

An assessment of the temporal dynamics of moral decisions

Gregory J. Koop*

Abstract

In the domain of moral decision making, models in which emotion and deliberation constitute competing dual-systems have become increasingly popular. Currently, the favored explanation of this interaction is what Evans (2008) termed a “default-interventionist” (DI) process where moral decisions are the result of a prepotent emotional response, which can be overridden with substantial deliberative effort. Although this “emotion-then-deliberation” sequence is often assumed, existing methods have lacked the requisite process resolution to clearly depict the nature of this interaction. The present work utilized continuous mouse tracking, or response dynamics, to develop and test predictions of these DI models of moral decision making. Study 1 utilized previously published moral dilemmas to validate the method for use with such complex stimuli. Although the data replicated typical choice and RT patterns, the process metrics provided by the response trajectories did not demonstrate the online preference reversals predicted by DI models. Study 2 utilized more rigorously constructed stimuli and an alternative presentation format to provide the strongest possible test of DI predictions, but again failed to show the predicted reversals. In summary, neither experiment provided data in accordance with the predictions of popular DI dual-systems models, which suggests that researchers should consider models allowing for concurrent activation of deliberative and emotional systems, or reconceptualize moral decisions within the typical multiattribute decision framework.

Keywords: moral decision making, methodology, process models, dual-systems models.

1 Introduction

Victor Hugo’s classic 19th century novel (as well as the modern musical adaptation) *Les Misérables* is dominated by the moral struggles of the main character, an ex-convicted turned philanthropist named Jean Valjean. In a particularly tense moment, Valjean must decide whether to kill his nemesis, thereby allowing him to continue a life of good works (thus benefiting many), or to obey his strongly held moral imperative not to kill, and thereby be unjustly incarcerated. Reflecting a certain degree of foresight on the part of Hugo, Valjean’s conflict between doing the most good (a utilitarian consideration) and obeying moral imperatives (a deontological consideration) closely reflects popular contemporary models of the moral decision-making process. However, whereas Hugo depicts this process as one of controlled deliberation, modern models of moral decision making also presume a role for an automatic, emotional system.

1.1 Dual-systems in moral decision making

Unlike earlier models that put the moral decision burden largely on either controlled deliberative processes (e.g.,

Kohlberg, 1969) or emotional processes (e.g., Haidt, 2001), the most popular contemporary model is a combination of these two perspectives (Greene et al., 2001, 2004, 2008; Haidt, 2007). This hybrid view associates the emotional system with moral laws and rules (i.e., deontological considerations) and the deliberative system with dispassionate utilitarian concerns. Furthermore, these models specify the nature in which these systems interact. Specifically, the emotional deontological system provides a prepotent response that can only be effortfully overridden by the slower, controlled deliberative system (Greene et al., 2001). This type of interaction, known as default-interventionist (DI; Evans, 2008), is similar to more general dual-systems models of decision making (e.g., Kahneman & Frederick, 2002; Loewenstein, Rick, & Cohen, 2008).

The data supporting this dual-systems model is largely from discrete choice behavior on sets of “small world” dilemmas where all possible outcomes and actions are known. Shortcomings in ecological validity aside (see Gigerenzer, 2010), this small world assumption allows researchers strict control over the degree to which the proposed systems conflict with one another. These dilemmas are often described as “personal” or “impersonal” dilemmas, partially as a function of this conflict (e.g., Greene et al., 2001; Moore, Clark, & Kane, 2008; Moore et al., 2011; Baron et al., 2012). Specifically, personal dilemmas are thought to be more emotionally aversive because they require “up close and personal” action that

Copyright: © 2013. The author licenses this article under the terms of the Creative Commons Attribution 3.0 License.

*Miami University. Now at Syracuse University. Address: Department of Psychology, 430 Huntington Hall, Syracuse, NY 13244. E-mail: gjkoop@syr.edu.

leads to serious harm to a specific person or group of people (Greene et al., 2001). Impersonal dilemmas maintain more psychological distance and do not elicit as strong of an emotional response.

Functional imaging has supported the claim that personal dilemmas elicit a heightened emotional reaction relative to impersonal dilemmas (Greene et al., 2001), and that choosing against these emotional considerations is associated with increased activation in cognitive control areas like dorsolateral prefrontal cortex (dlPFC) and anterior cingulate cortex (ACC; Greene et al., 2004). Case studies have suggested that the ventromedial prefrontal cortex (vmPFC) is integral to properly utilizing this emotional signal in moral judgment because individuals with damage to this area show an increased rate of utilitarian responding (Koenigs et al., 2007; Mendez, Anderson, & Shapira, 2005; Neary et al., 1998). Within the DI dual-systems framework, it is assumed that vmPFC, orbitofrontal cortex (OFC), and amygdala constitute part of the emotional circuit that generally dominates moral decision making (Mendez, 2009), but this circuit can be overridden with substantial effort via dlPFC.

Behavioral studies have also lent support to the notion that personal dilemmas create more conflict than impersonal dilemmas, and that this is especially true for utilitarian responses (Greene et al., 2001; Moore et al., 2011; Greene, 2009; but see McGuire et al., 2009). Greene and colleagues' (2008) use of a cognitive load manipulation is a prime example. Under cognitive load, participants took longer to make utilitarian choices relative to a control (no load) condition, whereas deontological judgments were unaffected (Greene et al., 2008). These authors concluded that there is a unique cognitive component involved in moral decision making; one that does not merely provide a post hoc rationalization of an emotion-driven response. Conway and Gawronski (2013; Experiment 2) used a process dissociation method in concert with a cognitive load task and drew similar conclusions about unique deliberative and emotional inclinations in moral decision making. Although these data are in accordance with predictions derived from the dual-systems model, both analyses rely on discrete choice outcomes and are thus somewhat limited in depicting the process that produces those choices (a point to which I will return shortly).

There are obviously other considerations that affect moral judgments like the doctrine of double-effect (Moore et al., 2008; Suter & Hertwig, 2011), physical proximity (Greene et al., 2009), intention (Moore et al., 2008), and action/inaction (e.g., Cushman, Young, & Hauser, 2006) among many others, yet the present work largely focuses on the personal-impersonal distinction given its ubiquity, empirical support, and critical predictions for dual-systems models. As foreshadowed above,

despite this body of research there remain questions surrounding specific aspects of the dual-systems account.

1.2 Questions of temporal dynamics

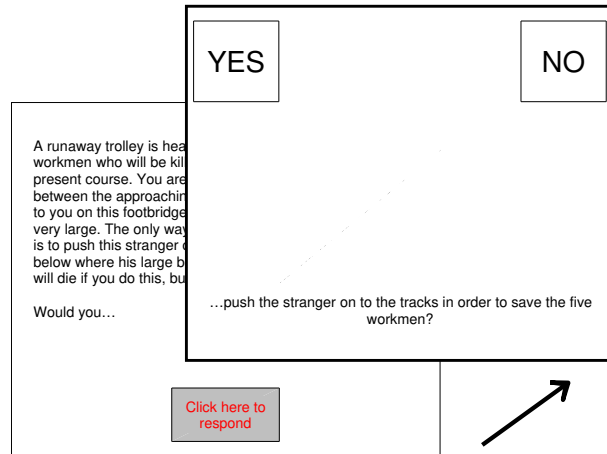
The intent of the personal-impersonal distinction is to allow researchers to design dilemmas in such a controlled way that the analysis of choice proportions or mean RTs is theoretically meaningful. However, the discrete nature of the response (usually a single key press) makes testing the *dynamic* aspects of these models difficult. For example, although it is commonly assumed that the emotional/deontological system and the deliberative/utilitarian system interact sequentially (i.e., the deliberative system overrides a prepotent emotional response), there is scant direct empirical evidence to support this specific claim. The behavioral and neuroscientific methods discussed above lack the temporal resolution to discriminate between sequential and concurrent interactions—a fact that is acknowledged by model proponents (Greene et al., 2004; Greene et al., 2008). It is not surprising, then, that there have been calls to better explicate the time course of this interaction (Huebner, Dwyer, & Hauser, 2009).

Recently, Suter & Hertwig (2011) directly examined these temporal predictions by manipulating decision time through time pressure or with instructions to either decide intuitively or deliberately. In line with the deliberative override prediction, participants made fewer utilitarian choices when under time pressure or when instructed to decide intuitively, yet this was true only on a subset of three “high-competition” personal dilemmas. Given the control generally afforded by the use of small world dilemmas, the fact that this result was only found on a very select subset of dilemmas suggests the need for further replication or converging evidence. At present, this study remains the lone example of response time manipulations affecting choice outcomes, so questions as to whether the two systems interact sequentially or concurrently would be well served by an assessment that can better depict the decision process itself. In short, the uncertainty surrounding the time course of the interaction between systems is classic example of why, to echo Johnson and colleagues' (2008) proclamation, “process models deserve process data.”

1.3 Response dynamics

Response dynamics have proven adept at providing the type of process data that has thus far been lacking in moral psychology. At their simplest, response dynamics experiments continuously track the mouse response as participants move from a central location to one of two disparately spaced onscreen choice options. The typ-

Figure 1: Typical response dynamics trial as used in Experiment 1. After reading the dilemma and clicking the “start” box, participants saw the proposed action and response options.



ical trial begins in the bottom-center of the screen with response options located in the upper left and right corners (Figure 1). Curvature in the mouse response when making a choice is interpreted as “competitive pull” from the non-chosen option (Spivey et al., 2005). Thus, even though participants make discrete choices, the mouse data provide a real-time portrait of how preference for a choice option develops over the course of a trial.

Researchers have used the method in a variety of domains, including the evaluation of statement veracity (McKinstry, Dale, & Spivey, 2008), categorization of atypical exemplars (Dale, Kehoe, & Spivey, 2007), stereotype activation (Freeman & Ambady, 2009), a metric of task learning (Dale et al., 2008, Koop & Johnson, 2011), and self-reported strength of recognition memory (Papesh & Goldinger, 2012). The prior application of response dynamics to questions of sequential versus concurrent processing in the domain of phonological processing (Spivey et al., 2005) indicates that the method is particularly relevant for use in moral psychology. Critically for testing default-interventionist predictions, there are examples of “changes of mind” using response dynamics tasks, both from externally directed target changes (Farmer, Anderson & Spivey, 2007), and from internally motivated preference reversals that arise naturally during the course of evidence evaluation (Resulaj et al., 2009; Koop & Johnson, in press).

Resulaj and colleagues (2009) used a random dot motion paradigm in combination with a directed reaching task to examine the ability of evidence accumulation models (e.g., Ratcliff & Rouder, 1998) to explain changes of mind. In this perceptual decision task, the stimulus was extinguished as soon as response movement was ini-

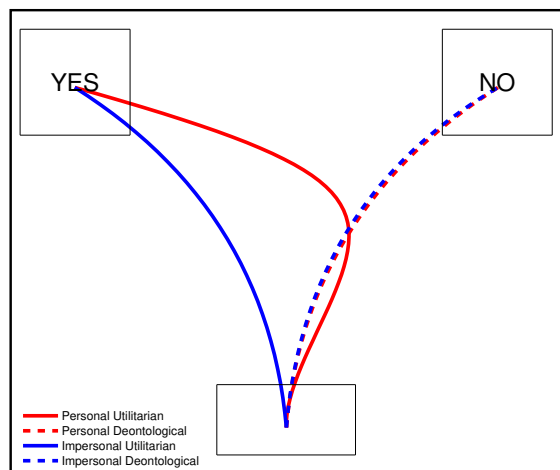
tiated, yet occasionally participants started in one direction before course correcting to the other alternative. The authors posited continued evidence sampling *after* movement initiation, and if sufficient evidence existed during this period a preference reversal could occur.

Koop and Johnson (in press) uncovered similar changes of mind in a higher-order preference task using economic gambles. Participants most directly selected safe gambles in the realm of gains, and risky gambles when in the realm of losses, which followed the classic maxim in behavioral economics of risk seeking in losses and risk aversion in gains (Kahneman & Tversky, 1979; Tversky & Kahneman, 1981). Crucially, choices contrary to this maxim were generally the product of an online preference reversal. For example, when participants selected risky gambles in the realm of gains, they first proceeded towards the safe gamble before reversing course and ultimately selecting the risky option. Like Resulaj et al. (2009), Koop and Johnson showed that a simple attention-driven evidence-accumulation model provided a good fit to the change-of-mind data.

These results are significant for moral psychology because they demonstrate the ability of response dynamics to capture the very behavior that is predicted by DI dual-systems models, yet has proven difficult to uncover. The critical conditions for these models are those combinations of dilemma and response that ostensibly require deliberative override. Specifically, utilitarian responses to personal dilemmas should begin towards the emotionally preferred deontological option before reversing course and ultimately selecting the utilitarian option (Figure 2, solid red line). Deontological responses to personal dilemmas, on the other hand, predict fairly direct response trajectories because the initial emotional impulse “wins” and is never overridden (Figure 2, dashed red line). To round out the response predictions, recall that impersonal dilemmas are not intended to produce conflict between the systems and so there should not be a difference in directness between responses (blue lines). Thus, utilitarian responses to personal dilemmas are critical to testing DI dual-systems models, which predict the same online preference reversals demonstrated in other domains by Farmer and colleagues (2007), Resulaj and colleagues (2009), and Koop and Johnson (in press).

In Experiment 1, I use classic dilemmas from the seminal studies in moral psychology (Greene et al., 2001; Koenigs et al., 2007) with the intent of replicating critical choice and RT results in order to show that the addition of mouse tracking does not produce idiosyncratic results. In Experiment 2, I slightly modify the experimental presentation and use more rigorously constructed stimuli (Moore et al., 2008) in order to provide the best possible environment for uncovering the predicted preference reversals.

Figure 2: Response format and predictions for default-interventionist dual-systems response trajectories. An online preference reversal is uniquely predicted for personal-utilitarian choices. “YES” responses indicate acceptance of the proposed utilitarian action, whereas “NO” responses indicate a deontological preference.



2 Experiment 1

2.1 Method

2.1.1 Participants

I recruited 91 participants from an introductory psychology course via an online sign-up tool, where this experiment was listed among many others. For their participation, students received course credit.

2.1.2 Stimuli

To assess the validity of response dynamics in the realm of moral decision making, I used a set of previously published dilemmas (Greene et al., 2001; Koenigs et al., 2007). Although concerns have been raised about the structure of some of these dilemmas (e.g., McGuire et al., 2009; discussed more fully below), their merit lies primarily in the abundance of empirical data to which the present findings can be compared. Participants read and responded to 29 dilemmas that were deemed acceptable for use with this particular participant population (see Supplemental Materials). Of these dilemmas, 15 were personal dilemmas, 9 were impersonal dilemmas, and the remaining 5 were non-moral “filler” dilemmas.

2.1.3 Procedure

Participants were seated in a group testing room where up to six individuals were tested per session. At least one

empty seat separated participants at all times. After providing informed consent, participants proceeded through self-paced instruction slides that described the nature of the task and provided animated examples of the response process. Participants were instructed that there were no right or wrong answers to these dilemmas, only that they were to make the choice that seemed best to them.

On each trial, participants first read the dilemma text, although the final proposed action was hidden until they clicked the “Start” box located in the bottom middle of a 640 x 480 pixel screen (Figure 1). After clicking the “Start” box, the proposed action appeared along with “YES” (the utilitarian response) and “NO” (the deontological response) response boxes located in the upper corners of the screen. From the time participants clicked the start box until they clicked on their chosen response, I recorded the (x,y) coordinates of the mouse response at a rate of 100 Hz. All participants completed all 29 dilemmas. Dilemma order was randomized for each participant, and the left-right order of the response boxes was counterbalanced across participants. After completing all dilemmas, participants rated each dilemma on the basis of difficulty and emotionality on a 1–9 scale. After completing the survey portion of the experiment, participants were thanked for their participation and dismissed.

2.2 Results

The primary aim of Experiment 1 was to ensure that the data acquired via response dynamics fit with previous studies, as well as to provide an initial descriptive test of DI dual-systems models of moral judgment. Each of these aims provides different analytic constraints. For the former, one must rely on the measures and metrics that have been used previously: choice proportions and RT. For the latter, only those dilemmas that manipulate emotional reactions are necessary to assay the supposed interaction between emotion and deliberation. With these concerns in mind, I first include all dilemmas in order to replicate classic choice and RT effects, before focusing solely on personal and impersonal dilemmas (in accordance with Greene et al., 2001, and Koenigs et al., 2007). Finally, in order to prepare for analysis of response trajectories, trials with RTs more than three standard deviations above the mean were excluded from all analyses.

2.2.1 Analysis of outcome-based metrics

As an initial test for effects of the response method, I compared the proportion of utilitarian choices for each dilemma type (Table 1, “All Dilemmas”) to the normal control data presented by Koenigs and colleagues (2007). In order to provide a fair comparison, I removed three impersonal dilemmas (Supplemental Materials, 22–24) that

Table 1: Choice proportions and response times in Experiment 1

Condition		Pr	Mean RT (SE)
a. All Dilemmas			
Personal	Utilitarian	0.38	4220.06 (157.17)
	Deontological	0.62	3604.62 (117.86)
Impersonal	Utilitarian	0.53	3091.67 (122.28)
	Deontological	0.47	3362.41 (151.14)
b. Non-Unanimous Dilemmas			
Personal	Utilitarian	0.52	4213.74 (161.47)
	Deontological	0.48	4173.20 (179.59)
Impersonal	Utilitarian	0.53	3741.69 (183.28)
	Deontological	0.47	3415.09 (161.27)

were exclusive to Greene et al. (2001). For consistency, these dilemmas are not included in subsequent analyses. Participants were most likely to provide a utilitarian (“yes”) response for non-moral dilemmas, Pr(UTL) = .76, followed by impersonal dilemmas, Pr(UTL) = .53, and were least likely to advocate a utilitarian response in personal dilemmas Pr(UTL) = .39, as evidenced by a statistically significant linear contrast, $F(1,90) = 360.64, p < .001$. Importantly, the overall trend replicates choice data from previous studies (Koenigs et al., 2007); however more theoretically meaningful comparisons can be performed using RT data.

The classic finding from Greene et al. (2001) was a strong Dilemma by Response interaction. Specifically, in personal dilemmas people were faster to make deontological responses relative to utilitarian responses, yet there was no difference in RTs for impersonal dilemmas. As shown in Table 1 (“All Dilemmas”), the present data replicate this finding using Experiment 1’s subset of moral dilemmas. A 2 (personal, impersonal) by 2 (utilitarian, deontological) repeated-measures ANOVA showed that participants were faster on impersonal dilemmas, $F(1,86) = 43.74, p < .001$, and that the difference between utilitarian and deontological responses was dependent on Dilemma, $F(1,86) = 19.74, p < .001$. Specifically, on personal dilemmas participants were faster for deontological responses than utilitarian responses, $t(86) = 4.42, p < .001$, which was not the case with impersonal dilemmas, $t(86) = -1.68, p = .096$.

Although these data mirror the patterns seen in Greene et al. (2001), subsequent research has questioned the cause of this interaction. Specifically, near unanimous deontological responses to a few dilemmas drove this effect (Greene et al., 2008; McGuire et al., 2009). With this in mind, I reevaluated the data while only focus-

ing on those dilemmas that were endorsed by more than 5% of respondents (McGuire et al., 2009). Again confirming that there were no idiosyncratic effects of the response method, RTs from non-unanimous dilemmas (Table 1; “Non-Unanimous Dilemmas”) were in accord with the results presented by McGuire and colleagues (2009; Analysis 3). Effects of Response, $F(1,78) = 1.49, p = .226$, and the Response-Dilemma interaction, $F(1,78) = 1.51, p = .223$, disappeared.¹ The lone remaining statistically significant finding was that participants responded faster to impersonal than personal dilemmas, $F(1,78) = 20.52, p < .001$. In accordance with these data, this 5% exclusion criterion will be used for all subsequent analyses on these stimuli.


An important caveat when performing RT analyses on response dynamics experiments is the possibility most processing is done “offline”—that is, prior to response initiation. If this were the case, any RT differences should be present in a pre-movement window but not while the mouse was actually in motion. To address this concern, I divided total RT into two separate measures: RT_{latency} and RT_{motion} (cf. Dale et al., 2008). RT_{latency} represents the time between clicking the “Start” button and moving the mouse outside of a 50 pixel radius. RT_{motion}, then, represents the remainder of the trial until a response is made. When the analyses just discussed were conducted on these variables, RT_{latency} showed no significant main effects or interactions (F 's < 1.5, p 's > .20). For RT_{motion}, personal dilemmas took longer than impersonal dilemmas, $F(1,78) = 14.19, p < .001$, thus replicating the total RT analyses discussed above. Given the lack of differences in RT_{latency}, it is appropriate to conclude that the data capture online processing.

2.2.2 Analysis of the mouse response

Having replicated traditional choice patterns and shown that the mouse response captured online processing, it is possible to shift focus to the central questions motivating the present research. Specifically, when making a utilitarian response, do individuals have to override an initial emotionally driven preference for the non-utilitarian action, and is this reversal uniquely present in personal dilemmas? Figure 3 presents the aggregate response trajectories for utilitarian and deontological responses in personal and impersonal dilemmas. In order to produce

¹McGuire and colleagues (2009) also performed an item analysis to examine whether the interaction found in Greene et al. (2001) was the product of a few aberrant dilemmas rather than a difference in the general characteristics of personal and impersonal dilemmas. This sort of analysis is beyond the scope of the present work, but suffice it to say that such an analysis would include the variability between dilemmas within a dilemma type (e.g., variability within personal dilemmas), and thus make finding the predicted Response-Dilemma interaction even less likely.

Figure 3: Aggregate response trajectories for non-unanimous dilemmas in Experiment 1. All responses are flipped to the upper left for ease of comparison. Dotted black line represents midpoint between response options; crossings of this axis represent absolute preference reversals.



these trajectories, I first time-normalized each response for each participant into 101 time steps, as is typical in the literature (Spivey et al., 2005, set this precedent). I next created aggregate trajectories for each Dilemma-Response condition for each participant, before collapsing across all participants. At first glance, it is readily apparent that utilitarian choices (solid lines) on personal dilemmas (red lines) do not demonstrate the large preference reversals that would be produced by a deliberative override. In fact, there does not seem to be an effect of either Dilemma or Response. However, to fully uncover any possible differences, one must appeal to individual-level analyses.

Given the dangers inherent in analyzing heavily aggregated data like the response trajectories, I utilized more refined metrics calculated at the level of individual participants to better uncover processing differences. Average absolute deviation (AAD) represents the average deviation of each response trajectory from a direct path between the beginning and end of each response. In general, AAD indicates the curvature of each response, and thus the amount of “competitive pull” being exerted by the non-chosen alternative. A second metric, X_{flips} , simply counts the number of directional changes along the x -axis during a response (perhaps most easily described as “uncertainty” in the response). Finally, I also calculated the number of global preference reversals (Reversals); here operationalized as the number of times a response crosses the y -axis (which is at the midpoint between the two response options) on a given trial. X_{flips} and Reversals may

be the most applicable metrics because they provide the number of momentary valence and absolute preference reversals during a response. Any choice that requires deliberative override (e.g., utilitarian choices on personal dilemmas) should produce more X_{flips} and Reversals than those that do not (e.g., deontological choices).

The response dynamics metrics calculated at the individual level are largely in accord with the pattern seen in the aggregate trajectories (Table 2). 2 (personal, impersonal) by 2 (utilitarian, deontological) repeated-measures ANOVAs run for the dependent measures of AAD, X_{flips} , and Reversals failed to show any main effects or interactions (p 's > .10). This finding was also true for X_{flips} calculated solely outside of the latency radius.

Because this is the first application of response dynamics to moral dilemmas, it is possible that the lack of any observable trend is due to a failure of the method. Two sets of analyses can help to rule out this possibility. First, as shown above, statistically significant differences in RT were present *outside* the latency radius, which suggests the method captured online data. A more convincing argument, however, can be made by using the difficulty ratings that participants provided after completing all dilemmas. Instead of aggregating by dilemma type, I grouped trajectories via a median split on these ratings (Table 2b). The results indicate that high-difficulty dilemmas evinced more curvature than did low-difficulty dilemmas, $F(1,78) = 6.76$, $p = .011$, though there was not an effect of Response. This same effect of Difficulty held true for X_{flips} , $F(1,78) = 13.97$, $p < .001$, and Reversals, $F(1,78) = 4.64$, $p = .034$, as well. Thus, the difficulty-based analysis demonstrates the ability of response dynamics to uncover meaningful differences using moral dilemmas.

Whereas the above analyses completely ignore the personal-impersonal distinction in favor of a difficulty-based distinction, recent work has especially focused on “high-conflict” personal dilemmas (e.g., Koenigs et al., 2007; Greene et al., 2008; Suter & Hertwig, 2011). I therefore repeated the difficulty-based analyses for personal dilemmas alone, which produced largely identical results to the overall difficulty-based analyses. Personal high-difficulty dilemmas showed more Reversals, $F(1,67) = 8.07$, $p = .006$, and X_{flips} , $F(1,67) = 7.39$, $p = .008$, than personal low-difficulty dilemmas, but there were no effects of Response (F 's < 1, p 's > .60) nor interactions (F 's < 1, p 's > .70). There were no significant effects in AAD (F 's < 2.10, p 's > .15).

2.3 Discussion

The goals of Experiment 1 were twofold: to demonstrate the ability of response dynamics to accommodate complex stimuli like moral dilemmas, and to test the specific temporal claims made by dual-systems models of

Table 2: Individual-level analyses for Experiment 1.

	Condition	AAD (SE)	X_{flips} (SE)	Reversals (SE)	N
a. Dilemma-based analyses					
Personal	Utilitarian	48.85 (3.71)	2.84 (0.20)	0.94 (0.06)	90
	Deontological	48.33 (3.83)	2.98 (0.23)	1.03 (0.07)	89
Impersonal	Utilitarian	49.41 (4.62)	3.01 (0.23)	1.07 (0.10)	84
	Deontological	41.90 (3.65)	2.73 (0.21)	0.95 (0.09)	88
b. Difficulty-based analyses.					
High difficulty	Utilitarian	54.13 (4.03)	3.12 (0.22)	1.05 (0.06)	90
	Deontological	54.33 (5.40)	3.36 (0.28)	1.15 (0.08)	83
Low difficulty	Utilitarian	49.15 (3.97)	2.85 (0.19)	1.04 (0.09)	88
	Deontological	41.56 (3.40)	2.65 (0.19)	0.89 (0.06)	90

moral decision making (Greene et al., 2001; Greene et al., 2004). In order to assess the first claim, I utilized a subset of dilemmas that have been widely used in the field (Greene et al., 2001; Greene et al., 2004; Koenigs et al., 2007). Using typical outcome-based measures, I demonstrated that the response dynamics method did not uniquely alter the frequency of utilitarian responses across dilemma types (as compared to Koenigs et al., 2007). Furthermore, RT data from the present study replicated the classic Response by Dilemma interaction described by Greene and colleagues (2001). However, these differences largely disappeared once poorly endorsed dilemmas were excluded from analysis, which is consistent with the reanalysis performed by McGuire and colleagues (2009). Collectively, these results replicated classic findings in the domain, and suggested that the response dynamics method did not have an idiosyncratic effect on decisions in moral dilemmas.

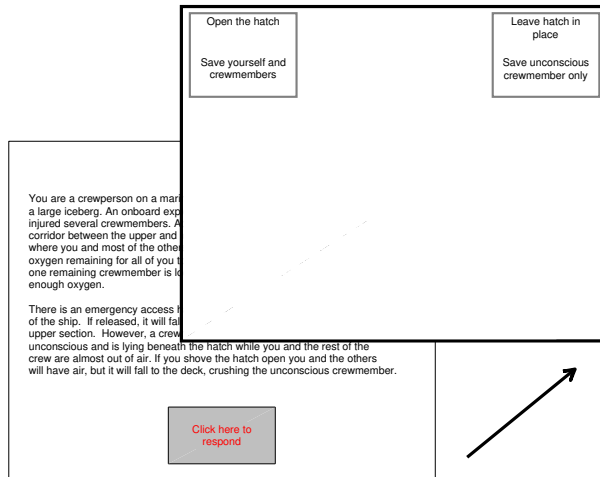
Following these replications, I used the unique analyses afforded by response dynamics to assess specific temporal predictions of DI dual-systems models of moral decision making. Analyses of AAD, X_{flips} , and Reversals failed to support the prediction that utilitarian choices on personal dilemmas require deliberative override. None of the metrics showed differences based on Response or Dilemma. The trajectories did, however, show meaningful differences when grouped by self-reported difficulty rather than by dilemma-type, which demonstrates the ability of the method to depict processing differences for complex stimuli like these dilemmas. Finally, when the difficulty-based analyses were restricted to personal dilemmas there was no effect of Response nor was there an interaction between Response and Difficulty. These findings mirror those of Baron and colleagues (2012), who specifically examined the impact of conflict and similarly failed to show differences between response types

(see the general discussion for lengthier treatment). It is most likely, then, that the failure to show the predicted interaction between Dilemma and Response (or Difficulty and Response for personal dilemmas) represents shortcomings of the personal-impersonal distinction in this dilemma set, rather than deficiencies in the response dynamics method.

3 Experiment 2

Given the difficulty in interpreting null effects, it is important to explore other possible explanations for the lack of support for DI dual-systems models in Experiment 1. To this end, I implemented a few changes in Experiment 2 in order to provide the strongest possible test of DI dual-systems models. First, a new set of stimuli (Moore et al., 2008) were used that better defined personal and impersonal dilemmas by controlling for severity of injury, and systematically varied whether proposed actions were self-beneficial (Benefit; self versus other) and whether harm to another was inevitable (Inevitability; inevitable versus avoidable). Second, in order to ensure that the mouse response captured as much online processing as possible, participants were asked to choose between two courses of action rather than just accepting or rejecting a single course of action. The structure of response presentation, then, was more similar to earlier work demonstrating online preference reversals (Koop & Johnson, in press). Finally, I increased screen resolution to 1280 x 768 pixels in order to require more response movement, and thus greater distance over which response differences could appear. With these changes, Experiment 2 aimed to build on the validation provided in Experiment 1 and create the best possible setting in which to test the predictions of DI dual-systems models of moral decision making.

Figure 4: Trial presentation in Experiment 2. Participants clicked a “start” box in order to populate the two response boxes with possible courses of action and associated outcomes.



3.1 Method

3.1.1 Participants

Participants were 96 introductory psychology students. Recruitment and compensation were identical to Experiment 1.

3.1.2 Stimuli

Participants completed 38 dilemmas that were based on previously published materials (Moore et al., 2008; Supplemental Materials). These 38 dilemmas were comprised of 24 critical dilemmas and 14 “filler” dilemmas. Each of the 24 critical dilemmas had a personal version where the actor must directly kill someone and an impersonal version where indirect actions resulted in another’s death. Whether a participant saw the personal or impersonal version of each dilemma was randomized between participants. Each participant also completed the same 14 “filler” dilemmas. These fillers served to make the stimulus set similar to that of Greene and colleagues (2001), and also made the patterns seen in the critical dilemmas (e.g., “kill one to save many”) slightly less apparent (Moore et al., 2008).

Recall that in Experiment 1 participants were asked whether they would perform a single proposed action. This format is similar to that which is typically used in studies of moral decision making, where participants are asked to judge whether a specific action is “appropriate” or “inappropriate”. The adaptation of this format for Experiment 1 raised a few concerns. First, it is possible that participants could occasionally guess the proposed action based on the partial stem. Secondly, the “yes” and “no”

responses introduced the possibility of an affirmation bias (e.g., Gilbert et al., 1993; McKinstry et al., 2008). In order to better encourage online processing and remove the confounds inherent in “yes/no” responses, dilemmas in Experiment 2 were slightly altered so that each response box was populated with a specific action and that action’s consequence. For example, take the dilemma entitled “Modified Submarine”:

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.

Rather than asking participants “Is it appropriate for you to open the hatch and crush the crewmember below to save yourself and the other crewmembers?”, participants were presented with two options: (a) “Open the hatch” to “save yourself and crewmembers” and (b) “Leave hatch in place” to “save unconscious crewmember only” (Figure 4).


3.1.3 Procedure

The procedure was virtually identical to that of Experiment 1, however participants were not asked to complete a rating block following the dilemmas.

3.2 Results

Although every attempt was made to remain faithful to the original dilemmas used by Moore et al. (2008), the “dual-stem” presentation required minor editing. Thus, in addition to descriptively testing model predictions, it is also important to compare choice results to those reported in previous applications of these dilemmas.


Figure 5: Choice proportions for Experiment 2 across the dimensions of Dilemma (personal, impersonal), Benefit (self, other), and Inevitability (inevitable, avoidable). ± 1 SE.



3.2.1 Analysis of outcome-based metrics

Unlike Experiment 1, dilemmas varied on dimensions of Benefit and Inevitability in addition to the crucial personal-impersonal (Dilemma) distinction. Although these additional factors are not central to the aims of this experiment, they present more opportunities to assess whether the modified stimuli utilized herein affected participants' choice proportions (Figure 5). A 2 (Dilemma) \times 2 (Benefit) \times 2 (Inevitability) repeated-measures ANOVA was used to analyze the average proportion of utilitarian choices for each participant (Figure 5). As predicted, participants were more likely to adopt utilitarian actions for impersonal dilemmas, $F(1,77) = 56.57, p < .001$, for dilemmas where they were saving themselves and others rather than only saving others, $F(1,77) = 45.13, p < .001$, and in instances where death was inevitable rather than avoidable, $F(1,77) = 53.82, p < .001$. Importantly, these are the same main effects described by Moore and colleagues (2008) in the original application of these stimuli, which suggests that the dual-system response paradigm did not substantially affect choice patterns. Unlike the original Moore article, the three-way interaction was not statistically significant, but the two-way interaction between Benefit and Inevitability was statistically significant, $F(1,77) = 9.40, p = .003$. Although participants made utilitarian responses more frequently in self-beneficial situations for both avoidable, $t(77) = 6.69, p < .001$, and inevitable dilemmas, $t(77) = 3.53, p = .001$, utilitarian decisions were particularly lessened when avoidable deaths were not self-beneficial. Overall, the choice behavior largely replicates the findings present in previous work using these stimuli, which allows for greater confidence when analyzing the response trajectories.

Figure 6: Aggregate response trajectories for Experiment 2. All responses are flipped to the upper left for ease of comparison. Dotted black line represents midpoint between response options; crossings of this axis represent absolute preference reversals.



3.2.2 Analysis of the mouse response

Although analyses based on specific dilemma dimensions (i.e., those in addition to the personal-impersonal distinction) are important insofar as they facilitate comparing data collected in response dynamics to previous studies, they are not central to the theoretical considerations driving the present study. The distinction between personal and impersonal dilemmas, however, remains crucial to testing the predictions of DI dual-systems models. The central question in Experiment 2 is whether fast emotional responses must be overridden in order to make a utilitarian choice. Figure 6 shows the time-normalized trajectories for all four Dilemma (personal, impersonal) by Response (utilitarian, deontological) combinations. Unlike Experiment 1, the aggregate response trajectories suggested an effect of Dilemma such that personal dilemmas generally elicited more direct paths than did impersonal dilemmas. More importantly, utilitarian choices in personal dilemmas, far from exhibiting the online preference reversal predicted by default interventionist models, actually seemed to take the most direct path. Finally, these trajectories did not show a meaningful difference between utilitarian and deontological choices. To unpack these results more fully, I again analyzed response trajectories on the level of individual participants.

As described in Experiment 1, AAD, X_{flips} , and Reversals can directly test the critical predictions of DI dual-systems models (Table 3). A 2 (Response) \times 2 (Dilemma) repeated-measures ANOVA on AAD showed that responses to personal dilemmas were subject to less competitive pull than impersonal dilemmas, $F(1,83) =$

Table 3: Individual-level dilemma-based analyses for Experiment 2.

Condition		AAD (SE)	X_{flips} (SE)	Reversals (SE)	N	Pr
Personal	Utilitarian	113.52 (6.37)	3.82 (0.15)	1.31 (0.06)	95	0.58
	Deontological	115.18 (7.23)	4.16 (0.29)	1.57 (0.09)	93	0.42
Impersonal	Utilitarian	124.30 (7.59)	3.83 (0.17)	1.35 (0.06)	96	0.77
	Deontological	138.08 (11.98)	4.33 (0.31)	1.52 (0.10)	87	0.23

4.72, $p = .033$. There was not a statistically significant effect of Response, nor was there a Dilemma-Response interaction (F 's < 1, p 's > .30).

While AAD gives a pretty good picture of the general characteristics of participants' responses, X_{flips} and Reversals remain the most appropriate metrics by which predictions regarding online preference reversals can be tested (Table 3). A 2 (Response) \times 2 (Dilemma) repeated-measures ANOVA revealed that participants showed a greater number of valence reversals (X_{flips}) when making deontological choices relative to utilitarian choices, $F(1,83) = 6.37$, $p = .013$, but that there was no statistically significant difference in X_{flips} between personal and impersonal dilemmas ($F < 1$, $p = .65$). Next, to ensure that these metrics captured online processing differences, I re-analyzed X_{flips} using only data from outside the latency radius. Given the increased screen resolution in Experiment 2, I expanded this latency radius from 50 to 100 pixels. Consistent with the notion that the trajectories reflect online processing, the same pattern of effects was seen for X_{flips} calculated on mouse movements outside of the latency radius. Finally, analysis of Reversals demonstrated the same pattern: participants showed more preference reversals for deontological responses than for utilitarian responses, $F(1,83) = 8.95$, $p = .004$, and this pattern did not differ by dilemma type, ($F < 1$, $p = .881$). Neither X_{flips} nor Reversals showed a significant interaction (F 's < 1, p 's > .45).

3.3 Discussion

The goal of Experiment 2 was to use dilemmas that have rigorously defined the personal-impersonal distinction within a two-alternative response dynamics task in order to test the specific temporal predictions of DI dual-systems models. Prior to testing the critical Dilemma-Response interaction it was important to ensure the modified dilemmas and response method did not have an idiosyncratic effect. The additional dimensions provided by Moore and colleagues (2008) in their dilemmas, although not theoretically critical to this study, offered extra opportunities for comparison and validation. Across the Benefit, Inevitability, and Dilemma dimensions, the

present data replicated the main effects described by Moore et al. (2008): participants were more likely to advocate for utilitarian action when the dilemma was impersonal, when it involved saving oneself, and when the affected person's death was inevitable. In summary, based on the analyses of choice data, the modified response format and slightly altered dilemmas did not have a systematic impact on participants' responses.

After confirming that the novel dilemma presentation did not affect participants' choice behavior, I directly tested the deliberative override assumption of DI dual-systems models. At a descriptive level, the aggregate response trajectories failed to show any hint of an online reversal of preference. In fact, responses to personal dilemmas actually showed less competitive pull from the non-chosen alternative than did impersonal dilemmas regardless of response. X_{flips} and Reversals, which provide the best indication of "changes of mind", also failed to show the predicted interaction. Although there was a main effect of Response, this effect was in the opposite direction as predicted by a DI dual-systems model. Participants showed more reversals when making deontological choices (i.e., those that should not require overcoming an emotional impulse) than when making utilitarian choices.

Although these data produced statistically significant differences between response types, the results should be treated with care. The finding that deontological responses were more direct than utilitarian responses is an unanticipated result, and one that runs contrary to apparent implications of several prior studies. As noted above, the original work by Moore and colleagues (2008) did not show this sort of pattern of effects, and additional studies using similar stimuli (Gürçay & Baron, in preparation; Baron et al., 2012) also failed to show such a dissociation. One possible explanation for this aberrant effect is that a few dilemmas elicited overwhelmingly utilitarian responses; if these dilemmas are removed from analyses the effect largely disappears. The main point of Experiment 2 remains, however, that the discrete choice data largely replicated earlier work, but the continuous response data did not support the decision process that is posited by DI dual-systems models of moral decision making.

4 General discussion

The two experiments presented above utilized response dynamics in order to test the temporal predictions of dual-systems models. Specifically, contemporary models of moral decision making assume that decisions are the product of a fast emotional system and a slower, controlled deliberate system. Furthermore, the emotional system is generally thought to be a prepotent response that must be subsequently overridden by the deliberative system (Greene et al., 2001, 2004, 2008; Haidt, 2007). This sequential interaction is a hallmark of default-interventionist models (DI; Evans, 2008), and has been resistant to direct empirical assessment (Greene et al., 2004; Greene et al., 2008; Huebner et al., 2009). Although the choice data from both experiments closely matched prior work, the mouse trajectories gave no indication of the online preference reversal predicted by DI dual-systems models.

Although these data fail to support the preference reversal predicted by DI dual-systems models, they do not address whether emotion and deliberation actually constitute dual-systems. This more fundamental debate is ongoing (see Greene et al., 2004, for a brief discussion of the general emotion-cognition distinction) but it is beyond the scope of the data presented here. These experiments simply address whether the presumed emotional and deliberative components are active serially or concurrently, and the data discussed above favor the latter explanation. Although these studies contribute the most applicable data to this debate, the idea that these systems operate concurrently is not a new one. Greene and colleagues (2004, 2008) discussed this possibility, and the neuroscientific evidence previously interpreted in light of a serial assumption can also easily accommodate concurrent systems. In short, the anterior cingulate cortex may indicate conflict between concurrently active components rather than actively initiating cognitive override (Greene et al., 2004).

Baron and colleagues (2012) provided an account of the moral decision making process that does not require two competing systems (though dual systems can be accommodated if concurrent activity is assumed). These authors reexamined existing dilemma data using a Rasch analysis and showed that the longest RTs were most likely when dilemma “difficulty” (i.e., the dilemma’s tendency to elicit a utilitarian response) matched a participant’s “ability” (i.e., the individual’s tendency to make a utilitarian response). At this point, the two responses were equally strong in theory, and equally likely in fact. The dual-systems prediction would hold that utilitarian responses would be slower at this point, because some of them would result from mind changing after an initial tendency to make a deontological response. In sum, contrary

to the predictions of a DI dual-systems model, RTs were best predicted by the fit between participant and dilemma, and there were no asymmetries between utilitarian and deontological responses.

Further complicating the depiction of the moral decision process is recent work by Kahane and colleagues, which suggests that the dilemmas most often used to provide support for the dual-systems account may be structured such that they are confounded with intuitiveness (Kahane et al., 2012; Kahane, 2012). That is, rather than reflecting unique utilitarian or deontological systems, it is more likely that any differences between responses are produced by the intuitiveness of the response option. For example, the utilitarian choice can actually be made intuitive if the deontological duty is trivial and the consequence is large (e.g., lying to prevent a murder). Thus, the authors acknowledge that the relative weight assigned to deontological and utilitarian considerations is critical to how this decision process unfolds.

In a similar vein, a significant amount of research has examined which factors play into this relative weighting of duty and consequence. That is, what are those factors that could make the utilitarian action a more intuitive response? The Moore et al. (2008) stimuli used in Experiment 2 provide three such considerations: Benefit, Inevitability, and the personal-impersonal distinction. The more numerous the factors influencing utilitarian tendencies become, the more these dilemmas begin to resemble traditional multiattribute choice experiments. Thus, perhaps the same methods used to study choices between innocuous items like laundry detergent or consumer electronics can be used to examine choices in moral dilemmas. For example, just as selecting a laundry detergent can be described by price, stain-removing quality, and scent, dilemmas can be described by the personal force required (Greene et al., 2009), intentionality (Moore et al., 2008), or the deontological rule being violated (Kahane et al., 2012), to name a few.

Reconceptualizing dilemmas in this fashion opens the possibility of utilizing established process methods like process tracing (e.g., Payne, Bettman, & Johnson, 1993) eye-tracking, or response dynamics to help develop more precise computational models. For example, recent work linked eye-tracking with response dynamics to test the ability of a simple attention-driven accumulation model to predict the mouse response between risky and safe economic gambles (Koop & Johnson, in pres). Future work could similarly utilize a dilemma presentation akin to Experiment 2 in order to examine the degree to which attention to proposed actions and outcomes predicts preference development. Not only would this further allay concerns that the method does not capture online cognition, but could provide additional evidence for the plausibility of an evidence accumulation account, which has been

extraordinarily successful in modeling choice behavior in other domains (e.g., Ratcliff & Rouder, 1998; Busemeyer & Townsend, 1993; Usher & McClelland, 2001).

4.1 Conclusion

In sum, when the moral decision process is allowed to unfold naturally (i.e., no secondary tasks or time pressure), there is nothing in the present data that indicates a default-interventionist interaction between emotional and deliberative systems. Thus, the next steps in the development of process models of moral psychology should involve accommodating the possibility of concurrent interaction within the dual-systems framework, or more fully developing alternative accounts of the moral decision making process. That the data presented here provide a picture of the decision process that diverges from popular models exemplifies the need to test process models with process data, rather than with outcome measures as has traditionally been done.

References

- Baron, J., Gürçay, B., Moore, A. B., & Starcke, K. (2012). Use of a Rasch model to predict response times to utilitarian moral dilemmas. *Synthese*, *189* (15), 107–117.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision Field-Theory—A dynamic cognitive approach to decision-making in an uncertain environment. *Psychological Review*, *100*, 432–459.
- Conway, P. & Gawronski, B. (2013). Deontological and utilitarian inclinations in moral decision making: A process dissociation approach. *Journal of Personality and Social Psychology*, *104*, 216–235.
- Cushman, F. A., Young, L., & Hauser, M. D. (2006). The role of reasoning and intuition in moral judgments: Testing three principles of harm. *Psychological Science*, *17*, 1082–1089.
- Dale, R., Kehoe, C. E. & Spivey, M. J. (2007). Graded motor responses in the time course of categorizing atypical exemplars. *Memory and Cognition*, *35*, 15–28.
- Dale, R., Roche, J., Snyder, K., McCall, R. (2008). Exploring Action Dynamics as an Index of Paired-Associate Learning. *PLoS ONE* *3*(3): e1728. <http://dx.doi.org/10.1371/journal.pone.0001728>.
- Evans, J. St. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, *59*, 255–278.
- Farmer, T. A., Anderson, S. E., & Spivey, M. J. (2007). Gradiency and visual context in syntactic garden-paths. *Journal of Memory and Language*, *57*, 570–595.
- Freeman, J. B. & Ambady, N. (2009). Motions of the hand expose the partial and parallel activation of stereotypes. *Psychological Science*, *20*, 1183–1188.
- Gigerenzer, G. (2010). Moral satisficing: Rethinking moral behavioral as bounded rationality. *Topics in Cognitive Science*, *2*, 528–554.
- Gilbert, D. T., Tafarodi, R. W., & Malone, P. S. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology*, *65*, 221–233.
- 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.
- Gürçay, B., & Baron, J. (n.d.). *New challenges for the two-systems model in moral judgment: Do two-systems models explain what is really going on?* Manuscript in preparation.
- Haidt, J. (2001). The emotional dog and its rational tail: A social-intuitionist approach to moral judgment. *Psychological Review*, *108*, 814–834.
- Haidt, J. (2007). The new synthesis on moral psychology. *Science*, *316*, 998–1002.
- Huebner, B., Dwyer, S., & Hauser, M. (2009). The role of emotion in moral psychology. *Trends in Cognitive Sciences*, *13*, 1–6.
- Johnson, E.J., Schulte-Mecklenbeck, M., & Willemsen, M. C. (2008). Process models deserve process data: Comment on Brandstätter, Gigerenzer, and Hertwig (2006). *Psychological Review*, *115*, 263–272.
- Kahane, G. (2012). On the wrong track: Process and content in moral psychology. *Mind and Language*, *25*, 519–545.
- Kahane, G., Wiech, K., Shackel, N., Farias, M., Savulescu, J., Tracey, I. (2012). The neural basis of intuitive and counterintuitive moral judgment. *Social, Cognitive and Affective Neuroscience*, *7*, 393–402.

- Kahneman, D., & Frederick, S. (2002). Representative-ness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin & D. Kahneman (Eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press, New York.
- Kahneman, D. & Tversky, A. (1979). Prospect theory—Analysis of decision under risk. *Econometrica*, *47*, 263–291.
- Koenigs, M., Young, L., Adolphs, R., Tranel, D., Cushman, F., Hauser, M., & Damasio, A. (2007). Damage to the prefrontal cortex increases utilitarian moral judgments. *Nature*, *446*, 908–911.
- Kohlberg, L., (1969). Stage and sequence: The cognitive-developmental approach to socialization. In D.A. Goslin (Ed.), *Handbook of socialization theory and research* (pp. 347-480). Chicago: Rand McNally.
- Koop, G. J. & Johnson, J. G. (2011). Response dynamics: A new window on the decision process. *Judgment and Decision Making*, *6*, 749–757.
- Koop, G. J. & Johnson, J. G. (in press). The response dynamics of preferential choice. *Cognitive Psychology*.
- Loewenstein, G., Rick, S., & Cohen, J. D. (2008). Neuroeconomics. *Annual Review of Psychology*, *59*, 647–672.
- 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.
- McKinstry, C., Dale, R., & Spivey, M. J. (2008). Action dynamics reveal parallel competition in decision making. *Psychological Science*, *19*, 22–24.
- Mendez, M. F. (2009). The neurobiology of moral behavior: Review and neuropsychiatric implications. *CNS Spectrums*, *14*, 608–620.
- Mendez, M. F., Anderson, E., & Shapira, J. S. (2005). An investigation of moral judgment in frontotemporal dementia. *Cognitive and Behavioral Neurology*, *18*, 193–197.
- 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.
- Moore, A. B., Lee, N. Y. L., Clark, B. A. M., Conway, A. R. A. (2011). In defense of the personal/impersonal distinction in moral psychology research: Cross-cultural validation of the dual process model of moral judgment. *Judgment and Decision Making*, *6*(1), 186–195.
- Neary, D., Snowden, J., Gustafson, L., Passant, U., Stuss, D., Black, S., Freedman, M., Kertesz, A., Robert, P. H., Albert, M., Boone, K., Miller, B. L., Cummings, J., & Benson, D. F. (1998). Frontotemporal lobar degeneration—A consensus on clinical diagnostic criteria. *Neurology*, *51*, 1546-1554.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.
- Papesh, M. H. & Goldinger, S. D. (2012). Memory in motion: Movement dynamics reveal memory strength. *Psychonomic Bulletin & Review*, *19*, 906-913.
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*, 347–356.
- Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, *461*, 263–266.
- Spivey, M. J., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 10393–10398.
- Suter, R. S. & Hertwig, R. (2011). Time and moral judgment. *Cognition*, *119*, 454–458.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and psychology of choice. *Science*, *211*, 453–458.
- Usher, M., and McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, *108*, 550–592.