In defense of the personal/impersonal distinction in moral psychology research: Cross-cultural validation of the dual process model of moral judgment

Adam B. Moore*   N. Y. Louis Lee†   Brian A. M. Clark‡   Andrew R. A. Conway§

Abstract

The dual process model of moral judgment (DPM; Greene et al., 2004) argues that such judgments are influenced by both emotion-laden intuition and controlled reasoning. These influences are associated with distinct neural circuitries and different response tendencies. After reanalyzing data from an earlier study, McGuire et al. (2009) questioned the level of support for the dual process model and asserted that the distinction between emotion evoking moral dilemmas (personal dilemmas) and those that do not trigger such intuitions (impersonal dilemmas) is spurious. Using similar reanalysis methods on data reported by Moore, Clark, & Kane (2008), we show that the personal/impersonal distinction is reliable. Furthermore, new data show that this distinction is fundamental to moral judgment across widely different cultures (U.S. and China) and supports claims made by the DPM.

Keywords: moral judgment, dual process model, Ex-Gaussian method, culture.

1 Introduction

The dual process model of moral judgment (DPM; Greene, Morelli, Lowenberg, Nystrom, & Cohen, 2008; Greene, Nystrom, Engle, Darley, & Cohen, 2004; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001) contends that moral judgments are the product of two partially separable neural systems—one fast, automatic and affective, the other slower, effortful, and more abstract. A significant body of research has emerged from this proposal, with researchers employing a variety of approaches to characterize the nature and function of each system, as well as their interaction. These efforts span experimental and individual difference methods (Cushman, Young, & Hauser, 2006; Greene, Cushman, Stewart, Lowenberg, Nystrom, & Cohen, 2009; Greene et al., 2008; Hauser, Cushman, Young, Jin, & Mikhail, 2007; Moore, Clark, & Kane, 2008; Valdesolo & DeSteno, 2006; 2008) as well as neuroimaging and lesion studies (Ciaramelli, Muccioli, Ladavas, & Di Pellegrino, 2007; Greene et al., 2004; Koenigs et al., 2007; Moll, Eslinger, Oliveira-Souza, 2001). On the whole, the balance of this work has supported the DPM in its broad strokes while also providing for revision and clarification of the details of its elaboration (but see Moll & de Oliveira-Souza,

2007, and Greene, 2007).

The original empirical foundation of the DPM (Greene et al., 2001) depended on the distinction between personal and impersonal moral dilemmas. Personal dilemmas are those that, according to Greene et al. (2001), require the harming (or killing) of another person or persons by an agent to achieve some goal, specifically when that harm is not simply redirected from one person or group onto another (that is, the agent must generate the harm themselves). Impersonal dilemmas are those that do not satisfy some aspect of these requirements. The key piece of behavioral evidence from Greene et al. (2001) providing support for the DPM was a response time (RT) interaction between response and dilemma type, such that responses affirming the moral appropriateness of harming or killing in personal dilemmas were associated with significantly longer RTs than disavowing such actions, but there were no differences in RTs for responses to impersonal dilemmas. This effect, in conjunction with consonant neuroimaging data, led Greene et al. to infer that personal dilemmas triggered an automatic/affective processing system, which biased responding toward the negative and decreased RTs for these responses. However, cognitive control mechanisms were sometimes recruited to support more abstract, outcome-oriented processing of personal dilemmas, described in the paper as utilitarian responding. When prepotent negative responses were overridden, subjects took longer to respond to the scenario and tended to affirm the moral appropriateness of killing to save others, thereby producing the RT interaction between response and dilemma type. A similar inter-

*Adam B. Moore, Psychology Dept., Green Hall, Princeton University, Princeton, NJ 08540. Email: amtwo@princeton.edu.
†Chinese University of Hong Kong
‡University of Oregon
§Princeton University

action emerged in a later study (Greene et al., 2004), and neuroimaging data again supported this interpretation.

Recently, McGuire and colleagues (2009) pointed out a potentially serious flaw in the original work of Greene and colleagues (2001). Specifically, their criticism rested on an item analysis performed to test the psychometric properties of the dilemmas used in that experiment and thereby the theoretical assumptions based on effects observed when using them as stimuli. The results of that reanalysis showed that the RT interaction was spurious, driven entirely by a subset of dilemmas that could reasonably be considered stimulus outliers or non-dilemmas. Upon elimination of the poorly constructed materials, repeating the original RT analysis of Greene et al. (2001) yielded an effect very different from that reported— personal moral dilemmas simply took longer to respond to than did impersonal ones. This appears to be a critical problem, since the DPM asserts that the personal dilemmas are preferentially triggering automatic, affective, and *fast* processing, whereas the impersonal dilemmas do not, but rather are processed in a more voluntary, abstract, and *slow* fashion. The reanalysis yielded a pattern of RT effects that contradicted this interpretation. McGuire et al. drew two conclusions from this result. They argued that there is not sufficient evidence to support either the strong distinction between personal and impersonal moral dilemmas or a dual process model of moral judgment more generally.

In a response to McGuire et al., Greene (2009) argued that their criticism conflates the personal/impersonal distinction with the DPM more generally and that McGuire et al. ignore substantial evidence for the DPM that has been gathered since the original 2001 study, evidence that does not depend on the personal/impersonal distinction as originally articulated (e.g., Greene et al., 2004; Greene et al., 2008). On our reading of this exchange, the key question up for debate now is whether or not there is evidence that utilitarian responses to personal moral dilemmas engender cognitive conflict (compared to non-utilitarian responses) whereas such responding to impersonal dilemmas does not. Greene (2009) reviewed several neuroimaging studies that reported results consistent with the increased recruitment of cognitive control structures (e.g., lateral and dorsolateral prefrontal cortex) when subjects approved harming in a personal moral dilemma, but in order to avoid redundancy, we will pass over these. Instead, we offer additional and independent evidence on this point by reanalyzing a large, previously published data set (Moore et al., 2008) and demonstrating novel findings in support of the conflict claim on the part of the DPM. This new evidence is strengthened by a new empirical study that replicates the results across disparate cultures, suggesting that the DPM is tracking a fundamental fact about human moral cognition.

2 Analysis 1

Moore et al. (2008) published a complex dataset that addressed several factors that affect moral judgment. Briefly, they demonstrated that moral approval of killing one to save many: (1) decreases as a function of physical directness (vs. environmentally mediated harm); (2) increases as a function of self interest (i.e., subjects were more likely to judge the killing as morally acceptable if it saved themselves as well as others vs. saving only others when their own life was not at risk); and (3) increases when the one to-be-sacrificed person would die regardless of the subjects' choice. Moreover, the authors demonstrated that, regardless of the type of response, response times were significantly faster for personal dilemmas than those for impersonal ones. This latter finding is the opposite of the results of McGuire et al.'s reanalysis. Moore et al. (2008) argued that this result was the product of having matched their personal and impersonal dilemmas for complexity and length. They pointed out that Greene et al. (2001) had personal dilemmas that were longer and more complex than their impersonal ones, and that this confounded the interpretation of their evidence. We will return to the RT results below with new data.

The useful contribution of McGuire et al. was to point out that observed results such as these could be driven by outlier examples within a stimulus category. Here we conduct an item analysis on data from Moore et al. (2008) in an effort to test the validity and robustness of the reported personal/impersonal effect. A key aspect of McGuire et al.'s strategy was to eliminate dilemmas from their analysis if less than 5% of subjects approved of the proposed action. It was this criterion that demonstrated the artifactual nature of the RT effect originally reported by Greene et al. (2001). However, this strategy depends on the fact that Greene et al. used different dilemmas in the personal and impersonal categories, and therefore does not apply to Moore et al.'s (2008) materials. This is because, with respect to the personal/impersonal effect, each scenario served in both categories, though no subject saw both versions of any scenario. Thus, if a particular scenario was approved less than 5% in one category but more than that in the other, this effect is properly attributable to the experimentally manipulated variable since all other details of the dilemmas were kept constant.[1]

---

[1]Only four of 24 dilemmas failed to show a difference in response preference in the predicted direction (i.e. Impersonal > Personal). Leave one out bootstrapping reproduced the direction and significance of the reported effect of personal vs. impersonal 100% of the time for responses and 75% of the time for the RT results.

2.1 Method

We performed a between item ANOVA where our dependent variable was the across-subjects proportion of moral approval for each dilemma and the cases were individual dilemmas (i.e., we treated dilemmas as subjects). The task for the subjects was to indicate whether or not they found a suggested course of action in a moral dilemma to be morally appropriate, which they indicated via button press on a standard keyboard (as done by Greene et al., 2001). For our purposes we will focus on the personal/impersonal effect and ignore the others, as they were by design orthogonal to it.[2]

2.2 Results and discussion

When killing one to save many, there was significantly less moral approval for personal compared to impersonal killing, $F(1, 46) = 5.45$, $p = .02$, $\eta_p^2 = .11$. Similarly, RTs were faster for personal dilemmas compared to impersonal, $F(1, 46) = 4.62$, $p = .037$, $\eta_p^2 = .09$. This replicates the previously reported effects of Moore et al. and demonstrates some measure of reliability for the personal/impersonal distinction, operationalized as relative directness of action.

3 Analysis 2

Both the item analysis reported here and Moore et al. showed that responses to personal moral dilemmas were faster overall than those to impersonal dilemmas. However, it has widely been recognized that RT data contain more information than is captured by simple summary statistics such as the mean (Ratcliff, 1979). Perhaps there is evidence in the available RT data that supports the idea of a greater amount of conflict in utilitarian responding to personal moral dilemmas. One method of interrogating this possibility is to examine the distribution of RTs directly and estimate the underlying generative parameters using maximum likelihood methods. One such approach that has been used for exactly the purpose that we propose is the Ex-Gaussian method (Brown & Heathcote, 2003; Heathcote, Brown, & Mewhort, 1991; McGill, 1963). This approach acknowledges the positive skew of RT data by proposing that such distributions are generated by the convolvement of a Gaussian and an Exponential distribution. The former has the standard parameters μ and σ, and the Exponential has one parameter, τ. The mean of the Ex-Gaussian distribution equals $\mu + \tau$, and the variance is $\sigma^2 + \tau^2$. In this case, it is the exponential parameter that produces the skewed tail of the RT distribution and

contributes to the observed mean. While the μ parameter reflects the decision component in the pure sense, the τ parameter is theorized to reflect a response transduction component of cognitive processing and accordingly has been shown to systematically relate to the increased demand for conflict resolution (Heathcote, Popiel, & Mewhort, 1991; Spieler, Balota, & Faust, 1996; Spieler, Balota, & Faust, 2000; Steinhauser & Hubner, 2009). If the DPM is accurate and utilitarian responding to personal moral dilemmas involves response conflict requiring cognitive control to overcome while non-utilitarian responding does not, then at minimum we should expect a greater estimate for the exponential parameter τ for this condition relative to the others.

3.1 Method

We averaged RTs over all subjects into four categories formed from the factorial crossing of the factors dilemma type (personal/impersonal, collapsing over self/other and inevitable/avoidable) and response type (approval/disapproval) while treating dilemmas as cases. In order to reliably determine the variability of the parameter estimates from the Ex-Gaussian estimation we used leave-one-out bootstrapping to create new data sets for each category. We sampled with replacement 100 times for each of 24 iterations (one iteration for each censored dilemma pair). Then we estimated the best fitting Ex-Gaussian parameters for each of those RT distributions using QMPE v2.18, a freely available quantile maximum likelihood estimator software package (Heathcote, Brown, & Cossineau, 2004; Heathcote, Brown, & Mewhort, 2002).[3]

3.2 Results and discussion

The average parameter values are reported in Table 1.

It is reassuring that, much like the aggregate statistical results reported above, the μ estimates for personal moral dilemmas were smaller than for impersonal, $t(22) = -9.21$, $p < .001$, $d = -1.88$, and there was no effect of response for this parameter, $t < 1$. The τ parameter estimate was larger across the personal dilemmas ($M = 1713.4$, $SE = 172.55$) relative to the impersonal ones ($M = 969.7$, $SE = 161.6$), $t(22) = 4.45$, $p < .001$, $d = 0.91$, suggesting greater conflict in personal dilemmas than in impersonal dilemmas, regardless of response. This may be due to the fact that in personal dilemmas one must weigh the value of saving lives against sacrificing another in a fashion that is emotionally provocative, where in the impersonal dilemmas this emotional provocation

---

[2]The other main effects reported by Moore et al. also replicate in an item analysis. Results are available on request.

[3]Available from `http://newcl.org/?q=node/10`. Note that we calculated estimates on the raw data, not the computed quantiles.

Table 1: Parameter estimates for RT distributions by personal/impersonal and response, from Moore, Clark, and Kane (2008). Values in parentheses are standard errors of the estimates.

| Parameter | Dilemma type | | | |
| | Personal | | Impersonal | |
	Appropriate	Inappropriate	Appropriate	Inappropriate
μ	3062.0 (117.0)	3115.1 (112.0)	4540.2 (215.3)	4656.0 (211.0)
σ	686.1 (141.9)	1001.8 (130.4)	1845.9 (193.5)	2194.3 (200.1)
τ	1924.9 (202.5)	1501.9 (142.6)	917.0 (220.0)	1022.4 (103.2)

is (relatively) absent. As predicted, the τ parameter estimate for moral approval of killing in a personal dilemma was larger than for disapproval in the same category, one tailed $t(22) = 2.45$, $p < .025$, $d = 0.50$.

These results suggest that there is evidence both for increased conflict involved in the processing of personal moral dilemmas, compared to impersonal, and that this conflict is selectively greater in those cases when subjects approve moral harm in personal dilemmas than when they disapprove of such harm. Moreover, this effect may be masked by aggregate statistics. However, reanalysis can only take us so far, since it is based on the assumption that the original sample is truly representative and random. A stronger test of the assertions that personal moral dilemmas involve additional conflict relative to impersonal ones and that conflict is greater when responding in a utilitarian fashion than in a non-utilitarian fashion would be replication of these results across disparate groups of subjects. To make this test as strong as possible, as well as to rigorously replicate the effects reported by Moore et al. (2008), we chose to compare groups of subjects who have been shown to reliably differ in their decision making styles in similar paradigms. If they nevertheless produce effects similar to those discussed above then this will be strong evidence for the reliability of the conclusions that form the core of the DPM.

4   Experiment 1

Since Markus and Kitayama's (1991) seminal proposal that Asians follow an interdependent self-construal style whereas Westerners follow an independent self-construal style, a large body of research has reported corresponding cross-cultural differences in both social cognition (see, e.g., Oyserman, Coon, & Kemmelmeier, 2002; Oyserman & Lee, 2008); and cognition in general (see, e.g., Nisbett, Peng, Choi, & Norenzayan, 2001; Nisbett, 2003). One reported cognitive difference is that Asians tend to reason holistically, considering more contextual factors, and intuitively, by recalling prior experience or existing

beliefs, thereby being more likely to succumb to the so-called "belief bias" where one's judgment of an argument's deductive validity is affected by whether the conclusion matches the reasoner's real world belief (see, e.g., Evans, Barston, & Pollard, 1983). On the other hand, Westerners tend to reason deliberately by applying logical rules (Norenzayan, Smith, Kim, & Nisbett, 2002; but see Unsworth & Medin, 2005, for a challenge to the belief bias finding). A direct cross-cultural comparison of responses to the moral dilemmas we have investigated thus far would therefore be relevant, as it could be argued either that (1) Asians, who tend to have an interdependent construal of self, may have a stronger bias towards saving more people at the expense of a single individual when compared to Westerners, or that (2) Asians, who tend to favor intuitive reasoning, would show a higher level of disapproval for personal killing than would Westerners. As moral judgment, at least insofar as it is tapped by these kinds of hypothetical dilemmas, is clearly a species of social cognition, the possibility that these groups might not treat the personal/impersonal distinction in the same way represents a test of the claims of the DPM.

A substantial proportion of previous work on cultural differences in moral judgment has dealt with the question of whether human moral psychology is fundamentally concerned with harms (psychological, political, or physical) or whether cultural influences extend the domain of moral concern to other areas, depending on the idiosyncrasies of the cultures in question (Haidt, Koller, & Dias, 1993; Miller, Bersoff, & Harwood, 1990; Shweder, Mahapatra, & Miller, 1987; Turiel, Killen, & Helwig, 1987). O'Neill and Petrinovich (1998) have shown many aspects of moral judgment to be similar in Taiwanese and U.S. subjects, but they did not investigate any distinction in the type of action undertaken, but rather focused on action vs. inaction. Likewise, Baron and Miller (2001) found that both Indian and American subjects found action of harming to be worse than inaction resulting in harm. More recent work, using a large online sample, has investigated moral judgment while attempting to control for the type of action taken, but did

not directly address the personal/impersonal distinction (Cushman, Young, & Hauser, 2006; Hauser, Cushman, Young, Jin, & Mikhail, 2007). Moreover, these latter studies did not directly test cultural differences (this information was used only as a covariate), and data were gathered only from English speaking subjects who logged onto the relevant website. Our goal here was to test the robustness of the personal/impersonal effect by directly comparing samples of subjects from two disparate cultures. To our knowledge, this is the first attempt to gather moral judgments about personal vs. impersonal dilemmas from two cultures in their native languages. To this end, we tested subjects in the U.S. and in China using the same battery of moral dilemmas reanalyzed above, appropriately translated. We predicted that the patterns of judgments reported by Moore et al. (2008) would replicate in both and that the Ex-Gaussian analysis of their RT distributions would produce the same pattern of estimated parameters. If so, then that would support the contentions that 1) the personal/impersonal distinction represents something fundamental about human moral cognition and 2) this distinction is at least partially characterized by the idea that utilitarian responding to personal moral dilemmas involves cognitive conflict in a manner that non-utilitarian responding does not.

4.1 Method

4.1.1 Subjects

The U.S. sample consisted of 35 college students at Princeton University who participated in exchange for partial credit toward a course requirement (20 females). The Chinese sample consisted of 41 college students (22 females) at the Chinese University of Hong Kong who participated in exchange for HK $50 (approximately U.S. $6.50). Both samples were the control groups for an unrelated experiment.

4.1.2 Materials

We used the battery of moral dilemmas used by Moore et al. (2008) and reanalyzed above. For the Chinese sample, two independent bilingual translators translated the stimuli into Chinese. They were then back-translated by the second author and compared to the original English versions. Any discrepancy in translation was resolved via discussion among the translators and the second author.

The battery of moral dilemmas consisted of 24 separate critical scenarios, all of which involved sacrificing one person to save multiple others from death. Each scenario had two possible resolutions; one personal, which involved relatively direct physical harm, and one impersonal, in which harm could be brought about by less direct, environmentally mediated means. These two res-

olutions were matched to within ±2 words in length. Subjects saw only one resolution (personal or impersonal) to each dilemma, counterbalanced across subjects. All dilemmas were presented in second person perspective. Half of all dilemmas involved killing one person to save oneself and others (self) and the other half involved killing to save only others (other). We crossed the self/other and personal/impersonal variables and orthogonally manipulated whether the to-be-sacrificed person would die regardless of the choice of the subject (inevitable) or the to-be-sacrificed person would survive unharmed unless the subject chose to sacrifice them to save others (avoidable). There were thus 48 possible dilemmas (a personal and impersonal version of each of 24), and each subject saw 24 of them, 6 from each of the 4 conditions defined by personal/impersonal and inevitable/avoidable. We also included 14 filler dilemmas that involved other social decisions of a moral nature (e.g. cheating, stealing, or lying) that did not require harming or killing.

4.1.3 Procedure

After giving informed consent, all subjects were presented with the text of the dilemmas as black text on a grey computer screen. Dilemmas were presented individually in a different randomized order for each subject. Each dilemma was presented in two paragraphs. After reading the first paragraph of the problem, subjects pressed the spacebar to cause each sentence of the resolution to appear, one at a time. All text remained visible on the screen until the subject responded by pressing one of two allowable buttons on the keyboard to indicate that the proposed action was "morally appropriate" or "morally inappropriate". This was the procedure used by Moore et al. (2008). The computer collected responses and RTs for the dilemmas. Afterward, subjects were debriefed and thanked.

4.2 Results of Experiment 1

All null hypothesis significance tests are non-directional with $\alpha = .05$ unless otherwise indicated. Effect sizes are reported as partial eta squared (η_p^2) or Cohen's *d*. The assumption of equality of group covariance matrices was not violated for the response data, so we present univariate results. Two subjects were dropped from the Chinese sample for responding randomly and a third whose RTs were trimmed as outliers, leaving 39 subjects for analysis of responses and 38 for analysis of times. For ease of presentation and continuity with the above discussion, we will present only those results that were statistically significant, while focusing on the personal/impersonal effect for the Ex-Gaussian RT analysis.

Table 2: Mean proportion (standard deviation) of "appropriate" responses by country and dilemma type.

	U.S.		China	
	Inevitable	Avoidable	Inevitable	Avoidable
Personal-Self	.71 (.35)	.51 (.41)	.53 (.37)	.43 (.33)
Personal-Other	.57 (.30)	.40 (.27)	.47 (.30)	.31 (.35)
Impersonal-Self	.78 (35)	.67 (.31)	.62 (.38)	.57 (.35)
Impersonal-Other	.69 (.30)	.45 (.32)	.61 (.29)	.44 (.32)

Table 3: Parameter estimates for RT distributions by personal/impersonal and response, for Chinese and U.S. samples. Values in parentheses are standard errors of the estimates.

		Dilemma type			
		Personal		Impersonal	
	Parameter	Appropriate	Inappropriate	Appropriate	Inappropriate
China	μ	3311.7 (144.1)	3427.3 (162.0)	3792.0 (184.9)	3820.2 (166.3)
	σ	731.2 (151.4)	483.9 (205.5)	631.4 (180.9)	497.8 (202.3)
	τ	2214.2 (155.0)	1976.6 (177.6)	1465.1 (274.4)	1550.2 (297.1)
U.S.	μ	3017.0 (126.3)	2887.3 (210.1)	3507.5 (213.0)	3703.9 (254.8)
	σ	526.9 (120.1)	550.3 (153.9)	762.9 (106.2)	1224.5 (207.4)
	τ	1874.5 (130.3)	1531.0 (159.6)	1029.6 (167.5)	1279.0 (206.1)

4.2.1 Responses

We used the proportion of dilemmas in each category for which subjects morally approved of the action as the dependent measure in a 2 (U.S./China) x 2 (personal/impersonal) x 2 (self/other) x 2 (inevitable/avoidable death) mixed factorial ANOVA where the dilemma variables were manipulated within subjects. The mean levels of moral approval by dilemma type and country are presented in Table 2.

Across both samples, subjects found physically indirect killing (impersonal) to be more morally appropriate than direct (personal), $F(1, 72) = 24.14$, $p < .001$, $\eta_p^2 = .25$. Killing to save oneself and others (self) was more morally appropriate than killing to save only others (other), $F(1, 72) = 32.93$, $p < .001$, $\eta_p^2 = .31$. Killing someone who would die anyway (inevitable death) was more morally appropriate than killing someone who would otherwise survive (avoidable death), $F(1, 72) = 47.81$, $p < .001$, $\eta_p^2 = .40$. There was not a significant effect of culture, $F(1, 71) = 3.1$, $p > .08$, $\eta_p^2 = .04$. There were no interactions between any of the dilemma factors and country/culture, all Fs < 2.3, all ps > .14. Thus, the overall pattern of conditional means was not significantly different across these samples.

4.2.2 Response times

As articulated above, the rationale behind conducting the Ex-Gaussian analysis is to use the three inferred parameters for significance tests instead of only the observed mean and standard deviation. Thus, we conducted separate Ex-Gaussian analyses on the U.S. and the Chinese samples to determine if the predicted pattern of parameter estimates would emerge from both groups. For each sample we created four categories by crossing the personal/impersonal factor and response type. We then sampled with replacement 100 times to create 24 new datasets leaving out one dilemma in each set.

4.2.3 Chinese RT parameter estimates

As can be seen in Table 3, the estimates for μ were significantly smaller for personal compared to impersonal dilemmas, paired samples $t(22) = -2.66$, $p < .02$, $d = -0.54$. There was no significant difference between the estimates for moral approval vs. disapproval in personal or impersonal dilemmas, ts < 1.

As we predicted, and in accord with the results from the reanalysis, the τ parameter is significantly larger for personal dilemmas than for impersonal, $t(22) = 2.60$, $p < .02$, $d = 0.53$. However, there was only a trend for larger

estimates for utilitarian responses to personal dilemmas than for non-utilitarian responses, one tailed $t(22) = 2.78$, $p = .074$, $d = 0.29$. The τ parameter was not significantly different for utilitarian compared to non-utilitarian responses within impersonal dilemmas, $t < 1$.

4.2.4 U.S. RT parameter estimates

Referring to Table 3, again we find that the estimates for μ were significantly smaller for personal compared to impersonal dilemmas, $t(22) = -3.25$, $p < .005$, $d = -0.66$. There was no significant difference as a function of response for either personal or impersonal dilemmas, ts < 1. Once again, the τ estimate for personal dilemmas is larger than for impersonal, $t(22) = 3.31$, $p < .005$, $d = 0.67$. The τ estimate for utilitarian responses to personal dilemmas is significantly larger than that for non-utilitarian responses, $t(22) = 2.37$, one-tailed $p < .015$, $d = .48$, but the estimates for utilitarian responses to impersonal dilemmas was not significantly larger than for non-utilitarian responses, $t(22) = -1.34$, $p > .15$, $d = -0.27$.

4.3 Discussion of Experiment 1

Both U.S. and Chinese subjects found relatively indirect, environmentally mediated killing to be more morally acceptable than the same death caused by more direct, physical means. Similarly, both samples showed higher moral approval of killing to save oneself and others than of killing to save only others and higher moral approval of killing when death was inevitable than when death was avoidable. There were no significant differences between the U.S. and Chinese subjects in the pattern of their responses across the eight different dilemma types used here. Ex-Gaussian analysis of RT distributions indicated that utilitarian responding to personal moral dilemmas involves more conflict than non-utilitarian responding, and that personal moral dilemmas generally involve more conflict than impersonal dilemmas regardless of response.

Our goal in this experiment was twofold. First, we sought to determine whether the factors we manipulated genuinely tap fundamental aspects of human moral cognition, as we have previously argued. It was possible that, although these materials survived an item analysis based on a large U.S. sample, they might nevertheless fail to generalize to another language and culture or to another U.S. sample of different demographic makeup. Second, we investigated the replicability of the Ex-Gaussian results from our reanalysis in order to gauge the strength of the finding as evidence for the DPM. The DPM claims that utilitarian responding to personal moral dilemmas involves cognitive conflict to a greater degree than does non-utilitarian responding, a claim that receives support

here in the form of Ex-Gaussian τ parameter estimates. These estimates were larger for personal dilemmas than for impersonal, and larger for utilitarian responses than for non-utilitarian responses to personal dilemmas. This replicates the results from the previous reanalysis. The overall replication of previous results across very different languages and socio-cultural samples reinforces support for the personal/impersonal distinction as well as introducing novel evidence that favors the DPM's claims of increased conflict in personal dilemmas.

5 General discussion

The DPM has served quite well to organize thinking and research on human moral judgment, particularly with respect to connecting empirical behavioral results to neural structures and function. McGuire et al. raised the valid point that methodological flaws may have produced the illusion of a pattern where none truly existed, thus possibly undermining the DPM's attempt to connect automatic moral intuitions and consequentialist moral judgments to separable neural circuitries. The original definition of personal vs. impersonal moral dilemmas roughly tracked those problems that evoke automatic intuitions and those that do not. Subsequent research from a variety of sources has shown that this definition can be narrowed, and methods refined, to isolate particular factors responsible for the evocation of such automatic responses, including direct physical contact (Cushman et al., 2006; Moore et al., 2008), self interest (Moore et al., 2008), and intentionality (Cushman et al., 2006). The current results demonstrate the importance of creating carefully constructed experimental materials. As mentioned in the Introduction, the original dilemmas constructed by Greene et al. (2001) were not matched for length, causal structure, or complexity. This clearly confounds the interpretation of RT data and lead to a larger debate about generalizability. In contrast, Moore et al. (2008) constructed dilemmas that were well matched on these dimensions, and these materials survived the item analysis and a cross-cultural replication reported above.

As research on the possible mechanisms of moral judgment has progressed, support for the DPM has increased. The current results have shown that factors previously identified as playing a strong role in moral judgment, particularly the personal/impersonal effect, seem to be robust across cultures. Perhaps more importantly, we have shown that there is intriguing new evidence that personal moral dilemmas involve greater conflict than impersonal ones, and this is particularly true when subjects respond in a utilitarian, or consequentialist, fashion within personal dilemmas. Taken in conjunction with other studies that have shown selective impairment of

such judgments by cognitive load (Greene et al., 2008) and greater recruitment of cognitive control regions of the prefrontal cortex when giving utilitarian vs. non-utilitarian responses to personal moral dilemmas (Greene et al., 2004), this evidence reinforces support for the DPM.

The DPM, in addition to being compatible with a wide variety of empirical work, has generated several questions for future research. Among the most interesting of these questions, yet to be answered, are: When and why is more effortful, abstract processing recruited for moral judgment? What is the nature of this processing in algorithmic terms? Are there stable differences among individuals as to their likelihood to engage in such effortful cognition in the moral domain? As we move toward answers to these questions, the form and content of the DPM will adapt to accommodate these developments if it is to survive as a useful characterization.

References

Baron, J. & Miller, J. G. (2001). Limiting the scope of moral obligation to help: A cross-cultural investigation. *Journal of Cross-Cultural Psychology, 31*, 705–727.

Ciaramelli, E., Muccioli, M., Ladavas, E., & di Pellegrino, G. (2007). Selective deficit in personal moral judgment following damage to ventromedial prefrontal cortex. *Social Cognitive and Affective Neuroscience*, *2*, 84.

Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment. *Psychological Science*, *17*, 1082–1089.

Evans, J. St. B. T., Barston, J. L., & Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory and Cognition, 11*, 295–306.

Greene, J. D. (2007). Why are VMPFC patients more utilitarian? A dual-process theory of moral judgment explains. *Trends in Cognitive Sciences*, *11*, 322–323.

Greene, J. D. (2009). Dual-process morality and the personal/impersonal distinction: A reply to McGuire, Langdon, Coltheart, and Mackenzie. *Journal of Experimental Social Psychology, 45*, 581–584.

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*, 364–371.

Greene, J. D., Morelli, S. A., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2008). Cognitive load selectively interferes with utilitarian moral judgment. *Cognition*, *107*, 1144–1154.

Greene, J. D., Nystrom, L. E., Engell, A. D., Darley, J. M., & Cohen, J. D. (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron, 44*, 389–400.

Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science, 293*, 2105–2108.

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

Heathcote, A., Brown, S., & Cousineau, D. (2004). QMPE: estimating lognormal, Wald, and Weibull RT distributions with a parameter dependent lower bound. *Behavior Research Methods, 36*, 277–290.

Heathcote, A., Brown, S., & Mewhort, D. (2002). Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin and Review, 9*, 394–401.

Heathcote, A., Popiel, S. J., & Mewhort, D. (1991). Analysis of response time distributions: an example using the Stroop task. *Psychological Bulletin, 109*, 340–347.

Koenigs, M., Young, L., Adolphs, R., Tranel, D., Cushman, F., Hauser, M., et al. (2007). Damage to the prefrontal cortex increases utilitarian moral judgements. *Nature*, *446*, 908–911.

Markus, H. R., & Kitayama, S. (1991). Culture and the self: implications for cognition, emotion, and motivation. *Psychological Review, 98* 224–253.

McGill, W. J. Stochastic latency mechanisms. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of mathematical psychology*, Vol. 1, New York: Wiley, 1963.

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*, 577–580.

Moll, J., Eslinger, P. J., & Oliveira-Souza, R. (2001). Frontopolar and anterior temporal cortex activation in a moral judgment task: preliminary functional MRI results in normal subjects. *Arquivos de neuropsiquiatria, 59*, 657–664.

Moll, J., & de Oliveira-Souza, R. (2007). Moral judgments, emotions and the utilitarian brain. *Trends in Cognitive Sciences*, *11*, 319–321.

Moore, A. B., Clark, B. A., & Kane, M. J. (2008). Who shalt not kill? Individual differences in working memory capacity, executive control, and moral judgment. *Psychological Science*, *19*, 549–557.

Nisbett, R. E. (2003). *The geography of thought: How asians and westerners think differently... and why*. New York: Free Press.

Norenzayan, A., Smith, E. E., Kim, B. J., & Nisbett, R. E. (2002). Cultural preferences for formal versus intuitive reasoning. *Cognitive Science, 26*, 653–684.

Oyserman, D., Coon, H. M., & Kemmelmeier, M. (2002). Rethinking individualism and collectivism: evaluation of theoretical assumptions and meta-analyses. *Psychological Bulletin, 128*, 3–72.

Oyserman, D., & Lee, S. W. S. (2008). Does culture influence what and how we think? Effects of priming individualism and collectivism. *Psychological Bulletin, 134*, 311–342.

Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin, 86*, 446–461.

Royzman, E. B., & Baron, J. (2002). The preference for indirect harm. *Social Justice Research*, *15*, 165–184.

Spieler, D. H., Balota, D. A., & Faust, M. E. (1996). Stroop performance in healthy younger and older adults and in individuals with dementia of the Alzheimer's type. *Journal of Experimental Psychology: Human Perception and Performance, 22*, 461–479.

Spieler, D. H., Balota, D. A., & Faust, M. E. (2000). Levels of selective attention revealed through analyses of response time distributions. *Journal of Experimental Psychology: Human Perception and Performance, 26*, 506–526.

Steinhauser, M. & Hubner, R. (2009). Distinguishing response conflict and task conflict in the Stroop task: Evidence from Ex-Gaussian distribution analysis. *Journal of Experimental Psychology: Human Perception and Performance, 35* 1398–1412.

Unsworth, S. J., & Medin, D. L. (2005). Cultural differences in belief bias associated with deductive reasoning? *Cognitive Science, 29*, 525–529.

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

Valdesolo, P. & DeSteno, D. (2008). The duality of virtue: Deconstructing the moral hypocrite. *Journal of Experimental Social Psychology, 44*, 1334–1338.

Appendix

Modified submarine

Personal version

You are a crewperson on a marine-research submarine traveling underneath a large iceberg. An onboard explosion has damaged the ship, killed and injured several crewmembers. Additionally, it has collapsed the only access corridor between the upper and lower parts of the ship. The upper section, where you and most of the others are located, does not have enough oxygen remaining for all of you to survive until you reach the surface. Only one remaining crewmember is located in the lower section, where there is enough oxygen.

There is an emergency access hatch between the upper and lower sections of the ship. If released, it will fall to the deck and allow oxygen to reach the upper section. However, a crewmember in the lower section was knocked unconscious and is lying beneath the hatch while you and the rest of the crew are almost out of air. If you shove the hatch open you and the others will have air, but it will fall to the deck, crushing the unconscious crewmember.

Is it appropriate for you to open the hatch and crush the crewmember below to save yourself and the other crewmembers?

Impersonal version

You are a crewperson on a marine-research submarine traveling underneath a large iceberg. An onboard explosion has damaged the ship, killed and injured several crewmembers. Additionally, it has collapsed the only access corridor between the upper and lower parts of the ship. The upper section, where you and most of the others are located, does not have enough oxygen remaining for all of you to survive until you reach the surface. Only one remaining crewmember is located in the lower section, where there is enough oxygen.

There is an emergency access hatch between the upper and lower sections of the ship. If released by an emergency switch, it will fall to the deck and allow oxygen to reach the area where you and the others are. However, the hatch will crush the crewmember below, since he was knocked unconscious and is lying beneath it. You and the rest of the crew are almost out of air though, and you will all die if you do not do this.

Is it appropriate for you to release the hatch and crush the crewmember below to save yourself and the other crew members?

Modified rowboat

Personal version

You are in a rowboat with a tour guide while sight-seeing on a lake in Alaska. You notice that three children have overturned their boat nearby. They are now in danger of quickly freezing to death in the icy water. You begin to row over to rescue them when you realize that your boat will not hold you, the tour guide, and the three children.

Two of the children are closer to your boat than the third child. As you and the tour guide pull in the first two it is obvious that one of them is too heavy and if you keep him on board your boat will sink and all of the children will die. Neither you nor the guide can get out because

you are rowing and the guide is performing CPR. The only way to save two children is to throw the large child off the boat and rescue the third one.

Is it appropriate to throw one child off the boat to die in order to save the other two?

Impersonal version

You are in a rowboat with a tour guide while sight-seeing on a lake in Alaska. You notice that three children have overturned their boat nearby. They are now in danger of quickly freezing to death in the icy water. You begin to row over to rescue them when you realize that your boat will not hold you, the tour guide, and the three children.

Two of the children are closer to your boat than the third child. As you and the tour guide pull in the first two it is obvious that they are too heavy and if you try to rescue the third child your boat will sink and all of the children will die. Neither you nor the guide can get out because you are rowing and the guide is performing CPR. The only way to save the two children is to row quickly away from the third one, leaving him to die.

Is it appropriate to leave the one child behind to die in order to save the other two?