Personality and Individual Differences 50 (2011) 621–625

Contents lists available at ScienceDirect

Personality and Individual Differences

journal homepage: www.elsevier.com/locate/paid

Individual differences in sensitivity to reward and punishment predict moral judgment

Adam B. Moore *, Julianna Stevens, Andrew R.A. Conway

Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08540, USA

ARTICLE INFO

Article history:
Received 8 September 2010
Received in revised form 23 November 2010
Accepted 3 December 2010
Available online 30 December 2010

Keywords:
Moral judgment
Individual differences
Reward sensitivity
Dual process model

ABSTRACT

Dual process models of moral judgment propose that such judgments are produced by interacting neural systems: a controlled cognitive system and an automatic affective system. Individual differences in moral judgment may therefore arise from variation in cognitive control ability and/or from variation in affective sensitivity. Previous research indicates that individual differences in cognitive control, indexed by working memory capacity, predict moral judgment (Moore, Clark, & Kane, 2008). Here we replicate group level findings from Moore et al. (2008) and demonstrate that individual differences in sensitivity to reward and punishment are strong predictors of moral judgment. Higher reward sensitivity positively correlates with willingness to sacrifice one life to save multiple others and moderates the impact of self-interest on participants' judgments. Higher punishment sensitivity negatively correlates with willingness to kill, particularly when negative affective information is present. These results help to revise current dual process models of moral judgment.

1. Introduction

On Wednesday, May 11, 2005 an unidentified aircraft violated the no-fly zone established around Washington, DC in the wake of the terror attacks on September 11, 2001. Alarmingly, the aircraft was flying directly at the White House and the pilot was responding neither to repeated attempts at radio contact by the Secret Service, nor to repeated fly-bys of F-16 jets. The fighter pilots dropped flares in the path of the small plane to get the pilot's attention, but to no avail. Security officials were faced with a high-stakes moral dilemma: Whether or not to order the F-16s to shoot down the Cessna, killing everyone inside and possibly innocents on the ground or risk a possible terrorist strike at key government centers. Fortunately the Cessna's pilot altered course, reportedly only seconds before the decision was made to shoot it down.

Recently, cognitive scientists have been studying how people make such decisions. Of course we do not ask participants to destroy planes, but we can ask people whether or not they find it morally acceptable to perform a harmful action in the service of some greater good (e.g. Greene, Nystrom, Engell, Darley, & Cohen, 2004; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Hauser, Cushman, Young, Jin, & Mikhail, 2007; Moore et al., 2008). The aircraft example highlights an important characteristic of the kinds of dilemmas used in moral psychology research, namely that there exists a tension between costs and benefits that subjects must evaluate in order to form a judgment. Much can be inferred from the patterns of subjects' responses, and neuroimaging methods have begun to elucidate brain mechanisms that underpin the processing that goes into considering these types of problems (Borg, Hynes, Van Horn, Grafton, & Sinnott-Armstrong, 2006; Greene et al., 2001, 2004). A leading account of the neuroimaging findings associated with moral judgment is the Dual Process Model (DPM; Greene et al., 2004), which proposes that moral judgments can be the product of either automatically elicited emotional responses or controlled cognitive processing. Each of these components has been linked to dissociable neural circuitries and separable impacts on the responses that subjects give to certain kinds of moral dilemmas (Greene et al., 2004, 2001). A core feature of this account is that moral judgments tend to systematically vary as a function of the emotional response elicited by the problem at hand. Concomitantly, experimental manipulations have shown that moral judgments are reliably affected by emotional priming manipulations (Valdesolo & DeSteno, 2006; Wheatley & Haidt, 2005).

Though considerable progress has been made in determining what experimental factors impact moral judgment, there has been very little work on what traits, intrinsic to the subjects themselves, may influence how people evaluate the costs and benefits involved in moral dilemmas. This is interesting since a key, perhaps defining, feature of human moral judgment is the large variability between individuals' judgments. The logic of the DPM suggests that those participants with greater cognitive control should be more utilitarian in their moral judgments. Moore et al. (2008) tested this by examining the relationship between moral judgment

* Corresponding author. Tel.: +1 609 258 5108.
E-mail address: amtwo@princeton.edu (A.B. Moore).

0191-8869/$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.paid.2010.12.006

and individual differences in cognitive control. They found that participants with greater cognitive control integrated abstract factors into their judgments more often, which resulted in more utilitarian responding in select cases. However, this effect was weak, and that experiment did not address individual differences in the affective component of the DPM. This latter issue is even more interesting because it may shed light on the larger question of when (or why) people engage cognitive control processes during moral judgment at all, given that there is a relatively fast, automatic affective system capable of generating such judgments independently. We address that question here by examining the role of individual differences in affective sensitivity to rewards and punishments (or, alternatively, gains and losses) and moral judgment.

1.1. Approach, avoidance, and the tradeoffs of moral judgment

The proposal that personality constructs can be formulated along functional and behavioral lines has engendered a view of brain systems as broadly organized into approach and avoidance related networks (Carver, Sutton, & Scheier, 2000; Carver & White, 1994; Gray, 2004). The overall function of these networks is thought to reflect, among other things, differences in the processing of emotional information (Gray, 1994; Gross, Sutton, & Ketelaar, 1998; Larsen & Ketelaar, 1991; Zelenski & Larsen, 1999; but see Harmon-Jones, 2003) and to underpin individual differences on two personality measures: *behavioral approach sensitivity* (BAS) and *behavioral inhibition sensitivity* (BIS; Carver et al., 2000; Carver & White, 1994; Gray, 1994). These constructs are psychometrically reliable and factor analysis supports their dissociation into two separate, but related, systems responsible for processing of (or sensitivity to) reward related/positive information (BAS) and punishment related/negative information (BIS; Carver & White, 1994; Jorm et al., 1998).

The logic of the current experiment is to determine if individual differences in the sensitivity of the BAS and BIS systems predict moral judgment in theoretically meaningful ways. Specifically, in moral dilemmas where harm to one person yields benefits to many others, the harm may be seen as a negative component of a complex stimulus (the dilemma) and the benefit as a positive component. The tradeoff of one against the other must be made within the context of the individual's sensitivity to each, and the key issue is that subjects are not equally sensitive to both kinds of information. Thus, we predict that a participant differentially more sensitive to positive gains than negative losses (i.e. high BAS but low BIS scores) should answer in a more utilitarian fashion, approving killing one to save many more often compared to one who has the reverse asymmetry. This effect should be particularly pronounced when comparing responses across personal and impersonal moral dilemmas; the former being those in which direct personal force (e.g. direct touching) must be applied to bring about the harm while the latter lack this requirement (Greene et al., 2009; Moore et al., 2008). Since the former generally have a lower moral approval rating, presumably due to the greater negativity associated with the use of personal force, the difference between moral approval of killing one to save many across personal and impersonal dilemmas should be much greater for those subjects with greater relative sensitivity to negative information (i.e. high BIS, low BAS scores).

2. Methods

2.1. Subjects

Participants were 77 undergraduates (48 female) from Princeton University who participated in exchange for partial fulfillment of course requirements.

2.2. Materials

2.2.1. BIS–BAS scales

Participants completed paper and pencil versions of Carver and White (1994) BIS–BAS scales. These consisted of 24 questions answered on a four point Likert scale and scores were computed in accordance with standard practice (see http://www.psy.miami.edu/faculty/ccarver/sclBISBAS.html for the full instrument, instructions, and scoring methods).

2.2.2. Moral dilemmas

We used the moral dilemma stimuli generated by Moore et al. (2008). Briefly, 24 dilemmas each had two resolutions, one personal and the other impersonal, all requiring the consideration of killing one person to save multiple others. These resolutions differed only in that personal resolutions required direct physical contact and impersonal resolutions enabled killing via a less direct, environmentally mediated route. In all other respects they were matched, including consequences of action and inaction, wording, number and length of sentences, and number of words (±2). Half of these critical dilemmas involved threat to the life of the participant as well as hypothetical others (self) while the other half involved threat only to others (other). Factorially crossed with this was the inevitability of death for the to-be-sacrificed person (inevitable) or the dependence on the participant's choice to kill them (avoidable). Thus, the dilemmas reflected a $2 \times 2 \times 2$ design, with four types of dilemmas (self-inevitable, self-avoidable, other-inevitable, other-avoidable) divided into an introductory paragraph and two possible conclusions (personal or impersonal).

Additionally, participants also judged 14 filler dilemmas. These were similar to the critical dilemmas except that two involved killing multiple people for the sake of only one other person, and 12 did not involve killing. Two lists were created, with every dilemma appearing on both lists but with only the personal or the impersonal conclusion. We randomly assigned participants to one of the two lists. All participants saw the same resolutions to the filler problems.

2.3. Procedure

We tested all participants individually. After giving informed consent, participants filled out the BIS–BAS instrument. The experimenter then read aloud on-screen instructions directing the participants to answer each scenario seriously, to disregard legal issues, and to respond only on the basis of the moral appropriateness of the described action. Participants completed two practice dilemmas and the experimenter then answered any questions. Dilemmas were presented as black text on a grey background via computer monitor. Presentation order was randomized across participants. Each dilemma appeared as an introductory paragraph that participants read silently. When finished, a key press caused the resolution to appear one sentence at a time; each successive key press revealed another resolution sentence until the final question asked about moral appropriateness (all text remained in view until participants made their judgment). Participants responded via two keys on a standard keyboard. The screen blanked for 1 s between dilemmas. The computer recorded final responses and response times for the final resolution sentence.

3. Results

All null-hypothesis significance tests were non-directional. Alpha was set at .05. Effect sizes are reported as η_p^2.

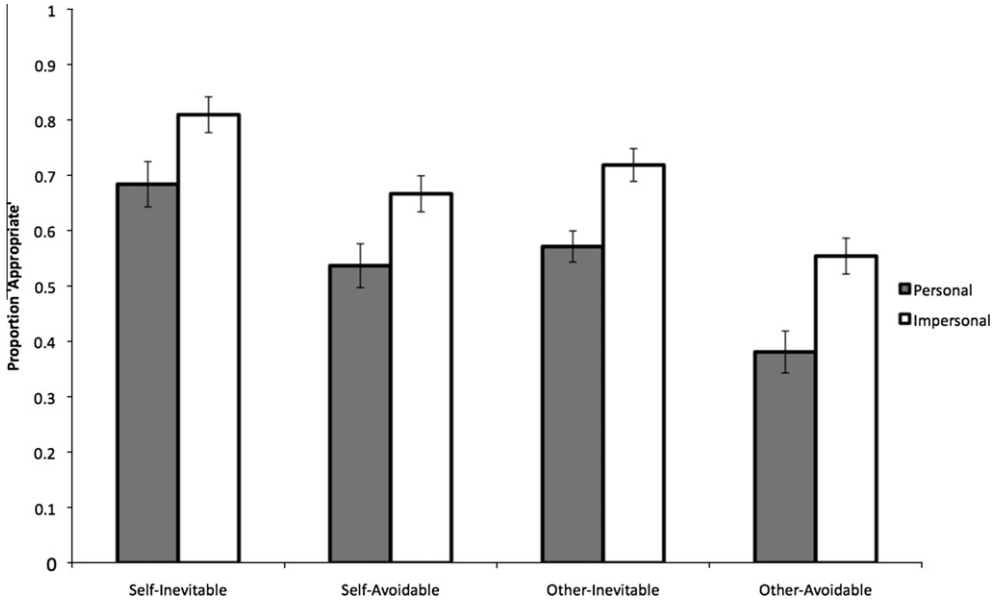

Fig. 1. Mean proportion of 'Appropriate' judgments by dilemma condition. Error bars are standard error of the mean.

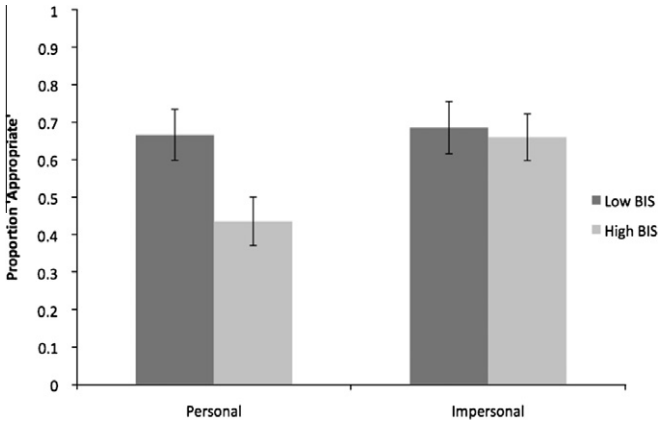

Fig. 2. Mean proportion of 'Appropriate' responses to moral dilemmas as a function of personal vs. impersonal killing. Extreme groups (defined as participants more than one standard deviation away from mean on BIS, $n = 13$ per group) are depicted to illustrate the interaction. The analysis was conducted on all participants. Error bars are standard error of the mean for each group shown.

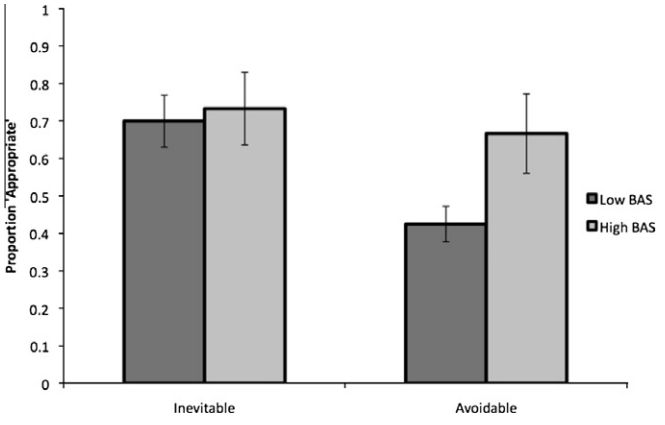

Fig. 3. Mean proportion of 'Appropriate' responses as a function of inevitable vs. avoidable dilemmas and BAS scores. Extreme groups (defined as the 20 participants with a BAS score one standard deviation or greater from the mean BAS score, $n = 10$ per group) are depicted to illustrate the pattern of the interaction. The analysis was conducted on all participants. Error bars are standard error of the mean of each group shown.

Table 1
Descriptive statistics and correlations for BIS–BAS scores.

Factor	Mean	Standard deviation
BIS	22.25	3.57
Drive	11.39	2.18
Fun seeking	11.70	1.84
Reward sensitivity	17.92	1.66
BAS	13.67	1.37

Correlations

	BIS	Drive	Fun	Reward	BAS
BIS	1				
Drive	.155	1			
Fun seeking	−.033	.29*	1		
Reward sensitivity	.49**	.35**	.16	1	
BAS	.27*	.81**	.67**	.67**	1

* $p = .01$.
** $p < .01$ BAS was formed as the average of the subscales Drive, Fun seeking, and Reward sensitivity (see results).

3.1. BIS–BAS scores

BIS–BAS scores were normally distributed and correlated within the normal range (see Table 1; Jorm et al., 1998). As the subscales of BAS are subordinate factors (Jorm et al., 1998), we treated BAS as a unitary construct in our analyses by averaging the subscale scores together.[1] We chose this approach because we had no a priori predictions regarding the subscales themselves. Cronbach's alpha was .77 for BIS and .63 for BAS, which accord well with previously published results (e.g. Jorm et al., 1998).

[1] The higher order BAS factor formed from the subscales by use of principal components analysis was statistically indistinguishable from the averaged score, $r = .998$.

Table 2
Mean response times (in milliseconds) by dilemma type and response.

Dilemma type	Response Appropriate	Inappropriate
Personal	6,600 (5,219)	6,190 (3998)
Impersonal	6,807 (5,730)	7,719 (7403)
Self	6,593 (6,453)	8,343 (12,916)
Other	6,933 (5,218)	5,967 (2880)
Inevitable	6,166 (5,134)	5,989 (3920)
Avoidable	6,852 (4,837)	8,089 (12,634)

Note: Values in parentheses are standard deviations of the means. All values were calculated from only the 75 subjects who were included in the response time analysis.

3.2. Responses

Mean proportions of moral endorsement (i.e. killing is morally appropriate) by condition are shown in Fig. 1. Using a 2(personal/impersonal) × 2(self/other) × 2(inevitable/avoidable) repeated measures analysis of covariance (ANCOVA), with mean-centered BAS and BIS scores as covariates, we replicate the main effects reported by Moore et al. (2008). Personal killing was judged less morally appropriate than impersonal killing, $F(1,74) = 34.55$, $p < .001$, $\eta_p^2 = .31$; killing to save oneself and others was judged more appropriate than to save only others, $F(1,74) = 43.00$, $p < .001$, $\eta_p^2 = .36$; and killing someone who would inevitably die was judged more appropriate than someone who would otherwise be spared, $F(1,74) = 60.42$, $p < .001$, $\eta_p^2 = .44$. There was a significant main effect of BIS, $F(1,74) = 6.83$, $p = .01$, $\eta_p^2 = .09$; participants with higher BIS responded with lower overall levels of moral approval of killing. There was also a significant main effect of BAS, $F(1,74) = 4.24$, $p = .04$, $\eta_p^2 = .05$, participants with higher BAS scores showed higher levels of approval of killing. After statistically removing these latter two main effects, BIS continued to significantly interact with personal harm, $F(1,74) = 11.69$, $p = .001$, $\eta_p^2 = .14$; participants with higher BIS scores were significantly more likely to disavow personal killing than impersonal killing, but those with lower BIS scores showed no effect of personal vs. impersonal killing (see Fig. 2). There was also a significant interaction of BAS and personal/impersonal, $F(1,74) = 10.88$, $p = .001$, $\eta_p^2 = .13$. Here, participants with higher BAS scores showed no difference in their approval of personal vs. impersonal killing, remaining very high on both. Lower BAS scores, however, correlated with a significant decrease in moral approval for personal killing, relative to impersonal.

The self/other factor did not interact with either BIS or BAS (F's < 1). Given that we considered saving oneself (in addition to others) as a type of benefit in the context of the tradeoffs participants made in the course of their judgments, we expected an interaction between self/other and BAS. To this end, we re-conducted this analysis with the sub-factors of BAS as covariates, instead of the global BAS factor, with $\alpha = .025$ to guard against Type 1 error. There was a significant interaction between self/other and the BAS sub-factor Fun Seeking, $F(1,72) = 6.10$, $p = .016$, $\eta_p^2 = .08$. Participants with lower Fun Seeking scores showed the largest effect of self vs. other, while those with higher scores showed little effect, responding to both types of dilemma with high levels of moral approval. The other sub-factors of BAS did not significantly interact with self/other (F's < 1.83, p's > .18).

The inevitable/avoidable factor interacted with BAS, $F(1,74) = 6.34$, $p = .01$, $\eta_p^2 = .08$; participants with higher BAS scores showed no significant effect of inevitable/avoidable. They morally approved of killing one to save many equally, and at high rates, in these two dilemma types. Those with lower BAS scores, however, showed a significant difference, finding killing when death was inevitable more morally appropriate than when avoidable (Fig. 3). There was a marginally significant interaction between BIS and inevitable/avoidable, $F(1,74) = 3.77$, $p = .056$, $\eta_p^2 = .05$; here, as BIS scores increased, the impact of the inevitable/avoidable factor also increased (i.e. the discrepancy between approval in the two cases grew larger). Participants with low BIS scores showed similarly (and relatively) higher rates of moral approval in both cases.

3.3. Response times

Mean response times (RT) (and standard deviations) are reported in Table 2. As with responses, a 2 × 2 × 2 repeated measures ANCOVA showed that participants responded more quickly to personal dilemmas than to impersonal dilemmas, $F(1,74) = 4.24$, $p = .043$, $\eta_p^2 = .05$.[2] There were no significant main effects of either self interest, $F < 1$, inevitability of death, $F(1,74) = 2.78$, $p = .10$, or BAS, $F(1,74) = 1.88$, $p = .17$, though there was a marginally significant main effect of BIS, $F(1,74) = 3.50$, $p = .065$, $\eta_p^2 = .05$; participants more sensitive to negative information took longer to make a response. An interaction between self/other and inevitable/avoidable resulted from participants responding more quickly to self-inevitable dilemmas (M = 5596 ms, SE = 434) than self-avoidable ones (M = 6856 ms, SE = 528), but showing no significant difference between RTs for other-inevitable (M = 6095 ms, SE = 393) and other-avoidable (M = 6017 ms, SE = 369). BIS and self/other interacted significantly, $F(1,74) = 4.76$, $p = .03$, $\eta_p^2 = .06$ (all other F's < 1.43, p's > .24). Exploration of the correlations between RTs for self and other dilemmas and BIS revealed a significant positive correlation, $r = .22$, $p = .05$, between RTs to self dilemmas and BIS, but no such relationship for other dilemmas, $r = .08$, $p = .47$. To clarify this interaction, we added response as a factor, which revealed that high BIS participants took selectively longer to approve killing, $r = .25$, $p = .03$, but not to disapprove of it, $r = .04$, $p = .76$.

Re-exploring the previous results with response as a factor showed no significant interactions as a function of response (all F's < 1.01) with the exception of a trend toward faster RTs to approve killing in self dilemmas than to disapprove, and the reverse for other dilemmas, $F(1,64) = 3.20$, $p = .078$, $\eta_p^2 = .05$. This replicates an identical crossover interaction pattern reported by Moore et al. (2008).

4. Discussion

Analysis of the group level effects replicated the previously reported impact of the personal/impersonal, self/other, and inevitable/avoidable factors on moral judgment (Moore et al., 2008). We also replicated the pattern of response by RT interaction for self/other dilemmas; participants were more likely to morally approve of killing one to save many if their own life was also in danger, and they were faster to make this decision than they were to morally disapprove. The opposite pattern characterized their responses and RTs when only others' lives were in danger. These effects, all replications, provided the context within which we examined the role of differential sensitivity to positive and negative aspects of the tradeoffs inherent in these types of moral dilemmas.

As we predicted, there were differential patterns of moral judgment as a function of individual differences in both BIS and BAS. Higher BIS scores predicted an overall decrease in moral approval of killing one to save many, likely reflecting aversion to killing another person, regardless of whether physical contact was required. However, even when accounting for this effect, the negative impact of personal harm on moral judgment was strongly related to sensitivity to negative information, with those highest on the BIS scale

[2] Response time analyses were conducted on log-transformed RTs as well, and produced qualitatively identical results.

showing the largest discrepancy in their judgments as a function of this factor, and those lowest on the BIS scale showing almost no effect at all. After controlling for this effect, there remained an interaction between BAS and personal harm such that participants most sensitive to positive information were indifferent between personal and impersonal killing, but those less sensitive to the gains represented by saving the lives of others found personal killing to be significantly less acceptable than impersonal killing. We interpret this straightforwardly: subjects who were highly sensitive to the positive gains evaluate the tradeoff against killing, either personally or impersonally, as being morally acceptable. Those less sensitive to the gains find the more negative personal harm to outweigh the gain of saving others more often than their counterparts.

BAS interacted with self-interest, but in an interesting way. The global BAS factor did not interact with self/other, but a sub-factor of BAS did. The theoretical distinction between these sub-factors is unclear,[3] and so we restrain our discussion to the global BAS system while recognizing that it may not be a homogenous construct. This interaction between BAS and self threat reveals that, as sensitivity to positive information decreases, the effect of self vs. other magnifies. In other words, those with low BAS scores do not find saving the lives of others sufficient motivation to endorse killing as morally appropriate very often, but when their own life is involved, the additional benefit increases their moral approval. This effect is attenuated in highly positive/reward sensitive individuals because their baseline rate of moral approval of killing to save others is already quite high. The BAS $\times$ inevitable/avoidable interaction follows the same pattern. Specifically, as BAS sensitivity declines, the inevitability of death for the potentially sacrificed person shows an increased effect. Put differently, participants with lower sensitivity to positive gains show a larger decrement in their moral approval in response to causing an avoidable death, presumably because the gains involved in saving other lives do not outweigh the cost of such a sacrifice.

These effects are both novel and substantially stronger than those reported by Moore et al. (2008). Individual differences in sensitivity to costs and benefits (or reward and punishment) predict moral judgment far more successfully than any other individual measure to date. These results are clearly interpretable within the DPM and add an additional level of specificity to this model. Participants' affective responsiveness to *both* positive and negative information impacts their likelihood of producing a utilitarian moral judgment; sensitivity to positive information, in the form lives saved, promotes utilitarian responding to personal moral dilemmas, presumably because the overall gain outweighs the cost of taking a life using personal force. Sensitivity to negative information (i.e. the use of personal force and lives lost) drives the personal-impersonal distinction as well as predicting overall levels of utilitarian moral judgment.

There is substantial variation in how people respond to moral dilemmas and such dilemmas vary in important ways. As the current research suggests, individual differences variables may be predictive of main effects, i.e., an overall tendency to respond in a more utilitarian fashion, and/or interaction effects, i.e., a tendency to respond in a more utilitarian fashion only in certain contexts. Indeed, these individual differences constrain whether or not such main effects will be observed at all. These results inform the larger debate over the relative roles of emotion vs. cognition in moral judgment insofar as they clearly demonstrate that sensitivity to the affectively eliciting factor(s) need to be evaluated separately for positive (i.e. gains) and negative (i.e. costs) elements of complex moral stimuli. Future research could investigate even more complicated interactions, considering multiple individual differences variables and multiple dilemma types with particular regard to the possibility that positive and negative information have differing impacts on the manner in which emotion and cognition combine to form judgments.

Acknowledgement

This work was partially funded by a National Science Foundation Graduate Research Fellowship awarded to the first author.

References

Borg, J. S., Hynes, C., Van Horn, J., Grafton, S., & Sinnott-Armstrong, W. (2006). Consequences, action, and intention as factors in moral judgments: An fMRI investigation. *Journal of Cognitive Neuroscience, 18*(5), 803–817.

Carver, C. S., Sutton, S. K., & Scheier, M. F. (2000). Action, emotion, and personality: Emerging conceptual integration. *Personality and Social Psychology Bulletin, 26*(6), 741.

Carver, C., & White, T. (1994). Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS/BAS scales. *Journal of Personality and Social Psychology, 67*(2), 319–333.

Gray, J. A. (1994). Personality dimensions and emotion systems. In P. Ekman & R. J. Davidson (Eds.), *The Nature of Emotion: Fundamental Questions* (pp. 329–331). New York: Oxford University Press.

Gray, J. R. (2004). Integration of emotion and cognitive control. *Current Directions in Psychological Science, 13*(2), 46–48.

Greene, J. D., Cushman, F. A., Stewart, L. E., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2009). Pushing moral buttons: The interaction between personal force and intention in moral judgment. *Cognition, 111*(3), 364–371.

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.

Gross, J. J., Sutton, S. K., & Ketelaar, T. (1998). Affective-reactivity views. *Personality and Social Psychology Bulletin, 24*(i3), 279.

Harmon-Jones, E. (2003). Clarifying the emotive functions of asymmetrical frontal cortical activity. *Psychophysiology, 40*(6), 838–848.

Hauser, M. D., Cushman, F., Young, L., Jin, R. K. X., & Mikhail, J. (2007). A dissociation between moral judgments and justifications. *Mind and Language, 22*(1), 1–21.

Jorm, A. F., Christensen, H., Henderson, A. S., Jacomb, P. A., Korten, A. E., & Rodgers, B. (1998). Using the BIS/BAS scales to measure behavioural inhibition and behavioural activation: Factor structure, validity and norms in a large community sample. *Personality and Individual Differences, 26*(1), 49–58.

Larsen, R. J., & Ketelaar (1991). Personality and susceptibility to positive and negative emotional states. *Journal of Personality and Social Psychology, 61*(1), 132–140.

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*(6), 549–557.

Valdesolo, P., & DeSteno, D. (2006). Manipulations of emotional context shape moral judgment. *Psychological Science, 17*(6), 476–477.

Wheatley, T., & Haidt, J. (2005). Hypnotic disgust makes moral judgments more severe. *Psychological Science, 16*(10), 780–784.

Zelenski, J. M., & Larsen, R. J. (1999). Susceptibility to affect: A comparison of three personality taxonomies. *Journal of Personality, 67*(5), 761–791.

[3] The names of these factors are for convenience only, being derived from the measurement items on Carver and White's scales (1994).