were asked to read and respond to our battery of 24 moral dilemmas, using the same fixed random order and the action framing of Studies 3a and 3b. After completion of the moral dilemmas in the second session, participants were asked to complete the same instructional attention check as in the Pilot Study.

Results

Moral dilemma responses were aggregated in line with the procedures of Study 1a. The cut-off for the identification of participants with low levels of psychopathy on the PPS was 1.51; the cut-off for the identification of participants with high levels of psychopathy was 2.30. Means and 95% confidence intervals of moral dilemma responses for the two groups are presented in Table 2.

Traditional analysis. In line with the traditional approach, we first investigated participants' responses on moral dilemmas involving a proscriptive norm that prohibits action in cases where the benefits of action outweigh its costs to well-being. Replicating the pattern found in previous studies, participants with high levels of psychopathy showed a significantly stronger preference for action over inaction in this type of dilemma than participants low levels of psychopathy, $t(196) = 2.70$, $p = .008$, $d = 0.384$ (see Table 2).

PD analysis. PD scores were calculated in line with the procedures of Study 1a.²² Analyses revealed that participants low in psychopathy showed significantly higher scores than participants high in psychopathy on both the D parameter, $t(196) = 3.90$, $p < .001$, $d = 0.855$ (see Table 3), as well as the U parameter, $t(196) = 5.97$, $p < .001$, $d = 0.554$ (see Table 3).

CNI model. The CNI model fit the data well, $G^2(2) = 0.18$, $p = .916$. Replicating the findings of Study 4a, participants with high levels of psychopathy showed significantly lower scores on the N parameter than participants with low levels of psychopathy, $\Delta G^2(1) = 111.80$, $p < .001$, $d = 1.48$ (see Figure 9). In addition, participants with high levels of psychopathy showed lower scores on the C parameter, $\Delta G^2(1) = 23.13$, $p <$

$.001$, $d = 0.695$, and the I parameter, $\Delta G^2(1) = 8.90$, $p = .003$, $d = 0.406$, compared to participants with low levels of psychopathy (see Figure 9). Together, these findings corroborate the conclusion that psychopaths show (1) a weaker sensitivity to consequences, (2) a weaker sensitivity to moral norms, and (3) a weaker general preference for inaction compared to non-psychopaths.

Discussion

Study 4b replicated the main findings of Study 4a with a different measure of psychopathy. In line with the pattern obtained in Study 4a, participants with high levels of psychopathy showed a lower sensitivity to moral norms on the N parameter, a lower sensitivity to consequences on the C parameter, and a weaker preference for inaction on I parameter compared to participants with low levels of psychopathy. Although the obtained differences on the C parameter and the I parameter were only marginally significant in Study 4a, the differences on all three parameters were statistically significant in the current study. These findings shed new light on earlier research showing that participants with sub-clinical levels of psychopathy are more willing to accept harmful action in the trolley paradigm compared to non-psychopathic participants (e.g., Bartels & Pizzaro, 2011; Kahane et al., 2015; Patil, 2015; Pletti et al., in press). In terms of the traditional interpretation of moral dilemma responses, this finding would suggest that psychopaths show a stronger preference for utilitarian over deontological judgments compared to non-psychopaths. The CNI model provides a more nuanced understanding of this finding, suggesting that psychopaths, compared to non-psychopaths, are (1) less sensitive to morally relevant consequences of their actions, (2) less sensitive to proscriptive and prescriptive norms, and (3) and less reluctant to engage in action irrespective of consequences and norms.

General Discussion

The distinction between utilitarianism and deontology has become a prevailing framework for conceptualizing moral judgment. According to the principle of utilitarianism, the moral status of a behavioral option depends on its consequences (*consequentialist morality*); the principle of deontology states that the moral status of a behavioral option depends on its consistency with moral norms (*rule-based morality*). To identify the processes underlying utilitarian and deontological judgments, research has investigated responses to moral dilemmas that pit one principle against the other (e.g., trolley problem). In the current article, we argued that the conceptual meaning of responses in this paradigm is ambiguous, because the defining aspects of utilitarianism and deontology, consequences and norms, are not manipulated (see Gawronski & Beer, in press). Conceptually, utilitarian judgments are defined by the property of being sensitive

²² One participant showed a U score of 1.00, which prevented the calculation a D score for this participant, because it would require a division by zero (see Conway & Gawronski, 2013). Data from this participant were excluded from the PD analysis.

to morally relevant consequences; deontological judgments are defined by the property of being sensitive to moral norms. Thus, to categorize a given judgment as utilitarian or deontological, it is essential to confirm their defining properties, which requires experimental manipulations of consequences and norms. Without such manipulations, theoretical interpretations of responses in the traditional paradigm remain ambiguous and prone to inadequate conclusions about the psychological processes underlying moral judgments.

To overcome this limitation, we proposed an alternative approach in which utilitarian judgments are inferred from their sensitivity to morally relevant consequences and deontological judgments are inferred from their sensitivity to moral norms (Gawronski & Beer, in press). Expanding on this approach, we presented a multinomial model that allows researchers to quantify sensitivity to consequences (*C*), sensitivity to moral norms (*N*), and general preference for inaction versus action irrespective of consequences and norms (*I*) in responses to moral dilemmas. To illustrate the more nuanced insights that can be gained from our CNI model, we presented 8 studies that used this model to investigate the effects of gender, cognitive load, question framing, and psychopathy on moral judgments.

Expanding on the finding that men, compared to women, show a stronger tendency to accept harmful action in the traditional trolley paradigm (e.g., Arutyunova et al., 2016; Friesdorf et al., 2015; Fumagalli et al., 2010), Studies 1a and 1b showed that this gender difference is rooted in a stronger sensitivity to norms and a stronger general preference for inaction among women. There was no evidence for gender differences in the sensitivity to consequences.

Addressing the widespread assumption that utilitarian judgments result from effortful cognitive processes (e.g., Conway & Gawronski, 2013; Greene et al., 2008; Suter & Hertwig, 2011), Studies 2a and 2b investigated the effect of low versus high cognitive load on moral dilemma judgments using our CNI model. The only significant effect in these studies was a significant increase in participants' general preference for inaction as a result of cognitive load. Cognitive load did not affect participants' sensitivity to morally relevant consequences. There was also no effect on participants' sensitivity to moral norms. A potential interpretation of this finding is that participants under high cognitive load feel that they do not have the cognitive capacity to make a well-informed decision. As a result, they may prefer not to engage in any action, given that harm caused by action is typically perceived as more severe than the same amount of harm caused by inaction (i.e., omission bias; see Cushman et al., 2006). From this perspective, cognitive load influences moral dilemma judgments by enhancing the omission bias, not by reducing sensitivity to consequences in a utilitarian sense.

Studies 3a and 3b demonstrated the usefulness of the CNI model in providing deeper insights into framing effects on moral dilemma judgments. Expanding on earlier evidence for framing effects on moral dilemma judgments (e.g., Greene et al., 2001; Pletti et al., in press), we manipulated the framing of our dilemma questions while keeping the moral nature of the described actions as well as their consequences identical across conditions. Toward this end, participants were asked to indicate either (1) if it is morally acceptable to perform the described action (i.e., moral judgment) or (2) if they would perform the described action (i.e., moral action). Results showed a stronger general preference for inaction and a weaker sensitivity to moral norms for participants in the moral action condition compared to participants in the moral judgment condition. To the extent that personal involvement is greater for action decisions compared to judgments of moral acceptability, our findings suggest that personal involvement may reduce sensitivity to moral norms and increase general preference for inaction.

Finally, Studies 4a and 4b helped to reconcile some counterintuitive findings obtained with the traditional paradigm, focusing particularly on the role of psychopathy. A common finding in the moral dilemma literature is that participants with sub-clinical levels of psychopathy show a stronger willingness to accept harmful action in the traditional dilemma paradigm compared to non-psychopaths (e.g., Bartels & Pizarro, 2011; Kahane et al., 2015; Patil, 2015; Pletti et al., in press). In terms of the traditional approach, this finding would suggest the counterintuitive conclusion that psychopaths are more concerned with maximizing well-being than non-psychopaths. The current findings resolve this paradox, showing that the difference between psychopaths and non-psychopaths is primarily driven by a weaker sensitivity to moral norms among psychopaths. Additionally, our findings suggest that psychopaths are less (not more) sensitive to morally relevant consequences than non-psychopaths. In the current research, participants with high levels of psychopathy also showed a weaker general preference for inaction than participants with low levels of psychopathy, in that participants with high levels of psychopathy were less reluctant to engage in action regardless of consequences and norms.

Benefits of the CNI Model

A major benefit of the CNI model is that it is more sensitive in identifying influences on moral dilemma judgments compared to earlier approaches. Because the traditional approach conflates sensitivity to consequences, sensitivity to norms, and general preference for inaction in a single outcome measure, it is unable to detect actually existing effects of a given factor when this factor influences the three components in a manner that compensates effects in the traditional approach. An illustrative example is provided by Studies

3a and 3b, which did not show any effect of question framing when moral dilemma judgments were analyzed with the traditional approach. Yet, counter to the conclusion that moral dilemma judgments were unaffected by the question framing, results obtained with the CNI model suggest that a focus on personal action had two simultaneous effects that cancelled each other out in the traditional approach. Specifically, a focus on personal action, compared to a focus on moral acceptability, decreased sensitivity to moral norms and increased general preference for inaction. Within the traditional approach, the former leads to acceptance of harmful action whereas the latter leads to rejection of harmful action. These simultaneous effects cancel each other out, making the traditional approach insensitive to actually existing effects of question framing.

Similar concerns apply to Conway and Gawronski's (2013) PD approach. Although the PD model acknowledges the conceptual and empirical independence of utilitarian and deontological inclinations in moral dilemma judgments, the two PD parameters are confounded with general preference for inaction versus action. Whereas the *U* parameter confounds sensitivity to consequences with general preference for action, the *D* parameter confounds sensitivity to moral norms with general preference for inaction (see Gawronski et al., 2016; Hütter & Klauer, 2016). As a result of these confounds, the PD approach is unable to detect actually existing effects of a given factor when this factor influences moral judgments in a manner that has compensatory effects on either of the two parameters. For example, in Studies 3a and 3b, the PD model did not show any effects of question framing. Yet, results obtained with the CNI model suggest that a focus on personal action, compared to a focus on moral acceptability, decreased sensitivity to moral norms and increased general preference for inaction. Whereas the former leads to lower scores on the *D* parameter, the latter leads to higher scores on the *D* parameter. These effects cancel each other out, making the PD approach insensitive to actually existing effects of question framing.

Another benefit of the more fine-grained approach of the CNI model is that it increases the reproducibility of empirical findings. Because the traditional and the PD approach conflate different sources of variance in moral dilemma judgments, their measurement outcomes tend to be noisier compared to the more differentiated outcomes of the CNI model. Such measurement noise can contribute to replication failures in at least two ways. First, noisy measurement can contribute to inflated estimates of effect sizes in initial demonstrations of an effect, which can undermine the success of replication studies that use these inflated estimates in power analyses for appropriate sample sizes (Loken & Gelman, 2017). Second, multiple sources of measurement variance can contribute to replication failures when the effect sizes of

simultaneous influences of a given factor vary across studies as a result of sampling error (see Cumming, 2014; Stanley & Spence, 2014). For example, in Studies 4a and 4b, analyses with the CNI model revealed replicable effects of psychopathy on sensitivity to consequences, sensitivity to moral norms, and general preference for inaction. Yet, analyses with the traditional approach replicated previous findings on the effect of psychopathy only in Study 4b, but not in Study 4a. Similarly, in Studies 2a and 2b, analyses with the CNI model revealed a replicable effect of cognitive load on general preference for inaction. Yet, counter to earlier findings by Conway and Gawronski (2013), neither of the two studies revealed a significant effect of cognitive load when moral dilemma judgments were analyzed with the PD approach. Thus, in addition to providing a more nuanced approach to studying the determinants of moral dilemma judgments, the CNI model contributes to the reproducibility of empirical findings by (1) resolving some of the problems that can undermine successful replications of actually existing effects, and thereby (2) enhancing researchers' ability to distinguish between false positives and actually existing effects (see Maxwell, Lau, & Howard, 2015).²³

Theoretical Implications

An important question concerns the implications of the current work for extant theories of moral dilemma judgment. One of the most influential theories in this area is Greene's dual-process model of moral judgment (Greene et al., 2001, 2004, 2008), which suggests that deontological judgments result from automatic emotional reactions to the idea of causing harm. In contrast, utilitarian judgments are assumed to result from cognitively effortful processes involving a deliberate analysis of costs and benefits and a suppression of automatic emotional responses favoring a deontological judgment.

When assessing the implications of the current work for the dual-process model of moral judgment, it is important to note that the dual-process model is based on the premise that a single judgment can be identified as either utilitarian or deontological by using moral dilemmas that pit one principle against the other (e.g., trolley problem). In this conceptualization, the moral

²³ An illustrative example are two studies reported in the Supplementary Materials, showing inconsistent effects of a manipulation designed to influence salience of harm. This manipulation was based on previous research by Conway and Gawronski (2013) aimed to manipulate salience of harm through pictures of the focal targets. Although the CNI model fit the data well in both studies, the effects on the three parameters were inconsistent across the two studies. The two studies also produced inconsistent effects with the traditional approach and the PD approach, both of which failed to replicate earlier findings by Conway and Gawronski (2013). Based on the inconsistency of findings across studies and data analytic approaches, we interpret these replication failures as evidence for the low reliability of Conway and Gawronski's (2013) picture manipulation.

nature of a given judgment is inferred from its mere consistency with either the utilitarian or the deontological principle (see Greene, 2007). Thus, it is *sufficient* for a judgment to qualify as utilitarian if it supports action in dilemmas where a proscriptive norm prohibits action and the benefits of action are greater than the costs. Conversely, it is *sufficient* for a judgment to qualify as deontological if it supports inaction in dilemmas where a proscriptive norm prohibits action and the benefits of action are greater than the costs. This conceptualization is different from the one proposed in the current work, which treats these criteria as *necessary*, but *insufficient*, for the identification of utilitarian and deontological judgments. According to our conceptualization, a given judgment cannot be categorized as utilitarian without confirming its property of being sensitive to consequences, which requires a comparison of judgments across dilemmas with different consequences. Similarly, a given judgment cannot be categorized as deontological without confirming its property of being sensitive to moral norms, which requires a comparison of judgments across dilemmas with different moral norms. Moreover, either of the two kinds of judgments have to be distinguished from a general preference for inaction regardless of consequences and moral norms.

The implications of the two conceptualizations for the dual-process model can be illustrated with the results of Studies 2a and 2b. In these studies, cognitive load reduced participants' acceptance of action in moral dilemmas where a proscriptive norm prohibited action and action led to better consequences for a larger number of people. From the perspective of the traditional approach, these findings reflect a reduced tendency for utilitarian judgments under cognitive load, which is consistent with the dual-process hypothesis that utilitarian judgments are the result of cognitive effortful processes. Yet, analyses with the CNI model suggested a different conclusion. Instead of reducing participants' sensitivity to consequences in a utilitarian sense, cognitive load increased participants' general preference for inaction. A potential interpretation of this finding is that participants under cognitive load feel that they do not have the capacity to make a well-informed decision. Thus, they may prefer not to engage in any action regardless of consequences and norms, because harm caused by action is typically perceived as more severe than the same amount of harm caused by inaction (Cushman et al., 2006). From this perspective, the current findings suggest that cognitive load influences moral dilemma judgments by enhancing the omission bias, not by reducing sensitivity to consequences in a utilitarian sense.

A potential way to reconcile this conclusion with the dual-process model is to interpret general preference for inaction as an instance of deontological responding. In line with this idea, the *doctrine of doing and allowing* (DDA) states that actively causing harm is morally worse than merely allowing harm, which is consistent with the

finding that harm caused by action is perceived as worse than equivalent harm caused by inaction (cf. Cushman, Knobe, & Sinnott-Armstrong, 2008; Cushman, Murray, Gordon-McKeon, Wharton, & Greene, 2012; Cushman et al., 2006). Conceptually, the DDA can be regarded as a deontological principle in the sense that the moral status of a behavioral option depends on its consistency with a general rule. From this perspective, the findings of Studies 2a and 2b could be reinterpreted in a manner that is consistent with the dual-process model. Specifically, the absence of a cognitive load effect on the *C* parameter suggests that cognitive load does not necessarily interfere with the deliberate analyses of costs and benefits. Instead, the obtained effect of cognitive load on the *I* parameter suggests that cognitive load reduces utilitarian judgments by interfering with the effortful suppression of automatic emotional responses favoring a deontological judgment.

Although a reinterpretation of the *I* parameter as an instance of deontological responding reconciles the dual-process model with the findings of the CNI model in Studies 2a and 2b, we deem such a reinterpretation problematic for at least three reasons. First, equating a general preference for inaction with deontological responding conflates sensitivity to moral norms with a general preference for inaction irrespective of moral norms. Yet, the two ways of responding to moral dilemmas are fundamentally different, in that general preference for inaction can lead to judgments that are either congruent or incongruent with the judgments suggested by moral norms. In dilemmas involving a proscriptive norm, both general preference for inaction and sensitivity to moral norms lead to inaction. Yet, in dilemmas involving a prescriptive norm, general preference for inaction leads to inaction whereas sensitivity to moral norms leads to action. Thus, if general preference for inaction is interpreted as an instance of deontological responding, one would still have to specify the particular way in which a judgment is deemed deontological: is the judgment deontological in the sense that it is sensitive to proscriptive and prescriptive norms or is it deontological in the sense that it is congruent with the DDA principle?

Second, although the two ways of responding may be deemed deontological in a philosophical sense, they should not be conflated in a psychological theory about the mechanisms underlying moral dilemma judgment. After all, sensitivity to moral norms and general preference for inaction are functionally distinct in terms of their psychological antecedents and their behavioral outcomes. Their distinct outcomes are reflected in the fact that the two ways of responding lead to different judgments in moral dilemmas involving a prescriptive norm. Their distinct antecedents are reflected in the current finding that a given factor can simultaneously strengthen one way of "deontological" responding while weakening the other way of "deontological" responding. For example, in Studies 3a and 3b, a focus on personal

action, compared to a focus on moral acceptability, decreased sensitivity to moral norms and increased general preference for inaction. Thus, although some variables may influence the two ways of responding in the same direction (e.g., gender, psychopathy), the antagonistic effects of question framing suggest that they cannot be treated as conceptually equivalent instances of deontological responding.

Third, if sensitivity to moral norms and general preference for inaction are treated as functionally distinct instances of deontological responding, any theory about the processes underlying moral dilemma judgments would have to specify these processes for each of the two ways of deontological responding. For example, if automatic emotional responses are claimed to influence moral dilemma judgments via general preference for inaction, the theory would have to make additional assumptions about the processes underlying sensitivity to moral norms. In fact, a clear distinction between the two ways of deontological responding may even resolve some debates about the processes underlying deontological judgments. For example, some researchers have argued that deontological judgments are the product of coherence-based reasoning processes, involving pragmatic considerations of rights and duties (e.g., Holyoak & Powell, 2016). Although these assumptions may seem in conflict with the hypothesis that deontological judgments are the result of automatic emotional responses, any such conflict would be spurious to the extent that (1) pragmatic considerations of rights and duties influence moral dilemma judgments via sensitivity to moral norms and (2) automatic emotional responses influence moral dilemma judgments via general preference for inaction. From the perspective of the CNI model, either of the two competing accounts seems limited in the sense that they focus on only one of the two ways of deontological responding while ignoring the other. The current findings suggest that any comprehensive theory of moral dilemma judgment needs to capture the conceptual and psychological differences between three distinct determinants of moral dilemma judgments: (1) sensitivity to consequences, (2) sensitivity to moral norms, and (3) general preference for inaction versus action.

Potential Objections

A potential objection against the approach underlying the CNI model is that, in contrast to the unambiguous manipulations of consequences, manipulations of moral norms tend to be much more ambiguous, in that participants may not understand the moral norm that is supposed to be reflected in a given dilemma. Whereas consequences can be explicitly described in objective terms (e.g., by the number of people who may be affected by a given action), the relevance of specific moral norms is implicit in the sense that these norms have to be inferred from a given

scenario. Although we agree that the two kinds of manipulations differ in terms of their relative ambiguity, this concern can be ruled out for the current research, which consistently showed estimates on the N parameter that (1) were significantly larger than zero and (2) systematically varied across groups and experimental conditions. If participants did not understand the moral norms that were supposed to be captured by our dilemmas, estimates on the N parameter should be close to zero overall and too noisy to show meaningful variation across groups and experimental conditions.

A related objection is that proscriptive and prescriptive norms tend to be psychologically different, in that violations of a prescriptive norm may often feel less immoral compared to violations of a proscriptive norm. Thus, estimates on the N parameter may be distorted by the psychological asymmetry between proscriptive and prescriptive norms. In response to this objection, we would argue that the proposed psychological difference is conceptually equivalent to the omission bias, which refers to the phenomenon that harm caused by inaction feels less immoral compared to the same amount of harm caused by action (Cushman et al., 2006). As a result, violations of prescriptive norms should feel less immoral compared to violations of proscriptive norms. In this sense, the criticism is actually well-captured by the CNI model which distinguishes between (1) norm-congruent responses that are in line with the proscriptive and prescriptive norms in a given dilemma (N parameter), and (2) general preference for inaction that may result from the psychological asymmetry between action and inaction known as the omission bias (I parameter).

Another objection concerns the specific operationalization of prescriptive norms in our moral dilemmas. Whereas some dilemmas described direct interactions between participants and a target of their actions (e.g., participants as doctors providing patients with drugs), other dilemmas involved a third agent who engaged in potentially immoral actions against a target (e.g., participants witnessing another person using illegal interrogation techniques). In the latter type of dilemma, participants were asked to judge the behavioral option of interfering with the immoral action of the third agent (e.g., stopping the person from using illegal interrogation techniques by reporting their actions to their supervisor). The use of moral dilemmas with third agents may be deemed problematic, because interference with the actions of another person may be regarded as morally questionable even when the person's actions are deemed immoral. To test whether our findings were influenced by the use of moral dilemmas with a third agent, we reran all of the CNI analyses with a reduced set of moral dilemmas that included only those that did not involve a third agent (i.e., assisted suicide dilemma, immune deficiency dilemma). Of the 30 significance tests reported for the 8 studies in the main text and the 2 studies in the Supplementary Materials, 27 tests produced identical

results. For the 3 tests that did not produce identical results, a formerly significant effect turned non-significant in the analyses with the reduced sample. Because (1) the mean level differences in parameter scores in these cases were equivalent to the ones in the original analyses and (2) the smaller number of observations increases the confidence intervals for parameter estimates with the CNI model, we consider these deviations as being due to lower statistical power rather than genuine effects of different kinds of moral dilemmas. Nevertheless, future research may provide deeper insights into this issue by directly manipulating the involvement of third agents.

Limitations

Although we deem the CNI model as superior to both the traditional approach and Conway and Gawronski's (2013) PD approach, it is important to acknowledge some limitations. The most significant limitation is that, with the current set of 24 moral dilemmas, the CNI model is not suitable for correlational designs. In the current work, we exclusively presented group-level comparisons that involved aggregate data from all participants in a given group. This approach is very common and widely accepted among researchers using multinomial modeling (Klauer, 2015). It is also the standard data analytic approach in the multiTree software that was used to analyze the current data (Moshagen, 2010). To obtain data that are suitable for correlational analyses, the CNI model would have to be fit to the responses of each individual participant (rather than groups of participants). However, with the small number of trials provided by our 24 moral dilemmas, such individual estimates tend to be unreliable.

First, a small number of observations often leads to poor model fit at the individual level. Thus, even when the CNI model fits the data well at the aggregate level, there are usually several participants for whom the model does not fit when the model is applied to a small number of observations at the individual level. Second, to the extent that a participant shows highly consistent responses across dilemmas, the number of trials that are available to estimate the respective other response patterns becomes rather small. For example, if a participant shows a utilitarian response pattern on 5 of the 6 basic dilemmas, estimations of the N and the I parameters are based on the four variants of a single basic dilemma, making such estimates even less reliable.

Although both of these problems can be resolved in group-level analyses by aggregating data for groups of participants, they undermine the suitability of the CNI model for correlational designs, at least with the small number of dilemmas employed in the current studies. To address this limitation, we are currently testing a larger set of moral dilemmas that could be used to increase the number of observations for the model. Yet, a potential drawback is that an increase in the number of dilemmas

may lead to fatigue over the course of the study. Thus, any such extension will have to balance the trade-off between unreliable parameter estimations with small numbers of dilemmas and the increased likelihood of fatigue with large numbers of dilemmas.

A more specific concern is that the CNI model showed suboptimal fit in two of the eight studies. In Study 2a, the probabilities predicted by the model showed a marginally significant deviation from the empirically observed probabilities. In Study 3a, the deviation between actual and predicted probabilities was statistically significant. However, in both studies the effect size of the deviation indicated negligible effects (Cohen, 1988; see Footnote 6). Moreover, we conducted replications for each of the two studies using the same procedures and materials. In both replication studies, the CNI model fit the data well and the obtained effects on the three parameters replicated in both cases. Based on these findings, we are confident that our findings in the two sets of studies are indeed reliable.

From a broader perspective, this issue illustrates the trade-offs that have to be considered when using multinomial models in studies with large samples. On the one hand, large samples are essential to ensure sufficient statistical power in testing experimental effects on the three parameters. On the other hand, large samples increase the likelihood that the predicted probabilities of the model significantly deviate from the empirically observed probabilities, because large samples also increase statistical power for detecting even small deviations. Based on the overall performance of the CNI model in the current studies, we would argue that higher priority should be given to the concern about sufficient statistical power.

Future Directions

The current work also raises a number of interesting questions for future research on moral judgment. One example is the use of the CNI model to gain deeper insights into the psychological underpinnings of counterintuitive findings that do not fit well to current theories of moral judgment. In Studies 4a and 4b, we already demonstrated the value of the CNI model for this purpose by providing a more nuanced picture of how psychopathy influences moral dilemma judgments. Counter to the paradoxical conclusion that psychopaths are 'more utilitarian' than non-psychopaths (e.g., Bartels & Pizarro, 2011; Kahane et al., 2015; Patil, 2015; Pletti et al., in press), our findings suggest that psychopaths are (1) less sensitive to the consequences of their actions, (2) less sensitive to moral norms, and (3) less reluctant to engage in action than non-psychopaths.

A similarly puzzling finding is that higher levels of blood alcohol are associated with a greater willingness to accept harmful action in the trolley paradigm (Duke & Bègue, 2015). This finding stands in contrast to the notion that (1) alcohol impairs cognitive function and (2)

impaired cognitive function should reduce utilitarian judgments according to the dual-process model of moral judgment (Greene et al., 2001, 2004, 2008). From the perspective of the CNI model, this paradox can be reconciled with the hypothesis that alcohol does not increase the sensitivity to morally relevant consequences. Instead, alcohol may (1) reduce the sensitivity to moral norms, or (2) increase the willingness to engage in action irrespective of consequences and norms (or both). Although the traditional approach is unable to distinguish between these possibilities, the CNI model provides a straightforward means to test these hypotheses.

At a broader level, the current work poses a challenge to research that has used the traditional approach to studying the mental underpinnings moral dilemma judgments. Similar concerns could be raised about research on the cognitive and emotional determinants of moral judgments (e.g., Bartels, 2008; Strohminger et al., 2011; Ugazio, Lamm, & Singer, 2012; Valdesolo & DeSteno, 2006) as well as their neural correlates (e.g., Bernhard et al., in press; Ciaramelli, Muccioli, Ladavas, & di Pellegrino, 2007; Hutcherson, Montaser-Kouhsari, Woodward, & Rangel, 2015; Koenigs et al., 2007). A major limitation of this research is that it provides little information on whether a given factor influenced moral dilemma judgments via (1) sensitivity to consequences, (2) sensitivity to moral norms, and (3) general preference for inaction. Because this ambiguity undermines straightforward inferences about the mental underpinnings of moral dilemma judgments, the current work calls for a reassessment of dominant theoretical assumptions. The CNI model provides a valuable tool in this endeavor by allowing researchers to quantify the three patterns of responding to moral dilemmas.

To facilitate research progress along these lines, we have prepared a zip-file with a multiTree template and hands-on instructions on how to analyze moral dilemma judgment data with the CNI model. The zip-file also includes copies of our moral dilemmas (see Appendix A), a template file for lab studies with our moral dilemmas using the psychological lab software MediaLab by Empirisoft, and an SPSS syntax file for the aggregation of data obtained with our MediaLab template file. The zip-file with these materials can be freely downloaded at: http://www.bertramgawronski.com/documents/CNI-Model_Materials.zip. We hope that psychologists who are interested in moral dilemma judgment will find these materials useful for their own research, and we are looking forward to seeing the products of future work using the CNI model.

Conclusion

To identify the processes underlying utilitarian and deontological judgments, research has investigated responses to moral dilemmas that pit one principle against the other (e.g., trolley problem). We argued that

the conceptual meaning of responses in this paradigm is ambiguous, because the defining aspects of utilitarianism and deontology, consequences and norms, are not manipulated. To overcome potential distortions in the interpretation of empirical findings, we proposed an alternative approach in which utilitarian judgments are inferred from the sensitivity to consequences and deontological judgments are inferred from the sensitivity to moral norms. Expanding on this approach, we presented a multinomial model that allows researchers to quantify sensitivity to consequences, sensitivity to moral norms, and general preference for inaction versus action in responses to moral dilemmas. The findings obtained with our CNI model offer more nuanced insights into the determinants of moral judgments, calling for a reassessment of dominant theoretical assumptions.

References

- Amit, E., & Greene, J. D. (2012). You see, the ends don't justify the means: Visual imagery and moral judgment. *Psychological Science, 23*, 861-868.
- Arutyunova, K. R., Alexandrov, Y. I., & Hauser, M. D. (2016). Sociocultural influences on moral judgments: East-West, male-female, and young-old. *Frontiers in Psychology, 7*:1334.
- Baron, J. (1994). Nonconsequentialist decisions. *Behavioral and Brain Sciences, 17*, 1-10.
- Bartels, D. (2008). Principled moral sentiment and the flexibility of moral judgment and decision making. *Cognition, 108*, 381-417.
- Bartels, D. M., & Pizarro, D. A. (2011). The mismeasure of morals: Antisocial personality traits predict utilitarian responses to moral dilemmas. *Cognition, 121*, 154-161.
- Batchelder, W. H., & Riefer, D. M. (1999). Theoretical and empirical review of multinomial process tree modeling. *Psychonomic Bulletin & Review, 6*, 57-86.
- Bauman, C. W., McGraw, A. P., Bartels, D. M., & Warren, C. (2014). Revisiting external validity: Concerns about trolley problems and other sacrificial dilemmas in moral psychology. *Social and Personality Psychology Compass, 8*, 536-554.
- Bernhard, R. M., Chaponis, J., Siburian, R., Gallagher, P., Ransohoff, K., Wikler, D., Perlis, R., & Greene, J. D. (in press). Variation in the oxytocin receptor gene (OXTR) is associated with differences in moral judgment. *Social Cognitive and Affective Neuroscience*.
- Blinder, A., & Grady, D. (2014, August 2). American doctor with Ebola arrives in U.S. for treatment. *New York Times*. Retrieved from <http://www.nytimes.com>
- Broeders, R., van den Bos, K., Müller, P. A., & Ham, J. (2011). Should I save or should I not kill? How people solve moral dilemmas depends on which rule is most accessible. *Journal of Experimental Social Psychology, 47*, 923-934.
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011).

- Amazon's mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6, 3-5.
- Cale, E. M., & Lilienfeld, S. O. (2002). Sex differences in psychopathy and antisocial personality disorder: A review and integration. *Clinical Psychology Review*, 22, 1179-1207.
- Cameron, C. D., Payne, B. K., Sinnott-Armstrong, W., Scheffer, J. A., & Inzlicht, M. (2017). Implicit moral evaluations: A multinomial modeling approach. *Cognition*, 158, 224-241.
- Cameron, C. D., Spring, V. L., & Todd, A. R. (in press). The empathy impulse: A multinomial model of intentional and unintentional empathy for pain. *Emotion*.
- Cao, F., Zhang, J., Song, L., Wang, S., Miao, D., & Peng, J. (in press). Framing effect in the trolley problem and the footbridge dilemma: Number of lives saved matters. *Psychological Reports*.
- Ciamarelli, E., Muccioli, M., Ladavas, E., & di Pellegrino, G. (2007). Selective deficit in personal moral judgment following damage to ventromedial prefrontal cortex. *Social Cognitive and Affective Neuroscience*, 2, 84-92.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. (2005). Separating multiple processes in implicit social cognition: The quad-model of implicit task performance. *Journal of Personality and Social Psychology*, 89, 469-487.
- 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.
- Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: Using ethnicity to disambiguate threatening individuals. *Journal of Personality and Social Psychology*, 83, 1314-1329.
- Crone, D. L., & Laham, S. M. (2017). Utilitarian preferences or action preferences? De-confounding action and moral code in sacrificial dilemmas. *Personality and Individual Differences*, 104, 476-481.
- Cumming, G. (2014). The new statistics: Why and how. *Psychological Science*, 25, 7-29.
- 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.
- De Houwer, J. (2011). Why the cognitive approach in psychology would profit from a functional approach and vice versa. *Perspectives on Psychological Science*, 6, 202-209.
- Duke, A. A., & Bègue, L. (2015). The drunk utilitarian: Blood alcohol concentration predicts utilitarian responses in sacrificial moral dilemmas. *Cognition*, 134, 121-127.
- Foot, P. (1967). The problem of abortion and the doctrine of double effect. *Oxford Review*, 5, 5-15.
- Friesdorf, R., Conway, P., & Gawronski, B. (2015). Gender differences in responses to moral dilemmas: A process dissociation analysis. *Personality and Social Psychology Bulletin*, 41, 696-713.
- Fumagalli, M., Ferrucci, R., Mameli, F., Marcegaglia, S., Mrakic-Sposta, S., Zago, S., Lucchiari, C., Consonni, D., Nordio, F., Pravettoni, G., Cappa, S., & Priori, A. (2010). Gender-related differences in moral judgments. *Cognitive Processing*, 11, 219-226.
- Gast, A., Gawronski, B., & De Houwer, J. (2012). Evaluative conditioning: Recent developments and future directions. *Learning and Motivation*, 43, 79-88.
- Gawronski, B., & Beer, J. S. (in press). What makes moral dilemma judgments "utilitarian" or "deontological"? *Social Neuroscience*.
- Gawronski, B., Conway, P., Armstrong, J., Friesdorf, R., & Hütter, M. (2016). Understanding responses to moral dilemmas: Deontological inclinations, utilitarian inclinations, and general action tendencies. In J. P. Forgas, L. Jussim, & P. A. M. Van Lange (Eds.), *Social psychology of morality* (pp. 91-110). New York, NY: Psychology Press.
- Gawronski, B., Ehrenberg, K., Banse, R., Zukova, J., & Klauer, K. C. (2003). It's in the mind of the beholder: The impact of stereotypic associations on category-based and individuating impression formation. *Journal of Experimental Social Psychology*, 39, 16-30.
- Greene, J. D. (2007). The secret joke of Kant's soul. In W. Sinnott-Armstrong (Ed.), *Moral psychology: Vol. 2. The cognitive science of morality*. Cambridge, MA: MIT Press.
- Greene, J. D., & Haidt, J. (2002). How (and where) does moral judgment work? *Trends in Cognitive Sciences*, 6, 517-523.
- Greene, J. D., Morelli, S. A., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2008). Cognitive load selectively interferes with utilitarian moral judgment. *Cognition*, 107, 1144-1154.
- 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.
- Hofmann, W., Wisneski, D. C., Brandt, M. J., & Skitka, L. J. (2014). Morality in everyday life. *Science*, 345, 1340-1343.

- Holyoak, K. J., & Powell, D. (2016). Deontological coherence: A framework for commonsense moral reasoning. *Psychological Bulletin, 142*, 1179-1203.
- Hütter, M., & Klauer, K. C. (2016). Applying processing trees in social psychology. *European Review of Social Psychology, 27*, 116-159.
- Hütter, M., & Sweldens, S. (2013). Implicit misattribution of evaluative responses: Contingency-unaware evaluative conditioning requires simultaneous stimulus presentations. *Journal of Experimental Psychology: General, 142*, 638-643.
- Hütter, M., Sweldens, S., Stahl, C., Unkelbach, C., & Klauer, K. C. (2012). Dissociating contingency awareness and conditioned attitudes: Evidence of contingency-unaware evaluative conditioning. *Journal of Experimental Psychology: General, 141*, 539-557.
- Hutcherson, C. A., Montaser-Kouhsari, L., Woodward, J., & Rangel, A. (2015). Emotional and utilitarian appraisals of moral dilemmas are encoded in separate areas and integrated in the ventromedial prefrontal cortex. *Journal of Neuroscience, 35*, 12593-12605.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30*, 513-541.
- Janoff-Bulman, R., Sheikh, S., & Hepp, S. (2009). Proscriptive versus prescriptive morality: Two faces of moral regulation. *Journal of Personality and Social Psychology, 96*, 521-537.
- Kahane, G., Everett, J. A. C., Earp, B. D. Farias, M., & Savulescu, J. (2015). 'Utilitarian' judgments in sacrificial moral dilemmas do not reflect impartial concern for the greater good. *Cognition, 134*, 193-209.
- Kahneman, D., & Tversky, A. (1982). The psychology of preferences. *Scientific American, 246*, 160-173.
- Klauer, K. C. (2015). Mathematical modeling. In B. Gawronski, & G. V. Bodenhausen (Eds.), *Theory and explanation in social psychology* (pp. 371-389). New York: Guilford Press.
- Klauer, K. C. & Wegener, I. (1998). Unraveling social categorization in the "Who said what?" paradigm. *Journal of Personality and Social Psychology, 75*, 1155-1178.
- Koenigs, M., Young, L., Adolphs, R., Tranel, D., Cushman, F., Hauser, M., & Damasio, A. (2007). Damage to the prefrontal cortex increases utilitarian moral judgments. *Nature, 446*, 908-911.
- Kohlberg, L. (1969). Stage and sequence: The cognitive-developmental approach to socialization. In D. A. Goslin (Ed.), *Handbook of socialization theory and research*. (pp. 347-480). Chicago: Rand McNally.
- Krieglmeyer, R., & Sherman, J. W. (2012). Disentangling stereotype activation and stereotype application in the Stereotype Misperception Task. *Journal of Personality and Social Psychology, 103*, 205-224.
- Levenson, M. R., Kiehl, K. A., & Fitzpatrick, C. M. (1995). Assessing psychopathic attributes in a noninstitutionalized population. *Journal of Personality and Social Psychology, 68*, 151-158.
- Lipsey, M. W., & Wilson, D. S. (2001). *Practical meta-analysis*. Thousand Oaks, CA: Sage.
- Loken, E., & Gleman, A. (2017). Measurement error and the replication crisis. *Science, 355*, 584-585.
- Maxwell, S. E., Lau, M. Y., & Howard, G. S. (2015). Is psychology suffering from a replication crisis? What does "failure to replicate" really mean? *American Psychologist, 70*, 487-498.
- Meissner, F., & Rothermund, K. (2013). Estimating the contributions of associations and recoding in the Implicit Association Test: The ReAL model for the IAT. *Journal of Personality and Social Psychology, 104*, 45-69.
- Moshagen, M. (2010). multiTree: A computer program for the analysis of multinomial processing tree models. *Behavioral Research Methods, 42*, 42-54.
- Moore, A. B., Clark, B. A., & Kane, M. J. (2008). Who shalt not kill? Individual differences in working memory capacity, executive control, and moral judgment. *Psychological Science, 19*, 549-57.
- Nadarevic, L., & Erdfelder, E. (2011). Cognitive processes in implicit attitude tasks: An experimental validation of the Trip Model. *European Journal of Social Psychology, 41*, 254-268.
- Nichols, S., & Mallon, R. (2006). Moral dilemmas and moral rules. *Cognition, 100*, 530-542.
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science, 349*(6251), aac4716.
- 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.
- 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., Neumann, C. F., & Hare, R. D. (2009). *Manual for the self-report psychopathy scale*. Toronto: Multi-Health Systems.
- Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology, 81*, 181-192.
- Payne, B. K., & Bishara, A. J. (2009). An integrative review of process dissociation and related models in social cognition. *European Review of Social Psychology, 20*, 272-314.
- Piazza, J., Sousa, P., Holbrook, C. (2013). Authority dependence and judgments of utilitarian harm. *Cognition, 128*, 261-270.
- Pletti, C., Lotto, L., Buodo, G., & Sarlo, M. (in press). It's immoral, but I'd do it: Psychopathy traits affect decision-making in sacrificial dilemmas and in

- everyday moral situations. *British Journal of Psychology*.
- Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, *1*, 7-59.
- Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, *27*, 76-105.
- Stahl, C., & Degner, J. (2007). Assessing automatic activation of valence: A multinomial model of EAST performance. *Experimental Psychology*, *54*, 99-112.
- Stanley, D. J., & Spence, J. R. (2014). Expectations for replications: Are yours realistic? *Perspectives on Psychological Science*, *9*, 305-318.
- Starcke, K., Ludwig, A., & Brand, M. (2012). Anticipatory stress interferes with utilitarian moral judgment. *Judgment and Decision Making*, *7*, 61-68.
- Strohlinger, N., Lewis, R. L., & Meyer, D. E. (2011). Divergent effects of different positive emotions on moral judgment. *Cognition*, *119*, 295-300.
- 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.
- Trémolière, B., & Bonnefon, J.-F. (2014). Efficient kill-save rations ease up the cognitive demands on counterintuitive moral utilitarianism. *Personality and Social Psychology Bulletin*, *40*, 923-930.
- Ugazio, G., Lamm, C. & Singer, T. (2012). The role of emotions for moral judgments depends on the type of emotion and moral scenario. *Emotion*, *12*, 579-590.
- Valdesolo, P. & DeSteno, D. (2006). Manipulations of emotional context shape moral judgment. *Psychological Science*, *17*, 476-477.
- van den Bos, K., Müller, P. A., & Damen, T. (2011). A behavioral disinhibition hypothesis of interventions in moral dilemmas. *Emotion Review*, *3*, 281-283.
- Whitlock, C. (2006, February 16). German court overturns law allowing hijacked airliners to be shot down. *Washington Post Foreign Service*. Retrieved from <http://www.washingtonpost.com>
- Yzerbyt, V. Y., Coull, A., & Rocher, S. J. (1999). Fencing off the deviant: The role of cognitive resources in the maintenance of stereotypes. *Journal of Personality and Social Psychology*, *77*, 449-462.

Acknowledgements

This research was supported by National Science Foundation Grant # 1449620 to the first author. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Appendix A

Moral dilemmas with proscriptive versus prescriptive norms and benefits of action for overall well-being that are either greater or smaller than the costs of action for well-being.

	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
Immune Deficiency Dilemma	<p>You are the director of a hospital in a developing country. A foreign student who is volunteering in the country got infected with a rare virus. The virus is highly contagious and deadly to seniors and children. The only medication that can effectively stop the virus from spreading has severe side-effects. Although the virus will not kill her, the student suffers from a chronic immune deficiency that will make her die from these side-effects.</p> <p>Is it acceptable in this case to give the student the medication?</p>	<p>You are the director of a hospital in a developing country. A foreign student who is volunteering in the country got infected with a rare virus. The virus is highly contagious and can cause severe stomach cramps. The only medication that can effectively stop the virus from spreading has severe side-effects. Although the virus will not kill her, the student suffers from a chronic immune deficiency that will make her die from these side-effects.</p> <p>Is it acceptable in this case to give the student the medication?</p>	<p>You are the director of a hospital in a developing country. A foreign student who is volunteering in the country got infected with a rare virus. The virus is highly contagious and can cause severe stomach cramps. The student suffers from a chronic immune deficiency that will make her die from the virus if she is not returned to her home country for special treatment. However, taking her out of quarantine involves a considerable risk that the virus will spread.</p> <p>Is it acceptable in this case to take the student out of quarantine to return her to her home country for treatment?</p>	<p>You are the director of a hospital in a developing country. A foreign student who is volunteering in the country got infected with a rare virus. The virus is highly contagious and deadly to seniors and children. The student suffers from a chronic immune deficiency that will make her die from the virus if she is not returned to her home country for special treatment. However, taking her out of quarantine involves a considerable risk that the virus will spread.</p> <p>Is it acceptable in this case to take the student out of quarantine to return her to her home country for treatment?</p>
Assisted Suicide Dilemma	<p>You are a doctor and are treating a seriously ill, suffering patient. Even the strongest drugs do not relieve him from his pain anymore.</p> <p>He feels terrible agony and you know that his condition will never improve. For days, he has been waiting for his death.</p>	<p>You are a doctor and are treating a seriously ill, suffering patient. Even the strongest drugs do not relieve him from his pain anymore.</p> <p>He feels terrible agony, but you know that he will get better again soon. For days, he has been waiting to see improvements.</p>	<p>You are a doctor and are treating a seriously ill, suffering patient. Even the strongest drugs do not relieve him from his pain anymore.</p> <p>He feels terrible agony, but you know that he will get better again soon. For days, he has been waiting to see improvements.</p>	<p>You are a doctor and are treating a seriously ill, suffering patient. Even the strongest drugs do not relieve him from his pain anymore.</p> <p>He feels terrible agony and you know that his condition will never improve. For days, he has been waiting for his death.</p>

Since he no longer wants to endure his pain, he repeatedly asks you to end his life. You could give him a drug causing his death, which will release him from his pain.

Is it acceptable in this case to provide your patient with this drug?

Abduction Dilemma

You are the president of your country. A guerilla group operating in a conflict zone has abducted a journalist from your country and threatens to behead him if your government does not pay a ransom of one million dollars. The group will use the ransom money to buy weapons for their guerilla war, which will cause the deaths of many people. Congress has approved payment of the ransom, but you have the power to veto the payment.

Is it acceptable in this case to veto the ransom payment?

Transplant Dilemma

You are a surgeon in a small hospital. One day, your hospital receives five badly hurt patients from a car accident. The patients all need organ transplants or they will die. You have no spare organs, but there is a patient who has been in a coma for several weeks and it seems unlikely that he will wake up again. You could terminate his life support and take his organs for the five

Since he no longer wants to endure his pain, he repeatedly asks you to end his life. You could give him a drug causing his death, which will release him from his pain.

Is it acceptable in this case to provide your patient with this drug?

You are the president of your country. A guerilla group operating in a conflict zone has abducted a journalist from your country and threatens to behead him if your government does not pay a ransom of one million dollars. The group will use the ransom money to buy food for their families, who live in an area that has been plagued by several droughts. Congress has approved payment of the ransom, but you have the power to veto the payment.

Is it acceptable in this case to veto the ransom payment?

You are a surgeon in a small hospital. One day, your hospital receives five badly hurt patients from a car accident. These patients all need organ transplants or they will have serious health problems for the rest of their lives. You have no spare organs, but there is a patient who has been in a coma for several weeks and it seems unlikely that he will wake up again. You could

Since he no longer wants to endure his pain, he repeatedly asks you to end his life. Suddenly, he has a severe heart attack. You could give him a drug to save him from dying.

Is it acceptable in this case to provide your patient with this drug?

You are the president of your country. A guerilla group operating in a conflict zone has abducted a journalist from your country and threatens to behead him if your government does not pay a ransom of one million dollars. The group will use the ransom money to buy food for their families, who live in an area that has been plagued by several droughts. As the president, you have the power to approve payment of the ransom.

Is it acceptable in this case to approve the ransom payment?

You are a surgeon in a small hospital. One day, your hospital receives five badly hurt patients from a car accident. These patients all need organ transplants or they will have serious health problems for the rest of their lives. You have no spare organs, but there is a patient who has been in a coma for several weeks and it seems unlikely that he will wake up again. One of your co-

Since he no longer wants to endure his pain, he repeatedly asks you to end his life. Suddenly, he has a severe heart attack. You could give him a drug to save him from dying.

Is it acceptable in this case to provide your patient with this drug?

You are the president of your country. A guerilla group operating in a conflict zone has abducted a journalist from your country and threatens to behead him if your government does not pay a ransom of one million dollars. The group will use the ransom money to buy weapons for their guerilla war, which will cause the deaths of many people. As the president, you have the power to approve payment of the ransom.

Is it acceptable in this case to approve the ransom payment?

You are a surgeon in a small hospital. One day, your hospital receives five badly hurt patients from a car accident. The patients all need organ transplants or they will die. You have no spare organs, but there is a patient who has been in a coma for several weeks and it seems unlikely that he will wake up again. One of your co-workers plans to terminate his life support and

accident victims, so that their lives will be saved.

Is it acceptable in this case to terminate the patient's life support to take his organs?

terminate his life support and take his organs for the five accident victims, so that they won't suffer from health problems.

Is it acceptable in this case to terminate the patient's life support to take his organs?

workers plans to terminate his life support and take his organs for the five accident victims, so that they won't suffer from health problems. You could stop your co-worker by informing the director of the hospital.

Is it acceptable in this case to stop your co-worker from terminating the patient's life support to take his organs?

take his organs for the five accident victims, so that their lives will be saved. You could stop your co-worker by informing the director of the hospital.

Is it acceptable in this case to stop your co-worker from terminating the patient's life support to take his organs?

Torture Dilemma

You are a member of a special police department who is trained to obtain information in particularly difficult cases. You are dealing with a case involving a male adult who is accused of having abducted several children. You don't know where he is hiding the children, and he refuses to tell you where they are. The children will likely die from dehydration if they are not found within the next 24 hours. You have tried every legal interrogation technique, but none of them were successful. To get information on where the children are, you consider the use of illegal techniques that are deemed torture.

Is it acceptable in this case to use illegal interrogation techniques?

You are a member of a special police department who is trained to obtain information in particularly difficult cases. You are dealing with a case involving a male adult who is accused of having stolen several paintings. You don't know where he is hiding the paintings, and he refuses to tell you where they are. The paintings will likely be shipped to a different country if they are not found within the next 24 hours. You have tried every legal interrogation technique, but none of them were successful. To get information on where the paintings are, you consider the use of illegal techniques that are deemed torture.

Is it acceptable in this case to use illegal interrogation techniques?

You are a member of a special police department who is trained to obtain information in particularly difficult cases. You are dealing with a case involving a male adult who is accused of having stolen several paintings. You don't know where he is hiding the paintings, and he refuses to tell you where they are. The paintings will likely be shipped to a different country if they are not found within the next 24 hours. You have tried every legal interrogation technique, but none of them were successful. To get information on where the paintings are, your partner starts using illegal techniques that are deemed torture, and you consider stopping him by reporting him to your supervisor.

Is it acceptable in this case to stop your partner from using illegal interrogation techniques?

You are a member of a special police department who is trained to obtain information in particularly difficult cases. You are dealing with a case involving a male adult who is accused of having abducted several children. You don't know where he is hiding the children, and he refuses to tell you where they are. The children will likely die from dehydration if they are not found within the next 24 hours. You have tried every legal interrogation technique, but none of them were successful. To get information on where the children are, your partner starts using illegal techniques that are deemed torture, and you consider stopping him by reporting him to your supervisor.

Is it acceptable in this case to stop your partner from using illegal interrogation techniques?

Vaccine

You are a doctor in an area that suffers from an outbreak of a

You are a doctor in an area that suffers from an outbreak of a

You are a doctor in an area that suffers from an outbreak of a

You are a doctor in an area that suffers from an outbreak of a

Dilemma

highly contagious disease. Preliminary tests have shown the success of a new vaccine that is not approved by the health department of your country, because of its severe side-effects. The side-effects of the vaccine will likely cause the death of dozens of people who are not infected, but the vaccine will save hundreds of lives by preventing spread of the virus.

Is it acceptable in this case to use the vaccine?

highly contagious disease. Preliminary tests have shown the success of a new vaccine that is not approved by the health department of your country, because of its severe side-effects. The side-effects of the vaccine will likely cause the death of dozens of people who are not infected, but the vaccine will save about the same number of lives by preventing spread of the virus.

Is it acceptable in this case to use the vaccine?

highly contagious disease. Preliminary tests have shown the success of a new vaccine that is not approved by the health department of your country, because of its severe side-effects. The side-effects of the vaccine will likely cause the death of dozens of people who are not infected, but the vaccine will save about the same number of lives by preventing spread of the virus. One of your colleagues plans to use the vaccine, but you could stop him by reporting his plans to the health department.

Is it acceptable in this case to report your colleague to the health department?

highly contagious disease. Preliminary tests have shown the success of a new vaccine that is not approved by the health department of your country, because of its severe side-effects. The side-effects of the vaccine will likely cause the death of dozens of people who are not infected, but the vaccine will save hundreds of lives by preventing spread of the virus. One of your colleagues plans to use the vaccine, but you could stop him by reporting his plans to the health department.

Is it acceptable in this case to report your colleague to the health department?

Appendix B

Model equations for the estimation of sensitivity to consequences (C), sensitivity to moral norms (N), and general preference for inaction versus action irrespective of consequences and norms (I) in responses to moral dilemmas with proscriptive versus prescriptive norms and benefits of action for overall well-being that are either greater or smaller than the costs of action for well-being.

$$\begin{aligned}
 p(\text{inaction} \mid \text{proscriptive norm, benefits} > \text{costs}) &= [(1 - C) \times N] + [(1 - C) \times (1 - N) \times I] \\
 p(\text{inaction} \mid \text{proscriptive norm, benefits} < \text{costs}) &= C + [(1 - C) \times N] + [(1 - C) \times (1 - N) \times I] \\
 p(\text{inaction} \mid \text{prescriptive norm, benefits} > \text{costs}) &= (1 - C) \times (1 - N) \times I \\
 p(\text{inaction} \mid \text{prescriptive norm, benefits} < \text{costs}) &= C + [(1 - C) \times (1 - N) \times I]
 \end{aligned}$$

$$\begin{aligned}
 p(\text{action} \mid \text{proscriptive norm, benefits} > \text{costs}) &= C + [(1 - C) \times (1 - N) \times (1 - I)] \\
 p(\text{action} \mid \text{proscriptive norm, benefits} < \text{costs}) &= (1 - C) \times (1 - N) \times (1 - I) \\
 p(\text{action} \mid \text{prescriptive norm, benefits} > \text{costs}) &= C + [(1 - C) \times N] + [(1 - C) \times (1 - N) \times (1 - I)] \\
 p(\text{action} \mid \text{prescriptive norm, benefits} < \text{costs}) &= [(1 - C) \times N] + [(1 - C) \times (1 - N) \times (1 - I)]
 \end{aligned}$$

Appendix C

Descriptions of behaviors in the Pilot Study. Participants were asked to rate each behavior for its moral relevance on 5-point scales ranging from 1 (*not at all*) to 5 (*very much*).

Behaviors in Moral Dilemmas of the Current Studies

Immune Deficiency Dilemma

Giving a patient a drug that will prevent the spread of a highly contagious virus when the patient has an immune deficiency that will make her die from the side-effects of the drug.

Bringing a patient who is infected with a highly contagious virus to her home country when taking her out of quarantine involves a considerable risk that the virus will spread.

Assisted Suicide Dilemma

Following the request of a terminally ill patient to give him a deadly drug to release him from his pain.

Providing a life-saving drug to a terminally ill patient who has a heart attack but no longer wants to live.

Abduction Dilemma

Vetoing the payment of a ransom to free a journalist abducted by terrorists who threaten to behead him.

Approving the payment of a ransom to free a journalist abducted by terrorists who threaten to behead him.

Transplant Dilemma

Terminating a patient's life support to take his organs for other patients.

Informing the hospital director about a co-worker who plans to terminate a patient's life support to take his organs for other patients.

Torture Dilemma

Using illegal interrogation techniques that are deemed to torture to obtain information from a criminal offender.

Reporting a co-worker who plans to use illegal interrogation techniques that are deemed to torture to obtain information from a criminal offender.

Vaccine Dilemma

Using a vaccine to fight a highly contagious disease when the vaccine has not been approved for use by the health department.

Reporting a colleague who plans to use a vaccine to fight a highly contagious disease when the vaccine has not been approved for use by the health department.

Moral Behaviors Adapted from Hofmann et al. (2014)

Care/Harm

Assisting a tourist with directions because he is looking lost.

Giving a homeless man an extra sandwich you have.

Smoking a cigarette with small children in the car.

Hiring someone to kill a muskrat that's ultimately not causing any harm.

Fairness/Unfairness

Talking to someone about treating others equally.

Reminding a waitress that you did not pay for your bill when she thought you did.

Congress making cuts across the board and not solving debt problems for the country.

Stealing your co-worker's nice balsamic vinegar while he was off shift.

Loyalty/Disloyalty

Paying tribute to veterans and families that have lost a loved one on Memorial Day.

Putting your family before your own fun.

Giving up on your team.

Arranging adulterous encounter.

Authority/Subversion

Enforcing a rule.

Appropriately disciplining a youth not your own.

Disrespecting your mother.

Having drinks with a colleague during work hours without the boss knowing.

Sanctity/Degradation

Talking about God with a family member.

Yoga Nidra meditation class.

Catching your teenage son looking at hard core porn.

Making your 3 year old eat her feces for having an accident.

Liberty/Oppression

Arguing on behalf an oppressed population in a public setting.
Freeing Beagles that had never seen daylight or felt grass, due to a life of captivity for animal testing.
Denying a girl a ticket to prom because she wants to take another girl as her date.
Kidnapping schoolgirls in Nigeria.

Honesty/Dishonesty

Being honest about a sales initiative.
Finding a lost cell phone and returning it to its owner.
Lying to someone by saying their suggestion was good when it wasn't.
Faking a bomb threat to cover up not graduating.

Self-Discipline/Lack of Self-Discipline

Following through on completing a work commitment by completing it after hours.
Not reading confidential info that is not for you even though you have access.
Being sloppy in your work because you didn't want to do it
Getting fast food although you promised someone you wouldn't have it.

Behaviors in Non-moral Decision Problems Adapted from Greene et al. (2001)

Plant Transport

Ruin the fine leather upholstery of your car by putting plants on the back seat.
Making two trips to avoid ruining the fine leather upholstery of your car by putting plants on the back seat.

Generic Brand

Purchasing a generic headache medicine when the brand-name product you planned to purchase is sold out.
Searching for the brand-name headache medicine you planned to purchase elsewhere when it is sold out at your pharmacy.

Brownies

Substituting macadamia nuts for walnuts in a recipe for brownies in order to avoid eating walnuts.
Using walnuts in a recipe for brownies although you do not like eating walnuts.

Computer

Waiting to buy a new computer in order to get it for a lower price.
Paying more for a new computer in order to get it right away.

Survey

Interrupting a quiet dinner by yourself in order to earn \$200 for participating in a national phone survey.
Refusing to participate in a national phone survey because you would have to interrupt a quiet dinner by yourself.

Choosing Classes

Taking a history class in order to help you fulfill your graduation requirements although it conflicts with a science you would like to take.
Taking a science class you would like to take although it conflicts with a history class that you need to fulfill your graduation requirements.

Table 1. Means and 95% confidence intervals of perceived moral relevance of behaviors in Pilot Study (see Appendix C). Scores can range from 1 to 5 with higher scores indicate higher perceived moral relevance.

	<i>M</i>	95% CI
Behaviors in Moral Dilemmas of Current Studies	3.85	[3.73, 3.97]
Immune Deficiency Dilemma	3.95	[3.78, 4.12]
Assisted Suicide Dilemma	3.77	[3.61, 3.92]
Abduction Dilemma	3.78	[3.63, 3.93]
Transplant Dilemma	4.00	[3.86, 4.14]
Torture Dilemma	4.04	[3.90, 4.18]
Vaccine Dilemma	3.61	[3.47, 3.76]
Moral Behaviors of Experience Sampling (Hofmann et al., 2014)	3.23	[3.15, 3.31]
Care/Harm	3.19	[3.09, 3.30]
Fairness/Unfairness	3.44	[3.32, 3.55]
Loyalty/Disloyalty	3.31	[3.20, 3.42]
Authority/Subversion	2.98	[2.87, 3.09]
Sanctity/Degradation	2.78	[2.68, 2.88]
Liberty/Oppression	3.71	[3.57, 3.84]
Honesty/Dishonesty	3.48	[3.38, 3.59]
Self-Discipline/Lack of Self-Discipline	2.96	[2.86, 3.07]
Behaviors in Non-moral Decision Problems (Greene et al., 2001)	1.88	[1.75, 2.02]
Plant Transport	1.86	[1.72, 2.01]
Generic Brand	1.81	[1.66, 1.97]
Brownies	1.64	[1.51, 1.77]
Computer	1.91	[1.75, 2.07]
Survey	1.97	[1.81, 2.13]
Choosing Classes	2.10	[1.94, 2.25]

Table 2. 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. Scores can range from 0 to 6. The neutral reference value of equal numbers of action and inaction responses is 3.

	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	
	<i>M</i>	95% CI	<i>M</i>	95% CI	<i>M</i>	95% CI	<i>M</i>	95% CI
Study 1a								
men	3.35	[3.06, 3.64]	2.10	[1.82, 2.39]	3.95	[3.66, 4.25]	2.94	[2.66, 3.22]
women	2.74	[2.43, 3.04]	1.36	[1.06, 1.66]	4.20	[3.89, 4.51]	2.95	[2.65, 3.25]
Study 1b								
men	3.01	[2.73, 3.29]	2.03	[1.73, 2.33]	4.29	[4.05, 4.54]	3.34	[3.05, 3.62]
women	2.75	[2.49, 3.02]	1.24	[0.95, 1.52]	4.47	[4.24, 4.71]	3.22	[2.95, 3.50]
Study 2a								
low load	3.24	[2.97, 3.51]	1.79	[1.50, 2.08]	4.24	[3.98, 4.50]	3.17	[2.91, 3.43]
high load	2.91	[2.64, 3.19]	1.67	[1.38, 1.96]	3.98	[3.71, 4.25]	3.11	[2.84, 3.38]
Study 2b								
low load	3.17	[2.89, 3.45]	1.81	[1.55, 2.07]	4.39	[4.12, 4.65]	3.23	[2.96, 3.49]
high load	2.72	[2.41, 3.02]	1.61	[1.32, 1.90]	3.97	[3.67, 4.26]	3.05	[2.76, 3.34]
Study 3a								
moral judgment	2.97	[2.69, 3.24]	1.34	[1.07, 1.61]	4.47	[4.16, 4.73]	3.63	[3.56, 3.90]
moral action	2.65	[2.37, 2.93]	1.18	[0.91, 1.46]	4.14	[3.85, 4.43]	2.65	[2.38, 2.93]
Study 3b								
moral judgment	3.27	[2.96, 3.57]	1.66	[1.35, 1.96]	4.52	[4.25, 4.80]	3.33	[3.05, 3.61]
moral action	2.93	[2.64, 3.22]	1.46	[1.17, 1.75]	3.85	[3.59, 4.11]	2.64	[2.37, 2.90]
Study 4a								
low psychopathy	2.94	[2.66, 3.21]	1.42	[1.13, 1.71]	4.04	[3.76, 4.32]	2.60	[2.32, 2.88]
high psychopathy	3.19	[2.91, 3.47]	1.94	[1.65, 2.24]	3.72	[3.43, 4.01]	2.58	[2.30, 2.87]
Study 4b								
low psychopathy	2.74	[2.48, 3.00]	1.13	[0.87, 1.39]	4.32	[4.07, 4.57]	2.75	[2.48, 3.03]
high psychopathy	3.27	[2.98, 3.56]	2.49	[2.21, 2.78]	3.24	[2.96, 3.52]	2.44	[2.14, 2.74]

Table 3. Means and 95% confidence intervals of *U* and *D* process dissociation parameters. Scores can range from 0 to 1. The neutral reference value for the *U* parameter is 0. The neutral reference value for the *D* parameter is 0.5.

	<i>U</i> Parameter		<i>D</i> Parameter	
	<i>M</i>	95% CI	<i>M</i>	95% CI
Study 1a				
men	.21	[0.16, 0.25]	.58	[0.53, 0.63]
women	.23	[0.18, 0.28]	.71	[0.66, 0.77]
Study 1b				
men	.16	[0.11, 0.21]	.61	[0.56, 0.66]
women	.25	[0.21, 0.30]	.76	[0.71, 0.81]
Study 2a				
low load	.24	[0.19, 0.29]	.63	[0.58, 0.69]
high load	.21	[0.16, 0.26]	.65	[0.60, 0.71]
Study 2b				
low load	.23	[0.18, 0.27]	.62	[0.57, 0.67]
high load	.18	[0.14, 0.23]	.67	[0.62, 0.73]
Study 3a				
moral judgment	.27	[0.22, 0.32]	.71	[0.65, 0.76]
moral action	.24	[0.20, 0.29]	.76	[0.70, 0.81]
Study 3b				
moral judgment	.27	[0.22, 0.32]	.65	[0.59, 0.71]
moral action	.24	[0.20, 0.29]	.68	[0.62, 0.74]
Study 4a				
low psychopathy	.26	[0.20, 0.30]	.69	[0.64, 0.74]
high psychopathy	.21	[0.16, 0.26]	.63	[0.57, 0.68]
Study 4b				
low psychopathy	.27	[0.22, 0.32]	.75	[0.70, 0.80]
high psychopathy	.13	[0.08, 0.18]	.53	[0.48, 0.58]

Figure 1. Multinomial processing tree 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.

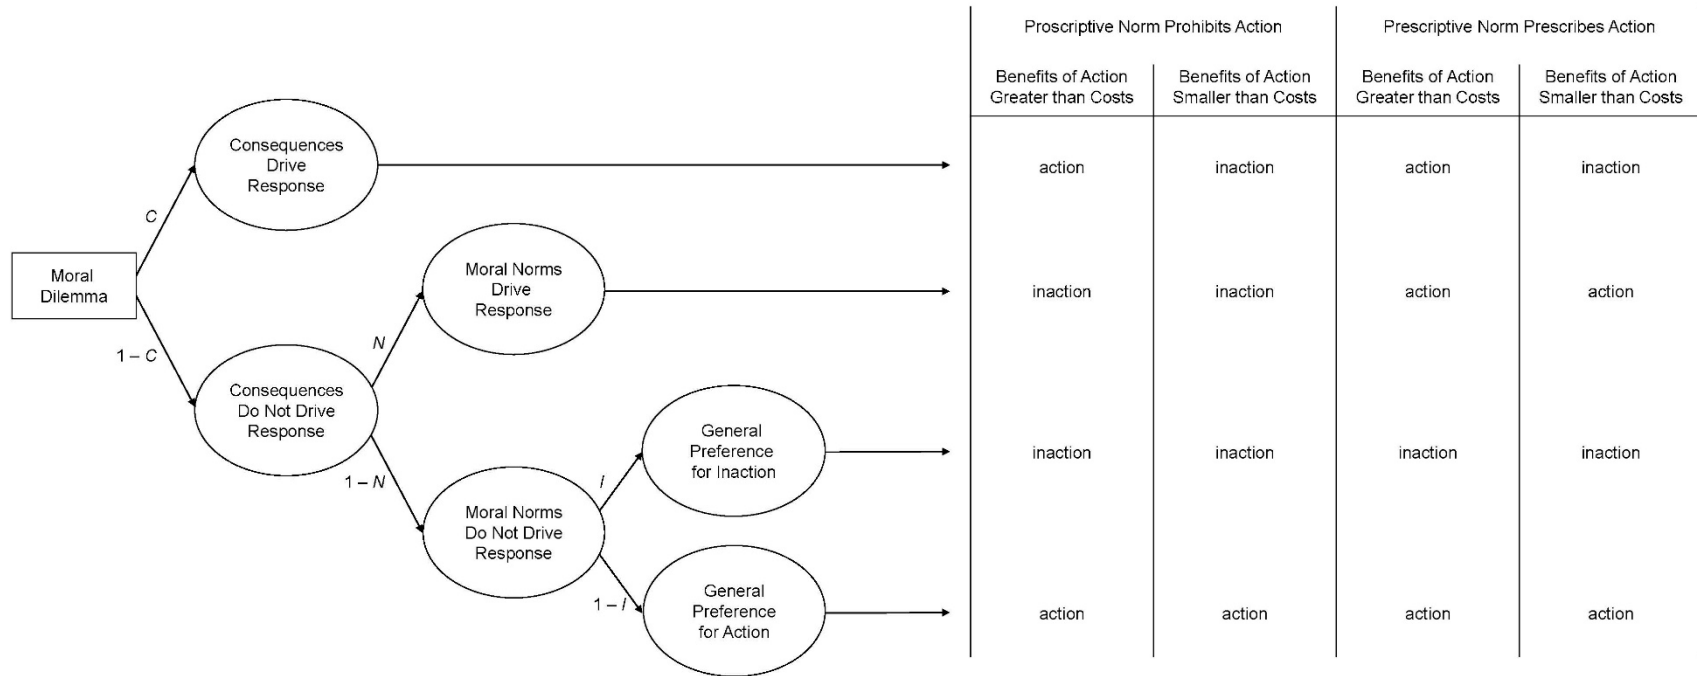


Figure 2. Parameter estimates of sensitivity to consequences (*C*), sensitivity to norms (*N*), and general preference for inaction versus action (*I*) for men and women, Study 1a. Error bars depict 95% confidence intervals.

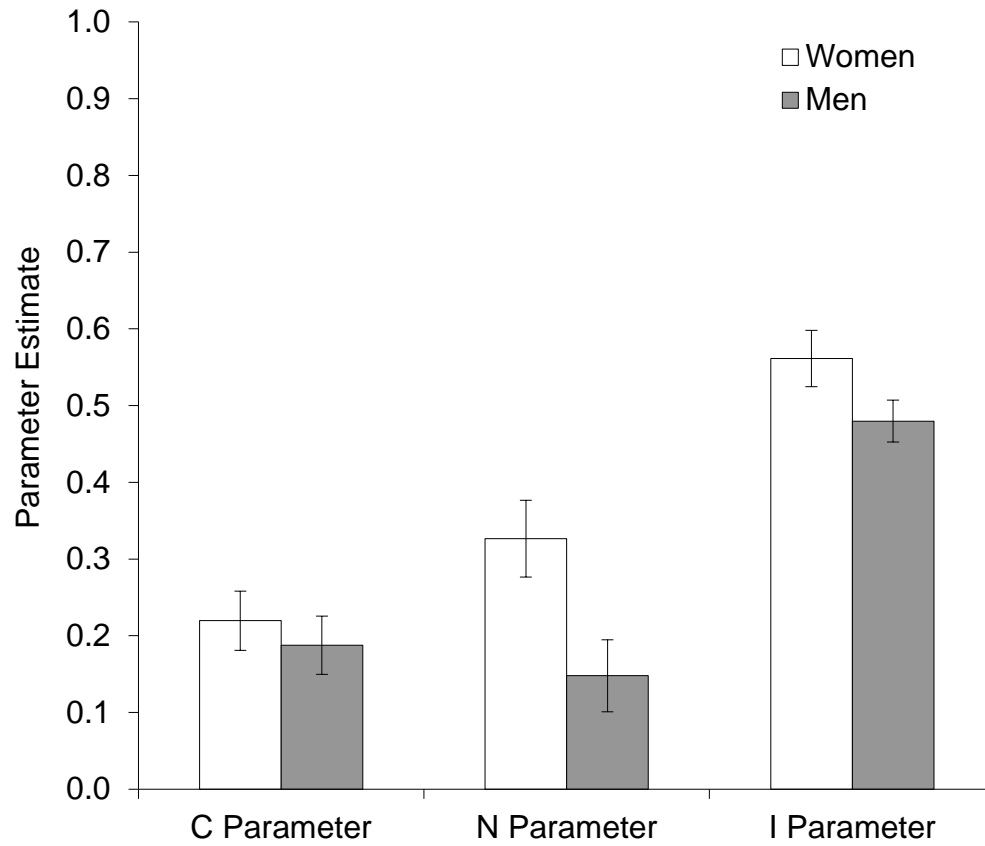


Figure 3. Parameter estimates of sensitivity to consequences (*C*), sensitivity to norms (*N*), and general preference for inaction versus action (*I*) for men and women, Study 1b. Error bars depict 95% confidence intervals.




Figure 4. Parameter estimates of sensitivity to consequences (*C*), sensitivity to norms (*N*), and general preference for inaction versus action (*I*) as a function of cognitive load (low load vs. high load), Study 2a. Error bars depict 95% confidence intervals.




Figure 5. Parameter estimates of sensitivity to consequences (*C*), sensitivity to norms (*N*), and general preference for inaction versus action (*I*) as a function of cognitive load (low load vs. high load), Study 2b. Error bars depict 95% confidence intervals.




Figure 6. Parameter estimates of sensitivity to consequences (*C*), sensitivity to norms (*N*), and general preference for inaction versus action (*I*) as a function of question framing (moral judgment vs. moral action), Study 3a. Error bars depict 95% confidence intervals.




Figure 7. Parameter estimates of sensitivity to consequences (*C*), sensitivity to norms (*N*), and general preference for inaction versus action (*I*) as a function of question framing (moral judgment vs. moral action), Study 3b. Error bars depict 95% confidence intervals.

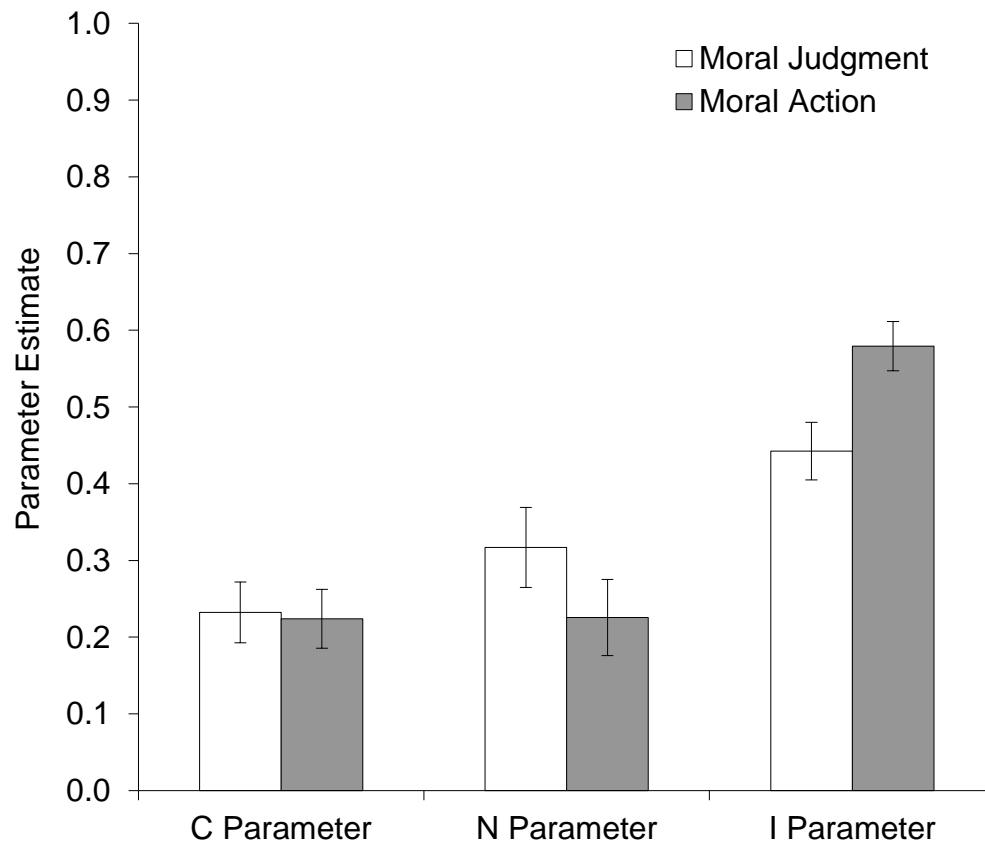


Figure 8. Parameter estimates of sensitivity to consequences (*C*), sensitivity to norms (*N*), and general preference for inaction versus action (*I*) as a function of psychopathy (low vs. high), Study 4a. Error bars depict 95% confidence intervals.

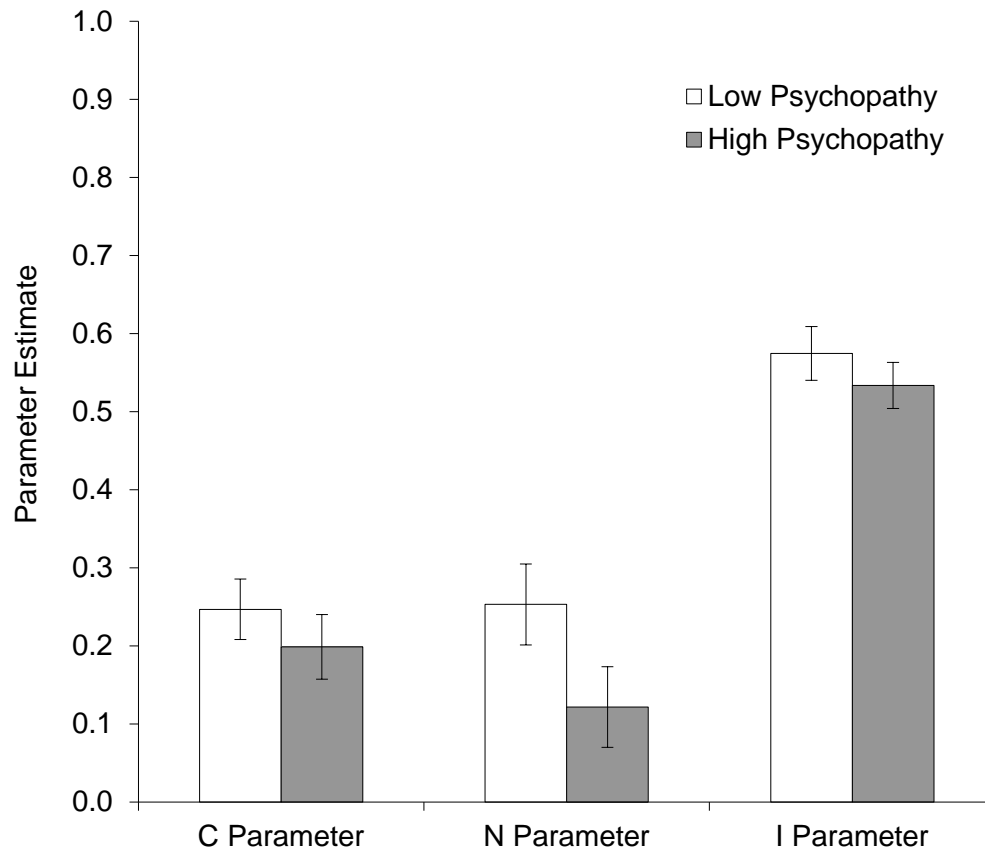


Figure 9. Parameter estimates of sensitivity to consequences (*C*), sensitivity to norms (*N*), and general preference for inaction versus action (*I*) as a function of psychopathy (low vs. high), Study 4b. Error bars depict 95% confidence intervals.

