

Correlations of trait and state emotions with utilitarian moral judgements

Jonathan Baron^a, Burcu Gürçay^a and Mary Frances Luce^b

^aDepartment of Psychology, University of Pennsylvania, Pennsylvania, PA, USA; ^bThe Fuqua School of Business, Duke University, Durham, NC, USA

ABSTRACT

In four experiments, we asked subjects for judgements about scenarios that pit utilitarian outcomes against deontological moral rules, for example, saving more lives vs. a rule against active killing. We measured trait emotions of anger, disgust, sympathy and empathy (the last two in both specific and general forms, the latter referring to large groups of people), asked about the same emotions after each scenario (state emotions). We found that utilitarian responding to the scenarios, and higher scores on a utilitarianism scale, were correlated negatively with disgust, positively (but weakly and inconsistently) with anger, positively with specific sympathy and state sympathy, and less so with general sympathy or empathy. In a fifth experiment, we asked about anger and sympathy for specific outcomes, and we found that these are consistently predictive of utilitarian responding.

ARTICLE HISTORY

Received 31 March 2016
Accepted 4 February 2017

KEYWORDS

Moral judgement;
utilitarianism; deontology;
emotion; anger; disgust;
sympathy; empathy

The relation between emotions and moral judgements have been widely discussed (e.g. Bechara, Damasio, & Damasio 2000; Bechara, Damasio, Damasio, & Lee 1999), particularly with respect to judgements in moral dilemmas that pit utilitarian considerations based on consequences against (roughly) deontological rules that define right actions somewhat independently of their consequences (e.g. Greene, Sommerville, Nystrom, Darley, & Cohen 2001). Our concern here is with individual differences in these utilitarian dilemmas, and particularly with emotions of empathy and sympathy, although we also examine anger and disgust.

Our usage of the terms empathy and sympathy is based on the sort of analysis described by Decety and Cowell (2014), who distinguish three facets of empathy: emotional sharing, empathic concern, and perspective taking. Decety and Cowell's concept of perspective taking is equivalent to what we here call sympathy, while the first two facets are what we call empathy.¹

Utilitarians are stereotypically "cold", because the "sacrificial dilemmas" often used in research require the sacrifice of one person's interest, or life, for the sake of the greater good of others. This stereotype is

supported by studies showing negative correlations between empathetic concern and utilitarian moral judgement (Choe & Min 2011; Conway & Gawronski 2013; Gleichgerrcht & Young 2013; Patil & Silani 2014; Robinson, Joel, & Plaks 2015; Royzman, Landy, & Leeman 2015), although some studies seem to find little or no relation with empathy (Miller, Hannikainen, & Cushman 2014). Note, however, that these sacrificial dilemmas are unusual in pitting a repugnant action against a more abstract outcome involving a larger number of victims (Baron 2011; Bartels and Pizarro 2011).

Decety and Cowell (2014) conclude that "utilitarian judgments are facilitated by a lack of empathic concern" (p. 526), yet they also argue that "perspective taking is a strategy that can be successfully used to ... extend the circle of empathic concern from the tribe to all humanity" (p. 526). We know of only two studies that examined perspective taking and utilitarian judgement: Gleichgerrcht and Young (2013) found no correlation with utilitarian responding in sacrificial dilemmas; and Conway and Gawronski (2013) found results that are unclear.²

We think that perspective taking should correlate positively with utilitarian judgement, other things

being equal. Utilitarianism has as its basis the assessment of consequences for everyone affected, and their aggregation across individuals (Hare 1981). Judgements must be based on an understanding of what it is like to be those affected, given their own values and desires. We call this “sympathy”. However, in sacrificial dilemmas, this effect would happen only when subjects consider all those affected, not just the person who must bear the sacrifice.

The association of all forms of empathy with deontological responding is also suspect. This seems to be the result of the attention given to a few cases in the psychology literature in which the utilitarian response is joined with something emotionally disturbing. It is not at all clear that real-world dilemmas are predominantly of this type, as opposed to the opposite type in which the utilitarian response is the one supported by emotion, particularly emotions such as empathy, while the deontological response is the result of rigid application of a rule, such as Kant’s famous example in which he argues that it would be wrong to lie in order to save someone’s life (Baron 2011). Very few people today would agree with Kant on this point, but people often apply rules of their religion or culture even when they must force themselves, against obvious human passions, to do so.³ One example might be following a rule against abortion, even when it saves the mother’s life and when the fetus would die anyway.

We also examine the role of anger and disgust. Anger has been found to correlate positively with utilitarian judgements, and disgust negatively (Choe & Min 2011). It is possible that these correlations are related to action vs. inaction. As we shall explain, we examine cases in which utilitarian responding is not confounded with acting in the usual simple way. Most of the literature on disgust shows that disgust is negatively correlated with utilitarian judgement (e.g. Choe & Min 2011).⁴

Our studies are about correlations, not causality. Although much research implicitly assumes that individual differences arise because of causal effects of emotion on moral judgement, they may arise for many different reasons (as discussed by Avramova & Inbar 2013; Horne & Powell 2016; Huebner 2015; Huebner, Dwyer, & Hauser 2009; Landy & Goodwin 2015; Teper, Zhong, & Inzlicht 2015). For example, moral beliefs and principles may arise in development in part as a result of experiences that involve emotion. Or emotional predispositions can, over time, lead to

development of moral views. Or certain kinds of education and child rearing could affect both emotions and judgements. Our primary concern here is with correlations between (trait and, to a lesser extent, state) emotions and judgements.

Similarly, we are not concerned with the relation between emotions and moral behaviour (as discussed, for example, by Teper et al. 2015).

1. Overview of experiments

The first four studies attempt to determine the correlation between emotion and utilitarian moral judgement. We were interested in the emotions of empathy, sympathy, anger, and disgust. Each of these emotions was assessed as a trait and, insofar as possible, as a state. The trait measures were self-report questionnaires designed to assess individual differences in susceptibility to each emotion. The state measures were one-item measures of responses to each of several scenarios designed to assess utilitarian responding.

Many researchers suppose that sympathy and empathy play similar roles in determining moral judgement, but we thought they might play different roles, so we tried to distinguish them. We think (following Hare 1981) that sympathy is necessary for utilitarian reasoning. Utilitarians must be able to imagine the preferences and desires of everyone affected in a situation, and weigh these against each other as if they were conflicting goals of a single person. This kind of simultaneous role taking, or some substitute for it, is required to arrive at a utilitarian judgment. Empathy is not required for utilitarian reasoning, since the resulting judgement is just a judgement, a conclusion about the relative value of alternative options for everyone affected by them. Empathy might help motivate people to act on their judgements, especially when action is costly, or to think more thoroughly about them. However, when the task involves hurting one person in order to help others, empathy might be greater for that single person, leading to opposition to utilitarian judgements.

We also tried to distinguish two kinds of sympathy and two kinds of empathy: specific and general. Specific responses are those that are given to particular people who are the focus of attention, such as the person who is the target of the action in a sacrificial dilemma. General responses are those that are given to those outside of the focus of attention. General

sympathy is likely required for utilitarian thought about the problems of the world, now and in the future. We thought it might be a better indication of utilitarian thinking in general. Our hypothesis was not supported.

The four studies we did yielded inconsistent results by the criterion of statistical significance. Effects of interest were sometimes significant and sometimes not significant. Although we shall report these results, we rely more heavily on an analysis of all four experiments combined. For this analysis, results are clearer.

The results we found apply to trait measures. As expected disgust is correlated negatively with utilitarian judgements, and anger positively (but very little). Specific sympathy is correlated positively with utilitarian judgements, which might be surprising given past research. General sympathy is not correlated with utilitarian reasoning. Nor, for the most part, is empathy of either type.

Correlational results like these can be of interest even if they do not represent any particular population. A correlation indicates that some causal factor affects both variables. Our concern should be that this causal factor is not just the experimenter's sampling procedure. Our subjects came from a group who volunteered to complete questionnaires for pay. We cannot see how this sampling procedure can induce the sorts of correlations that we find.

We summarise each of the four main studies briefly, then report an analysis of all four studies combined, and then report a fifth study concerned with state emotions.

2. Experiment 1

To test utilitarian judgement against deontological rules, we need to pit the two against each other. The "standard" sacrificial dilemmas used in most recent research on moral judgement do this by asking about harming or killing (hence sacrificing) one person (usually) in order to prevent greater harm/death to others. They always involve act/omission, and they always involve numbers. That is, they ask subjects to compare killing one person, and thus violating some deontological rule, to doing nothing and letting several others die. The deontological rule forbids killing the one, but, to balance that and create a conflict, numbers are used to make the other option more attractive. Moreover, the standard items always put emotion on the side of doing nothing by making the killing vivid in various ways.

In Experiment 1 and subsequent experiments, we use the standard items, which we call "number" items, but we supplement them with a set of new items, which we call "rule" items. These pit deontological rules against various outcomes that are obviously worse in ways that do not involve numbers. To make up such items, we needed to find fairly strong deontological rules. (A rule against lying in general would probably not suffice, but a rule against lying under oath might work, as it is a composite of a rule against lying with a rule against breaking a promise.) The rule items, in principle, allow emotion to be on either side. For most of them, we tried to put emotion at least a little on the utilitarian side, unlike the role of emotion in the standard items. The online supplement shows the items we used in Experiment 1 and all other studies.

2.1. Method

The subjects were from a panel of about 1200 people who volunteered to do studies for pay on the World Wide Web over the last 15 years, through advertising, links from various web sites (such as those dealing with "how to make money on the Web"), and word-of-mouth. These were mostly Americans, varying considerably in age, income, politics, and educational level, but with women over-represented. Subjects who did not take previous studies seriously had been removed over the years. The panel was divided into three groups, so that closely related studies could use different samples. However, it was necessary to give Experiments 2 and 4 to the same group, and Experiments 3 and 5 to another group. Overlap was not complete even when studies were given to the same group; only a subset of each group of 400 actually did the study while it was available. We aimed for 100 subjects in each study but could not stop exactly at that number because the experiments were programmed in JavaScript, which runs entirely on the subject's computer. Subjects were paid a fixed amount for each study, ranging from \$5 to \$7, depending on the number of questions.

In each experiment, subjects responded to 15–20 moral dilemmas, presented in a different random order to each subject. The dilemmas all involved the choice of another person (such as X), and the first question was of the form "should X [choose one of the options]". In most studies, we used two types of dilemmas: sacrificial dilemmas in which one person

is sacrificed in order to save others; and “rule dilemmas” in which a deontological rule must be violated in order to achieve a superior outcome. A number of other questions, different for each experiment, followed on the same page. After these scenarios, subjects answered a number of questions about trait emotions and moral opinions, as we shall describe.

The 104 subjects in Experiment 1 had a median age of 45 (range 22–79); 32% were male. We used 21 scenarios (see online supplement), each followed by a yes/no question and a number of other questions. They were presented in a different random order to each subject. The utilitarian (U) answer was always “yes”. Eleven dilemmas were the new ones with rules. Ten were from Gürçay and Baron (2016), slightly modified.

Each dilemma was followed by a yes/no question asking about what the agent should do. This was followed by other questions, on the same page. For example, the first rule item in Experiment 1 read as follows:

X is a doctor and medical researcher who believes that he has a method to undo much of the brain damage caused by serious head injuries. It has worked on animals. The treatment is under review for testing on humans, but it is taking a long time, because the review board is worried about the fact that unconscious patients cannot give informed consent.

Meanwhile Joan is brought to the emergency room, unconscious, and alone, with exactly the kind of injury that X thinks he can treat. Without treatment, Joan has no chance of meaningful recovery and will be severely disabled. The treatment cannot make her worse, and it might allow her a nearly complete recovery. Nobody will know whether X has given his treatment to Joan.

Should X try the treatment on Joan? yes/no
How much did you experience the feelings of those involved in this case?

Not at all A little Strongly As if I were in the situation
As you read this case, how much did you feel as if you understood how those involved would feel?

Not at all A little Strongly As if I were in the situation
How much of each the following would you experience if you had to make the decision?

anger: None A little A great deal
sadness: None A little A great deal
conflict: None A little A great deal
sympathy: None A little A great deal

When you made your choice about what X should do, did you base your choice on what the story said about the effects of each option?

Yes.

Mostly, but I took into account the possibility that unexpected things could happen.

No, I thought the story could never happen, and I made my choice on the basis of what I thought would really happen.

The first two follow-up questions were designed to assess empathy and sympathy, respectively. We had no hypothesis about sadness and conflict and, in retrospect, should have omitted those questions. We do not discuss them further.⁵

2.1.1. Scales

Following the scenarios, subjects completed several scales, all shown in the supplement if not in the text. The first (supplement) was an attempt to assess utilitarian morality, constructed by us. We modify this in subsequent studies, but we call it *U*-scale throughout.

Next, following Choe and Min (2011), we used the revised disgust scale of Olatunji et al. (2007).

Our sympathy and empathy items were modifications of those used by Escalas and Stern (2003). We used two versions of each, one referring to specific individuals and one referring to groups. We thought that the group measures might be more correlated with utilitarian responses, but this did not happen:

Sympathy items

Specific

When I read or hear news stories about specific other people:

I understand what the people were feeling.

I understand what was bothering people.

I try to understand the events as they occurred.

I try to understand people’s motivation.

I am able to recognize the problems that the people in the story had.

General When I read or hear news articles about groups of other people:

[same questions]

Empathy items

When I read or hear news stories about other people:

I experience feeling as if the events were really happening to me.

I feel as though I were one of the characters.

I feel as though the events in the story were happening to me.

I experience many of the same feelings as those in the story.

I feel as if the people’s feelings were my own.

Similar changes for general items.

The trait anger scale was that of Spielberg, Jacobs, Russel, & Crane (1983), using the original 4-point response scale (Almost never, Sometimes, Often, Almost always) although Choe and Min (2011) used a 5-point scale.

2.2. Results

We called the yes/no response Act (1 for yes, 0 for no). The individual differences in Act across the 21 items were quite consistent, and, interestingly, the rule items did not form a separate factor. (If we forced 2 factors, each one loads on items from both types.) Although a single-factor model can be rejected, the first principal component accounts for more than twice as much variance as all the others.

The trait emotion measures were highly reliable. Coefficient α was .87 for Anger, .87 for Disgust, .89 for Empathy, and .90 for Sympathy. The two versions of the sympathy measure (specific and general) were so highly correlated with each other ($r = .80$) that they seemed to be measuring the same thing; their individual reliabilities were .84 for specific and .83 for general. (We created new measures for subsequent studies.)

Table 1 shows the correlations of the two moral-judgement scales with *U*-scale, the four trait emotions and the state emotions of interest. The Anger and Disgust results replicate Choe and Min, who found a positive correlation of utilitarian responding with

trait anger and a negative correlation with trait disgust. Interestingly, both of these correlations were just as strong for rule items as for number items. The empathy and sympathy measures, while highly reliable, had non-significant positive correlations with moral judgement. Choe and Min found a negative correlation between trait empathy and utilitarian responding, and we did not replicate this result.

Table 2 shows the results for the state emotions of interest. Experience (a measure of empathy) and Understand (sympathy) correlate clearly and positively with utilitarian responding (Act). These results are based on logistic regression (using the `glmer()` function in the `lme4` package of R: Bates, Maechler, Bolker, & Walker 2015) with random effects for subjects and items, and random slopes for both, testing one measure at a time. There was no interaction with type of item (Rule vs. Number). The main effect is the opposite of the result for empathy reported by Choe and Min. No measures correlate with *U*-scale (assessed by Pearson correlations across subjects, using the subject means of Understand and

Table 1. Correlations of utilitarianism and trait emotions with moral judgements in each of the four experiments. *U*-scale is the utilitarianism scale. *P*-values are for correlations with Act (all scenarios) and are two-tailed and rounded.

	<i>U</i> -scale	Act-rule	Act-number	Act, combined	<i>p</i> -value
Experiment 1					
<i>U</i> -scale	1.00	0.18	0.23	0.24	.00
Disgust	−0.21	−0.19	−0.18	−0.22	.03
Specific sympathy	0.10	0.17	0.13	0.16	.10
General sympathy	0.14	0.15	0.08	0.13	.17
Empathy	0.06	0.16	0.16	0.18	.06
Anger	0.09	0.24	0.26	0.29	.00
Experiment 2					
<i>U</i> -scale	1.00	0.19	0.36	0.35	.61
Disgust	−0.01	−0.04	−0.26	−0.22	.04
Specific sympathy	0.07	0.21	0.15	0.21	.04
General sympathy	−0.05	0.14	0.06	0.12	.24
Specific empathy	−0.04	0.09	0.17	0.17	.10
General empathy	−0.05	0.08	0.02	0.06	.58
Anger	0.15	0.10	0.04	0.07	.48
Experiment 3					
<i>U</i> -scale	1.00	0.37	0.32	0.42	.00
Disgust	−0.02	−0.13	−0.11	−0.15	.13
Specific sympathy	0.27	−0.01	0.24	0.14	.16
General sympathy	0.24	0.04	0.12	0.10	.31
Specific empathy	0.08	0.00	0.06	0.04	.70
General empathy	0.05	−0.09	−0.04	−0.08	.42
Anger	0.10	0.05	−0.01	0.03	.79
Experiment 4					
<i>U</i> -scale	1.00	0.41	0.55	0.59	.00
Divine	−0.42	−0.19	−0.23	−0.26	.01
Specific sympathy	0.15	0.11	0.07	0.11	.28
General sympathy	−0.01	−0.02	−0.06	−0.05	.64
Specific empathy	−0.01	0.09	0.08	0.11	.31
General empathy	−0.07	−0.06	−0.09	−0.09	.39
Anger	0.15	0.06	0.01	0.04	.69

Table 2. Standardized regression coefficients of state emotions with moral judgements (Act) and their correlations with the utilitarianism scale (*U*-scale) in each of the four experiments. *P*-values are two-tailed.

	Act		<i>U</i> -scale	
	Coefficient	<i>p</i> -value	Correlation	<i>p</i> -value
Experiment 1				
Experience	0.486	.0000	0.065	.5569
Understand	0.547	.0000	0.065	.5565
Anger	−0.025	.7890	−0.049	.6630
Experiment 2				
Experience	0.366	.0000	0.243	.0164
Understand	0.205	.0079	0.282	.0051
Anger	−0.073	.4540	0.191	.0610
Experiment 3				
Empathy	0.152	.2761	0.079	.2761
Sympathy	0.359	.0117	0.246	.0107
Anger	−0.074	.4988	0.046	.6358
Experiment 4				
Empathy	−0.122	.1787	−0.152	.1417
Sympathy	−0.055	.5566	−0.080	.4402
Anger	−0.166	.0867	−0.217	.0347

Empathy), but the current version of *U*-scale was replaced in subsequent studies with better versions.

The final question about Belief in the story, treated as a 1–3 scale, correlated strongly with Act. We analysed these data, and other trial-by-trial data reported here, using a mixed-effects model in which both subjects and scenarios were treated as random effects (Bates et al. 2015). Here we also included random slopes for both subjects and scenarios. We used a logistic regression, since Act is coded as 1 or 0. The effect of Belief (standardised) on Act was highly significant ($p < .001$) with a coefficient of .65 (which corresponds to an odds ratio of 1.9). The direction of the effect was that a lower tendency to accept the story as stated led to a lower probability of taking the action. This result is difficult to interpret. It could mean that subjects who have trouble accepting the story tend to think that the truth would argue against the utilitarian choice. But why wouldn't disbelief in the story also allow subjects to favour that choice, if they were so disposed? We suspect that disbelief in the story was at least in part an effect, rather than a cause, of the decision, that is, a *post hoc* justification of a deontological response that would, on the face of it, lead to worse outcomes.

Another major purpose of Experiment 1 was to test the new utilitarian belief scale, *U*-scale. That scale did correlate with Act (Table 1), but it was not very reliable because of the reverse-scored items. It did not correlate significantly with any of the trait emotions. It also was incorrectly designed.⁶

For Experiment 2, we refined the scale by checking the individual items. Item 10 (CAST, see supplement) correlated negatively with Act, so we removed it. We checked the dilemmas the same way and also found one to eliminate because it had a big negative correlation with *U*-scale. The correlation between these refined scales was more substantial (.336, $p = .0019$). This is, of course, data snooping. But note that the non-snooping correlation was significant too, so the result is real, and the purpose of the snooping is to select items for the future. Without CAST, we ended up with 10 dilemmas of each type.

3. Experiment 2

3.1. Method

Experiment 2 was an attempt to improve Experiment 1 in several ways.

We completely re-wrote *U*-scale, breaking it into two parts with different answer types (supplement). On the basis of preliminary analysis (before looking at correlations with Act), we decided to create composite items from the paired items, subtracting one from the other and treating the result as a single item. Again subjects often gave the same or similar answers to both items. (Experiment 3 will have yet another rewrite.)

We re-wrote the trait sympathy and empathy scales. The personal sympathy scale was written to be more such as standard empathy scales, dealing with close social relationships, and the empathy scale had additional items:

Sympathy items:

When something happens to **someone I know**:

I understand what he/she is feeling.

I understand what is bothering him/her.

I try to understand the events as they occurred.

I try to understand the motivation of those involved.

I am able to recognize the problems of those involved.

When I hear or read news stories about **other people whom I don't know**:

I understand what the people were feeling.

I understand what was bothering them.

I try to understand the events as they occurred.

I try to understand people's motivation.

I am able to recognize the problems that the people in the story had.

Empathy items:

When something happens to **someone I know**:

I feel as if the events were happening to me.

I feel as though I were the person affected.

I experience many of the same feelings as the person affected.

I feel as if the person's feelings were my own.

When I hear or read news stories about **other people whom I don't know**:

- I feel as if the events were happening to me.
- I feel as though I were one of the characters.
- I experience many of the same feelings as those in the story.
- I feel as if the people's feelings were my own.

We rewrote the rule dilemmas so that the act/omission distinction was not confounded with the utilitarian/deontological distinction. Almost all scenarios involved a choice of two acts (supplement). We omitted HIKER (because it was difficult to rewrite), leaving 10 items of each type. We also edited the scenarios on the basis of subjects' comments, especially ORPHANAGE, CINDERBLOCK, BIKE WEEK, NUN, and VACCINE, with a minor change in TYCOON.

We revised the Belief question as follows:

- Some people suppose that X has other options, or that the consequences might be different from what the story says. How did you think about this story?
- I assumed that the story was exactly as stated.
- I thought about other options or consequences, but they did not change my answer.
- I thought about other options or consequences, and my answer would change if I assumed that the story was exactly as stated.

We changed "those involved" to "those affected" in the state emotion questions after each scenario. Finally, we re-wrote the rule items so as to unconfound omission from deontological responding (supplement), as well as unconfounding emotion from deontological.

The 97 subjects had a median age of 47 (range 24–69); 32% were male. Because of a programming error, 38 subjects did not do the FIRING scenario, but these subjects' data from the other scenarios, and the data for FIRING scenario from the remaining subjects are included in all analyses. One subject did not do the rule items, but we included data from the number items.

3.2. Results

Table 1 shows the major results. Act-rule (utilitarian responding in the rule scenarios) did not correlate so highly with Act-number (utilitarian responding in the number scenarios) as in Experiment 1 ($r = .31, p = .002$), suggesting that part of the correlation was in fact due to confounding of act/omission with utilitarian/deontological. However, the reliability of the Act-rule scenarios was low ($\alpha = .45$, compared to .75 for Act-

number), which would limit the possible correlation between the two sets of scenarios. Thus, we continue to report correlations with Act, the overall mean of utilitarian responding for each subject.

The new *U*-scale had very low reliability ($\alpha = .19$), possibly because the items measured a variety of utilitarian elements and the correlations among them were no longer confounded by a tendency to agree or disagree overall. As shown in Table 1, it correlated reasonably well with Act, particularly Act-number, but not significantly with any of the trait emotions except specific sympathy.

The four measures of general and specific empathy and sympathy all had reliabilities greater than .80. The two sympathy measures correlated .61 with each other, and the two empathy measures correlated .63, but the cross correlations between sympathy and empathy measures were never greater than .30. The specific/general distinction did not seem to matter much. As shown in Table 1, the correlations of trait emotions with responses to dilemmas were in the same direction as in Experiment 1, but only specific sympathy was significant.

The results for state measures are shown in Table 2.

Again, Understand and Experience correlated positively with Act (and did not interact with type of item). This time, these measures correlated with the somewhat improved version of *U*-scale.

The belief question no longer predicted Act responding at all, whether we treated it as a scale or as a binary measure taking the third response only (indicating that the subject's answer would change except for her additional assumptions, an answer given for 10% of all responses). Compared to Experiment 1, the question was directly about whether added assumptions changed the answer to the question. These results support our earlier interpretation, which is that assumptions are added to bolster a response that would have been chosen anyway.

4. Experiment 3

We made several modifications in Experiment 3. In Experiment 2, we asked about state sympathy and empathy for "those affected" on each page. But those affected are different in the rule and number cases. In the rule cases, breaking a rule usually helps one person. There is no trade-off between different people. The trade-off is between the good of one person and an abstract rule. In the number cases, there is a trade-off between one person and many

other people. So “those affected”, if it means “those potentially harmed”, are different for the two options. The sympathy and empathy questions for the number cases could focus the subject on the harm to the person *directly* affected, thus pushing against the utilitarian choice. In the rule cases, asking about the one person affected pushes *for* the utilitarian choice. We thus eliminated these questions and added sympathy and empathy to the list of emotions. As it happens, these questions did not show any differences between the Rule and Number items, but that was our concern.

We made up an entirely new utilitarian belief scale (*U*-scale, supplement). We presented all the rule items together and all the number items together. The order of the two blocks was counterbalanced across subjects. We also increased the stakes for the affected person in some of the rule dilemmas and eliminated dilemmas that did not correlate well with other dilemmas in Experiment 2, leaving 8 rule cases and 7 number cases (supplement).

4.1. Method

The state-emotion questions now had the form: “How much of each the following would you experience if you had to make the decision?” followed by “anger”, “sadness”, “conflict”, “empathy (experiencing negative feelings that others would have)”, and “sympathy (understanding how others would feel)”. The answer options for each of these were “None, A little, A great deal.”

We used a new version of *U*-scale, refined on the basis of other studies (supplement).

The 107 subjects had a median age of 44 (range 23–69); 35% were male.

4.2. Results

U-scale had a reliability (α) of .63. Reliabilities of Act-rule and Act-number were .64 and .72, respectively. As shown in Table 1, Act did not correlate significantly with any trait emotions, but *U*-scale correlated significantly with trait sympathy ($p = .01$). *U*-scale also correlated with Act, and its correlations with Act-number and Act-rule were both significant at $p < .001$. It did not correlate significantly with any trait emotions (Table 1).

The counterbalancing of rule and number items allowed us to ask about possible order effects between these two-item types. When the rule items

came first, Act (utilitarian responding) was lower overall, with means of .58 (rule first) and .67 (number first; $t_{105} = 2.39$, $p = .019$; the interaction with number vs. rule was near zero and not close to significant). In addition, the correlation between rule and number items, while significant overall ($r = .32$, $p = .00$) was present only when the rule items came first ($r = .49$ for these, and $r = .05$ when number items came first). We do not know why these results occurred or whether they are even real (given that it was a *post hoc* discovery), and we did not explore it further.

As shown in Table 2 state Sympathy correlated with both Act and *U*-scale, but Empathy did not correlate with either.

In sum, we now seem to have a useful version of *U*-scale, and it does correlate with utilitarian responding in both kinds of scenarios. This solidifies our conclusion that the different kinds are measuring a common source of individual differences, despite the fact that most rule items do not involve action vs. omission, or here, emotional bias.

5. Experiment 4

The number items used in Experiments 1–3 were all based on fantastic “sacrificial dilemmas” of the sort that philosophers make up, while the rule items were designed to be more realistic. The number items almost invite subjects to fill in missing details, as noted in Experiment 1. These items have historical interest. They originated with Joshua Greene (e.g. Greene et al. 2001) and have been used subsequently by many other investigators interested in following up Greene’s work. This was an ongoing line of research, where it was relevant and important to use the dilemmas that others had used. However, this rationale is largely irrelevant for our purposes.

Another type of number item was designed to be more realistic. First used by Ritov and Baron (1999), these items and others like them have also been used in subsequent studies. In Experiment 4, we replaced the number items with a set of items derived from these (see supplement).

5.1. Method

The format of each page was identical to that of Experiment 3, except that, for half of the items of each type, picked at random for each subject, the agent “X” was replaced with “you”, throughout the text. We thought that this manipulation might make the state emotion

questions more salient, but it had no effect on either moral judgements or state emotion judgements, so we ignore it henceforth. Order of the items was completely randomised for each subject; we did not block number and rule items as in Experiment 3. The rule dilemmas were edited and supplemented with additional dilemmas from another study (supplement).

We used the state-emotion scales as Experiment 3 for sympathy, empathy, and anger, but not disgust (which is designed for a different sort of dilemma, in the literature). At the end, we used the same trait-emotion scales, but not disgust.

U-scale was expanded to include the last five items from the “Morality scale” shown in the supplement (based on items from Piazza & Landy 2013; Piazza & Sousa 2014). These items improved the reliability (α) of the overall scale, from .55 to .60. The first 6 items of the Morality scale were the measure of divine-command theory used in Baron, Scott, Fincher, and Metz (2015), the belief that all morality comes from God and cannot be fully understood by people. Finally, we included a version of the Actively Open-minded Thinking scale used by Baron et al. (2015), which assesses beliefs about whether good thinking aims at reducing myside bias, the tendency to search for reasons favouring pet conclusions and to give these reasons more weight than reasons on the other side.

The 95 subjects had a median age of 42 (range 19–71); 38% were male.

5.2. Results

Table 1 shows some of the major correlational results. *U*-scale again correlated strongly with both Act-rule and Act-number, and these two scales correlated with each other ($r = .34, p = .001$), despite the replacement of Act-number with new items. Reliabilities (α) of both Act-rule and Act-num were .73. The measure of belief in divine-command theory correlated negatively with *U*-scale and with Act, as expected on the basis of the results of Baron et al. (2015). None of the trait emotion correlations was significant. The measure of actively open-minded thinking, not shown in Table 1, did not correlate significantly with anything except the divine-command scale ($r = -.36, p < .001$), as found by Baron et al. (2015).

As shown in Table 2 correlations with state empathy and sympathy were absent this time, and correlations with anger were in the opposite direction of other findings.

6. Joint analysis of Experiments 1–4

Experiments 1–4 all addressed the question of the relation between emotion and utilitarian judgement. Each experiment attempted to improve on the last. Each involved state and trait emotions.

In general, the results of the four studies are variable, weak, and somewhat inconsistent. The difficulty of replicating results like these becomes apparent when we cannot even replicate our own results with slightly different measures. The most interesting results are that some of the emotion measures, especially sympathy, previously found to correlate negatively with utilitarian responses, sometimes correlate positively. We replicated previous results concerning anger and disgust in some studies but not others.

However, it is possible to combine the results of the four studies so as to get one large data set. Many of the scenarios were common to the different experiment, sometimes with small modifications, and the emotions measures were often identical, or similar. We used the following trait measures: anger, disgust (3 studies), sympathy, empathy, separately for general and specific sympathy and empathy. For state emotions, we used Understand and Experience for Experiments 1 and 2, and sympathy and empathy for Experiments 3 and 4. These were different measures, and they behaved differently.

We also had two types of scenarios: number and rule. We continue to call the utilitarian response “Act”, even though the Rule scenarios had two acts as the options.

We also attempted to develop a scale of utilitarian beliefs, *U*-scale. This got progressively better over the four studies, both in terms of reliability and of correlating with the Act responses, so we used this scale only in Experiment 3 and 4.

6.1. Method

We gave each scenario a name that could be used across experiments in which that scenario was repeated (even if modified). Thus, we could see if a given trait emotion correlated with each scenario, using all possible data. We merged all the data into one matrix with 7819 rows. Each row represented a single response of one subject to one item. Each row included the subject’s response (Act, i.e. choice of the utilitarian option: 1 or 0), the type of item (rule or numer), the name of the scenario, the trait scores for Anger, specific sympathy (Sympspec), general

sympathy (Sympgen), specific empathy (Empspec), general empathy (Empgen), and Disgust, and finally, the experiment (1–4). For sympathy and empathy, we ignored Experiment 1, as the general and specific measures there correlated very highly, and we changed them after that. A subject identifier was not included, because we made sure to include only one response to each scenario from each subject. Note that Experiments 2 and 4 were sent to the same group of subjects, so there was considerable overlap (52 subjects out of about 95 in each experiment). In these cases, when a scenario was used in both experiments, we used the subject's Act and emotion responses from Study 4, except for Disgust, which was omitted from Study 4. We assumed that Study 4 was better designed.

6.2. Results

Table 3 shows the results of the overall analysis. In all analyses of the predictors of Act, we tested the interaction with type (Rule vs. Number) and none

Table 3. Combined data from Experiments 1–4. (2) indicates that subjects from Experiment 2 were included for this measure only if they did not do Experiment 4.

Measure	Experiments	Coefficient	<i>p</i> -value
State emotions predicting Act			
Anger	1,(2),3,4	−0.090	.1470
Understand	1,2	0.372	.0000
Experience	1,2	0.410	.0000
Sympathy	3,4	0.090	.2984
Empathy	3,4	−0.025	.7637
Trait emotions predicting Act			
Anger	1,(2),3,4	0.135	.0402
Disgust	1,(2),3,4	−0.193	.0121
General sympathy	(2),3,4	0.014	.8641
Specific sympathy	(2),3,4	0.224	.0056
General empathy	(2),3,4	−0.126	.1396
Specific empathy	(2),3,4	0.101	.2122
State emotions predicting <i>U</i> -scale			
Anger	3,4	−0.062	.3791
Sympathy	3,4	0.139	.0478
Empathy	3,4	0.014	.8486
Trait emotions predicting <i>U</i> -scale			
Anger	3,4	0.133	.0599
Disgust	3	−0.017	.8617
General sympathy	3,4	0.123	.0800
Specific sympathy	3,4	0.195	.0055
General empathy	3,4	−0.039	.5790
Specific empathy	3,4	0.012	.8674

Notes: For prediction of Act, regressions are logistic using the standardised score as the predictor, including random effects (and random slopes) for subjects and items, and *p*-values are based on a large-sample approximation. Prediction of *U*-scale (Experiments 3 and 4 only) are based on Pearson correlations. For state emotions, we used the mean across items.

of these interactions was significant. Trait Disgust predicts Act (negatively, as found by others) but does not predict *U*-scale. If the tendency to choose a utilitarian option differs from pre-existing utilitarian beliefs, then it is the former that is affected by disgust.

Trait anger had the expected but very weak positive correlation with Act and *U*-scale. In Experiments 2–4, where rule items were designed not to confound utilitarian responding with action, anger correlates at least as highly with Act-rule as with Act-number (Table 1), so it is unlikely to result from the correlation between anger and the tendency to act.

The most interesting effects are those for sympathy and empathy. In contrast to some previous research, we found several positive correlations with utilitarian responses for both Act and *U*-scale. The state measures Understand and Experience correlated positively with Act (Table 3), and with *U*-scale in Experiment 2 (Table 2, where the scale was improved compared to Experiment 1, which showed no correlation). In Experiments 3 and 4 combined, state sympathy as well as trait sympathy (both specific and general) correlated with *U*-scale. Specific sympathy and specific empathy correlated positively with Act, although general sympathy did not. General empathy correlated negatively with Act. Additionally, specific sympathy correlated with Act in the combined analysis (Table 3). In general, then correlations with utilitarian responding are more positive for specific than general and more positive for sympathy than for empathy.⁷

It is unclear why understand and experience predicted Act so well but state sympathy and empathy did not predict Act at all, even though the latter questions were intended as substitutes for the former. The former asked subjects what they thought or felt as they read the case. The latter asked subjects to predict their feelings if they were the decision-maker, thus perhaps leading subjects to focus on the decision rather than its outcomes. The results thus suggest that, in future studies, questions about state emotions should focus on reactions to the story and not the decision.

7. Experiment 5

To understand better the roles of state sympathy and anger, we did one more experiment in which asked about each option specifically in the state measures. We report here only the results from this new

measure. This study was sent to the same subjects as those who did Experiment 3, so there was considerable overlap. We felt that the questions we were asking could be answered even by those who had seen the same scenarios before. We report here only the results for sympathy and anger, although (as before) we also asked about sadness.

7.1. Method

The method was the same as Experiment 4, except that we used only the following items: BABY, BRAIN-DRUG, BDI, BIKEWEEK, CINDER, EBOLA, ER, EUTHAN, NUKE, NUN, SUB, FLU, LYING, BURNBLDG. (These included both Rule and Number cases.) After each item, we presented the two options explicitly, so that we could refer to them in the questions on the same page, e.g.:

You are a nun and an administrator in a Catholic hospital, which forbids abortion under all circumstances. Without an abortion, a patient will die and the fetus will also be lost. You know the rules about abortion and you took an oath to follow all the rules, but you never thought you would face such a case. You could allow the abortion. If you did, you would have to quit your job and your order, but you could get a similar job in a lay hospital. What should you do?

Option 1: Break the rules and approve the abortion. The patient will be saved, and she will be able to get pregnant again.

Option 2: Follow the rules and disallow the abortion. It is too late for the patient to move to another hospital, so she would face death.

How much of each of the following would you experience if you had to make the decision?

Anger about the possible results of Option 1:

None, A little, More than a little, A great deal

Anger about the possible results of Option 2: [same options]

Anger about having to choose:

Anger about the whole situation:

Sadness about the possible results of Option 1:

Sadness about the possible results of Option 2:

Sadness about having to choose:

Sadness about the whole situation:

Sympathy (understanding how others would feel) for anyone harmed by Option 1:

Sympathy for anyone harmed by Option 2:

None (or nobody is harmed), A little, More than a little, A great deal

These were the only questions asked on the page.

We also removed the trait empathy scales but kept the other trait scales used in Experiment 4.

The 104 subjects had a median age of 44 (range 21–76) and 46% were male.

7.2. Results⁸

Table 4 shows the main correlational results. Of primary interest are the first two rows. “Sympathy difference” represents the rating for “Sympathy for anyone harmed by Option 2” minus the rating for “Sympathy ... for anyone harmed by Option 1” (the utilitarian option). Likewise, “Anger difference” also represents anger for Option 2 minus anger for Option 1. Option 1 is the utilitarian option. Thus, when subjects chose that option they were more sympathetic for those harmed by the foregone option (Option 2) and more angry about what would happen to them.

The results for the sum of the two sympathy ratings and the sum of the two anger ratings show no correlation with Act. Together with the results for differences, these results suggest that the state correlations are specific to the outcomes of each option and not the results of some overall non-specific emotional response. Likewise, we found no correlations with anger about choice difficulty, or about the entire situation. Supporting this interpretation, trait anger had no relation to state anger difference (coefficient of standardized variable -0.017 , with random effects) or with state anger sum (.059). Although the state sympathy difference was unrelated to trait specific sympathy (coefficient .051) or general sympathy ($-.001$), the state sympathy *sum* was related weakly to trait specific sympathy (.127, $p = .0786$) and more strongly to general sympathy (.211, $p = .0019$). Thus, if anything, trait emotions predict overall higher ratings of the same emotions

Table 4. Results from Experiment 5 concerning predictors of Act and prediction of state measures from *U*-scale.

	Act predictors		U-scale predictions	
	Coefficient	<i>p</i> -value	Coefficient	<i>p</i> -value
Sympathy difference	0.273	.0420	0.247	.0001
Anger difference	0.346	.0078	0.168	.0014
Sympathy sum	-0.104	.3870	0.091	.2257
Anger sum	-0.133	.2470	0.037	.5949
Choice anger	-0.299	.0147	0.019	.7829
Situation anger	-0.208	.0703	0.055	.3961
Trait anger	-0.216	.1880	-0.100	.3103
Trait general sympathy	-0.024	.8790	0.104	.2917
Trait specific sympathy	0.292	.0664	0.083	.4016

Note: Results for Act are based on logistic regression with standardised predictors, with random effects for subjects and items. Results for *U*-scale are standardised coefficients, with the same random effects (except for traits, which are Pearson correlations).

in state measures, regardless of the outcome, but they do not predict differences between outcomes, which are more highly correlated with choice.

8. General discussion

As described in the Introduction, past literature on correlations of utilitarian judgements with trait emotions are labile at best. We have replicated some of those correlations but not others. In particular, we replicated effects of disgust and anger. We found results opposite to previous research for our measures of empathy (although this was weak) and especially sympathy, and the latter occurred in measures of state emotions as well as trait emotions.

We do not know why correlations with anger and disgust occur. Differences among conditions such as rule/number are too small and inconsistent to support conclusions, but we would tentatively suggest that the confounding of utilitarian responding with action in previous studies does not play a role (Table 1⁹).

Anger may be associated with a utilitarian approach because it is largely anger about waste, about “leaving money on the table” (e.g. lives that could be saved). A possible account of correlations with disgust – to our knowledge completely speculative at this point – is that both disgust and rule-based thinking are related to an obsessive personality style, which is often concerned with cleanliness issues as well as rule-governed behaviour. Of course, this hypothesis just puts off the question about why that style incorporates both of these somewhat different traits.

We have suggested that sympathy in particular is associated with true utilitarian reasoning. It is indeed necessary for act-utilitarian analysis, as Hare (1981) and others have argued. (And, in our scenarios, rule utilitarianism could be indistinguishable from deontology.) It is unclear why others have not found the positive correlations that we found. The results for empathy were much more variable, but we suspect that empathy is somewhat dependent on sympathy. If you do not understand someone else’s suffering, it is more difficult to respond to it with a negative emotion of your own.

We have also found labile results within the experiments we did. This lability, however, is mostly the result of some effects being statistically significant in one experiment but not another.

These are mostly small effects, so this kind of variation is not surprising.

In Experiments 1–4, we asked about state emotions in general and found some correlations that were mostly consistent with those from trait emotions. In Experiment 5, we made a distinction concerning the target of the emotion, and the results for anger and sympathy were consistent with the general hypothesis that utilitarian choices are related to specific emotional responses.

In thinking about the role of emotions, we could distinguish three kinds of emotional responses to scenarios or to particular outcomes of one option or the other. First, subjects could have real emotional responses, analogous to the kinds of responses we have to fiction. We know that the scenarios and outcomes are not real, but we become involved enough in them to respond as if they were. Second, subjects could be making judgements about what emotions they *would* have if the scenarios or outcomes were real, or if the subjects were more involved in them. Third, subjects could have affective reactions that are not in themselves emotions. These reactions themselves would not affect such measures as skin conductance. They are evaluative feelings, which can still be characterised in terms of their corresponding emotions. (And “sympathy” is, as we use the term, not an emotion at all, but just such an evaluative feeling.) Of course, these three kinds of response are surely correlated with each other.

Our data suggest that affective reactions (in the general sense, which includes emotions) are related to subjects’ responses. Causality is unclear here. People with a utilitarian world view would tend to get angry at deontological responses *because* they are utilitarians. This causal direction seems more plausible than the opposite, for it is not clear what would cause the affective reaction except for people’s moral opinions. The role of trait emotions is difficult to understand in this context, unless moral violations are frequent enough triggers of emotion to affect how subjects answer questions about their trait emotions.

Our results also support the conclusion that “utilitarian thinking” is a general trait. We found correlations among three kinds of dilemmas: rule dilemmas that mostly did not involve an act/omission distinction, fantasy sacrificial dilemmas of the sort used in recent research, and more realistic number dilemmas. In turn, these all correlated with

a scale that measured (probably) pre-existing moral beliefs.

Notes

1. Our usage follows that of Escalas and Stern (2003) and others. The term “perspective taking” as used in the emotion literature is more descriptive but has a different meaning in the literature on cognitive development, referring to a cognitive ability.
2. Conway and Gawronski found negative correlations of utilitarian judgment with both empathic concern and perspective taking. But they measured utilitarian responding as the difference between judged appropriateness of acting in the usual sacrificial dilemmas and appropriateness in a set of dilemmas designed to make the consequences of acting worse than those of doing nothing. Subjects often said that acting was appropriate in the latter dilemmas, so it is unclear what this difference means. The correlation between perspective taking and appropriateness judgements was not significant for the standard dilemmas only, or for the set of all dilemmas of both types. (We thank Bertram Gawronski for providing the data.)
3. See Kahane et al. (2012), for a similar argument.
4. Ong, Mullette-Gillman, Kwok, and Lim (2014) found that disgust priming *increased* utilitarian responding in one experiment (and not significantly in a second experiment), contrary to earlier results, which they discuss. Moreover, this increase was greater in subjects with greater trait disgust. But the effects were small, and the tests were done across subjects with no attention to item variance in the effect of disgust.
5. The sadness item did not correlate with anything. The conflict item tended to correlate with utilitarian responding, but this is irrelevant to our interest here.
6. The questions had “always” in them, and the responses ranged from “almost never” to “almost always”. Some subjects were giving nonsensical answers, saying “almost always” to one item *and its opposite*. (We included the opposites on purpose, just for this reason, e.g. “1. Some moral rules should always be followed, even if they lead to outcomes that are worse than those from breaking the rules. 2. If a moral rule leads to outcomes that are worse than those from breaking the rule, we should break the rule.”) We then eliminated subjects who gave “almost always” answers to both questions.
7. The non-significant negative results for general empathy may be related to the fact that general empathy was very low. People do not often experience empathic emotions for abstract groups of distant others, although it is possible to understand their situation. The means for the four traits were (on a 1–4 scale): 2.50 for general sympathy; 2.95 for specific sympathy; 1.88 for general empathy; and 2.34 for specific empathy.
8. The supplement contains an additional analysis of sex differences.
9. Bear in mind that the rule cases in Experiment 1 still confounded utilitarian responding with action, although, by design, not with emotion.

Acknowledgments

We thank to Wes Hutchinson for helpful comments.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

- Avramova, Y. R., & Inbar, Y. (2013). Emotion and moral judgment. *Wiley Interdisciplinary Reviews: Cognitive Science*, 4(2), 169–178.
- Baron, J. (2011). Utilitarian emotions: Suggestions from introspection. *Emotion Review*, 3, 286–287 (special issue on “Morality and emotion” edited by Joshua Greene).
- Baron, J., Scott, S., Fincher, K., & Metz, S. E. (2015). Why does the Cognitive Reflection Test (sometimes) predict utilitarian moral judgment (and other things)? *Journal of Applied Research in Memory and Cognition*, 4(3), 265–284 (special issue on Modeling and Aiding Intuitions in Organizational Decision Making, edited by J. Marewski and U. Hoffrage. <http://dx.doi.org/10.1016/j.jarmac.2014.09.003>).
- Bartels, D. M., & Pizarro, D. A. (2011). The mismeasure of morals: Antisocial personality traits predict utilitarian responses to moral dilemmas. *Cognition*, 121, 154–161.
- Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. doi:10.18637/jss.v067.i01
- Bechara, A., Damasio, H., & Damasio, A. R. (2000). Emotion, decision making and the orbitofrontal cortex. *Cerebral Cortex*, 10(3), 295–307.
- Bechara, A., Damasio, H., Damasio, A. R., & Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. *The Journal of Neuroscience*, 19(13), 5473–5481.
- Choe, S. Y., & Min, K. (2011). Who makes utilitarian judgments? The influences of emotions on utilitarian judgments. *Judgment and Decision Making*, 6, 580–592.
- 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.
- Decety, J., & Cowell, J. M. (2014). Friends or foes: Is empathy necessary for moral behavior? *Perspectives in Psychological Science*, 9(5), 525–537.
- Escalas, J. E., & Stern, B. B. (2003). Sympathy and empathy: Emotional responses to advertising dramas. *Journal of Consumer Research*, 29(4), 566–578.
- Gleichgerricht, E., & Young, L. (2013). Low levels of empathic concern predict utilitarian moral judgment. *PLOS One*, 8(4), 1–9.
- 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. (2016). Challenges for the sequential two-systems model of moral judgement. *Thinking and Reasoning*, 23, 49–80. Retrieved from <http://www.tandfonline.com/doi/full/10.1080/13546783.2016.1216011>

- Hare, R. M. (1981). *Moral thinking: Its levels, method and point*. Oxford: Oxford University Press (Clarendon Press).
- Horne, Z., & Powell, D. (2016). How large is the role of emotion in judgments of moral dilemmas? *PLOS One*, *11*(7), e0154780.
- Huebner, B. (2015). Do emotions play a constitutive role in moral cognition? *Topoi*, *34*(2), 427–440.
- Huebner, B., Dwyer, S., & Hauser, M. (2009). The role of emotion in moral psychology. *Trends in Cognitive Sciences*, *13*(1), 1–6.
- 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*, 392–402.
- Landy, J. F., & Goodwin, G. P. (2015). Does incidental disgust amplify moral judgment? A meta-analytic review of experimental evidence. *Perspectives on Psychological Science*, *10*, 518–536.
- Miller, R. M., Hannikainen, I. A., & Cushman, F. A. (2014). Bad actions or bad outcomes? Differentiating affective contributions to the moral condemnation of harm. *Emotion*, *14*(3), 573–587.
- Olatunji, B. O., Williams, N. L., Tolin, D. F., Abramowitz, J. S., Sawchuk, C. N., Lohr, J. M., & Elwood, L. S. (2007). The disgust scale: Item analysis, factor structure, and suggestions for refinement. *Psychological Assessment*, *19*, 281–297.
- Ong, H. H., Mullette-Gillman, O. A., Kwok, K., & Lim, J. (2014). Moral judgment modulation by disgust is bi-directionally moderated by individual sensitivity. *Frontiers in Psychology*, *5*, Article 194.
- Patil, I., & Silani, G. (2014). Reduced empathic concern leads to utilitarian moral judgments in trait alexithymia. *Frontiers in Psychology*, *5*. doi:10.3389/fpsyg.2014.00501
- Piazza, J., & Landy, J. F. (2013). 'Lean not on your own understanding': Belief that morality is founded on divine authority and non-utilitarian moral judgments. *Judgment and Decision Making*, *8*, 639–661.
- Piazza, J., & Sousa, P. (2014). Religiosity, political orientation, and consequentialist moral thinking. *Social Psychological and Personality Science*, *5*(3), 334–342.
- Ritov, I., & Baron, J. (1999). Protected values and omission bias. *Organizational Behavior and Human Decision Processes*, *79*, 79–94.
- Robinson, J. S., Joel, S., & Plaks, J. E. (2015). Empathy for the group versus indifference toward the victim: Effects of anxious avoidant attachment on moral judgment. *Journal of Personality and Social Psychology*, *56*, 139–152.
- Royzman, E. B., Landy, J. F., & Leeman, R. F. (2015). Are thoughtful people more utilitarian? CRT as a unique predictor of moral minimalism in the dilemmatic context. *Cognitive Science*, *39*(2), 325–352.
- Spielberger, C. D., Jacobs, G., Russell, S., & Crane, R. S. (1983). Assessment of anger: The State-Trait Anger Scale. In J. N. Butcher & C. D. Spielberger (Eds.), *Advances in personality assessment* (pp. 159–187). Hillsdale, NJ: Lawrence Erlbaum.
- Teper, R., Zhong, C., & Inzlicht, M. (2015). How emotions shape moral behavior: Some answers (and questions) for the field of moral psychology. *Social and Personality Psychology Compass*, *9*(1), 1–14.