



Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

COGNITION

Cognition xx (2005) 1–13

www.elsevier.com/locate/COGNIT

Moral dilemmas and moral rules

Shaun Nichols*, Ron Mallon

Department of Philosophy, University of Utah, Salt Lake City, UT 84112, USA

Received 12 January 2005; accepted 29 July 2005

Abstract

Recent work shows an important asymmetry in lay intuitions about moral dilemmas. Most people think it is permissible to divert a train so that it will kill one innocent person instead of five, but most people think that it is not permissible to push a stranger in front of a train to save five innocents. We argue that recent emotion-based explanations of this asymmetry have neglected the contribution that rules make to reasoning about moral dilemmas. In two experiments, we find that participants show a parallel asymmetry about versions of the dilemmas that have minimized emotional force. In a third experiment, we find that people distinguish between whether an action violates a moral rule and whether it is, all things considered, wrong. We propose that judgments of whether an action is wrong, all things considered, implicate a complex set of psychological processes, including representations of rules, emotional responses, and assessments of costs and benefits.

© 2005 Published by Elsevier B.V.

Keywords: Moral judgement; Moral rules; Normative ethics; Trolley problems

1. Moral dilemmas and descriptive normative ethics

These are good times for moral psychology. There is now a long tradition of excellent research on the capacity to distinguish moral from conventional violations (e.g., Blair, 1995; Nucci, 2001; Smetana, 1993; Turiel, 1983), and a complementary new tradition is emerging on the psychological factors involved in assessing moral dilemmas. In both traditions, one of the most exciting findings has been that emotions play a critical role in moral judgment (e.g., Blair, 1995; Greene, Sommerville, Nystrom, Darley, & Cohen 2001). However, the enthusiasm for emotion-based explanations has led researchers to

* Corresponding author. Tel.: +1 801 581 5197.

E-mail address: snichols@philosophy.utah.edu (S. Nichols).

neglect the role of rules in moral judgment. Elsewhere, one of us has pressed this objection against Blair's account of moral judgment (Nichols, 2002a, 2004). Here, we will argue that the contribution of rules is also overlooked in the most prominent account of the psychological processing involved in judging moral dilemmas.

The tradition of work on moral dilemmas grows out of a large body of research in philosophy. The philosophical project is to consider our intuitions about a wide range of dilemmas and to determine a set of principles that captures our intuitions about the cases.¹ The most intensively studied dilemmas are the 'trolley cases', which serve to isolate different factors that might affect our intuitions. In the *bystander* case, we are asked to imagine that a person sees a train approaching that will kill five innocents on the track, and the only way to prevent the deaths of these five is to flip a switch that will divert the train to a side track. Diverting the train to the side track will lead to the death of the person on the side track, but it is the only way to save the five people on the main track. Philosophers have maintained that the intuitive position is that it is acceptable to flip the switch to divert the train, leading to the death of one instead of five. In the *footbridge* case, the situation is quite similar except that there is no side track, and the only way for the protagonist to save the five is to push a large stranger off of a footbridge in front of the oncoming train, which will kill the stranger. In this case, philosophers maintain that the intuitive position is that it is wrong to push the stranger (Foot, 1967; Quinn, 1989; Thomson, 1976). One goal of the philosophical investigations has been to develop a unified normative theory that will accommodate our intuitions about such cases. This goal has been exceedingly difficult to meet, and few would maintain that philosophers have succeeded in finding a unified normative theory that fits with all of our intuitions about moral dilemmas.

This work in philosophy was unapologetically *a priori*, but recently researchers have conducted interview and survey experiments with these sorts of cases (Greene et al., 2001; Hauser, Cushman, Young, Jin, & Mikhail forthcoming; Mikhail 2000). The results have largely confirmed what philosophers maintained about the bystander and footbridge cases: most people have the intuition that it is acceptable to flip the switch in *bystander* but that it is not acceptable to push the stranger in *footbridge*. The interesting subsequent question concerns the psychological underpinnings of these judgments. On a simple utilitarian calculus, the cases seem quite parallel: five lives can be saved for the price of one. So why do people judge pushing the stranger as inappropriate but turning the train as appropriate? What are the psychological factors that contribute to our judgments about moral dilemmas? Here the empirical work promises new insight into the psychology of morality.

2. Emotion-based accounts

Although philosophers have not produced a unified normative theory that accommodates all of our intuitions about moral dilemmas, the empirical work has reinvigorated the investigation of moral dilemmas. In an important recent discussion,

¹ In addition to this descriptive philosophical project, there is a related prescriptive project, which attempts to characterize the normative theory that should guide our judgments in these cases.

Joshua Greene proposes that the response in footbridge-style cases is generated by the fact that these actions are ‘personal’ and such actions generate greater emotional engagement than ‘impersonal’ actions. The personal/impersonal distinction is drawn as follows:

A moral violation is personal if it is: (i) likely to cause serious bodily harm, (ii) to a particular person, (iii) in such a way that the harm does not result from the deflection of an existing threat onto a different party... A moral violation is impersonal if it fails to meet these criteria. ...Pushing someone in front of a trolley meets all three criteria and is therefore ‘personal’, while diverting a trolley involves merely deflecting an existing threat, removing a crucial sense of ‘agency’ and therefore making this violation ‘impersonal’ (Greene forthcoming; Greene & Haidt 2002, 519).

Before we proceed with this distinction, we need to clarify one aspect of Greene’s description. What does ‘moral violation’ mean here? It does not seem to mean *transgression* because diverting the trolley in the bystander case is cast as an ‘impersonal moral violation’, but it is doubtful that that action is a transgression at all. Indeed, if transgressions are measured by judgments of permissibility, the available evidence indicates that diverting the train is not a transgression—in all of the extant studies, a majority of participants judge it to be permissible (Greene et al., 2001; Hauser et al. forthcoming; Mikhail 2002). Since ‘violation’ is easily confused with ‘transgression’, we want to set aside this terminology.

This leaves us with the following explanation of the responses on the footbridge case:

“If an act is manifestly personal, then it is judged impermissible.” We will call this the ‘Personal Hypothesis’. This hypothesis would explain why pushing the stranger elicits impermissibility judgments, for the action is clearly personal. Furthermore, as noted above, Greene maintains that a crucial feature of personal acts is that they elicit strong emotions.² And indeed, Greene et al. found that emotional processing plays a distinctive role when people consider personal acts like the footbridge dilemmas (Greene et al., 2001).

Despite the intriguing evidence in favor of the Personal Hypothesis, there are numerous *prima facie* counterexamples, cases in which manifestly personal (and emotionally salient) acts are not judged impermissible. Some acts of self-defense, war, and punishment are plausibly personal and emotional, but regarded as permissible nonetheless. For instance, many people think that spanking their own child is permissible, even though it is obviously personal and emotional. Similarly, there is cultural variation in the harms that are judged impermissible. Among Yanomamö men, wife beating is judged permissible, despite being personal and emotional (Chagnon, 1992). Closer to home, in much of Western culture, male circumcision is permissible though an informal survey of our colleagues suggests it is regarded as personal. So the Personal Hypothesis is threatened by a wide range of apparent

² Jesse Prinz offers a related account of the responses to moral dilemmas (forthcoming). Prinz maintains that emotional response plays the key role in generating the different responses to footbridge and bystander cases, but he does not rely on the personal/impersonal distinction.

counterexamples. Of course, that would hardly be a decisive problem in the absence of a promising competitor. But there is in fact a venerable competitor.

3. Traditional rule-based accounts

According to traditional rule-based accounts of morality, an action is wrong if it violates a moral rule. There are quite different versions of such rule-based accounts. Deontological accounts, according to which certain types of actions are intrinsically wrong, provide one familiar example (Kant, 1785/1964; Ross, 1930). But even some utilitarian theorists are committed to the claim that an action is wrong if it violates a moral rule. According to some ‘rule utilitarian’ theorists, it is morally wrong to violate a rule that is justified by its consequences (e.g., Brandt, 1985). These traditional rule-based accounts naturally suggest a psychological proposal: an action is judged to be morally impermissible if the action violates a moral rule that is embraced by the judge. This approach can give an obvious explanation for why personal acts like self-defense, punishment, and circumcision are not judged impermissible. The judge does not embrace a rule against them.

In addition to this rich philosophical tradition, there is a rich empirical tradition that adverts to rules as essential to certain normative judgments. In the literature on the moral/conventional distinction, it is widely agreed that at least for judgments of conventional violations (e.g., standing up during story time), these judgments depend on knowledge of local rules (see, e.g., Turiel, Killen, & Helwig 1987). Thus there is independent reason to think that people make at least some normative judgments by drawing on their knowledge of rules. In addition, by appealing to agents’ knowledge of local rules, we get an obvious explanation for cross-cultural differences in normative judgments. For example, people in the US but not people in China would think it wrong not to tip servers in local restaurants. The obvious explanation for this difference is that people in the US embrace a rule about tipping (in the US) and people in the China do not embrace that rule about tipping (in China). Thus, there is independent reason to think that rules do play a role in at least some normative judgments.

The traditional rule-based approach does owe an answer to the original challenge, however. Why do people judge that choosing five over one is acceptable in the bystander case but not in the footbridge case? The traditional answer is that the different status of these actions is explained by what the rules do and do not forbid. One important proposal is that a rule like “Do not kill persons” forbids one from intending to bring about a person’s death, but it does not necessarily forbid acting in a way that brings about a person’s death as an unintended but foreseen side effect. Hence, in the bystander case, it is permissible to divert the trolley even though it has an effect (the killing of an innocent) which it would be impermissible to intend. There has been protracted debate on how exactly to characterize what moral rules do and do not forbid (Foot, 1967; Quinn, 1989; Thomson, 1976). Sorting all this out has been enormously complicated and has not produced any consensus, but we need not delve into this debate. For the important point is simply that the traditionalist maintains that we can explain the intuitions about the trolley cases in terms of what the rules do and do not forbid.

The rule-based approach offers an alternative to the emotion-based explanation of the asymmetry between the footbridge and bystander cases. However, in light of the apparent failure of philosophers to achieve consensus on how to characterize what the rules do and do not forbid, the rule-based explanation of the asymmetry between the footbridge cases and bystander cases might seem ad hoc. Is there an independent way to support the claim that the asymmetry between the footbridge cases and the bystander cases is explained by what the rules do and do not forbid? Our hypothesis is that it is a common feature of many rules, not specific to personal contexts, that they exhibit the asymmetry reflected in the footbridge and bystander cases. The experiments below were designed to test this hypothesis. Our prediction was that in *impersonal* scenarios with minimized emotional content, the asymmetry between footbridge-style cases and bystander-style cases will still occur. This would support the traditionalist's interpretation of the results on the original trolley cases. However, there is an important possible complication for the traditional account. For even if an action is thought to violate a rule, it might also be regarded as acceptable, *all things considered*. As a result, we wanted to explore this possible complication in our experiments as well. Judgments that an action violated a rule will be called judgments of "weak impermissibility". Judgments that an action was wrong, all things considered, will be called judgments of "all-in impermissibility".

4. Moral dilemmas and rules: some empirical results

4.1. Experiment 1

This experiment investigated whether the footbridge/bystander distinction would be drawn in *impersonal* dilemmas, and also whether participants would treat some actions as weakly impermissible but not all-in impermissible.

4.1.1. Method

4.1.1.1. Participants. Thirty-nine students from an introductory philosophy course at the University of Utah participated in this study. Eleven participants were female; 28 were male.

4.1.1.2. Materials. Four scenarios were used in this study. The first two scenarios were impersonal analogues of the bystander and footbridge cases. The last two scenarios were versions of the personal bystander and footbridge cases. In the bystander case, the protagonist switches the track, which leads to the death of one instead of five. In the footbridge case, the protagonist throws the man off the footbridge, leading to the death of one instead of five. In the impersonal cases, teacups were substituted for people.

Impersonal bystander case: When Billy's mother leaves the house one day, she says "you are forbidden from breaking any of the teacups that are on the counter." Later that morning, Billy starts up his train set and goes to make a snack. When he returns, he finds that his 18 month old sister Ann has taken several of the teacups and placed

them on the train tracks. Billy sees that if the train continues on its present course, it will run through and break five cups. Billy cannot get to the cups or to the off-switch in time, but he can reach a lever, which will divert the train to a side track. There is only one cup on the side track. He knows that the only way to save the five cups is to divert the train to the side track, which will break the cup on the side track. Billy proceeds to pull the lever and the train is diverted down the side track, breaking one of the cups.

Impersonal footbridge case: When Susie’s mother leaves the house one day, she says “you are forbidden from breaking any of the teacups that are on the counter.” While Susie is playing in her bedroom, her 18 month old brother Fred has taken down several of the teacups and he has also turned on a mechanical toy truck, which is about to crush 5 of the cups. As Fred leaves the room, Susie walks in and sees that the truck is about to wreck the cups. She is standing next to the counter with the remaining teacups and she realizes that the only way to stop the truck in time is by throwing one of the teacups at the truck, which will break the cup she throws. Susie is in fact an excellent thrower and knows that if she throws the teacup at the truck she will save the five cups. Susie proceeds to throw the teacup, which breaks that cup, but it stops the truck and saves the five other teacups.

After each of the impersonal scenarios, two questions were asked, a weak permissibility question (e.g., “Did Susie break her mother’s rule?”) and an all-in permissibility question (e.g., “All things consider, was it okay for Susie to throw the teacup?”). After each of the personal cases, participants were asked whether the protagonist broke a moral rule and also whether all things considered, it was okay for the protagonist to act as he did.

4.1.1.3. Procedure. Participants were given questionnaires in a classroom. The impersonal cases were counterbalanced.

4.1.2. Results

As expected, in the personal cases, participants were more likely to say that the actor broke a moral rule in the footbridge case than in the bystander case (Sign test, $N=39$, $P<0.01$). But the real test comes with impersonal scenarios. Here too we found a distinction between responses on the footbridge-style case and the bystander-style case.

Participants were more likely to say that the rule was broken in the impersonal footbridge case than in the impersonal bystander case (Sign test, $N=39$, $P<0.001$ (two-tailed)). (The frequency of response patterns is shown in [Table 1](#).)

Table 1
Frequency of responses to impersonal cases in experiment 1

Footbridge	Bystander	
	Broke rule	Did not break rule
Broke rule	22	12
Did not break rule	0	5

Thus our prediction was confirmed. We also found that participants recognize a distinction between weak impermissibility and all-in permissibility. In the impersonal footbridge case, the vast majority maintained that it was all-in permissible for the protagonist to throw the teacup (92%). Indeed, 72% of the participants maintained both that the protagonist broke a rule and that it was all-in permissible, which differs significantly from what would be expected by chance alone (χ^2 goodness-of-fit (1, $N=39$)=7.410, $P<0.01$ two-tailed). However, we also found marginal effects of order. For instance, participants who got the impersonal bystander case second were marginally more likely to think that the person broke a rule in this case (χ^2 (1, $N=39$)=2.839, $P<0.1$ two-tailed).³ In light of the marginal order effects in this study, it was important to investigate the phenomena with a between subjects design.

4.2. Experiment 2

This experiment again investigated whether the impersonal footbridge and bystander cases would be treated differently.

4.2.1. Method

4.2.1.1. Participants. Fifty-two students from an introductory philosophy course at the University of Utah participated in this study. Thirteen participants were female; 39 were male.

4.2.1.2. Materials. There were two questionnaires, each consisting of a single scenario, either the impersonal bystander case or the impersonal footbridge case.⁴ After each scenario, the weak permissibility and all-in permissibility questions were asked.

4.2.1.3. Procedure. Participants were given the questionnaires in a classroom. Participants in condition 1 ($N=27$) were given the impersonal footbridge case. Participants in condition 2 ($N=25$) were given the bystander case.

4.2.2. Results

Both of the results on the impersonal cases were replicated. 96% of the participants said that a rule was broken in the footbridge case, but only 44% said that a rule was broken in the bystander case. This difference was quite significant (χ^2 (1, $N=52$)=17.296, $P<0.0001$ two-tailed, $\Phi^2=0.33$). See [Table 2](#) for frequency of response patterns.

Participants also recognized the distinction between weak and all-in impermissibility. Although most participants claimed that the rule was broken in the impersonal footbridge case, they also claimed that it was, all things considered, okay. 85% of the participants said both that the actor broke a rule and that it was, all things considered, okay. This differs

³ After running this experiment, we learned from Marc Hauser that he and his colleagues have found numerous order effects in their work on trolley cases.

⁴ The bystander scenario was exactly the same as in experiment 1, the footbridge case was revised slightly to make it shorter.

Table 2
Frequency of responses to impersonal cases in experiment 2

	Footbridge	Bystander
Broke rule	26	11
Did not break rule	1	14

significantly from what would be expected by chance (χ^2 goodness-of-fit (1, $N=27$) = 13.370, $P < 0.001$ two-tailed). And again, this suggests that participants appreciate the distinction between weak and all-in impermissibility.

4.2.3. Discussion

The foregoing experiments indicate that something like the asymmetry between footbridge and bystander is observed even in impersonal cases. This provides independent reason to think that traditional rule-based accounts illuminate the asymmetric responses on the personal bystander and footbridge cases. Namely, the judgment of impermissibility in the footbridge cases is guided by a moral rule that is not violated in the bystander cases. The experiments also indicate that participants recognize a distinction between weak impermissibility and all-in impermissibility. However, this distinction was only observed in the impersonal cases. In experiment 1, in the personal footbridge cases, most participants (70%) said both that a moral rule was broken and also that all things considered, it was not okay. So the experiments do not address whether people distinguish weak and all-in impermissibility in personal and moral contexts. Furthermore, in philosophical ethics, one important view, absolutist deontology, maintains that if an action violates a moral rule, it is thereby the wrong thing to do, all things considered (e.g., Fried, 1978). That is, when it comes to moral rules, there is no distinction between weak and all-in impermissibility. In addition, some rule-utilitarians also endorse the primacy of moral rules, even in the face of consequences that favor breaking the rule. On such views, weak impermissibility is (at least typically) sufficient for all-in impermissibility. If either account is a correct view of the role of moral rules in actual moral judgments, then we have a very plausible explanation of the original footbridge results. In those case, the action violates a moral rule, and violating a moral rule suffices for generating a judgment of all-in impermissibility. Experiments 1 and 2 fail to show that even in moral contexts, people can recognize a distinction between weak and all-in permissibility. The following experiment was designed to test this.

4.3. Experiment 3

This experiment draws on a traditional problem that pits absolutist deontology against utilitarianism. The problem emerges when we consider cases in which catastrophic consequences will ensue unless one breaks a moral rule (see e.g., Davis, 1991; Nagel, 1972). We developed such a catastrophe case modeled on the footbridge case. Unless one pushes the stranger, terrible consequences will follow. The existing results on the footbridge cases led us to expect that people would maintain that the action violated a

moral rule; but we also thought that with such dramatically bad consequences, participants might be less likely to maintain that the action was wrong, all things considered.

4.3.1. Method

4.3.1.1. *Participants.* Twenty-five students from an introductory philosophy course at the University of Utah participated in this study. Ten participants were female; 15 were male.

4.3.1.2. *Materials.* A new scenario was devised for this experiment, a “catastrophe” case, modeled on the footbridge case. In this case, catastrophically bad consequences will follow unless the protagonist pushes the stranger in front of the train.

Catastrophe case: A train is transporting an extremely dangerous artificially produced virus to a safe disposal site. The virus is profoundly contagious and nearly always leads to the death of the victim within a matter of weeks. If the virus were to be released into the atmosphere, billions of people would die from it. Indeed, there is a chance that it will wipe out more than half of the human population. Jonas is one of the scientists who was responsible for ensuring that the virus would be destroyed, and he is watching the train from a footbridge. As the train is approaching he sees through his binoculars that there is a powerful bomb planted on the tracks ahead, and there is no way for him to communicate with the train operators to get them to stop the train in time. If the train passes over the bomb, it will explode and the virus will be released into the environment with catastrophic consequences. There is a large stranger looking over the footbridge next to Jonas. Jonas knows that the stranger has nothing to do with the bomb, but the only way to stop the train from hitting the bomb is to push this stranger over the railing. For unlike Jonas’s body, the stranger’s body is big enough that it will bring the train to a halt, although this will kill the stranger. Jonas proceeds to push the stranger over the railing, which kills the stranger, but it prevents the explosion and saves billions of people from dying from the virus

After the scenario, participants were asked a weak impermissibility question, “Did Jonas break a moral rule by pushing the stranger over the railing?”. They are then asked an all-in impermissibility question, “All things considered, did Jonas do the wrong thing?”.

4.3.1.3. *Procedure.* Participants were given the questionnaires in a classroom.

4.3.2. Results

Sixty-eight percent of participants said that the actor broke a moral rule, but only 24% said that the action was, all things considered, the wrong thing to do. A sign test indicated that participants were significantly more likely to say that the actor broke a moral rule than that the action was all-in impermissible ($N=25$, $P=0.001$, two-tailed).

4.3.3. Discussion

These results reinforce the familiar problem posed by catastrophe cases: they indicate most people are not absolutist deontologists. People think that sometimes it is all-in permissible to do something that violates a moral rule, including the rule that forbids killing innocent people.

More importantly for our purposes, the results show that even in the case of moral rules, people appreciate a distinction between weak and all-in impermissibility.

The results also support the idea that there are two partly independent mechanisms underlying moral judgment. On the one hand, people have a general capacity to reason about how to minimize bad outcomes. On the other hand, people have a body of rules proscribing certain actions.⁵ This body of rules cannot be subsumed under the capacity to reason about how to minimize bad outcomes. For in the catastrophe experiment, participants judge that an action broke a moral rule despite the fact that the action obviously minimizes bad outcomes.

5. The psychology of judgments of all-in impermissibility

According to traditional rule-based accounts, people make different judgments in the footbridge and bystander cases because they embrace a rule that is violated in the footbridge case but not in the bystander case. As we saw in Section 3, traditional rule based accounts can easily accommodate cases like punishment and circumcision, which are personal but not impermissible (in some cultures). This provides one reason to favor the rule-based account over the Personal Hypothesis. Experiments 1 and 2 provide further support for the rule-based account. For even in the impersonal versions of those cases, participants say that a rule was broken in the footbridge-style case but not the bystander-style case.

However, the results on weak and all-in impermissibility pose a serious complication for the traditional rule-based theorist. For the connection between weakly morally impermissible and all-in morally impermissible is not straightforward. The psychological factors here are multifarious, but a judgment of weak impermissibility is obviously insufficient for generating a judgment of all-in impermissibility. For in the impersonal footbridge case and the catastrophe case, many participants maintained that the action was weakly but not all-in impermissible.

This leaves open difficult questions about the relationship between judgments of weak and all-in impermissibility and, more generally, about the psychological underpinnings of judgments of all-in impermissibility. In some cases, it is plausible that the judgments of all-in permissibility are generated by (roughly) utilitarian considerations. Thus, in the impersonal cases, the fact that more cups will be saved by breaking the rule makes it all-in permissible to throw the cup. Similarly, in the catastrophe cases presumably the overwhelmingly bad consequences lead people to judge it all-in acceptable to break the moral rule. But this brings us back to our original conundrum. In the personal footbridge case in experiment 1, most participants (about 70%) said both that pushing the stranger broke a moral rule and that it was, all things considered, not okay. Why does the moral rule trump the utilitarian considerations in these cases?

⁵ Although we think that rules make an independent contribution to reasoning about moral dilemmas, we also think that emotions play an important role in fixing the rules that we have (Nichols, 2002b, 2004).

Here is where we think the results of [Greene and colleagues \(2001\)](#) on affect and moral dilemmas are particularly revealing and instructive. They found that people show greater activation in brain areas associated with emotional processing when they consider footbridge cases than when they consider bystander cases. More importantly, Greene and colleagues found that the participants who judged it appropriate to push the stranger in the footbridge case showed longer reaction times, and the researchers offer the plausible interpretation that this is because judging it permissible to push the stranger is incongruous with the negative emotions elicited by the footbridge case. Thus, what the results indicate is that affect influences the all-in judgment in footbridge-style cases. But this does not mean that rules play no role. Rather, we suggest that judgments on the footbridge cases are guided by *affect-backed rules*: our all-in judgment to footbridge-style cases is a product of both rules and emotions. Emotional responses elevate the rule above the utilities. But the rules play an essential role in everyday moral judgment.

Thus, we propose that assessments of all-in impermissibility implicate three factors: cost/benefit analysis, checking for rule violations, and emotional activations. In both the personal and impersonal trolley cases, all the actions have favorable outcomes. The judgment of all-in impermissibility in such cases depends on both the presence of an emotion and the judgment that a rule has been violated. So Greene and colleagues are right to say that emotional activations play a key causal role. In the absence of emotion, the cost-benefit analysis typically wins. But the traditional rule-based theory is also right that rules play a key role. In the absence of proscriptive rules, the action is typically not treated as a candidate for being all-in impermissible. Typically, then, when a person judges an action as all-in impermissible despite its having favorable outcomes this depends on both emotional activation and on thinking that a rule has been violated. However, these two factors clearly do not necessitate the judgment that an action is all-in impermissible, for as illustrated by the catastrophe cases, when the cost-benefit ratio is sufficiently high, people tend not to judge the action as all-in impermissible.

This three factor account can explain all of the problem cases we have considered here. For instance, the account can explain why self-defense, punishment, circumcision, and other culturally permitted harms do not generate all-in judgments of inappropriateness despite the fact that they are personal and generate considerable emotional response. They do not count as impermissible even in the weak sense - there are no rules against them (in the culture). The account also explains why some rule violations e.g., the impersonal footbridge cases, only count as impermissible in the weak sense. That is because the prohibited actions do not engage the emotions. In such emotionally bland cases, cost/benefit considerations prevail in the all-in judgments. The account can also accommodate the catastrophe cases. When the consequences are overwhelmingly bad, the consequences can trump the moral rules.

We have proposed that judgments of all-in impermissibility implicate at least three separate psychological factors—representations of rules, assessments of costs and benefits, and emotional activations. It is probably unrealistic to expect a tidy processing account of how these factors interact to generate judgments of all-in impermissibility. But the fact that multifarious psychological factors impact judgments of all-in impermissibility brings us back to the difficulty philosophers have had in reaching a unified normative theory that captures our intuitions about moral dilemmas. If judgments of all-in

impermissibility arise from the interaction of a diverse collection of psychological mechanisms—representations of prohibitions, utilitarian assessments, and emotions—then it is probably misguided to expect that *there is* a single normative criterion that can capture our intuitions about moral dilemmas.

Acknowledgements

An earlier version of this paper was presented at the Fall 2004 workshop for the Moral Psychology Research Group. We would like to thank the participants for their useful feedback. We'd also like to thank Cynthia Stark and an anonymous referee for their helpful suggestions on an earlier version of the manuscript.

References

- Blair, R. (1995). A cognitive developmental approach to morality: Investigating the psychopath. *Cognition*, 57, 1–29.
- Brandt, R. (1985). *A theory of the good and the right*. Oxford: Clarendon Press.
- Chagnon, N. (1992). *Yanomamö* (4th ed.). New York: Harcourt Brace Jovanovich.
- Davis, N. (1991). Contemporary deontology. In P. Singer (Ed.), *A companion to ethics*, 205–218.
- Foot, P. (1967). *The problem of abortion and the doctrine of the double effect Oxford review* 5. Oxford: Oxford University Press (Reprinted in *virtues and vices*).
- Fried, C. (1978). *Right and wrong*. Cambridge, MA: Harvard University Press.
- Greene, J. (forthcoming). Cognitive neuroscience and the structure of the moral mind. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Structure and content*. New York: Oxford University Press.
- Greene, J., & Haidt, J. (2002). How (and where) does moral judgment work? *Trends in Cognitive Science*, 6(12).
- Greene, J., Sommerville, R., Nystrom, L., Darley, J., & Cohen, J. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, 293, 2105–2108.
- Hauser, M., Cushman, F., Young, L., Jin, R., & Mikhail, J. (forthcoming). A dissociation between moral judgments and justifications.
- Kant, I. (1785/1964). *Groundwork of the metaphysics of morals*. New York: Harper & Row (translated by H.J. Paton).
- Mikhail, J. (2000). *Rawls' linguistic analogy*. Unpublished PhD thesis, Cornell University.
- Nagel, T. (1972). War and massacre. *Philosophy and Public Affairs*, 1.
- Nichols, S. (2002a). Norms with feeling: Towards a psychological account of moral judgment. *Cognition*, 84, 221–236.
- Nichols, S. (2002b). On the genealogy of norms: A case for the role of emotion in cultural evolution. *Philosophy of Science*, 69, 234–255.
- Nichols, S. (2004). *Sentimental rules: On the natural foundations of moral judgment*. New York: Oxford University Press.
- Nucci, L. (2001). *Education in the moral domain*. Cambridge: Cambridge University Press.
- Prinz, J. (forthcoming). *The emotional construction of morals*. Oxford: Oxford University Press.
- Quinn, W. (1989). Actions, intentions, and consequences: The doctrine of double effect. *Philosophy and Public Affairs*, 18.
- Ross, W. (1930). *The Right and the good*. Oxford: Clarendon Press.
- Smetana, J. (1993). Understanding of social rules. In M. Bennett (Ed.), *The development of social cognition: The child as psychologist* (pp. 111–141). New York: Guilford Press.

- Thomson, J. (1976). Killing, letting die, and the trolley problem. *The Monist*, 59, 204–217.
- Turiel, E. (1983). *The development of social knowledge: Morality and convention*. Cambridge: Cambridge University Press.
- Turiel, E., Killen, M., & Helwig, C. (1987). Morality: Its structure, functions, and vagaries. In J. Kagan, & S. Lamb (Eds.), *The emergence of morality in young children* (pp. 155–244). Chicago: University of Chicago Press.