

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

More than a feeling: When emotional reactions don't predict moral judgments

Permalink

<https://escholarship.org/uc/item/7w29s0dt>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

ISSN

1069-7977

Authors

Horne, Zachary
Powell, Derek

Publication Date

2013

Peer reviewed

More than a feeling: When emotional reactions don't predict moral judgments

Zachary Horne (horne2@illinois.edu)
Department of Philosophy, 810 S. Wright Street
Urbana, IL 61801 USA

Derek Powell (derepowell@ucla.edu)
Department of Psychology, 1285 Franz Hall,
Los Angeles, CA 90095 USA

Abstract

Many moral psychologists have proposed that the difference between people's moral judgments about the Trolley and Footbridge dilemmas can be explained by their differing emotional responses to the dilemmas. In two experiments, we tested this explanation by presenting the dilemmas and measuring participants' reactions using a self-report emotion measure (PANAS-X). As might be expected, participants experienced more intense emotions after reading moral dilemmas when compared to a non-moral dilemma. However, participants' emotional reactions to the Trolley and Footbridge dilemmas did not differ. Our findings call the oft cited emotion explanation into question.

Keywords: moral psychology; emotion; PANAS-X; decision making.

Recently, research on moral dilemmas has crossed disciplines, branching out from ethics into psychology and cognitive neuroscience (e.g., Ciaramelli, Muccioli, Ládavas, & Di Pellegrino, 2007; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Greene, Nystrom, Engell, Darley, & Cohen, 2004; Koenigs, Young, Adolphs, Tranel, Cushman, Hauser, & Damasio, 2007). Within these lines of inquiry, there has been widespread interest in participants' judgments about two famous moral dilemmas: the Trolley dilemma and the Footbridge dilemma. In the Trolley dilemma, participants are told they are at the wheel of a runaway train that will cause the deaths of five workmen if it proceeds on its present course. The only way to avoid the deaths of these workmen is to hit a switch, redirecting the train to a side track where it will kill a single workman instead. This dilemma is often contrasted with another famous moral dilemma, the Footbridge dilemma, which asks participants to consider a situation in which they are on a footbridge, in between a runaway trolley and five workmen who will be killed if nothing is done. Participants are then told that the only way to save the lives of the five workmen is to push a stranger off the bridge and onto the tracks below where his large body will stop the trolley. Surprisingly, in the Trolley dilemma, approximately 80% of people judge that it is *appropriate* to take action (hit the switch, killing one person), yet in the Footbridge dilemma, about 80% judge that it is *inappropriate* to take action (push the man off the bridge, killing him). Given the apparent similarity of these moral situations, researchers have sought

to understand why people give such drastically different judgments about these cases.

Many researchers have suggested that these two dilemmas can be distinguished by the emotional reactions they elicit. Specifically, the Footbridge dilemma is thought to elicit strong negative emotional reactions, whereas the Trolley dilemma is thought to elicit weak negative emotional reactions. We will refer to this claim as the *Emotion Explanation*.

The Emotion Explanation has been interpreted as having wide-ranging implications for theories in psychology, cognitive neuroscience, and ethics. For instance, some psychologists have suggested that the Emotion Explanation is revealing of the psychological mechanisms recruited in moral judgment in general (e.g., Greene et al., 2001; Koenigs et al., 2007). Meanwhile, philosophers have employed the Emotion Explanation to diverse ends, using it to advance or undercut normative ethical theories, and to make arguments about the epistemic status of moral intuitions (e.g., Singer, 2005).

There are several reasons to take this explanation seriously. First, there is good evidence that emotions do importantly influence moral judgments (Haidt, 2001). Most directly, cognitive neuroscientists using functional Magnetic Resonance Imaging (fMRI) have reported finding activity in areas of the brain associated with emotional processing while participants made 'deontological' moral judgments, such as judging that it is inappropriate to push the man in the Footbridge dilemma (Greene et al., 2001, 2004). Researchers have also examined the moral judgments of participants with damage to the ventromedial prefrontal cortex (vmPFC), an area of the brain associated with affect. When given a battery of moral dilemmas, these participants gave 'utilitarian' moral judgments in the Footbridge dilemma, whereas healthy control participants tended to give deontological judgments (Koenigs et al., 2007).

Despite this evidence, there are a number of reasons to examine the Emotion Explanation more closely. First, it is not clear how the Emotion Explanation coheres with existing theories about morally relevant emotions, such as guilt. For example, social functionalist theories of emotion suggest that emotions like guilt serve both intrapersonal and interpersonal functions (Keltner, Haidt, & Shiota, 2006). On this view, guilt is typically elicited as a result of one's perceived transgressions against another, which can damage

relationships (Keltner, 1995). Upon contemplating some action, the experience of guilt might serve an intrapersonal function, deterring the potentially damaging action. Alternatively, if the action is ultimately taken, guilt can serve an interpersonal function, helping to repair whatever damage was done (Baumeister, 1994).

While it is clear that moral dilemmas can differ in the intensity of the emotions they evoke (e.g., dilemmas that ask participants to consider killing their child versus lying), the emotionally salient aspects of both the Trolley and Footbridge dilemmas are held constant. In both scenarios, participants are asked if they should *kill* a stranger in order to save five other strangers. Causing the death of another person seems likely to place a strain on interpersonal relationships, whatever one's motivation for doing so. Cast in this light, social functionalist theories of emotion seem to predict that people will experience feelings of guilt even as they approve of taking action in the Trolley dilemma.

Second, extant research in moral psychology provides only indirect support for the Emotion Explanation. The most direct support comes from fMRI investigations by Greene and colleagues (2001, 2004). In these studies, researchers presented participants with two large groups of moral dilemmas. One group was composed of the Trolley dilemma and 18 other dilemmas, and the other group was composed of the Footbridge dilemma and 24 other dilemmas. They called the dilemmas they considered Trolley-like 'impersonal', and the dilemmas they considered Footbridge-like 'personal', due to the 'closeness' of the action being performed in the dilemma. They found that the 'impersonal' dilemmas activated areas of the brain associated with deliberative cognition, whereas the 'personal' dilemmas activated areas of the brain associated with emotion. They regarded these results, in part, as evidence for the Emotion Explanation.

However, in order for the data from the studies by Greene and colleagues (2001, 2004) to provide compelling evidence for the Emotion Explanation per se, each of their 'impersonal' and 'personal' dilemmas would have to elicit emotions with valence and intensity similar to the Trolley and Footbridge dilemmas, respectively. It is doubtful that such parity was achieved. Dilemmas deemed 'impersonal' frequently did not involve physically harming someone (even though the Trolley dilemma clearly does), whereas those deemed 'personal' almost invariably did. In fact, only 10 of the 19 'impersonal' dilemmas involved harm, whereas 24 of the 25 'personal' dilemmas involved physically harming someone (Moore, Clark, & Kane, 2008). This difference is a problematic confound because moral situations concerning the bodily harm of another have been shown to elicit stronger negative emotional responses than those not involving any bodily harm (Heekeren, 2005). Misgivings about the representativeness of the dilemmas in these sets are further supported by an item analysis performed by McGuire and colleagues (2009) on Greene and colleagues' data. This analysis demonstrated that only a subset of the moral dilemmas was responsible for the

significantly different patterns of neural activation Greene and colleagues observed, again impugning evidence for the Emotion Explanation.

Given the confound in these materials, and the findings of the item analysis (McGuire, Langdon, Coltheart, & Mackenzie, 2009), we suspect that people experienced more intense emotions when considering 'personal' dilemmas because the dilemmas much more frequently involved physical harm, *not* because the Footbridge dilemma itself is more emotionally engaging than the Trolley dilemma.¹ In other words, Greene and colleagues' (2001) findings do not provide clear evidence for the Emotion Explanation, and thus corroboration of the Emotion Explanation requires a more direct test. These criticisms apply equally to a number of other studies by researchers who have used these same materials and have claimed to have found evidence for the Emotion Explanation (e.g., Ciaramelli et al., 2011; Koenigs et al., 2007). To our knowledge, there is no empirical data taken as evidence for the Emotion Explanation that avoids this criticism.

We sought to test the Emotion Explanation by examining people's emotional responses to the Trolley dilemma and Footbridge dilemma, *individually* and *specifically*. We measured participants' emotional responses to moral dilemmas using the PANAS-X, a comprehensive emotional state, trait, and mood self-report measure (Watson, Clark, & Tellegen, 1988) that has been shown to correlate with neural activation in the amygdala (Irwin, Davidson, Kalin, Sorenson, & Turski, 1998), as well as in the vmPFC (Zald, Mattson, & Pardo, 2002). Importantly, use of this self-report measure allowed us to investigate people's emotional responses to individual moral dilemmas, rather than to a battery of different dilemmas. More advanced methodologies such as GSR or fMRI, while in many ways superior to self-report, are ill-suited for investigating responses to individual stimuli.

In line with previous research, we hypothesized that considering the Footbridge dilemma would elicit negative emotions. In contrast to the Emotion Explanation, we predicted that considering the Trolley Dilemma would also elicit increased guilt, as well as other negative emotions. Moreover, because the Trolley and Footbridge dilemmas contain very similar emotionally-relevant content (they both call on participants to imagine killing another person), we expected to find very little difference in people's emotional reactions to these dilemmas. These last two predictions stand in contrast to the way moral psychologists have conceived of the Footbridge and Trolley dilemmas (e.g., Ciaramelli et al., 2007; Greene et al., 2001; Greene et al.,

¹ These concerns do not necessarily undermine the distinction between personal and impersonal dilemmas, for which there is independent evidence (Moore, Clark & Kane, 2008). The two claims are independent; one can accept the personal-impersonal distinction without thinking that people's moral judgments in these cases are driven by differences in the emotions the situations evoke.

2004; Koenigs et al., 2007); hence support for our predictions would cast doubt on the Emotion Explanation.

Experiment 1

Method

Participants The participants in Experiment 1 were 442 students enrolled in various undergraduate courses at Arizona State University. Approximately 54% of the participants were males. The mean age of participants was 20.1 years old.

Materials Four different scenarios were presented to participants in this study, with each participant reading one scenario. Three of these scenarios were moral vignettes, the 'Trolley,' the 'Footbridge,' and the 'Crying Baby' vignette, all taken verbatim from Greene et al. (2001).

One potential concern is that the PANAS-X subscales might not be sensitive enough to detect differences between moral dilemmas of different emotional intensity. To evaluate the sensitivity of our chosen PANAS-X subscales, we presented some participants with a highly emotional dilemma, the Crying Baby dilemma. In this dilemma, participants must consider smothering their own infant child to save the lives of their townspeople.

The Crying Baby dilemma also afforded us the opportunity to explore whether there was a relationship between emotions and moral judgments within an emotionally engaging case. Prior research shows that people tend to be evenly divided over whether or not it is appropriate to take action in the Crying Baby dilemma, permitting meaningful comparisons between those who approve and those who disapprove. The fourth dilemma was a non-moral control dilemma (Coupon) adapted from Greene et al. (2001).

In Experiment 1, we used selected sub-scales from the PANAS-X to measure the extent to which participants experienced several relevant emotions. Participants rated how strongly they felt certain feelings on a 1 to 5 Likert scale, where 1 corresponded to 'Very slightly or not at all' and 5 corresponded to 'Extremely.' Importantly, participants were asked to report on their current emotional state rather than to assess the vignette or to speculate as to how they might feel if placed in the situation they read. These individual ratings were then averaged to obtain a score for each emotion subscale. Data were collected for the Positive Affect, Negative Affect, Hostility, Guilt, Joviality and Self-Assurance sub-scales from the PANAS-X. The Guilt sub-scale included words like 'Guilty', 'Ashamed', and 'Disgusted With Self', and the Hostility sub-scale included words like 'Angry', 'Disgusted', and 'Hostile'. The Joviality subscale included words like 'Happy', 'Joyful', and 'Cheerful', and the Self-Assurance subscale included words like 'Proud', 'Strong', and 'Confident'. As the Positive and Negative Affect scales contain many of the same emotion words found in the Guilt, Joviality, Hostility and

Self-Assurance subscales, we do not discuss analyses of the Positive and Negative Affect scales.

Procedure Experiment 1 was administered and data were collected via pen-and-paper questionnaires. The questionnaires consisted of text describing a moral dilemma, followed by an emotion measure (i.e., the PANAS-X). Participants were asked to respond to the emotion measure after considering a statement of the form, 'You are thinking about (action) in order to (outcome).' For example, when considering the Trolley dilemma, participants would read the statement, 'You are thinking about hitting the switch in order to avoid the deaths of the five workmen.' The text indicating the outcome was italicized for all conditions. On the opposite side of the questionnaire participants were asked to make a judgment of appropriateness. In making their judgments of appropriateness, participants were asked, 'Is it appropriate for you to (action) in order to (outcome).' They were instructed to circle 'yes' or 'no.' Participants were also asked to describe the emotions they experienced while reading the vignettes, and were given several lines to write in their response. These descriptions were used to identify problematic responses or inconsistencies between participants' reported emotions on the PANAS-X. Finally, participants were asked to provide demographic information.

Experiment 1 was conducted during various lecture classes at Arizona State University. The questionnaires were passed out at the beginning of the class, and students were given approximately 10 minutes to complete them. Instructions were given both verbally and in writing. Participants were instructed not to talk to each other while completing the survey, and were observed for compliance.

Results

Of the 442 participants originally involved in the experiment, 12 gave written descriptions of their emotions that were primarily non sequitur or conflicted with the emotion ratings they had given on the PANAS-X. These 12 participants were removed from further analyses. Only participants who had complete data sets for all the emotion subscales were included in the final analysis. This resulted in 86 participants responding to the control vignette, 147 to the Crying Baby, 71 to the Footbridge, and 77 to the

Table 1: Participants' judgments across conditions (percent approval).

Control	Trolley	Footbridge	Crying Baby
92%	91%	25%	46%




Figure 1: Mean emotion ratings on each PANAS-X sub-scale for the four dilemmas in Experiment 1.

Trolley. Additional participants were included in the Crying Baby in order to achieve adequate statistical power for comparisons to be made within that condition. Mean subscale ratings for each vignette are shown in Figure 1.

Participants' judgments are summarized in table 1. Their responses accord with those obtained by previous investigations of these scenarios.

PANAS-X scores were examined with four ANOVA tests with vignette type as a between-subjects variable. Significance was obtained on the Guilt ($F(3, 377) = 62.57$, $MSE = 1.334$, $p < .001$), Hostility ($F(3,377) = 65.66$, $MSE = .953$, $p < .001$), Joviality ($F(3, 377) = 21.19$, $MSE = .453$, $p < .001$) and Self-Assurance ($F(3, 377) = 2.74$, $MSE = .737$, $p = .04$) subscales.

For each of these ANOVA tests, planned comparisons examined differences between the three moral vignettes and the control vignette in the emotion ratings on the Guilt, Hostility and Joviality subscales. The moral vignettes elicited significantly more Guilt ($t(377) = 8.69$, $p < .001$) and Hostility ($t(377) = 7.90$, $p < .001$), and significantly less Joviality ($t(377) = -6.83$, $p < .001$) than the control vignette. Though significance was obtained on the respective ANOVA analyses, no significant difference was found between the moral vignettes and control vignette for Self-Assurance ratings ($t(377) = -.475$, $p = .635$), suggesting that the Self-Assurance subscale was not relevant to moral judgments. No further analyses were conducted for Self-Assurance ratings.

As predicted, participants' ratings of Guilt, Hostility and Joviality were compared between the Trolley and Footbridge vignettes with three separate t-tests. These tests found no significant differences between the two groups for Guilt ($t(146) = -.09$, $p = .93$), Hostility ($t(148) = -.33$, $p = .74$) or Joviality ($t(148) = 1.24$, $p = .22$). This finding is consistent with our prediction that such highly similar moral

Table 2: Summary of Logistic Regression Models Predicting Moral Judgments in the Crying Baby Dilemma in Experiment 1.

	Model 1				Model 2			
	β	SE	e^{β}	p	β	SE	e^{β}	p
Hostility	-.21	.23	.81	.36	-.22	.24	.81	.37
Guilt	-.06	.19	.94	.76	-.02	.20	.98	.91
Joviality	.21	.40	1.24	.59	.41	.41	1.20	.66
Gender					1.19	.35	3.27	.001
χ^2	4.08				15.81			
df	3				4			
p	.253				.003			

scenarios were unlikely to elicit meaningfully different emotional reactions.

Participants who read the Trolley vignette reported greater levels of Guilt ($t(377) = 4.66$, $p < .001$) and Hostility ($t(377) = 4.22$, $p < .001$), and lower levels of Joviality ($t(377) = -5.39$, $p < .001$) than those who read the control vignette. Further planned comparisons tested whether the Crying Baby vignette was more emotionally salient than the Footbridge and Trolley vignettes. The Crying Baby vignette elicited more Guilt ($t(377) = 9.11$, $p < .001$) and Hostility ($t(377) = 10.23$, $p < .001$), and less Joviality ($t(377) = -2.79$, $p < .001$) than the Footbridge and Trolley vignettes. These differences indicate that the PANAS-X is sufficiently sensitive as an emotion measure to detect differences between moral dilemmas of different emotional intensity.

Finally, we used logistic regression analyses to assess the relationship between emotions and moral judgments in the Crying Baby dilemma. We evaluated this relationship using two different regression models. The resulting equations are summarized in Table 2. The first model uses the Guilt, Hostility, and Joviality subscale scores as predictors of people's moral judgments. No significant relationship was found between participants' emotion ratings and their moral judgments for the Crying Baby problem ($\chi^2(3) = 4.082$, $p = .25$). The second model adds gender to the three regressors from the original model. This model significantly predicted moral judgments ($\chi^2(4) = 15.813$, $p < .01$). Within this model, the gender coefficient was statistically significant ($\beta = 1.185$, $e^{\beta} = 3.271$, $p < .01$), indicating that women were less likely to judge it appropriate to smother the child. Importantly, this relationship between gender and moral judgments was significant over and above the emotion ratings.

Experiment 2

As the results of Experiment 1 stand in stark contrast to the Emotion Explanation, we sought to replicate the findings of Experiment 1 in a different population.

Method

The participants in Experiment 2 were 221 workers recruited using Amazon Mechanical Turk (mTurk). Approximately 53% of the participants were males. The mean age of participants was approximately 31 years old.

Materials The materials and measurements used in Experiment 2 were identical to those of Experiment 1, with two exceptions. The results of Experiment 1 indicated that participants' emotion ratings on the Self-Assurance subscale in moral conditions did not differ from participants in the control condition. Consequently, in Experiment 2, we only used the Guilt, Joviality and Hostility subscales to measure participants' emotional responses. In addition, rather than asking participants to give moral judgments as a binary 'yes/no' response, we asked them to rate the moral rightness or wrongness of the proposed act on a 1-6 Likert scale. The end points of the Likert scale were labeled 'Completely Inappropriate' and 'Completely Appropriate.' This change was made in order to increase our statistical power for detecting any potential relationship between emotions and moral judgments in the Crying Baby dilemma.

Collection Experiment 2 was administered and data was collected through the mTurk work-distribution website. Eligible workers were redirected to Qualtrics, where they completed the study. Afterwards, workers were directed back to mTurk, where they were compensated with \$.20.

Procedure The survey presented to participants in Experiment 2 was nearly identical to the survey in Experiment 1, save for being computerized rather than pen and paper based. It consisted of text describing a moral dilemma, followed by an emotion measure (PANAS-X). After rating their emotions, participants made a moral judgment about the dilemma presented to them.

Results

Experiment 2 replicated the findings from Experiment 1. In Experiment 2, 61 participants read the Coupon vignette, 51 read the Crying Baby, 57 read the Footbridge, and 52 read the Trolley. On average, participants judged it appropriate to use the coupon (mean = 5.51) and to switch the track in the Trolley dilemma (mean = 4.69). Participants were divided over whether to smother the baby (tending to disapprove somewhat, mean = 2.88), and in general disapproved of pushing the man (mean = 2.32).

An ANOVA revealed significant differences between the moral and control vignettes in terms of the emotions they elicited, with significance obtaining for the Guilt ($F(3, 217) = 31.86$, $MSE = 1.389$, $p < .001$), Hostility ($F(3, 217) = 52.57$, $MSE = .929$, $p < .001$), and Joviality ($F(3, 217) = 15.339$, $MSE = .824$, $p < .001$) subscales. Planned comparisons again revealed the sensitivity of the PANAS-X for measuring the different emotions elicited by different moral cases. The Crying Baby vignette elicited more Guilt

Table 3: Summary of Linear Regression Model Predicting Moral Judgments in the Crying Baby Dilemma in Experiment 2.

	<i>b</i>	SE	<i>t</i>	<i>p</i>
Guilt	-.020	.345	-.059	.954
Hostility	.298	.392	.758	.452
Joviality	.175	.336	-.521	.605
Intercept	2.216	1.038		
<hr/>				
F			.764	
df			47	
<i>p</i>			.52	
R ²			.047	

($t(217) = 4.08$, $p < .001$) and Hostility ($t(217) = 5.72$, $p < .001$) than did the Footbridge and Trolley vignettes. The same contrast for the joviality subscale approached significance ($t(217) = 1.81$, $p = .07$).

As was observed in Experiment 1, participants who read the Trolley vignette reported greater levels of Guilt ($t(111) = 6.05$, $p < .001$) and Hostility ($t(111) = 7.01$, $p < .001$), and lower levels of Joviality ($t(111) = -3.67$, $p < .001$) than those who read the non-moral control vignette. To provide a second test of the Emotion Explanation, participants' ratings of Guilt, Hostility and Joviality were compared between the Trolley and Footbridge vignettes with three separate *t*-tests. These tests again found no significant differences between the two groups for Guilt ($t(107) = -1.42$, $p = .16$), Hostility ($t(107) = -.34$, $p = .74$) or Joviality ($t(107) = 1.27$, $p = .21$).

The use of Likert scale ratings of moral judgments in Experiment 2 enabled us to assess the relationship between emotions and moral judgments in the Crying Baby dilemma using a simpler linear regression model. The model used the three PANAS-X sub-scale scores as predictors of people's moral judgments. The resulting equation is summarized in Table 3. Replicating our findings in Experiment 1, no significant relationship was found between participants' emotion ratings and their moral judgments for the Crying Baby dilemma ($F(3, 45) = .90$, $p = .45$). This model accounted for less than 5 percent of the variance in participants' Crying Baby judgments, indicating an *extremely* weak (and non-significant) relationship between emotions and moral judgments about the case.

Discussion

We hypothesized that the Trolley dilemma would be more emotionally engaging than many moral psychologists have claimed. Additionally, we suspected that because of the high degree of similarity between the Trolley and Footbridge dilemmas, any differences between the emotions elicited by the two dilemmas would not be sufficient to explain the large difference in participants' judgments about the dilemmas. The results of Experiments 1 and 2 confirm these two predictions. We found that participants reported significantly more guilt and hostility in response to the Trolley dilemma than to the control dilemma, and

significantly less joviality. Additionally, we found no difference in emotion ratings from participants who responded to the Footbridge dilemma and those who responded to the Trolley dilemma, despite the rather large sample size of our experiment.

To our knowledge, this is the first use of the PANAS-X to compare the emotions elicited by different moral dilemmas. Additionally, one of our primary hypotheses was that there would be no differences between the emotions elicited by the Footbridge and Trolley judgments. This made it important to confirm that the PANAS-X is sufficiently sensitive to detect emotional changes, not just between moral dilemmas and a non-moral dilemma, but also between different moral dilemmas. We found that participants who responded to the Crying Baby dilemma reported significantly stronger negative emotions and weaker positive emotions than did participants who responded to the Trolley and Footbridge dilemmas, demonstrating the sensitivity of the PANAS-X.

The Emotion Explanation, as stated, concerns the Trolley and Footbridge dilemmas specifically. However, this explanation has been employed to a number of different ends, not all of which necessarily hinge on facts about these specific dilemmas. In this way, some proponents of the Emotion Explanation might argue for a more general claim, viz., in general, deontological moral judgments recruit gut, emotional processes. In an emotionally salient case like the Crying Baby dilemma, proponents of this more general Emotion Explanation ought to predict a relationship between emotions moral judgments, such that the stronger the negative emotional reaction a participant experiences, the more likely they are to give a deontological judgment. Across two experiments, we observed no such relationship between emotions and moral judgments in the Crying Baby dilemma.

In sum, our findings indicate that the Trolley, Footbridge and Crying Baby dilemmas fail to conform to the more general Emotion Explanation. It is a question for future research whether these dilemmas are simply exceptions to the rule, or whether the more general Emotion Explanation should be rejected. Our findings cast doubt on the Emotion Explanation and may call for revision of the psychological, ethical, and epistemological theories in which it has been employed.

Acknowledgments

Preparation of this paper was supported by grant N000140810186 from the Office of Naval Research.

References

Baumeister, R. F., Stillwell, A. M., & Heatherton, T. F. (1994). Guilt: an interpersonal approach. *Psychological Bulletin*, *115*(2), 243-267.

Ciaramelli, E., Muccioli, M., Làdavas, E., & Di Pellegrino, G. (2007). Selective deficit in personal moral judgment following damage to ventromedial prefrontal cortex. *Social cognitive and affective neuroscience*, *2*(2), 84-92.

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*(2), 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*(5537), 2105-2108.

Haidt, J. (2001). The Emotional Dog and its Rational Tail: A Social Intuitionist Approach to Moral Judgment. *Psychological Review*, *108*(4), 814-834.

Heekeren, H. R., Wartenburger, I., Schmidt, H., Prehn, K., Schwintowski, H. P. et al. (2005). Influence of bodily harm on neural correlates of semantic and moral decision-making. *NeuroImage*, *24*(3), 887-897

Irwin, W., Davidson, R. J., Kalin, N. H., Sorenson, J. A. & Turski, P. A. 1998 Relations between human amygdala activation and self-reported dispositional affect. *J. Cogn. Neurosci. (Suppl. S)*, 109.

Keltner, D. (1995). Signs of appeasement: Evidence for the distinct displays of embarrassment, amusement, and shame. *Journal of Personality and Social Psychology*, *68*(3), 441-454.

Keltner, D., Haidt, J., & Shiota, M. N. (2006). Social Functionalism and the Evolution of Emotions. In Schaller, M., Simpson, J. A., & Kenrick, D. T. (Eds.) *Evolution and Social Psychology* (115-142). New York, NY: Psychology Press.

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*(7138), 908-911.

McGuire, J., Langdon, R., Coltheart, M., & Mackenzie, C. (2009). A reanalysis of the personal/impersonal distinction in moral psychology research. *Journal of Experimental Social Psychology*, *45*(3), 577-580.

Singer, P. (2005). Ethics and Intuitions. *The Journal of Ethics*, *9*.

Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect schedule. *Journal of Personality and Social Psychology*, *54*(6), 1063-1070.

Zald, D. H., Mattson, D. L., & Pardo, J. V. (2002). Brain activity in ventromedial prefrontal cortex correlates with individual differences in negative affect. Proceedings of the National Academy of Sciences of the United States of America, *99*(4), 2450-4.