

**Moral foundations vignettes: A standardized stimulus database of scenarios  
based on moral foundations theory**

Forthcoming at *Behavior Research Methods*

Scott Clifford<sup>1\*†</sup>, Vijeth Iyengar<sup>2\*</sup>, Roberto Cabeza<sup>2</sup>, and Walter Sinnott-Armstrong<sup>3</sup>

<sup>1</sup>Department of Political Science, University of Houston

<sup>2</sup>Center for Cognitive Neuroscience, Duke University

<sup>3</sup>Department of Philosophy, and Kenan Institute for Ethics, Duke University

\*Authors contributed equally to this manuscript

†Correspondence should be addressed to Scott Clifford, Department of Political Science,  
University of Houston, 447 Philip Guthrie Hoffman Hall, Houston, TX 77204. Phone: (713) 743-  
8799. Email: sclifford@uh.edu.

**Abstract**

Research on the emotional, cognitive, and social determinants of moral judgment has surged in recent years. The development of moral foundations theory (MFT) has played an important role, demonstrating the breadth of morality. Moral psychology has responded by investigating how different domains of moral judgment are shaped by a variety of psychological factors. Yet, the discipline lacks a validated set of moral violations that span the moral domain, creating a barrier to investigating influences on judgment and how their neural bases might vary across the moral domain. In this paper, we aim to fill this gap by developing and validating a large set of moral foundations vignettes (MFVs). Each vignette depicts a behavior violating a particular moral foundation and not others. The vignettes are controlled on many dimensions including syntactic structure and complexity making them suitable for neuroimaging research. We demonstrate the validity of our vignettes by examining respondents' classifications of moral violations, conducting exploratory and confirmatory factor analysis, and demonstrating the correspondence between the extracted factors and existing measures of the moral foundations. We expect that the MFVs will be beneficial for a wide variety of behavioral and neuroimaging investigations of moral cognition.

Moral psychology has experienced a “renaissance” in recent years, generating a large empirical literature emphasizing the intuitive and emotional aspects of moral judgment (Greene, 2011). The social intuitionism model (Haidt, 2001) and moral foundations theory (Haidt & Joseph, 2004) have been particularly influential in this burgeoning field. According to the social intuitionism model, moral judgment is an intuitive process, characterized by automatic, affective reactions to stimuli. Moral foundations theory builds on this model, categorizing our moral intuitions into “foundations.” Each foundation represents a set of intuitions that have evolved to solve certain social dilemmas. The current and most widely accepted draft of the theory posits five foundations, though proponents argue there are likely more (Graham et al., 2013). These foundations concern dislike for the suffering of others (Care/harm), proportional fairness (Fairness/cheating), group loyalty (Loyalty/betrayal), deference to authority and tradition (Authority/subversion), and concerns with purity and contamination (Sanctity/degradation). Researchers have recently proposed a sixth foundation (Liberty/oppression), focusing on concerns about domination and coercion (Haidt, 2013; Iyer, Koleva, Graham, Ditto, & Haidt, 2012).

Moral foundations theory has had a large impact in psychology and a variety of other disciplines (for a review, see Graham et al. 2013). Researchers have begun integrating moral foundations theory with research on personality (Hirsh, DeYoung, Xiaowen Xu, & Peterson, 2010; Lewis & Bates, 2011), psychological dispositions (Federico, Weber, Ergun, & Hunt, 2013), and life narratives (McAdams et al., 2008). In political science, MFT has been used to explain political attitudes and ideology (Graham, Haidt, & Nosek, 2009; Koleva et al., 2012; Weber & Federico, 2013; Kertzer et al., 2014), to classify moral rhetoric (Clifford & Jerit, 2013;

Sagi & Dehghani, 2013), and to understand how the public evaluates politicians' character (Clifford, 2014).

Cognitive neuroscientists have also investigated the neurological basis of moral judgment using neuroimaging techniques including functional magnetic resonance imaging (fMRI). Early studies compared the neural correlates of moral judgment by comparing evaluations of scenarios with or without moral content (Moll, Eslinger, & Oliveira-Souza, 2001; Moll, de Oliveira-Souza, Bramati, & Grafman, 2002; Moll, de Oliveira-Souza, Eslinger, et al., 2002). More recent studies have examined moral violations of bodily harm (Heekeren et al., 2005), different aspects of purity violations such as incest, iniquity, and infection (Schaich Borg, Lieberman, & Kiehl, 2008), and compared harm with purity violations (Parkinson et al., 2011). Others have even begun to explore the relationship between regional brain volumes and responses to the moral foundations questionnaire (Lewis, Kanai, Bates, & Rees, 2012). Although existing research has found differences in the brain regions evoked between moral domains, to our knowledge, no neuroimaging study has examined the full spectrum of morality as depicted by moral foundations theory.

Scholars conducting research relying on moral foundations theory have developed two widely used measures – the Moral Foundations Questionnaire (MFQ) and the Sacredness Scale (MFSS). The MFQ measures endorsement of abstract moral principles and self-theories, while the MFSS measures one's willingness to perform moral transgressions for money. Yet, the literature lacks a validated and comprehensive stimulus set consisting of vignettes designed to elicit moral *judgment*. Given the dramatic growth of research on moral judgment, the absence of a judgment scale that captures the full breadth of morality represents a substantial impediment to testing and developing theories of morality.

In this article, we introduce and validate the Moral Foundations Vignettes (MFVs). The MFVs consist of 132 scenarios, each of which is a short description of a behavior that violates a particular moral foundation. The MFVs provide a standardized set of scenarios that allows researchers across various disciplines to test a wide variety of theories about the nature of moral judgment. In Study 1, we show that respondents are able to correctly classify which moral foundation is being violated in each scenario. Additionally, we minimized differences between the foundations on a variety of parameters to ensure that subsequent findings could not simply be explained by differences in scenario structure or complexity. These stimulus restrictions should make the MFV particularly useful to researchers in the cognitive neuroscience community. In Study 2, we demonstrate the correspondence of the vignettes to existing moral foundations scales, and investigate the factor structure of our scenarios. Our results demonstrate that the vignettes tap into the intended foundations and correspond with existing measures, but promise to offer new insight into moral judgment.

### **Limitations of existing moral foundations stimuli**

Existing research relies primarily on the Moral Foundations Questionnaire (MFQ), which consists of two sections. The relevance section of the MFQ asks respondents to rate the relevance of 15 considerations to questions of right and wrong, such as “whether or not a person suffered emotionally” (Care), or “whether or not someone did something to betray their group” (Loyalty). The judgment section consists of 15 agree/disagree items, such as “justice is the most important requirement for a society” (Fairness) and “chastity is an important and valuable virtue” (Sanctity).

While the MFQ has been extensively validated (Graham et al., 2011), its design limits the types of questions that can be addressed with it. Most crucially, the MFQ largely relies on respondents' rating of abstract principles, rather than judgment of concrete scenarios. As Graham et al. argue (2009, 1031), moral relevance "does not necessarily measure how people actually make moral judgments," but these ratings are "best understood as self-theories about moral judgment." Yet, individuals' theories of morality (i.e., endorsement of moral principles) might diverge from their specific moral judgments (Haidt, 2001). For example, one might view harm or loyalty as highly relevant to morality, yet refrain from making harsh judgments about others' harmful or disloyal behavior. Moreover, many of these items include an "unstated, and ambiguous referent," such as an authority figure, yet people may "judge MFT issues differently depending on the referents" (Frimer, Biesanz, Walker, & MacKinlay, 2013). Furthermore, some have argued that the MFQ may be overstating ideological divides in morality by focusing on "points of disagreement between partisans—controversial issues (e.g., chastity)—issues unrepresentative of the full spectrum of moral judgments that people make" (Frimer, Biesanz, Walker, & MacKinlay, 2013). These authors call for a broader set of cases to explore judgments of right and wrong.

Some researchers have used the MFQ as a dependent variable (Lee, Sohn, & Fowler, 2013; Napier & Luguri, 2013; Wright & Baril, 2011), but it's unclear how these relationships would translate to moral judgment. Notably, some of these authors even refer to the MFQ as a measure of "moral judgment," highlighting the need for such a measure (Lee, Sohn, & Fowler, 2013). Moreover, neither the MFQ nor MFSS is ideally suited for techniques measuring neural activity. The brevity of the scales would lead to insufficient statistical power and additionally, the lack of control for stimulus length and complexity may introduce potential confounds.

Indeed, some prior evidence shows that simple variations to syntactic structure, including length, can lead to differences in neural activity (Baciu, Ans, & Carbonnel, 2002; Church, Balota, Petersen, & Schlaggar, 2011; Hauk & Pulvermüller, 2004).

Alternatively, some research has employed the Moral Foundations Sacredness Scale (MFSS), which was designed to test respondents' willingness to engage in taboo tradeoffs (Tetlock, Kristel, Elson, Green, & Lerner, 2000), such as kicking a dog in the head (Care) or renouncing one's citizenship (Loyalty) for money (Graham et al., 2009). However, the MFSS is designed to measure an individual's willingness to violate moral norms in exchange for money, as opposed to judgments of others' behaviors.

Finally, numerous papers have used variations of vignettes on a more ad hoc basis. For example, two scenarios involving incest and sex with a dead chicken have been used frequently in research on moral judgment (Feinberg, Willer, Antonenko, & John, 2012; Pennycook, Cheyne, Barr, Koehler, & Fugelsang, 2013). Other researchers have used vignettes including incest and eating a pet dog (Eskine, Kacirik, & Prinz, 2011; Schnall, Haidt, Clore, & Jordan, 2008; Wheatley & Haidt, 2005). Within the neuroimaging literature, researchers have typically constructed their own moral vignettes on an ad hoc basis with some devising scenarios corresponding to the harm and purity moral foundations (Heekeren et al., 2005; Parkinson et al., 2011; Schaich Borg et al., 2008; Schaich Borg et al., 2011). Other neuroimaging studies have used pictorial stimuli to depict certain moral violations but collapse across all forms of violations in their analyses, making it difficult to examine potential differences between specific moral foundations (Harenski, Antonenko, Shane, & Kiehl, 2008; Harenski & Hamann, 2006; Moll, de Oliveira-Souza, Eslinger, et al., 2002). Overall, researchers have relied on a variety of different moral vignettes, many of which have not been normed and none of which cover the full breadth

of the moral domain. Though we believe the MFV will be useful for a wide variety of research, below we discuss two areas of research in which we believe the MFV will be particularly useful.

### *Emotion and moral judgment*

A growing literature examines the effects of emotion on moral judgment. Much research has focused specifically on the role of disgust, with some arguing that disgust causes harsher moral judgments (Eskine et al., 2011; Pizarro, Inbar, & Helion, 2011; Schnall et al., 2008; Wheatley & Haidt, 2005) and others arguing that disgust uniquely affects judgments in the Sanctity domain (Horberg, Oveis, Keltner, & Cohen, 2009; Horberg, Oveis, & Keltner, 2011). More recently, some have argued that arousal, rather than specific emotions, are the driving force behind moral judgments (Cheng, Ottati, & Price, 2013). Yet, as noted by Horberg et al. (2009), the absence of a standardized set of scenarios that violate particular moral foundations makes it difficult to ascertain the specific effects of emotions across moral domains.

### *The neurological bases of moral judgment*

MFT holds that morality follows a five (or six) factor structure with each moral foundation corresponding to separate modules (Haidt, 2013). Whether morality is unified or can be deconstructed into five or fewer factors has been a question of debate (Sinnott-Armstrong, 2007) that could benefit from neuroimaging studies of moral judgment. Initial fMRI evidence points to morality as a non-unified construct, with distinct brain areas for judgments of disgust, honesty, and harm-based violations (Parkinson et al., 2011), though this study was not explicitly framed in terms of the moral foundations.<sup>1</sup> Furthermore, another study found different brain regions corresponding to the binding and individualizing superordinate moral foundations (Lewis et al., 2012). Neuroimaging techniques such as fMRI could also prove useful in testing cross-

---

<sup>1</sup> The Moral Foundations Dictionary (Graham et al., 2009) classifies honesty under the Fairness foundation and also as a “general” moral term, so the precise implications of these findings for moral foundations theory are unclear.



cultural differences found for judgments of moral violations (Graham et al., 2011), given evidence that cultural differences may influence both the location and levels of observed neural activity (for a review see Han & Northoff, 2008). While fMRI studies of morality are potentially useful for testing MFT, due to an absence of a standardized stimulus set, most have focused only on a subset of moral domains, such as purity (Moll et al., 2005; Schaich Borg et al., 2008), harm (Heekeren et al., 2005), fairness (Robertson et al., 2007), or a subset of the moral foundations (Parkinson et al., 2011), and consequently are unable to compare and contrast all of the moral domains.

### **Development of a standardized stimulus set of moral vignettes**

In order to aid researchers in addressing some of the theoretical questions discussed above, we sought to design a new stimulus set that would satisfy the following criteria: a) measure judgment of concrete behaviors, b) contain subsets mapping onto the moral foundations, c) contain a subset of social norm (i.e., non-moral) violations, and d) be suitable for use in behavioral and neuroimaging paradigms. Notably, we do not view these vignettes as measuring every aspect of morality. Our vignettes focuses specifically on judgment of third-party moral violations, as opposed to separable dimensions such as moral praise (Wiltermuth, Monin, & Chow, 2010) or moral character (Chadwick, Bromgard, Bromgard, & Trafimow, 2006).

We began by writing a large number of scenarios representing face valid violations of particular moral foundations, adapting previous stimuli whenever possible. Because our moral intuitions are argued to have evolved in response to social interactions in small group settings, we focus the content of our scenarios on events that could plausibly occur in everyday life. We made an effort to vary the content of scenarios *within* any given foundation in order to reduce redundancy, avoid interactions with memory, and ensure full conceptual coverage of the

foundation. It was also important to avoid overtly political content and reference to particular social groups in order to avoid tautological claims in political research. Furthermore, we tried to avoid scenarios that might require temporally or culturally bounded knowledge. Finally, we made an effort to eliminate any reference to other foundations to increase the likelihood of isolating the influence of a particular moral foundation.

We also took several steps to ensure that our vignettes are suitable for use in neuroimaging studies of moral judgment. First, we constrained the length of the scenarios (14-17 words, 60-70 characters) and maximized readability and comprehensibility by limiting the reading level to above 30 (ranging from 35.5 – 95.7 with an average of 70.8) and reading ease to below 12 (ranging from 3.6-11.7 with an average of 7.08) as measured by the Flesch-Kincaid reading level and reading ease indices.

Second, to encourage respondents to visualize themselves as third-party witnesses, all of our scenarios begin with a “You see...” formulation. For example, a Care violation reads: “You see a girl telling a boy that his older brother is much more attractive than him.” This structure ensures that respondents imagine a third party committing the violation and that any emotions evoked are the result of imagining witnessing the transgression (for a related approach, see Cannon, Schnall, & White, 2010).

Finally, we created a set of social norms violations that were intended to be unusual but not considered morally wrong (for example, drinking coffee with a spoon). The social norms will play an important role in serving as a control stimulus set in neuroimaging studies of moral judgment by allowing for a comparison between appraisals of scenarios that depict a moral violation and scenarios that depict a social, but not moral, violation. Additionally, the social norms violations prevent respondents from expecting a morally loaded transgression in every

scenario. Non-moral control conditions have been used in several fMRI studies to identify the neural circuitry involved when making moral judgments (Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Moll, de Oliveira-Souza, Bramati, et al., 2002; Moll, de Oliveira-Souza, Eslinger, et al., 2002; Parkinson et al., 2011). Below we detail the specific guidelines used to construct scenarios for each moral foundation.

For the Care foundation, we focused on three forms of harm that reflect the diversity of the original conception of Care: emotional harm to a human, physical harm to a human, and physical harm to a non-human animal. This division of Care also reflects evidence from an fMRI study showing that the introduction of bodily harm into either a moral or non-moral scenario can influence the levels of observed neural activity in certain brain regions (Heekeren et al., 2005). To avoid confounds with the Authority and Liberty moral foundations, we avoided any scenarios that invoked a social hierarchy (pretesting suggested that downward harm in a social hierarchy invoked Liberty, while upward harm invoked Authority). Additionally, we focused on scenarios involving strangers to avoid Loyalty considerations.

Fairness vignettes focused on instances of cheating or free riding (e.g., cheating on a test, lying about work hours). Again, we avoided scenarios involving close-knit groups that might invoke Loyalty. Additionally, we attempted to avoid scenarios involving disobedience towards a superior, which might invoke Authority. We avoided instances of unfairness involving race, gender, or structural inequality due to concerns that they would measure political and social attitudes more than moral concerns.

Loyalty violations consisted of individuals putting their own interests ahead of their group, with group defined as family, country, sports team, school, or company. Throughout our testing, we developed three guidelines for an effective Loyalty violation: the behavior occurs

publicly to threaten the reputation of the group, there is a clear out-group in competition with the actor's group, and the actor is perceived as a spokesperson or identifiable member of the group. This characterization of Loyalty is somewhat narrower than usual in the moral foundations literature, but we wanted to avoid scenarios with overt harm that might create a confound with the Care foundation.

Authority violations primarily consisted of disobedience or disrespect towards traditional authority figures (e.g., a boss, judge, teacher, or parent) or towards an institution or symbol of authority (e.g., courthouse, police department).

Sanctity violations include sexually deviant acts (promiscuity, incest), as well as behaviors that would be considered degrading (drunkenly making out with strangers on a bus) or raise contamination concerns (urinating in a public pool, using a stranger's toothbrush). Additionally, we included vignettes inspired by previous research, such as eating a dead pet dog (Haidt, Koller, & Dias, 1993). Notably, all of our scenarios include elements of physical disgust. We focused on physical disgust in order to avoid disputes about whether moral disgust (which might be directed towards political corruption or sick motives) really is the same emotion as physical disgust. Some have suggested that certain symbols, such as a church, cross, or flag, can become sacralized (Haidt, 2013). However, we chose not to include any violations of sacred objects due to concerns that a symbol may become sacralized for reasons related to other moral foundations, creating a confound in the scenario.

Liberty vignettes consisted of behaviors that are coercive or reduce freedom of choice, particularly actions by those in a position of power over another person (e.g., a man forcing his wife to change her religion). Agents in these scenarios include parents, husbands, bosses, and social leaders.

Finally, we created a set of scenarios that represented a violation of social norms, in that they would be seen as unusual, but not morally wrong. We also sought to avoid any content related to a particular moral foundation. Examples include lifting weights in business clothes and wearing a large sun hat indoors.

## **Study 1**

We first sought to validate our scenarios by asking respondents to rate the moral wrongness of each behavior, as well as *why* they believe each behavior is morally wrong. This first step ensured that people understood the scenarios as intended and further reduced the inter-subject variability that may arise due to differences in how respondents classified the particular moral violation. A similar procedure has been employed in a previous fMRI investigation of a subset of the moral foundations (Parkinson et al., 2011) as well as research on the impact of disgust on moral judgment (Horberg et al., 2009). Additionally, we asked respondents to rate the scenarios on imageability, vividness, arousal, frequency, and comprehension.

### *Method*

Respondents were recruited in three waves ( $n = 330, 192, 94$ ) from a national online panel by Qualtrics. After each wave, we discarded or modified scenarios that did not meet our requirements (described below), then fielded a new wave of the study to test new and modified scenarios. Respondents were limited to the age range of 18-40 ( $M = 35, 32, 33$ ), similar to the age range of respondents used in several previous fMRI investigations of moral judgments (Greene, Nystrom, Engell, Darley, & Cohen, 2004; Moll, de Oliveira-Souza, Bramati, et al., 2002; Moll, de Oliveira-Souza, Eslinger, et al., 2002; Schaich Borg, Sinnott-Armstrong, Calhoun, & Kiehl, 2011) and balanced on ideology (to maintain an equal number of liberals,

moderates, and conservatives). We also screened out respondents who failed an instructional manipulation check at the beginning of the survey (Berinsky, Margolis, & Sances, 2013; Oppenheimer, Meyvis, & Davidenko, 2009). The first sentence of the question asked respondents which sections of the newspaper they like to read, while the remaining three sentences instructed respondents to select “classifieds” and “none of the above.” Respondents who did not follow the instructions were considered inattentive and not allowed to complete the survey.

### *Measures*

Respondents were given a random subset (14-16) of the vignettes such that each vignette was rated by approximately 30 individuals. Respondents were first asked to rate how morally wrong the behavior is on a 5-point scale labeled *not at all wrong*, *not too wrong*, *somewhat wrong*, *very wrong*, *extremely wrong*. Next respondents were asked “Why is the action morally wrong? (Select the main reason).” Response options corresponded with each of the moral foundations:

- It violates norms of harm or care (e.g., unkindness, causing pain to another)*
- It violates norms of fairness or justice (e.g., cheating or reducing equality)*
- It violates norms of loyalty (e.g., betrayal of a group)*
- It violates norms of respecting authority (e.g., subversion, lack of respect for tradition)*
- It violates norms of purity (e.g., degrading or disgusting acts)*
- It violates norms of freedom (e.g., bullying, dominating)*
- It is not morally wrong and does not apply to any of the provided choices*

Crucially, we did not use any of the words from the descriptions of the foundations (e.g., betrayal, degrading) in the actual vignettes, minimizing concerns that classification is driven by shared language.<sup>2</sup> Next respondents were asked to rate the comprehensibility (“How easy is it for you to understand what is described in the scenario?”), imageability (“How easy is it for you to

---

<sup>2</sup> There is one exception – a Fairness vignette uses the word “cheating,” which also appears in the description of the Fairness foundation.

clearly imagine what is happening in the scenario?”), frequency (“How often do you see or hear about actions like the one described in this scenario in the media or your daily life?”), and the strength of their emotional response (“How strong was your emotional response to the behavior depicted in this scenario?”) all on 5-point fully-labeled scales.

### *Results and Discussion*

Table 1 displays the results of Study 1, retaining only scenarios that met two criteria. First, at least 60% of respondents must have classified the scenario as violating the intended moral foundation. Second, we excluded scenarios for which 20% or more of respondents classified the scenario as violating a particular unintended foundation (e.g., 20% selected Liberty as the reason why a Care violation is wrong).

<Table 1 here>

The average ratings within each foundation are shown in Figure 1. The top left panel shows the average wrongness ratings for each foundation, which range from 2.0 (Loyalty) to 2.8 (Sanctity), with social norms averaging 0.2 (scales range from 0 to 4). Notably, the Loyalty violations were rated the least wrong among the foundations, with respondents rating them as “somewhat wrong,” on average. The top right panel shows the average percentage of respondents correctly classifying why a behavior is wrong. Average within-foundation classification rates range from 69% (Care) to 94% (social norms). The middle row of the figure shows the average comprehension (left) and vividness ratings (right). As shown in the middle row of the figure, all foundations received high comprehension ratings (2.9-3.3) and high vividness ratings (2.8-3.3). Additionally, we collected ratings of stimulus frequency, as some have suggested that stimulus familiarity is a factor potentially driving the differences in evaluations of stimuli belonging to different categories (Somerville & Whalen, 2006). Across the

foundations, the vignettes were rated as fairly uncommon (bottom left panel), ranging from 0.7 to 1.7.

<Figure 1 here>

Finally, the foundations all induced moderately strong emotional responses (1.6-2.3), with the exception of social norms (0.5). This is unsurprising, as emotional arousal is highly correlated with moral wrongness ratings in our data. Overall, the results suggest that respondents shared a clear and common understanding of the scenarios, and that there are few differences between the foundations in terms of the frequency of the depicted behaviors. Our findings also support the validity of our social norms scenarios, as they were rated low on wrongness and arousal, yet also low on frequency.

## **Study 2**

As our next step we sought to validate the scenarios by establishing internal validity through factor analysis and criterion validity with existing measures of the moral foundations (MFQ and MFSS).

### *Method*

We recruited 510 respondents from a national online panel through Qualtrics. Respondents were limited to the age range of 18-40 and balanced on ideology (to maintain an equal number of liberals, moderates, and conservatives). We also screened out respondents who failed an instructional manipulation check at the beginning of the survey (Berinsky, Margolis, & Sances, 2013; Oppenheimer, Meyvis, & Davidenko, 2009). Respondents were asked what they were doing at that moment. Respondents who selected swimming or riding a bike were excluded, as were respondents who did *not* report using a computer or electronic device.



## *Measures*

Respondents were asked to rate the moral wrongness of all 132 scenarios that met the criteria established in Study 1, fill out the Moral Foundations Questionnaire and Sacredness Scale (Graham et al., 2009), and answer basic political and demographic questions. Respondents were randomly assigned to fill out the MFQ either before or after making the moral judgments. Respondents who failed either of two attention checks embedded in the MFQ were removed from analysis ( $n = 94$ ).<sup>3</sup> For respondents passing both checks, the correlation between ideology and partisan identification was  $r(416) = .62, p < .001$ , for respondents failing at least one check, the correlation was  $r(96) = .21, p = .04$ . This suggests that the attention checks were effectively picking out inattentive respondents.

## *Results and Discussion*

Although we have a predicted factor structure, we begin with an exploratory factor analysis to examine the extent to which individual scenarios load cleanly onto the expected factors, and not others. Respondents' wrongness ratings for all moral scenarios, with the exception of social norms violations, were entered as manifest variables in a maximum likelihood exploratory factor analysis.<sup>4</sup> A parallel analysis was then conducted, which indicated that nine factors should be retained. The nine factors were then submitted to promax rotation. Results are shown in Table 2 with factor loadings greater than or equal to .4 in bold and factor loadings less than .3 in gray.

In the first factor, all 16 Care scenarios representing emotional harm have factor loadings greater than .4. However, none of the Care scenarios representing physical harm load onto this

---

<sup>3</sup> Following previous work, we excluded participants who rated being "good at math" as "very" or "extremely" morally relevant and participants who disagreed that it is "better to do good than to do bad."

<sup>4</sup> We did not include the MFQ or MFSS items because they measure distinct concepts. Whereas we are interested in measuring judgment of others' behaviors, the MFQ is more of a measure of endorsement of general principles and the MFSS measures willingness to make taboo tradeoffs.

factor, with the exception of one modest loading (.33). This scenario involves a slapping a woman during an argument, which may have been construed as emotionally harmful in addition to being physically harmful. All nine of the Care scenarios involving harm to animals load strongly onto the fifth factor, and three of the seven scenarios involving harm to humans load onto this factor as well, though the factor loadings are less strong. The remaining three Care scenarios involving harm to humans load weakly onto the ninth factor along with two Authority items. Given the weak and inconsistent loadings on this ninth factor (only one scenario has a factor loading greater than .4), it likely is not substantively meaningful. Overall, the findings regarding our Care scenarios support research finding that moral violations may be processed differently when they involve bodily harm (Heekeren et al., 2005), though we did not find a clear separation between physical harm to humans and animals.

In the second factor, all 17 Authority scenarios have factor loadings greater than .4, though one has a modest cross-loading (.36) on Loyalty (refusing to stand for a judge). However, four Fairness scenarios and four Sanctity scenarios also have loadings greater than .3 on this factor. 14 of 17 Fairness scenarios load onto the fourth factor with loadings greater than .3 (10 greater than .4), although two of these scenarios cross-load onto the Authority factor. Of the four total Fairness scenarios that load onto Authority, three involve children or students deceiving an authority figure and one involves a greedy manager.

Turning to the third factor, all 16 Loyalty scenarios have factor loadings greater than 0.3. Overall, we find clear separation between the Authority and Loyalty foundations.

11 of the 17 Liberty scenarios load onto the sixth factor, while no other scenarios load onto this factor. Of the remaining six scenarios, one loads weakly onto Fairness, and one loads weakly onto the eighth factor. Given that only two scenarios (Liberty and Care – animals) load

on the eighth factor, it does not appear to be substantively meaningful. Our results regarding Liberty provide some initial support for this new foundation. However, it should be noted that all of the scenarios with factor loadings greater than .4 involved either coercion of children on the part of their parents, or coercion of a woman by a husband or boyfriend. The scenarios involving coercion by a teacher or boss did not load strongly onto the factor. Thus, this factor only seems to be picking up a particular aspect of Liberty involving coercion within family units.

Finally, 14 of the 17 Sanctity scenarios load onto the seventh factor. However, several of these scenarios modestly cross-load onto Authority, though it is not clear what features these scenarios share with Authority. Finally, one of the Sanctity scenarios involving eating a pet dog cross-loads weakly onto the physical harm factor, which includes harm to animals. Overall, the Sanctity scenarios form a coherent factor, including both sexual and non-sexual aspects of physical disgust.

In summary, our findings from the exploratory factor analysis show strong support for the expected divisions within the moral domain. We uncovered factors associated with each of the original moral foundations, as well as the newer Liberty foundation. We also found evidence of a division within the Care foundation depending on whether the violation involves emotional or physical harm. However, at this time it is unclear whether this latter aspect of the Care foundation should be interpreted as primarily representing harm to non-human animals, or as more broadly representing physical harm.

### *Confirmatory Factor Analysis*

Since prior research provides strong predictions about the factor structure of moral judgment, we also analyzed our data using a confirmatory factor analysis. We estimated an eight-factor model consisting of Care – emotional, Care – physical-human, Care – physical-animal,

Fairness, Liberty, Authority, Loyalty, and Sanctity. Our eight-factor model fit the data well ( $\chi^2(6,526) = 12,616.71$ ; RMSEA = .047). We also estimated several simpler models for comparison (shown in Table 3) and in each case our hypothesized model provided a better fit to the data as judged by improvements in the RMSEA and AIC, though the improvements are modest in moving from the seven-factor models to the eight-factor model. Overall, the results provide support for our hypothesized model.

< Table 3 here >

### *Criterion Validity*

We next demonstrate the criterion validity of our new scales by comparing them to two existing measures of the moral foundations—the Moral Foundations Questionnaire (MFQ) and the Moral Foundations Sacredness Scale (MFSS). Because the EFA demonstrated that several of the scenarios did not load on the expected factors, we utilize the seven factor scores from the EFA reported above. All other scales were standardized (correlations shown in Table A2 in the Appendix). Each factor of the MFVs was predicted as a function of each scale of the MFQ (Table 4, top panel) and the MFSS (Table 4, bottom panel) using OLS regression models.<sup>5</sup> Focusing first on the MFQ, in each case the MFQ criterion is a strong predictor of the MFVs subscale ( $\beta$ s = .24-.47; all  $p$ s < .001). Moreover, the criterion scale is always the strongest predictor, although we do not have the statistical power to distinguish between all of these coefficients. Although the MFQ does not contain a Liberty subscale, we examine the predictive power of the original five foundations included in the MFQ. We find that the MFQ Fairness

---

<sup>5</sup> We control for all moral foundations rather than examine bivariate correlations due to the moderate to high correlations between the MFQ and MFSS measures.

scale is a strong predictor of Liberty ( $\beta = .39$ ,  $t(406) = 6.14$ ,  $p < .001$ ), followed by Care ( $\beta = .15$ ,  $t(406) = 2.42$ ,  $p = .02$ ).

The bottom panel of Table 4 predicts our scales as a function of the MFSS subscales. We again find that the criterion subscales are strong predictors of our own scales ( $\beta$ s = .17 to .57; all  $ps < .05$ ). Though we cannot statistically distinguish the criterion coefficients from all other coefficients, we find that in every case, the MFSS criterion is the largest coefficient with one exception. In the case of our Authority vignettes, the MFSS Care coefficient ( $\beta = -.25$ ) is actually larger in magnitude than the MFSS Authority coefficient ( $\beta = .20$ ). Notably, this pattern did not occur with the more extensively validated MFQ scale. The MFSS also has no Liberty subscale, but we again predict our Liberty scale as a function of the original five foundations. In contrast to our MFQ analysis, we find that the MFSS Care foundation is the strongest predictor of the Liberty foundation ( $\beta = .22$ ,  $t(406) = 2.54$ ,  $p = .01$ ), while the Fairness foundation is small and statistically insignificant ( $\beta = .06$ ,  $t(406) = 0.66$ ,  $p = .51$ ).

<Table 4 about here>

Overall, we find strong evidence for the criterion validity of our new scales. Using both existing measures of the moral foundations, we observed the expected pattern in every case. Moreover, the MFVs correspond better with both the MFQ and MFSS than these two existing scales correspond with each other. As shown in the appendix (see Table A1), the correspondence between these two existing scales is only moderate ( $\beta$ s = .05 to .26,  $ps < .42$ ) and in only three out of five cases is the MFQ criterion a significant predictor of the corresponding MFSS scale. Thus, the MFVs demonstrate impressive correspondence with existing measures.

*Correlations with Political Ideology*

Next, we assess the relationships between the moral foundations and political ideology. Previous work finds consistent relationships between the moral foundations, with liberals tending to endorse Care and Fairness more strongly and conservatives tending to endorse Authority, Loyalty, and Sanctity more strongly (Graham et al., 2009). Table 5 shows the correlation between ideology and each foundation, with each column representing a different method for measuring the foundations (the MFQ, MFSS, and MFVs). Starting with the first column, the MFQ measures correlate with ideology (with higher values representing more conservative views) in the expected ways—Care and Fairness are negatively correlated with conservatism, while Authority, Loyalty, and Sanctity are positively correlated with conservatism (all  $ps < .01$ ). Turning to the second column, the MFSS correlations are mostly in the expected directions, with the exception of Fairness, which has an insignificant positive relationship with conservatism. However, only Loyalty and Sanctity are distinguishable from zero ( $ps < .05$ ). Finally, the last column shows that the MFV factors are generally correlated with ideology in the expected directions. The emotional Care factor is insignificant and negatively related to political conservatism, while the physical Care factor has a significant negative relationship ( $p < .001$ ). Similar to the MFSS, the Fairness factor appears to be unrelated to political ideology, in contrast to the MFQ. We suspect this result comes from our focus on proportional fairness, which concerns proportional rewards or compensation for contributions, as opposed to equal outcomes (consistent with Haidt, 2013). For the binding foundations, we find that each is positively associated with conservatism (all  $ps < .05$ ), consistent with the MFQ. Interestingly, Liberty shows a significant negative correlation with conservatism ( $r = -.14, p < .01$ ), suggesting that our scenarios tapped into more liberal aspects of Liberty, such as the autonomy of women and children (Haidt, 2013). Overall, we find largely the same relationships between the moral

foundations and political ideology, regardless of whether we rely on the MFQ or the MFV, with the exception of Fairness.

<Table 5 here>

### *Recommended Set of Vignettes*

Although the factor scores show strong correspondence with criterion measures, a number of vignettes either did not load strongly on the predicted factor, or cross-loaded onto another factor. Thus, in spite of the results from Study 1 and the face-validity of these vignettes, some of them may not be good measures of the intended concept, though we recognize that it is common to only rely on explicit classification ratings (e.g., Horberg et al., 2009; Parkinson et al., 2011). Below, Table 6 displays 90 vignettes that cleared both sets of criteria from Studies 1 and 2. Vignettes are retained only if they met our classification requirements (Study 1), demonstrated factor loadings  $\geq .3$  on the predicted factor, and did not have cross-loadings  $\geq .3$  (Study 2). The resulting set of 90 vignettes contains 10-16 vignettes per foundation, which should provide researchers with a sufficiently large and diverse set of stimuli.

<Insert Table 6 here>

## **General Discussion**

Moral psychology is a rapidly growing field, yet progress is limited by the quality and availability of existing measures. Resolving puzzles regarding the effects of emotion and cognition on moral judgment or investigating the neurological basis of moral judgment demands a validated, standardized stimuli set that covers the moral domain. In this paper, we attempt to contribute to the literature by creating such a stimuli set.

The Moral Foundations Vignettes (MFVs) go beyond existing measures of moral judgment by providing a large, diverse set of scenarios that represent concrete moral violations. Crucially, the MFVs have been carefully controlled in terms of the format and content of the vignettes. With respect to the factors governing the format of the vignettes, we carefully controlled for syntactic structure, word and character length, and comprehensibility as measured by reading ease and reading level. In terms of the factors governing how vignettes were evaluated, we took careful consideration in verifying the ease with which a vivid image could be formed from each vignette, as well as each vignette's classification into a particular moral foundation. Finally, we avoided including behaviors in the vignettes that required previous knowledge of a particular topic or endorsed overtly political attitudes. Apart from the methods by which we standardized the stimuli, we have constructed and validated a large number of vignettes for each foundation affording future researchers the ability to select subsets of stimuli matching the needs of their investigations.

These features of the MFVs will be particularly useful in driving future neuroimaging studies of moral judgment by providing a large set of normed stimuli that will address some of the statistical power and standardization issues arising from existing moral foundations scales. Neuroimaging investigations will be useful in informing questions critical to MFT, such as whether moral judgment is a distinct or unified process, whether five foundations best explain moral judgments, whether moral judgments are consistent across cultures, and whether we see differences in moral judgments based on political ideology.

Our scales correspond well with existing scales, but raise new questions about the moral domain. For instance, we find evidence that the Care foundation consists of at least two separate aspects (emotional and physical harm). Additionally, we find evidence in support of a separable



Liberty foundation, although our results do not perfectly support expectations from initial investigations, as only vignettes involving familial coercion loaded well on this factor. Further research will be needed to clarify the multiple aspects of the Care foundation and further test the validity of the Liberty foundation.

Finally, we found that our vignettes corresponded with with political ideology largely in the same way as previous scales, with the exception of Fairness, which fits with more recent theorizing (Haidt, 2013a). However, we were surprised to find that the Liberty foundation was negatively correlated with political conservatism. This may have been due to the particular scenarios that loaded on this factor, as discussed above, but again suggests that further research is needed on this foundation.

### **Acknowledgments**

This work was supported by internal funds awarded to W.S.A from Duke University. S.C. was supported by the Social Science Research Institute at Duke University. V.I. was supported by a National Science Foundation (NSF) Graduate Research Fellowship (grant number 110640). We thank Jeffrey Brooks for assisting in both the construction and piloting of the moral scenarios and Jonathan Haidt, Jennifer Jerit, Brad Jones, Christopher Johnston, and Matt Motyl for helpful comments. The authors declare that there are no conflicts of interest.

## References

- Baciu, M., Ans, B., & Carbonnel, S. (2002). Length effect during word and pseudo-word reading. An event-related fMRI study. *Neuroscience Research and Communications*, 30(3), 155–165.
- Berinsky, A. J., Margolis, M. F., & Sances, M. W. (2013). Separating the shirkers from the workers? Making sure respondents pay attention on self-administered surveys. *American Journal of Political Science*, 00(0), 1–15. doi:10.1111/ajps.12081
- Cannon, P. R., Schnall, S., & White, M. (2010). Transgressions and Expressions: Affective Facial Muscle Activity Predicts Moral Judgments. *Social Psychological and Personality Science*, 2(3), 325–331. doi:10.1177/1948550610390525
- Chadwick, R. A., Bromgard, G., Bromgard, I., & Trafimow, D. (2006). An index of specific behaviors in the moral domain. *Behavior Research Methods*, 38(4), 692–7.
- Cheng, J. S., Ottati, V. C., & Price, E. D. (2013). The arousal model of moral condemnation. *Journal of Experimental Social Psychology*, 49(6), 1012–1018. doi:10.1016/j.jesp.2013.06.006
- Church, J. A., Balota, D. A., Petersen, S. E., & Schlaggar, B. L. (2011). Manipulation of length and lexicality localizes the functional neuroanatomy of phonological processing in adult readers. *Journal of Cognitive Neuroscience*, 23(6), 1475–93. doi:10.1162/jocn.2010.21515
- Clifford, S. (2014). Linking issue stances and trait inferences: A theory of moral exemplification. *The Journal of Politics*, 76(03), 698–710. doi:10.1017/S0022381614000176
- Clifford, S., & Jerit, J. (2013). How words do the work of politics: Moral foundations theory and the debate over stem cell research. *The Journal of Politics*, 75(03), 659–671. doi:10.1017/S0022381613000492
- Eskine, K. J., Kacinik, N. A., & Prinz, J. J. (2011). A bad taste in the mouth: Gustatory disgust influences moral judgment. *Psychological Science*, 22(3), 295–9. doi:10.1177/0956797611398497
- Federico, C. M., Weber, C. R., Ergun, D., & Hunt, C. (2013). Mapping the connections between politics and morality: The multiple sociopolitical orientations involved in moral intuition. *Political Psychology*, 34(4), 589–610. doi:10.1111/pops.12006
- Feinberg, M., Willer, R., Antonenko, O., & John, O. P. (2012). Liberating reason from the passions: Overriding intuitionist moral judgments through emotion reappraisal. *Psychological Science*, 23(7), 788–95. doi:10.1177/0956797611434747

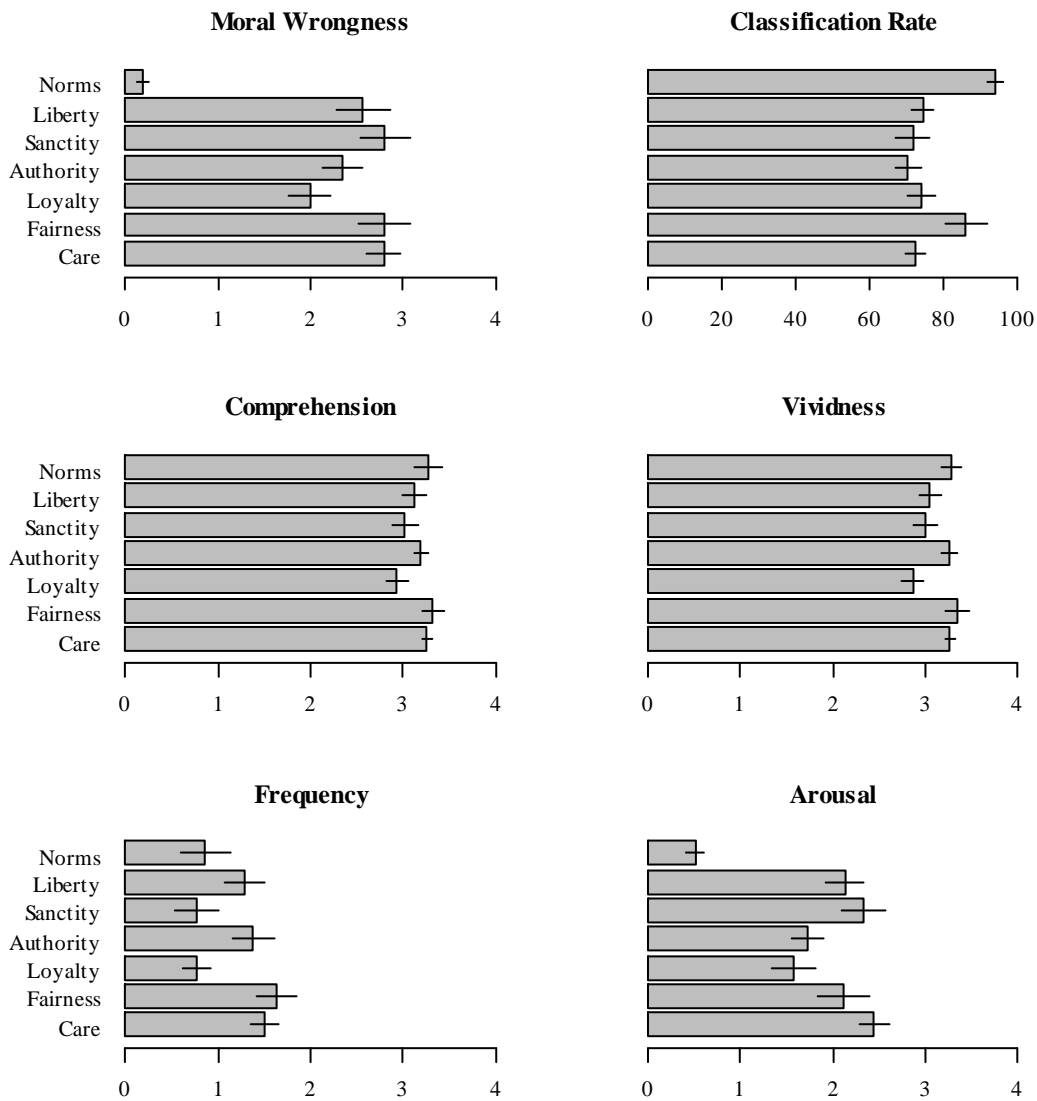
- Frimer, J. A., Biesanz, J. C., Walker, L. J., & MacKinlay, C. W. (2013). Liberals and conservatives rely on common moral foundations when making moral judgments about influential people. *Journal of Personality and Social Psychology, 104*(6), 1040–59. doi:10.1037/a0032277
- Graham, J., Haidt, J., Koleva, S. P., Motyl, M., Iyer, R., Wojcik, S. P., & Ditto, P. H. (2013). Moral Foundations Theory: The pragmatic validity of moral pluralism. *Advances in Experimental Social Psychology, 47*, 55–130.
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology, 96*(5), 1029–46. doi:10.1037/a0015141
- Graham, J., Nosek, B. A., Haidt, J., Iyer, R., Koleva, S., & Ditto, P. H. (2011). Mapping the moral domain. *Journal of Personality and Social Psychology, 101*(2), 366–85. doi:10.1037/a0021847
- Greene, J. D. (2011). Emotion and morality: A tasting menu. *Emotion Review, 3*(3), 227–229. doi:10.1177/1754073911409629
- Greene, J. D., Nystrom, L. E., Engell, A. D., Darley, J. M., & Cohen, J. D. (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron, 44*(2), 389–400. doi:10.1016/j.neuron.2004.09.027
- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science, 293*(5537), 2105–8. doi:10.1126/science.1062872
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review, 108*(4), 814–834. doi:10.1037//0033-295X.
- Haidt, J. (2013). *The Righteous Mind: Why Good People Are Divided by Politics and Religion* (Vintage) (p. 528). Vintage.
- Haidt, J., & Joseph, C. (2004). Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus, 55*–66.
- Haidt, J., Koller, S. H., & Dias, M. G. (1993). Affect, culture, and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology, 65*(4), 613–28.
- Han, S., & Northoff, G. (2008). Culture-sensitive neural substrates of human cognition: A transcultural neuroimaging approach. *Nature Reviews Neuroscience, 9*, 646–654.
- Harenski, C. L., Antonenko, O., Shane, M. S., & Kiehl, K. A. (2008). Gender differences in neural mechanisms underlying moral sensitivity. *Social Cognitive and Affective Neuroscience, 3*(4), 313–21. doi:10.1093/scan/nsn026

- Harenski, C. L., & Hamann, S. (2006). Neural correlates of regulating negative emotions related to moral violations. *NeuroImage*, *30*(1), 313–24. doi:10.1016/j.neuroimage.2005.09.034
- Hauk, O., & Pulvermüller, F. (2004). Effects of word length and frequency on the human event-related potential. *Clinical Neurophysiology*, *115*(5), 1090–103. doi:10.1016/j.clinph.2003.12.020
- Heekeren, H. R., Wartenburger, I., Schmidt, H., Prehn, K., Schwintowski, H.-P., & Villringer, A. (2005). Influence of bodily harm on neural correlates of semantic and moral decision-making. *NeuroImage*, *24*(3), 887–97. doi:10.1016/j.neuroimage.2004.09.026
- Hirsh, J. B., DeYoung, C. G., Xiaowen Xu, & Peterson, J. B. (2010). Compassionate liberals and polite conservatives: Associations of agreeableness with political ideology and moral values. *Personality & Social Psychology Bulletin*, *36*(5), 655–64. doi:10.1177/0146167210366854
- Horberg, E. J., Oveis, C., & Keltner, D. (2011). Emotions as moral amplifiers: An Appraisal tendency approach to the influences of distinct emotions upon moral judgment. *Emotion Review*, *3*(3), 237–244. doi:10.1177/1754073911402384
- Horberg, E. J., Oveis, C., Keltner, D., & Cohen, A. B. (2009). Disgust and the moralization of purity. *Journal of Personality and Social Psychology*, *97*(6), 963–76. doi:10.1037/a0017423
- Iyer, R., Koleva, S., Graham, J., Ditto, P., & Haidt, J. (2012). Understanding libertarian morality: The psychological dispositions of self-identified libertarians. *PloS One*, *7*(8), e42366. doi:10.1371/journal.pone.0042366
- Koleva, S. P., Graham, J., Iyer, R., Ditto, P. H., & Haidt, J. (2012). Tracing the threads: How five moral concerns (especially Purity) help explain culture war attitudes. *Journal of Research in Personality*, *46*(2), 184–194. doi:10.1016/j.jrp.2012.01.006
- Lee, J. J., Sohn, Y., & Fowler, J. H. (2013). Emotion regulation as the foundation of political attitudes: Does reappraisal decrease support for conservative policies? *PloS One*, *8*(12), e83143. doi:10.1371/journal.pone.0083143
- Lewis, G. J., & Bates, T. C. (2011). From left to right: How the personality system allows basic traits to influence politics via characteristic moral adaptations. *British Journal of Psychology*, *102*(3), 546–58. doi:10.1111/j.2044-8295.2011.02016.x
- Lewis, G. J., Kanai, R., Bates, T. C., & Rees, G. (2012). Moral values are associated with individual differences in regional brain volume. *Journal of Cognitive Neuroscience*, *24*(8), 1657–63. doi:10.1162/jocn\_a\_00239
- McAdams, D. P., Albaugh, M., Farber, E., Daniels, J., Logan, R. L., & Olson, B. (2008). Family metaphors and moral intuitions: How conservatives and liberals narrate their lives. *Journal of Personality and Social Psychology*, *95*(4), 978–90. doi:10.1037/a0012650

- Moll, J., de Oliveira-Souza, R., Bramati, I. E., & Grafman, J. (2002). Functional networks in emotional moral and nonmoral social judgments. *NeuroImage*, *16*(3), 696–703. doi:10.1006/nimg.2002.1118
- Moll, J., de Oliveira-Souza, R., Eslinger, P. J., Bramati, I. E., Mourão-Miranda, J., Andreiuolo, P. A., & Pessoa, L. (2002). The neural correlates of moral sensitivity: A functional magnetic resonance imaging investigation of basic and moral emotions. *Journal of Neuroscience*, *22*(7), 2730–6. doi:20026214
- Moll, J., de Oliveira-Souza, R., Moll, F. T., Ignácio, F. A., Bramati, I. E., Caparelli-Dáquer, E. M., & Eslinger, P. J. (2005). The moral affiliations of disgust: A functional MRI study. *Cognitive and Behavioral Neurology*, *18*(1), 68–78.
- Moll, J., Eslinger, P. J., & Oliveira-Souza, R. (2001). Frontopolar and anterior temporal cortex activation in a moral judgment task: Preliminary functional MRI results in normal subjects. *Arquivos de Neuro-Psiquiatria*, *59*(3-B), 657–64.
- Napier, J. L., & Luguri, J. B. (2013). Moral mind-sets: Abstract thinking increases a preference for “Individualizing” over “Binding” moral foundations. *Social Psychological and Personality Science*, *4*(6), 754–759. doi:10.1177/1948550612473783
- Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, *45*(4), 867–872. doi:10.1016/j.jesp.2009.03.009
- Parkinson, C., Sinnott-Armstrong, W., Koralus, P. E., Mendelovici, A., McGeer, V., & Wheatley, T. (2011). Is morality unified? Evidence that distinct neural systems underlie moral judgments of harm, dishonesty, and disgust. *Journal of Cognitive Neuroscience*, *23*(10), 3162–80. doi:10.1162/jocn\_a\_00017
- Pennycook, G., Cheyne, J. A., Barr, N., Koehler, D. J., & Fugelsang, J. A. (2013). The role of analytic thinking in moral judgements and values. *Thinking & Reasoning*, *0*(0), 1–27. doi:10.1080/13546783.2013.865000
- Pizarro, D., Inbar, Y., & Helion, C. (2011). On Disgust and Moral Judgment. *Emotion Review*, *3*(3), 267–268. doi:10.1177/1754073911402394
- Robertson, D., Snarey, J., Ousley, O., Harenski, K., DuBois Bowman, F., Gilkey, R., & Kilts, C. (2007). The neural processing of moral sensitivity to issues of justice and care. *Neuropsychologia*, *45*(4), 755–66. doi:10.1016/j.neuropsychologia.2006.08.014
- Sagi, E., & Dehghani, M. (2013). Measuring moral rhetoric in text. *Social Science Computer Review*. doi:10.1177/0894439313506837

- Schaich Borg, J., Lieberman, D., & Kiehl, K. A. (2008). Infection, incest, and iniquity: Investigating the neural correlates of disgust and morality. *Journal of Cognitive Neuroscience*, *20*(9), 1529–46. doi:10.1162/jocn.2008.20109
- Schaich Borg, J., Sinnott-Armstrong, W., Calhoun, V. D., & Kiehl, K. A. (2011). Neural basis of moral verdict and moral deliberation. *Social Neuroscience*, *6*(4), 398–413. doi:10.1080/17470919.2011.559363
- Schnall, S., Haidt, J., Clore, G. L., & Jordan, A. H. (2008). Disgust as embodied moral judgment. *Personality & Social Psychology Bulletin*, *34*(8), 1096–109. doi:10.1177/0146167208317771
- Sinnott-Armstrong, W. (2007). Is moral phenomenology unified? *Phenomenology and the Cognitive Sciences*, *7*(1), 85–97. doi:10.1007/s11097-007-9065-z
- Somerville, L. H., & Whalen, P. J. (2006). Prior experience as a stimulus category confound: An example using facial expressions of emotion. *Social Cognitive and Affective Neuroscience*, *1*(3), 271–4. doi:10.1093/scan/nsi040
- Tetlock, P. E., Kristel, O. V., Elson, S. B., Green, M. C., & Lerner, J. S. (2000). The psychology of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social Psychology*, *78*(5), 853–870. doi:10.1037//0022-3514.78.5.853
- Weber, C. R., & Federico, C. M. (2013). Moral Foundations and heterogeneity in ideological preferences. *Political Psychology*, *34*(1), 107–126. doi:10.1111/j.1467-9221.2012.00922.x
- Wheatley, T., & Haidt, J. (2005). Hypnotic disgust makes moral judgments more severe. *Psychological Science*, *16*(10), 780–4. doi:10.1111/j.1467-9280.2005.01614.x
- Wiltermuth, S. S., Monin, B., & Chow, R. M. (2010). The orthogonality of praise and condemnation in moral judgment. *Social Psychological and Personality Science*, *1*(4), 302–310. doi:10.1177/1948550610363162
- Wright, J. C., & Baril, G. (2011). The role of cognitive resources in determining our moral intuitions: Are we all liberals at heart? *Journal of Experimental Social Psychology*, *47*(5), 1007–1012. doi:10.1016/j.jesp.2011.03.014

Figure 1. Foundation-level ratings of moral vignettes.



Note: wrongness represents the average wrongness rating given to a foundation's scenarios, where 0=not wrong at all, 4=extremely wrong. Classification is the percentage of respondents correctly classifying the scenarios. Comprehension is how well respondents could understand the scenario, vividness is how clearly they could form a mental image of the scenario, frequency is how often they see the scenario portrayed and arousal is the strength of their emotional response to the scenario. Segments represent 95% confidence intervals.

If you would like the full details of the vignettes, please contact Scott Clifford.

*Table 3. Confirmatory factor analysis of wrongness ratings*

# of Factors	Description	$\chi^2$	<i>df</i>	$\chi^2/df$	RMSEA	90% C.I.	AIC
2	Individualizing and binding	16,156.29	6,553	2.47	.059	.058, .061	129,997.09
6	Original MFT, plus Liberty	13,607.23	6,539	2.08	.051	.050, .052	127,476.03
7	Division of physical and emotional harm	12,933.23	6,533	1.98	.049	.047, .050	126,814.03
7	Division of human and animal harm	12,893.93	6,533	1.97	.048	.047, .050	126,774.73
<b>8</b>	<b>Division of emotional, physical, animal harm</b>	<b>12,616.71</b>	<b>6,526</b>	1.93	<b>.047</b>	<b>.046,</b> <b>.049</b>	<b>126,511.51</b>



Table 4. Predicting MFV with existing moral foundations scales.

	Care (e)	Care (p)	Fairness	Liberty	Loyalty	Authority	Sanctity
<b>MFQ</b>							
Care	<b>0.34</b> *** (.06)	<b>0.47</b> *** (.06)	0.08 (.06)	0.15 * (.06)	0.06 (.06)	-0.06 (.06)	0.14 * (.06)
Fairness	0.21 ** (.06)	<b>0.23</b> *** (.06)	<b>0.34</b> *** (.06)	<b>0.39</b> *** (.06)	0.09 (.06)	0.18 * (.06)	-0.03 (.06)
Loyalty	0.15 * (.06)	-0.05 (.06)	0.00 (.07)	0.09 (.07)	<b>0.41</b> *** (.06)	0.22 *** (.06)	0.16 ** (.06)
Authority	-0.03 (.07)	-0.04 (.07)	0.05 (.07)	-0.07 (.07)	0.10 (.07)	<b>0.24</b> *** (.07)	0.10 (.06)
Sanctity	0.05 (.06)	-0.10 (.06)	0.06 (.07)	-0.10 (.06)	0.05 (.06)	0.14 * (.06)	<b>0.37</b> *** (.06)
Constant	-0.01 (.04)	-0.04 (.04)	-0.01 (.04)	-0.01 (.04)	0.03 (.04)	0.03 (.04)	0.03 (.04)
$R^2$	0.30	0.30	0.17	0.20	0.34	0.33	0.38
<b>MFSS</b>							
Care	<b>0.17</b> * (.09)	<b>0.57</b> *** (.08)	0.15 (.08)	<b>0.22</b> * (.09)	-0.22 ** (.09)	-0.25 ** (.08)	0.06 (.08)
Fairness	0.10 (.09)	-0.02 (.08)	<b>0.42</b> *** (.08)	0.06 (.09)	-0.11 (.09)	0.13 (.08)	-0.04 (.08)
Loyalty	0.04 (.08)	-0.08 (.08)	0.05 (.08)	0.04 (.08)	<b>0.35</b> *** (.08)	0.12 (.08)	0.11 (.08)
Authority	0.09 (.07)	-0.08 (.07)	-0.19 ** (.07)	-0.08 (.07)	0.03 (.07)	<b>0.20</b> ** (.07)	0.04 (.07)
Sanctity	-0.02 (.08)	-0.05 (.07)	0.07 (.04)	-0.03 (.08)	0.14 (.08)	0.15 * (.07)	<b>0.26</b> *** (.07)
Constant	-0.04 (.05)	-0.08 (.05)	-0.07 (.04)	-0.04 (.05)	0.00 (.05)	0.00 (.05)	-0.04 (.04)
$R^2$	0.08	0.13	0.15	0.03	0.09	0.14	0.12

Note: cell entries show OLS coefficients for MFQ (top) and MFSS (bottom) foundations predicting MFV scales. All scales are standardized. The largest coefficient in each row is bolded. Note that there are no existing criterion variables for the Liberty foundation. Care (e) refers to the emotional harm factor, while Care (p) refers to the physical harm factor.

\*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ , two-tailed. Standard errors in parentheses.

*Table 5. Correlations between the moral foundations and political ideology*

	MFQ	MFSS	MFV
Care	-0.16 **	-0.02	
Care (e)			-0.05
Care (p)			-0.20 ***
Fairness	-0.25 ***	0.04	-0.01
Loyalty	0.22 ***	0.12 *	0.12 *
Authority	0.20 ***	0.08	0.11 *
Sanctity	0.21 ***	0.11 *	0.23 ***
Liberty			-0.15 **

*Note:* cell entries are correlations between each foundation scale and political ideology (where higher values represent greater conservatism. Care (e) refers to emotional harm vignettes while Care (p) refers to vignettes involving physical harm to humans and animals.

\*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ , two-tailed.

Table A1. Predicting MFSS with MFQ

	MFSS				
	Care	Fairness	Loyalty	Authority	Sanctity
MFQ					
Care	<b>0.26</b> ***	0.03	0.09	0.13 *	0.02
	(.05)	(.06)	(.06)	(.06)	(.06)
Fairness	0.03	0.05	-0.08	-0.17 *	0.05
	(.05)	(.06)	(.06)	(.07)	(.06)
Loyalty	-0.05	0.07	<b>0.22</b> **	0.12	0.07
	(.06)	(.07)	(.06)	(.07)	(.07)
Authority	0.01	-0.02	-0.02	0.09	0.03
	(.06)	(.07)	(.07)	(.07)	(.07)
Sanctity	0.02	<b>0.13</b> *	0.11	<b>0.14</b> *	<b>0.19</b> **
	(.05)	(.07)	(.06)	(.07)	(.06)
Constant	0.14 ***	0.11 *	0.13 **	0.04	0.10 *
	(.04)	(.04)	(.04)	(.04)	(.04)
$R^2$	0.11	0.05	0.11	0.12	0.10

Note: cell entries show OLS coefficients for MFQ foundations predicting MFSS. The largest coefficient in each row is bolded. All measures are standardized.

\*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ , two-tailed. Standard errors in parentheses.

*Table A2. Correlations between moral foundations scales*

	MFV							MFQ					MFSS				
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1 MFV - Care (e)	-																
2 MFV - Care (p)	0.61	-															
3 MFV - Fairness	0.54	0.57	-														
4 MFV - Liberty	0.62	0.63	0.57	-													
5 MFV - Loyalty	0.57	0.35	0.43	0.46	-												
6 MFV - Authority	0.63	0.35	0.54	0.45	0.61	-											
7 MFV - Sanctity	0.49	0.40	0.43	0.28	0.54	0.50	-										
8 MFQ - Care	0.50	0.51	0.32	0.34	0.32	0.28	0.36	-									
9 MFQ - Fairness	0.45	0.44	0.40	0.43	0.31	0.33	0.27	0.65	-								
10 MFQ - Loyalty	0.34	0.10	0.21	0.15	0.56	0.50	0.50	0.37	0.35	-							
11 MFQ - Authority	0.30	0.12	0.24	0.11	0.47	0.51	0.50	0.39	0.38	0.70	-						
12 MFQ - Sanctity	0.31	0.10	0.23	0.08	0.43	0.47	0.58	0.41	0.32	0.65	0.70	-					
13 MFSS - Care	0.25	0.33	0.29	0.18	0.04	0.12	0.26	0.33	0.23	0.09	0.12	0.13	-				
14 MFSS - Fairness	0.25	0.14	0.36	0.12	0.10	0.29	0.24	0.14	0.14	0.18	0.16	0.21	0.64	-			
15 MFSS - Loyalty	0.23	0.10	0.25	0.11	0.24	0.29	0.28	0.17	0.10	0.31	0.24	0.27	0.59	0.70	-		
16 MFSS - Authority	0.23	0.05	0.15	0.05	0.15	0.33	0.24	0.17	0.05	0.28	0.28	0.30	0.47	0.68	0.65	-	
17 MFSS - Sanctity	0.20	0.14	0.22	0.10	0.17	0.27	0.34	0.17	0.17	0.25	0.25	0.30	0.64	0.63	0.63	0.58	-