Beyond good and bad: Varieties of moral judgment

William Jiménez-Leal¹, Samuel Murray², Santiago Amaya¹, Sergio Barbosa,^{1,3}

¹ Universidad de los Andes

² Duke University

³ Universidad del Rosario

Word count: ~13600 words

Author note

The first three authors contributed equally. Correspondence concerning this article should be addressed to William Jiménez-Leal, Universidad de los Andes. E-mail: w.jimenezleal@uniandes.edu.co. We thank Juan Pablo Bermudez for helpful discussions of the design of Experiments 1 and 2 and Joan Urriago for help coding analyses for Experiment 6

Abstract

We argue that people regularly encounter situations characterized by the presence of moral conflicts among permissible options. These scenarios, which some have called *morally charged situations*, reflect perceived tensions between moral expectations and moral rights. Studying responses to such situations marks a departure from the common emphasis on sacrificial dilemmas and widespread use of single-dimension measures. In 6 experiments (n=1607), we show that people use a wide conceptual arsenal when assessing actions that can be described as suberogatory (bad but permissible) or supererogatory (good but not required). In Experiment 1 we find that people freely describe actions as suberogatory or supererogatory. Experiment 2 shows that they differentially assess these actions in terms of how permissible, optional, and good they considered them. Experiment 3 tests the use of these evaluative dimensions with sacrificial dilemmas. We found that differences between these categories did not emerge when people respond to dilemmas, even when controlling for trait utilitarian tendencies. By including judgments of blameworthiness and sanction, Experiments 4 and 5 provided additional evidence of the richness sub/super erogatory evaluations. In Experiment 6 people offered their own explanations of their responses. Qualitative analyses revealed that they frequently appeal to character traits, the presence of rights, and the absence of explicit duties. Taken together these results suggest a richer spectrum of both situations and concepts relevant to characterize moral judgment than moral psychologists up to this point have generally recognized.

Keywords: moral judgment, moral dilemmas, erogation, supererogatory, duty, norms

Beyond good and bad: Varieties of moral judgment

The role of dilemmas in the science of morality

Dilemmas have traditionally played a prominent role in the study of moral cognition. While many have recently registered dissatisfaction with this state of affairs (Andrade, 2019; Bauman, McGraw, Bartels & Warren 2014; Dahl & Oftedal, 2019; Everett & Kahane, 2020; Kahane, 2020), there is considerably less said about how we got here. Dilemmas take on this crucial role because of the assumption that morality is a system of rules. But this assumption is problematic for reasons pertaining to both normative theories of morality and the psychology of moral judgment.

From its early days, psychologists interested in moral judgment understood morality as a structured set of rules prohibiting and prescribing certain behaviors. Piaget (1932), for instance, opened his landmark *The Moral Judgment of the Child* with the statement: “All morality consists in a system of rules, and the essence of all morality is to be sought for in the respect which the individual acquires for these rules” (1932, p. 1). Fauconnet echoes this sentiment in his *Responsibility*, stating that moral responsibility is a “quality belonging to those who must...in virtue of a rule be chosen as the passive subjects of a punishment” (1920, p. 11).

Thinking of morality as a system of rules has several theoretical pay-offs. For one, various deontic concepts can be inter-defined through rules. For example, acts are permissible if and only if those acts accord with the rules, and something counts as good only if it accords with the rules. Likewise, the impermissible is whatever constitutes violating a rule, and something is bad insofar as doing it violates a rule (see, for example, Kanger (1971) and Anderson (1958) for two formal attempts to explain inter-definability). This, as we shall see, results in a methodological

56 advantage. If true, asking whether something is impermissible, bad, or in violation of a duty
57 comes close to asking one and the same question.

58 Focusing on rules also provides a concrete way to measure moral development. Under the
59 assumption that maturing moral judgment consists in possessing a greater moral understanding, it
60 is possible to define this understanding in terms of an increasing ability to apply more
61 sophisticated evaluative rules (Kohlberg & Hersh, 1977). This insight famously informed the
62 stage theory of moral development (Kohlberg, Levine, & Hewer, 1983), which identified moral
63 stages according to the kinds of rules that informed judgments: from pre-conventional rules
64 (“Doing this is bad because I can get punished”) to universal, exceptionless principles of
65 impartial justice (“Killing is wrong”).

66 Moral dilemmas characteristically present actors with situations where they have to decide
67 among options that are impermissible, in the sense that each violates some plausible moral
68 imperative (Sinnott-Armstrong, 1988: 29-30). In this respect, they help isolate different sets of
69 rules (egocentric v. altruistic, in-group v. out-group, etc.) that are normally taken to inform moral
70 cognition. Thus, by looking at the choices people make in these situations or by studying how
71 they assess the decisions made by actors depicted in them, it seems possible to better understand
72 which rules inform their judgments of goodness and badness. For similar reasons, moral
73 dilemmas also seem to provide a good instrument to measure moral development and test
74 hypotheses about individual developmental trajectories.

75 Obviously, theories of moral cognition that fall under this rule-based tradition differ from
76 one another in important respects (see Darley & Schultz 1990 for discussion). A sign of this are
77 the well-known controversies that exist among them. Kohlberg’s research program, for example,

was heavily criticized for the use of biased samples (Gilligan, 1977) and the idea of a linear progression in moral development (Rest, 1979). There were some who also criticized the kind of dilemmas used by researchers working on this tradition, focusing on the artificial nature or the mundane character of them (Rosen, 1980; Bauman et al., 2014). At bottom, however, many of Kohlberg's critics agreed that morality was a system of rules that could be studied by means of situations where plausible moral rules would conflict with one another. They just disagreed about how to properly characterize these rules and the situations that best exemplified these conflicts.¹

Rules and Commonsense Morality

Many contemporary psychologists have sought to take distance from this early tradition by proposing dual models of moral cognition. Their proposals have been premised on the belief that commonsense morality reflects attitudes that lie on a continuum between full-blown utilitarian and deontological ethics (Bartels & Pizarro, 2011; Christensen & Gomila, 2012; Conway & Gawronski, 2013; Djeriouat & Trémolière, 2014; Holyoak & Powell, 2016; Lee & Gino, 2015). We believe, however, that these models are in an important respect a *continuation of the rule-based tradition* that has dominated the study of moral cognition. Let us explain.

Dual models are typically predicated on the idea that different systems encode different styles of moral cognition: reason-based systems produce characteristically utilitarian responses; emotional/intuitive systems produce characteristically deontological responses (Bialek & De Neys, 2017; Cushman, Young, & Greene, 2010; Moll & de Oliveira-Souza, 2007). This has led

¹ One notable exception comes from feminist critiques of ethical theory. Several feminist ethicists have noted that the central preoccupation of moral theorizing prioritizes abstract principles over the particularity of ethical life (Gilligan, 1982, p. 32-38). The emphasis on abstract generalization obscures the details that are crucially important to a well-lived life (Young, 1987, p. 61-62).

some researchers to question whether rules need to be explicitly represented in the process of forming moral judgments. On their view, when emotion and intuition drive moral judgment, the processes by which this happens are better characterized as non-symbolic and computed non-sequentially, in a way that is antithetical to rule-based processing (see Blair, 1995; Greene et al., 2001; Haidt, 2001; Mikhail, 2011).

There are considerable grounds for debating to what extent these processes are genuinely “rule-free.” One possibility, for example, is that emotion and intuition are actually guided by the unconscious representation of moral rules (for discussion Mallon & Nichols, 2010). Note, however, that independently of this issue, the picture of morality endorsed here is still one dominated by rules, whether these are utilitarian or deontological. In brief, the phenomena that these computational models are meant to capture, namely, the styles of thinking that dominate commonsense morality, are still conceptualized in terms of moral rules.

In general, rules can be defined as functions from relevant inputs to behavioral imperatives. For example, a rule against murder is a general function from some action’s being an instance of murder to an imperative against so acting. As Sidgwick (1981, p. 228) explains, “...rules of duty ought to admit of precise definition in a universal form.” Under this definition, many generalizations (but not all of them) count as moral rules: the prohibition not to harm innocents, the prescription to maximize some valued quantity (saved lives), etc.

Thus, as we see it, it might be an open question whether the processes that result in moral judgments (reason-based vs. intuitive/emotional) take as inputs explicit representations of these transitions from behavioral inputs to imperativial outputs. Be that as it may, the fact is that researchers working under this new paradigm still classify the *outputs* of these processes by

measuring conformity to the prescriptions of rule-based moral systems. The styles of thinking modeled by them, which are supposedly characteristic of commonsense morality, are still styles of thinking in *accordance with some distinctive moral rules*.

It is not surprising, then, that contemporary work on the psychology of moral judgment continues to be dominated by moral dilemmas, in particular sacrificial dilemmas. Observing people's choices or their evaluation of the available options when each choice or option is made impermissible by deontological or utilitarian rules seems a natural way to measure people's attitudes towards each of these moral frameworks. Whatever the underlying computational processes are, dilemmas seem to be recommended by the goal of understanding whether commonsense morality embodies deontological or utilitarian rules.

In sum, there is a long tradition, spanning from Piaget to contemporary dual models, that views morality as a rule-based system. Dilemmas appear useful for studying moral judgment in virtue of this underlying assumption. Hence, moving beyond dilemmas requires moving beyond the assumption that morality is a rule-based system. And moving beyond the assumption of morality as a rule-based system requires moving beyond the study of moral judgment in situations where rules conflict.

There are, as we shall see, numerous dimensions of the moral life that are not governed by rules. Hence, a science of morality that focuses only on dilemmas and rule-based judgments risks painting a picture of moral cognition that is overly narrow and stilted (Bauman et al., 2014). We want to challenge the use of dilemmas to study moral judgment because we reject the assumption that morality consists in a structured system of rules. While rules constitute part of the moral domain, there are various aspects of the moral life that are not governed by them. Seeing how

moral cognition operates beyond the rules provides a fresh perspective for the study of moral judgment. We examine moral judgment in cases of suberogatory behavior. To explain our approach, we begin by describing the nature of suberogatory behavior in more detail below.

The Supererogatory and the Suberogatory

From time to time, people are faced with the option of doing more than what they are required to do, for example, spending some of their free time volunteering at the local animal shelter or donating a large portion of money to charity. These actions are admirable, despite the fact that failing to do these things does not merit condemnation. Some classify these actions as ‘supererogatory’, going above and beyond the call of duty when failing to go above and beyond is perfectly morally permissible (Archer, 2018).

More controversially, people seem at times to under perform relative to some norm in a way that is permissible, for instance, tipping just a few dollars after having a pleasant meal that cost several hundred dollars. Doing this tends to be regarded as bad despite the fact that there is no rule that requires tipping above a certain exact amount or percentage. Some use the label ‘suberogatory’ to describe this kind of behavior (see Driver, 1992; Hurd, 1998). The behavior is morally objectionable but there is no well-defined duty that it violates.

In failing to do a supererogatory action, one usually does not do anything bad. It’s an admirable thing to donate money to charity, but failing to do so is not reprehensible. However, in some situations, failing to do a supererogatory action constitutes suberogatory behavior. If a tourist asks you for directions, you are completely within your rights to walk away without saying anything. Doing it, though permissible, is bad, whereas giving directions is good despite not being required. This possibility suggests a different kind of conflict that people encounter in

their day-to-day experiences of morality: conflicts between equally permissible good and bad options. Here, we refer to these as *morally charged situations* (Driver, 1992). We claim that these situations, along with the concepts used to evaluate them, offer a distinctive opportunity to study moral judgment.

Sacrificial dilemmas represent an interesting albeit limited subset of what people encounter in their everyday life. Using an ecological momentary assessment, Hofmann, Wisneski, Brandt and Skitka (2014) asked a large sample of people to report whether they had witnessed, committed, or heard about a moral situation during the last hour, five times a day for three days. While some situations resemble the kind of conflict expressed in sacrificial dilemmas (e.g., “Reminded waitress I did not pay for my bill when she thought I did”), many of the situations reported seemed more similar to morally charged situations. The choice between the competing options didn’t seem a matter of aligning oneself with some well-defined rule.

Over-relying on dilemmas risks papering over these distinctions and, more generally, the differences between moral categories that inform varieties of moral judgment. O’Hara, Sinnott-Armstrong, and Sinnott-Armstrong (2010) compared responses to 15 dilemmas where they asked people to rate how wrong, forbidden, inappropriate and blameworthy an action was. They found that “the influence of wording variations on moral judgments was negligible” (p. 552) and they analyzed the small differences found as a matter of magnitude. Likewise, many researchers treat terms like ‘forbidden’ or ‘blameworthy’ as linguistic variations of some homogenous moral judgment (Bjorklund, 2003; Cushman, Young, & Hauser, 2006; Greene et al., 2001b; Koenigs et al., 2012). The assumption is that morally evaluating a situation results in a monolithic type of cognitive product: an attribution of overall rightness or wrongness to some behavior.

While there have been calls to more carefully use these measures (Christensen & Gomila, 2012; Monin et al., 2007) it not always clear what the rationale for using different terms could be (but see Cushman, 2008; Barbosa & Jiménez-Leal, 2017). But, once the repertoire of moral conflicts is expanded to include morally charged situations, it is possible that nuances and variability in moral cognition will emerge. In sum, by expanding the kind of moral encounters used when empirically probing people's intuitions and by enlarging the dimensions along which these encounters are conceived we can more adequately study the variation of human moral judgment.

Here we present 6 experiments that study moral judgments in morally charged situations. One key feature of these situations is that the correct choice—if there is one—is not obviously settled by appealing to rules. Experiment 1 is an exploratory study that maps out the descriptions people offer of different situations. We find that people freely describe situations as suberogatory ('bad but permissible') or supererogatory ('good but not required'). Experiment 2 shows that judgments of good/bad, permissible/impermissible, and optional/obligatory dissociate when evaluating suberogatory and supererogatory situations. We also find that people's beliefs about duties negligibly correlate with judgments of goodness, permissibility, obligation, and blame. In Experiment 3, we compare judgments in sacrificial dilemmas to see whether the same distinctions appear. We do not find the same dissociations, which suggests that eliciting these patterns of judgment requires more than giving participants the options to judge along various dimensions. In Experiment 4, we included a measure of praise/blame to see whether it correlates with the main "erogatory" measures. In Experiment 5, we utilized different vignettes that describe more characteristically moral situations (adapted from classic philosophical thought experiments about abortion and property rights). We also measured whether beliefs about rights

predict any kind of judgment. We find the same pattern of dissociations in these different vignettes and again find a negligible correlation between judgments and individual beliefs about rights. In Experiment 6, we conducted a qualitative study to begin exploring the variety of factors that differentially drive different judgments. Using independent coders for a qualitative analysis, we find that people describe these situations using the language of character traits and rights rather than duties. In fact, in many cases people explicitly state the *absence* of duties to do anything in our scenarios

A methodological coda: Even though this research is mostly exploratory, its ideas are developed against a backdrop of well established findings in moral psychology. We decided to preregister Experiments 2 to 6 because we believe that clearly setting up a design and analysis plan can help distinguish the confirmatory and exploratory aspects of our research by clearly specifying our intent and thus reducing needless post hoc interpretations (Nosek et al, 2019).² Materials, data, preregistrations and code for all experiments are available in the OSF page of the project (<https://osf.io/cr8be/>).

² The only important deviation from preregistration plans occurred in Experiment 2, where the main statistical analysis proposed (a repeated measures ANOVA) was replaced by a mixed linear model, since it is better suited to model our data.

Experiment 1

The objective of this first study is exploratory. We presented participants with vignettes describing either suberogatory or supererogatory behavior. Participants were instructed to select words from a list to describe these scenarios and to offer a description in their own words. The goal of this is twofold. The first is to see what language people use to spontaneously describe a morally charged situation. The second is to see whether people recognize a distinction between different judgment categories that maps onto the complex category of suberogatory. This requires that people have distinct concepts of permissibility and goodness such that they can describe some behavior as bad but permissible. Hence, we decided to run an exploratory study using word selection and open response to see whether participants utilize the moral categories we aim to study without being prompted.

We expected people to always select always more than one word (e.g., “good”) and to give descriptions that characterize both the action and the person.

Method

Participants

95 participants (60 women and 35 men, mean age = 32.42, SD = 10.43), based in the United States and recruited through Prolific Academic, took part in the study in exchange for 40 pence. Participants were aware that their answers would be anonymous and were monetarily compensated for their participation. The average completion time was 5.2 minutes.

Materials and procedure

We constructed four scenarios, some of which were based on thought experiments by Driver (1992). Each scenario described an individual faced with a choice between a suberogatory

and a supererogatory option. Additionally, in order to account for possible asymmetries between actions and omissions, we created two versions of each scenario that describe either an action or an omission. This generated eight vignettes, described below (suberogatory versions of these vignettes are in brackets):

Newlyweds Two newlyweds are boarding a plane to go on their honeymoon. Because of a booking error by the airline, the couple does not have seats together. They ask someone, already seated, if they would switch seats so the couple could sit together. The passenger switches seats, and the newlyweds can sit together. [*The passenger does not switch seats, and the newlyweds have to sit separately.*]

Kidney Alex is suffering from severe kidney failure and Alex's only hope is to obtain a transplanted kidney. Alex's cousin, Jamie, is the only known compatible donor. Jamie offers to donate the kidney to Alex. [*Jamie does not offer to donate the kidney to Alex.*]

Mowing Early one Sunday morning when the neighbors are usually sleeping, Sam notices that the lawn needs to be mowed. Although it is his property and it would be inconvenient to do it later, he decides to not mow the lawn. He knows that starting the lawn mower will probably wake up the neighbors. [*Even though he knows that starting the lawn mower will probably wake up the neighbors, he does it anyway. It's his property and it will be inconvenient to mow the lawn later.*]

Raffle During the Christmas party, the secretary publicly announced the results of the office raffle: "Congratulations to Alex, who has won the trip for two to Disney World. She can come up front to claim her prize or she can let a cash equivalent go to a hurricane relief fund". After hearing the news, Alex looked excited: ". Even though I have the winning ticket and Disney World sounds fun, I am going to donate the prize to one of the charities." [*After hearing the news, Alex looked excited: "I have the winning ticket! Even though I don't really care much about Disney World, I am going to claim the prize anyway".*]

Participants were presented with four out of the eight possible variations, so each participant was presented with a suberogatory omission, a suberogatory action, a supererogatory

omission, and a supererogatory action. After reading the vignettes, participants completed a word selection task by selecting “the word(s) that you think best describe the situation.” The word choices were: “permissible”, “impermissible”, “required”, “good”, “obligatory”, “allowed”, “bad”, “optional” and “compulsory”. Participants also completed an open description task by offering a description of the situation in their own words. The vignette presentation order was counterbalanced and the tasks and words order was randomized.

Results

For the word selection task, participants selected 2.6 words on average (see Figure 1). For suberogatory behaviors, people most often chose the words “optional” (27%) and “allowed” (25%), followed by “permissible” (21%). For supererogatory behaviors, “good” (35%), “optional” (28%), and “allowed” (14%) were the most common choices. “Allowed” and “optional” are the most common pair of words used across all vignettes. “Allowed” and “permissible” are more strongly associated with evaluating suberogatory behaviors. Figure 1 summarizes these results.

283

284 *Figure 1. Mosaic plot for words selection.* Height is associated with selection frequency and

285 width is associated with prevalence (Hornik, Zeileis & Meyer, 2006). Darker shades and solid

286 lines represent positive associations, whereas lighter shades and dotted lines represent a negative

287 association within the sub/supererogatory categories. The plot represents a model contrasting

288 observed and expected frequencies.

290 In the open descriptions, people predominantly used character trait descriptors, like *selfish*
291 for the suberogatory vignettes and *thoughtful* for the supererogatory situation (see Figure 2).
292 People rarely used words like “permissible”. However, some of the more nuanced descriptions of
293 the suberogatory situations pointed out that “he is allowed to do it” (mowing), “Jamie has a right
294 to her own decision” (Kidney), “it is understandable” (kidney), comments that generally
295 highlighted the optionality of the suberogatory response. Other participants made similar
296 comments:

- 297 • “It is a little bit selfish, but then again she has the right to keep her organs”
- 298 • “While it was mean, the passenger has the right to refuse
- 299 • “Rude, but it is his lawn”
- 300 • “Giving an organ is a big thing to ask. It is something that is optional, and there is no
- 301 mention of Alex asking him to do it so he is not required to offer”

302 People were generally sensitive to the different aspects implied by the situations. Participants
303 were quick to describe the action/person negatively (mean, rude, etc) but they also recognized
304 how rights and expectations are at play in the decisions of the protagonists.

Figure 2. Wordcloud for suberogatory (left) and supererogatory (right) situations

While these results are exploratory, they provide evidence that moral thinking about a situation is complex. People recognize that bad behavior can be permissible, for example. We also find that thoughts about rights and character seem to be ingredients in moral judgment. The acceptance of the categories of suberogatory and supererogatory made possible the use of the measures and procedures described in the subsequent experiments.

Experiment 2

In Experiment 1, we found some evidence that people describe certain moral situations in ways that are sensitive to the categories of the suberogatory and the supererogatory. This suggests that people distinguish between permissibility, obligation, and goodness. Here we test for quantifiable differences between these evaluative categories.

We asked people to rate the sub– and supererogatory situations along three dimensions: good/bad, permissible/impermissible, and obligatory/optional. We manipulated the Erogation

category (within-subjects, sub- and supererogatory) and the Situation Type (between-subjects, action or omission). We hypothesized that people would judge suberogatory situations as worse than supererogatory situations, but that ratings of permissibility and obligatoriness between the two situations would not be significantly different. We expected ratings of permissibility to be higher than ratings of obligatoriness in both conditions, though the difference between permissibility and obligation would be greater in the suberogatory condition than in the supererogatory condition. We did not expect any differences between judging actions and omissions. We also collected data on attitudes towards duties, expecting that ratings along the Erogration category would be associated with these attitudes.

Participants

A power analysis for a mixed ANOVA suggested a sample size of 272. To account for exclusions, we recruited 311 participants (186 women and 125 men, mean age = 32.77, SD =11.18) through Prolific Academic. Each person voluntarily participated in the study and received 38 pence as compensation. The average completion time was 3.11 minutes.

Materials and procedure

Materials were the same as Experiment 1. Each participant saw four of eight scenarios. We manipulated the moral category within participants and action/omission between participants, so each participant saw two supererogatory situations and two suberogatory situations, where all of them were either actions or omissions. For each vignette, participants evaluated the situation along different dimensions with a 100-point sliding scale. The dimensions included degree of permissibility (impermissible = 0, neither impermissible nor permissible = 50, permissible =100), degree of goodness (bad = 0, neither good nor bad = 50, good = 100), and degree of obligation (optional=0, neither optional nor obligatory = 50, obligatory = 100). The dimension order was

randomized across trials. The scale appeared only with the anchors, the slider was always placed in the center of the scale, and participants were not given a numerical representation of where they placed the slider.

Participants viewed one vignette at a time and the sliders were placed on the same page as the vignette. After completing the study, participants indicated their age, gender, and political orientation (on a five-point Likert scale from *very liberal* to *very conservative*). We also collected information about personal sense of duty. Participants indicated agreement with a 7-point Likert scale (1 = strongly disagree, 7 = strongly agree) with the idea that there are duties to respect your neighbors, to help anybody who needs help, and to help one's family members.

Results

Pre-registered analyses were integrated into a set of linear random effect models fitted with the lme4 R library (Bates, Mächler, Bolker, & Walker, 2015; R Core Team, 2019), with participant id as a random intercept.³ Results are summarized in Table 1 and Figures 3 and 4. We fitted four nested models to contrast the interactive effect of the manipulated variables (see models 1 and 2) and to test the effect of differences between scenarios and endorsement of norms (models 3 and 4). Goodness of fit indicators (AIC, BIC and deviance) favor selection of model 2. Significant interactions between judgment type and condition suggest differences both between and within the erogation category. Supererogatory behaviors were judged as better and more

³ It can be argued that the data of this experiment could be considered a cross-classified data set. However, the items in each situation type are not completely equivalent, which makes the corresponding items nested within situation type. Additional models were fitted with additional random term but since results are equivalent, we restrict the presentation here to the different conditions as fixed effects.

360 permissible ($M = 86.4$, $SD = 18.7$, $n = 622$ and $M = 82.0$, $SD = 20.6$, $n = 622$) than
361 suberogatory behaviors ($M = 36.6$ $SD = 23.3$, $n = 622$ $M = 62.2$ $SD = 29.7$ $n = 622$) Welch's t -
362 test Good: $t(1105) = -41.6$, $p = .001$, $M_{diff} = -2.4$, 95% CI [-2.2, -2.5] and Welch's t -test
363 Permissible: $t(1067) = -13.7$, $p = .001$, $M_{diff} = -0.78$, 95% CI [-0.66, -0.89], regardless of whether
364 they were actions or omissions. Supererogatory behaviors were judged to be marginally more
365 obligatory (i.e., less optional) ($M = 30.9$, $SD = 32.7$, $n = 622$) than the corresponding
366 suberogatory responses ($M = 23.5$, $SD = 26.3$, $n = 622$) Welch's t -test: $t(1188) = -4.45$, $p = .001$,
367 $M_{diff} = -0.25$, 95% CI [-0.14, -0.36] though it is clearly a smaller effect. This is presumably
368 because there are no explicit duties in these situations.

369 Interestingly, the size of the differences between permissibility and goodness ratings are
370 vastly different when looking at sub and supererogatory responses. Within the suberogatory
371 responses, this difference amounts to 25.6 points (Cohen's $d = 0.96$) while for the supererogatory
372 category, this difference is only 4.46 points (Cohen's $d = 0.22$). That is, goodness and
373 permissibility judgements are very similar within the supererogatory condition, but not for the
374 suberogatory responses, where they are more clearly tracking different aspects of the situation.

375

	score			
	Model 1	Model 2	Model 3	Model 4
Omission	-2.74*** (-4.65, -0.83)	-6.411*** (-10.500, -2.322)	-7.33*** (-10.41, -4.24)	-7.26*** (-10.36, -4.17)
OBLIGATORY	-34.30*** (-36.41, -32.20)	-17.449*** (-21.389, -13.509)	-38.12*** (-41.07, -35.17)	-38.12*** (-41.07, -35.17)
PERMISSIBLE	10.57*** (8.47, 12.67)	21.758*** (17.818, 25.698)	7.58*** (4.63, 10.53)	7.57*** (4.62, 10.52)
SUPER EROGATORY	25.70*** (23.98, 27.41)	50.771*** (46.830, 54.711)	25.70*** (23.99, 27.41)	25.70*** (23.98, 27.41)
Mowing	-7.48*** (-9.90, -5.05)		-7.48*** (-9.89, -5.06)	-7.47*** (-9.89, -5.05)
Newlyweds	-3.97*** (-6.39, -1.54)		-3.97*** (-6.39, -1.55)	-3.96*** (-6.39, -1.54)
Raffle	-0.15 (-2.58, 2.26)		-0.16 (-2.58, 2.26)	-0.15 (-2.57, 2.26)
DUTY TO FAMILY				-0.21 (-0.99, 0.56)
DUTY TO HELP				-0.01 (-0.79, 0.77)
DUTY TO NEIGHBORS				-0.31 (-1.35, 0.71)
Omission: OBLIGATORY		8.75*** (3.15, 14.35)	7.72*** (3.53, 11.91)	7.718*** (3.526, 11.91)
Omission: PERMISSIBLE		7.75*** (2.15, 13.35)	6.04*** (1.85, 10.23)	6.042*** (1.85, 10.23)
Omission: SUPER EROGATORY		-1.84 (-7.44, 3.76)		
OBLIGATORY: SUPER EROGATORY		-41.34*** (-46.92, -35.77)		

PERMISSIBLE: SUPER EROGATORY		-28.36***		
		(-33.93, -22.78)		
Omission: OBLIGATORY: SUPER EROGATORY		-2.07		
		(-9.99, 5.84)		
Omission: PERMISSIBLE: SUPER EROGATORY		-3.419		
		(-11.34, 4.50)		
Constant	52.90*** (50.41, 55.40)	39.74*** (36.86, 42.62)	55.17*** (52.41, 57.94)	58.14*** (51.41, 64.87)
N	3732	3732	3732	3732
Log Likelihood	-17580.00	-17367.00	-17570.00	-17569.00
AIC	35180.00	34763.00	35163.00	35168.00
BIC	35242.00	34850.00	35238.00	35261.00

***p < .01; **p < .05; *p < .1

377 Our results do, however, show that people strongly distinguish between goodness,
378 permissibility, and obligation. People rated supererogatory behaviors as good, optional, and
379 permissible. People rated suberogatory behaviors as bad, optional, and permissible. Despite both
380 being rated permissible, supererogatory behaviors were rated as more permissible than
381 suberogatory behaviors.

Figure 3. Scores by type of judgment. Error bars represent 95% confidence intervals.

We did not find any significant differences of judgment type between action and omission across erogation conditions. That is, even though actions were rated as better ($M = 65.1$, $SD = 32.7$, $n = 628$) than omissions ($M = 57.8$, $SD = 32.3$, $n = 616$) Welch's t -test: $t(1242) = 3.98$, $p = .001$, $M_{diff} = 0.41$, 95% CI [0.30, 0.52], there were no differences between permissibility and obligatory ratings for actions and omissions. Similarly, actions and omissions were rated similarly for all judgment types in both erogation conditions.

We explored the association between ratings for each type of judgment and participant's endorsement of statements about personal duties. We did not find any consistent association between judgment scores and responses pertaining to personal duties (see Table 2). There are significant correlations between endorsing different statements of personal duty (ranging from 0.13 to 0.34) but most correlations between sense of personal duty and different judgment categories were negligible (from -.01 to 0.07) and non-significant. The one exception is that

beliefs about duties to help others significantly correlated with judgments of permissibility, but the correlation is miniscule.

Table 2. Bivariate correlations between personal norms and judgment scores

	Permissible	Good	Obligatory	Duty to neighbors	Duty to help
Permissible					
Good	0.55***				
Obligatory	-0.24**	-0.22***			
Duty to neighbors	-0.03	0.0	-0.02		
Duty to help	-0.07**	-0.02	0.07	0.33***	
Duty to family	-0.03	-0.02	0.01	0.30***	0.23***

*** $p < .01$; ** $p < .05$; * $p < .1$

Breaking down responses by scenario reveals some variability across vignettes (see Figure 4). For example, while donating a kidney to a cousin is judged to be better and more permissible than not donating a kidney, the same pattern does not hold in the raffle scenario. In this case, both situations are equally permissible, but donating the raffle prize is better than not. Notice however that these small differences do not account for significant variance, according to the model fitting presented in Table 1.

Figure 4. Scores by type of judgment and scenario. Error bars represent 95% confidence intervals

Discussion

These results reveal quantifiable differences between different evaluative categories that are employed in moral judgment. Similar to our results in Experiment 1, we found that judgments of permissibility, obligation, and goodness dissociate when people make judgments of suberogatory behavior. Including additional measures allows variability in moral judgment to emerge. Moreover, beliefs about personal duties are negligibly correlated with different kinds of judgments. This suggests that people's responses to these situations are not indicative of an underlying adherence to rules or generalized statements about what one ought to do, even though the duties we asked about specifically applied to the scenarios in the study.

We mentioned in the Introduction that the use of dilemmas to study moral judgment reinforces assumptions about the way to measure moral judgment. When assessing the relationships between different measures, researchers have found negligible differences between

different judgment categories. We see now that including different measures of moral judgment allows for variability, challenging a central assumption about the relationship between different measures of moral judgment. This raises a question: is the lack of variability between judgment categories a function of either the measures used to study moral judgment or the dilemmas commonly used to elicit judgments? It might be the case that some measurement devices function as self-reinforcing demand effects. Alternatively, there might be something about dilemmas themselves that make a difference to moral judgment independently of the measures used. To begin answering these questions, we applied our measures to sacrificial dilemmas in Experiment 3 to see how they align with the results of this experiment.

Experiment 3

In this experiment, we consider different kinds of judgments in cases of sacrificial dilemmas. In our previous experiments, we found dissociations between different kinds of judgments. However, the situations used to test these differences differ significantly from sacrificial dilemmas, the stimuli most consistently used to elicit moral judgments. First, harm is unavoidable in a sacrificial dilemma, so there is no possibility of supererogatory behavior. Second, despite being unrealistic, they elicit distinctive emotional reactions (Christensen et al., 2014). Third, dilemmas are situations where explicit rules come into conflict (e.g., prohibitions against causing harm and prohibitions against allowing easily preventable harm).

We chose sacrificial dilemmas that have been widely used in empirical studies of moral judgment. Based on an analysis from Christensen et al. (2014), we selected dilemmas that produce the largest differences in reactions (varying in the use of personal force and the inevitability of harm). We also excluded dilemmas where various interactive effects might

plausibly drive moral judgments (e.g., when causing harm is self-beneficial). Characteristics and full text of the dilemmas selected are presented in the supplemental materials.

Sacrificial dilemmas are sometimes thought to bring out the contrasts between deontological and utilitarian ethical theories, because each theory usually recommends different behaviors in the face of a sacrificial dilemma. This suggests that people committed to either utilitarianism or deontology might produce different moral judgments in reaction to sacrificial dilemmas. This presents a challenge. Because we are measuring for dissociations among different judgment categories, and different normative theories make different recommendations for navigating dilemmas, it is possible that each side will cancel the other out, thereby giving the appearance of similarity between judgment categories. To ensure that people with utilitarian tendencies do not cancel out people with deontological tendencies, we used the Oxford Utilitarianism Scale (Kahane et al., 2018) to assess trait utilitarianism across two dimensions of utilitarianism, impartial beneficence and instrumental harm.

The character of this study is primarily exploratory but we expected to replicate some of the main findings in this literature and some of the patterns observed by Christensen et al. (2014). We expected utilitarian responses in personal dilemmas to be associated with lower ratings of goodness and permissibility relative to impersonal dilemmas (Greene et al. (2001b)). We also expected behaviors that cause unavoidable deaths to be judged as better than situations where causing death was avoidable (Christensen et al., 2014).

Method

Participants

304 participants (153 women and 151 men, mean age = 33.21, SD =11.62), recruited through Prolific Academic, took part in the study in exchange for 50 pence. The average completion time was of 7.21 minutes and none of the participants had taken part in the previous experiments.

Materials and procedure

Participants saw the four dilemmas and were randomly presented with the version where the main character decides to cause harm (utilitarian response) or allow harm (deontological response). For each scenario, participants judged the main character's response along dimensions of permissibility, obligation, goodness, and blameworthiness using the same sliders as before. After judging the dilemmas, participants were presented with the items of the *Oxford Utilitarianism Scale* (OUS) in a random order. Finally, they are asked some basic demographic questions and were asked to rate their experience with this kind of dilemmas in a scale from 1 (not at all familiar) to 5 (extremely familiar).

Results

Results are summarized in Table 3 and Figures 5 and 6. As before, we fitted a set of random effect models with judgment type, personal/impersonal utilitarian/deontological, avoidable/unavoidable and the score in the OUS as fixed predictors allowing for a random effect of participant. The OUS had good reliability for its both subscales (Instrumental Harm Cronbachs' $\alpha = .65$ and Impartial Beneficence Cronbachs' $\alpha = .72$). The models differed in the specification of the interaction of the fixed effects. The best fitting model is presented in Table 3.

Table 3. Best fitting model for Experiment 3.

	score
BLAME	-5.15** (-9.56, -0.73)
OBLIGATORY	-2.98 (-7.39, 1.43)
PERMISSIBLE	12.78*** (8.37, 17.19)
Personal	-0.37 (-3.47, 2.74)
UNAVOIDABLE	0.15 (-2.95, 3.26)
Utilitarian	-4.65*** (-7.79, -1.52)
BLAME: Personal	4.49** (0.10, 8.89)
OBLIGATORY:Personal	-4.91** (-9.31, -0.52)
PERMISSIBLE:Personal	-4.95** (-9.34, -0.55)
BLAME: UNAVOIDABLE	0.23 (-4.16, 4.63)
OBLIGATORY: UNAVOIDABLE	-1.82 (-6.21, 2.57)
PERMISSIBLE: UNAVOIDABLE	-4.94** (-9.33, -0.54)
BLAME:Utilitarian	6.82*** (2.43, 11.22)
OBLIGATORY:Utilitarian	1.50 (-2.89, 5.89)
PERMISSIBLE:Utilitarian	-2.37 (-6.77, 2.02)
Constant	52.79*** (49.64, 55.95)
N	4864
Log Likelihood	-23061.00

AIC	46158.00
BIC	46275.00

*** $p < .01$; ** $p < .05$; * $p < .1$

Some of the participants (18%) reported having experience with sacrificial dilemmas but their responses do not differ from participants without experience. There are significant main effects of judgment type and the utilitarian status of the response judged, but not of its being personal (Personal $M = 51.64$, $SD = 25.8$, and Impersonal $M = 49.87$, $SD = 30.8$, $n = 2432$) or avoidable (Avoidable $M = 51.46$, $SD = 25.01$, and unavoidable $M = 50.04$, $SD = 31.4$, $n = 2432$). These variables only have interactive effects, which shows that there is not an overall effect of the manipulation across all judgment dimensions.

Figure 5 shows how personal deontological responses are judged as better ($M = 61.75$, $SD = 28.89$, $n = 303$), way more permissible ($M = 65.76$, $SD = 28.56$, $n = 303$), more obligatory ($M = 51.38$, $SD = 33.94$, $n = 303$) and less blameworthy ($M = 42.31$, $SD = 27.59$, $n = 303$) than the corresponding utilitarian responses (Good $M = 39.19$, $SD = 28.12$, $n = 305$, Permissible $M = 43.54$, $SD = 31.06$, $n = 305$, Obligatory $M = 36.53$, $SD = 31.23$, $n = 305$, and blame $M = 60.01$, $SD = 27.22$, $n = 305$) but only for the unavoidable outcomes. When the outcome is avoidable, the pattern is the opposite for both impersonal and personal responses. The main effect of personal contact consists in making the deontological response more acceptable (less blameworthy, more obligatory, better and more permissible) than the corresponding utilitarian ratings (See lower panel, Figure 5). Overall this picture is consistent with prior findings where the effect of these variables on dilemma responses is conditional on several factors (See Christensen et al. (2014)).

Figure 5. Mean ratings for each judgment type by utilitarian status, personal force and avoidable. Error bars represent 95% confidence intervals.

In the model fitting process, the overall score of the OUS (or its subscales) did not result in significant predictors of participant judgments. However, the utilitarian scores show near significant interaction effects, presumably not observable due to statistical power limitations. To further explore these interactions, we assigned participants to a low, medium and high utilitarianism group, by splitting participants scores by the 33rd and 66th percentile. Zooming in

on participants with the lowest and highest utilitarian scores (See Figure 6), we see that participants in the Low Utilitarianism Group (left panel) clearly judge causing harm as less permissible, less optional, more blameworthy than not causing harm. People high on utilitarianism show an opposite pattern, albeit not as distinctive (See full set of comparisons in the supplementary materials).

Figure 6. Mean ratings by of judgment, grouped by Utilitarianism level (panels) and Utilitarian response (bars)

Also notice than permissible judgments tend to always be higher than goodness judgments, suggesting that even though they are related, their difference in magnitude suggests that these judgments are tracking different features of the scenario. This difference is larger and statistically significant for participants with low utilitarian tendencies, who consider harm relatively more permissible ($M = 50.95$, $SD = 30.51$) than it is good ($M = 44.55$, $SD = 27.23$). With participants who score in the higher end of utilitarianism, there is virtually no difference

between these judgments (permissibility $M = 56.02$, $SD = 27.52$ vs obligatoriness $M = 55.35$, $SD = 26.71$). A similar pattern is found with the obligatoriness levels: participants who score low in utilitarianism are the ones who consider causing harm as optional (as revealed by a one-sample t test against the indifference point $M = 37.09$, $SD = 31.02$) and participants who score high in utilitarianism consider *not* causing harm as optional ($M = 44.20$, $SD = 29.81$). These tendencies are statistically significant when explored with a fixed effects model with these terms (see Supplementary materials). Bear in mind that this model should not be taken as testing a hypothesis but as exploring a tendency and as such, lacks confirmatory value.

Discussion

We wanted to see whether adding additional measures allowed variability in moral judgment to emerge, even when using dilemmas as stimuli. However, when we had people make judgments about dilemmas, the variability we observed in Experiment 2 nearly disappears. This is not an effect driven by Utilitarian members of the participant pool. Even when we separate people according to their utilitarian tendencies (which tendencies pull *against* making the distinctions outlined in Experiment 2), we do not find strong dissociations between the categories of goodness, permissibility, and obligation. The dissociations are weak enough that they might seemingly justify the claim that “the influence of wording variations on moral judgments was negligible” (O’Hara et al., 2010).

Our results suggest that changes in measurement alone are not sufficient to indicate the underlying complexity of moral judgment. The kind of situation being evaluated makes a difference to moral judgment. When we limit ourselves to using only dilemmas, we generate an

overly narrow view of moral thinking. Worse, this hinders the ability to make useful generalizations about the operations of moral cognition in general.

Thus, we need more than different measures. The situation being evaluated appears to have an influence on the shape of moral judgment. Our materials in Experiment 2 provide a first step toward expanding the range of situations used as stimuli to study moral judgment. In our next two experiments, we wanted to introduce different measures and scenarios, as well as explore potential connections between moral judgments and generalized rules.

Experiment 4

In Experiment 3 we found that judgment dimensions do not behave uniformly when used to judge responses in sacrificial dilemmas. That is, utilitarian and deontological responses are judged to be permissible, optional, blameworthy and good at different levels, depending on variables known to impact participant responses (Christensen et al. (2014)). However, unlike the pattern observed in Experiment 2, differences between categories were generally non-significant. Moreover, significant differences tended to be small, to arise from interactions, and to be driven mainly by individual differences in trait utilitarianism. This shows that sacrificial dilemmas do not elicit differences in judgment that are found when using different erogation conditions. Conversely, this suggests that not all situations of moral conflict could be assessed in terms of the dimensions proposed.

Additionally, in Experiment 3, participants made judgments of blameworthiness for the response of the person allowing, or causing harm. This marks an important difference with the procedure used in Experiment 2. In Experiment 2, we found that individuals rate scenarios in a way that reflects the categories of suberogatory and supererogatory. However, we did not find

that individuals rate suberogatory and supererogatory behaviors as equally permissible and optional, despite finding a difference in judgments of goodness.

One explanation of this failure might be that people disapprove of suberogatory behaviors, but the various judgment types (permissibility, obligation, and goodness) do not have any obvious place for people to register this disapproval. Hence, disapproval might be skewing judgments in a way that drives the differences between permissibility and obligatoriness across scenarios. To correct for this, and to replicate the findings of Experiment 2, we ran the experiment again with an additional judgment type of blameworthiness.

Participants

We recruited 311 participants (166 women and 145 men, mean age = 31.95, $SD = 11.00$) through Prolific Academic. Sample size was determined from the prior study (Experiment 2). All participants voluntarily agreed to participate in the study and were monetarily compensated (38 pence). The average completion time was 4 minutes.

Materials and procedure

We used the same materials and procedures from Experiment 2. We included an additional measure for participants to rate the blameworthiness of the character in the vignette (praiseworthy = 0, neither praiseworthy nor blameworthy = 50, blameworthy = 100).

Results

Overall, the results replicated the pattern observed in Experiment 2 (see Figure 7). We fitted the same random effect models as in Experiment 2. Model 2, which includes the interactive

587 effect of judgment type, omission and erogation condition, was the best model. Table 4 below
 588 presents only the two best fitting models.

589 Table 4. Summary of Models fitted for Experiment 4

	score	
	Model 1	Model 2
Omission	3.90** (0.19, 7.60)	5.70*** (2.40, 8.90)
GOOD	-19.00*** (-23.0, -16.0)	29.00*** (26.0, 32.0)
OBLIGATORY	-41.00*** (-45.00, -38.00)	-16.00*** (-19.0, -13.0)
PERMISSIBLE	4.20** (0.47, 7.90)	40.0*** (37.00, 43.00)
SUPER EROGATORY	-50.00*** (-53.00, -46.00)	5.20*** (3.60, 6.80)
Mowing		-1.0 (-4.20, 0.460)
Newlyweds		-1.400 (-3.70, 0.92)
Raffle		0.260 (-2.00, 2.60)
Omission:GOOD	-12.00*** (-18.00, -7.20)	-14.00*** (-19.00, -9.30)
Omission:OBLIGATORY	3.70 (-1.50, 9.00)	0.180 (-4.40, 4.80)
Omission:PERMISSIBLE	-7.90*** (-13.00, -2.70)	-10.00*** (-15.0, -5.70)
Omission:SUPER EROGATORY	3.50 (-1.70, 8.80)	
GOOD: SUPER EROGATORY	97.00*** (92.0, 102.0)	
OBLIGATORY: SUPER EROGATORY	51.00*** (46.0, 56.0)	
PERMISSIBLE: SUPER EROGATORY	72.00***	

	(67.00, 77.00)	
Omission: GOOD: SUPER EROGATORY	-3.00	
	(-10.00, 4.40)	
condiomi:Judg TypeOBLIGATORY:ErogationSUPER EROGATORY	-7.10*	
	(-14.0, 0.36)	
condiomi:Judg TypePERMISSIBLE:ErogationSUPER EROGATORY	-4.80	
	(-12.00, 2.60)	
Constant	62.00***	35.00***
	(59.00, 64.00)	(32.00, 38.00)
N	4976	4976
Log Likelihood	-22765.00	-23867.00
AIC	45566.00	47761.00
BIC	45683.00	47852.00

***p < .01; **p < .05; *p < .1

590 As before, both permissibility and goodness ratings were significantly higher for the
591 supererogatory (Permissibility $M = 85.4$, $SD = 19.5$, $n = 622$ and goodness $M = 85.6$, $SD = 20.9$,
592 $n = 622$) than for the suberogatory condition ($M = 63.9$, $SD = 29.9$, $n = 616$ and goodness M
593 $= 38.0$, $SD = 20.5$, $n = 622$) Welch's t -test permissibility: $t(1067) = -15.0$, $p = .001$, $M_{diff} = -21.5$,
594 95% CI [-24.3, -18.7], Welch's t -test goodness: $t(1242) = -40.6$, $p = .001$, $M_{diff} = -47.2$, 95% CI [-
595 49.9, -45.3]. However, the difference between obligatory ratings was not replicated (Welch's t -
596 test: $t(1228) = 0.21$, $p = .8$, $M_{diff} = 0.32$, 95% CI [-2.77, 3.41]). Blame ratings differed
597 significantly across the suberogatory and supererogatory conditions. Participants rated
598 suberogatory behaviors as more blameworthy ($M = 63.70$, $SD = 20.77$, $n = 622$) than
599 supererogatory behaviors ($M = 15.76$, $SD = 20.62$, $n = 622$) Welch's t -test: $t(1242) = 40.8$, $p =$
600 $.001$, $M_{diff} = 47.9$, 95% CI [45.6, 50.2]. As in Experiment 2, there are significant differences
601 between scenarios, with supererogatory responses in the *Kidney* scenario having the highest
602 positive scores (as better, more permissible, more praiseworthy) and suberogatory responses in

603 the *Mowing* scenario the most negative (as worse, less permissible, less optional, and more
604 blameworthy).

605 Blame ratings were not significantly correlated with statements about personal sense of
606 duty across either scenarios or conditions. The same overall pattern of correlations between the
607 other ratings and statements about duties was observed (See Table 2, Experiment 2).

608 Supererogatory behaviors are consistently rated as praiseworthy, while suberogatory
609 behaviors are consistently rated as blameworthy. However, the magnitude of praiseworthiness
610 judgments is much greater in the case of supererogatory behaviors. This asymmetry suggests that
611 supererogatory behaviors deserve more praise than the corresponding suberogatory behaviors
612 deserve blame. (see Figure 7). This asymmetry has been repeatedly observed in other studies
613 (Monroe et al., 2018, Pizarro et al., 2007).

614 It is worth noting that, again as in Experiment 2, there are no important differences
615 between the permissibility and goodness means in the supererogatory condition ($M = 85.61$, SD
616 $= 20.86$ vs $M = 85.43$, $SD = 19.46$ respectively) while for the suberogatory responses we find a
617 difference of 25.88 points between these (Cohen's $d = 1.01$), suggesting that participants are
618 considering different information when assessing permissibility and goodness in these situations.

Figure 7. Scores by type of judgment Experiment 4. Error bars represent 95% confidence intervals

Discussion

In this experiment, we replicated the key findings of Experiment 2. When comparing judgments about suberogatory and supererogatory behaviors, people distinguish between the goodness, permissibility, and blameworthiness of the action. However, people do recognize suberogatory and supererogatory behaviors as equally optional. Despite statistically significant differences between the permissibility of suberogatory and supererogatory actions, participants rated suberogatory behaviors as permissible.

Now that we have provided some evidence for distinctions between these evaluative categories, we turn to another question: what factors differentially affect these judgments? As we saw in Experiment 2, beliefs about personal duties do not relate significantly to people's judgments. The significant difference in blame judgments suggests that perhaps norms of social

sanctioning help to inform at least one kind of judgment. To understand the reasoning underlying these judgments, we conducted two additional experiments.

Experiment 5

In this experiment we explored how judgments on deserved social sanction are associated with the different kinds of moral judgment identified in previous experiments. Thus far, we have established that some situations can be understood in terms of suberogatory or supererogatory characterizations. As such, in order to understand people's responses to these situations, we need to deploy a richer array of judgment dimensions. In this experiment, we explored whether suberogatory responses are associated with potential behavioral consequences, such as social punishment. While our previous results establish that supererogatory behaviors merit praise and suberogatory responses merit blame, we have not yet explored any behavioral consequences associated with these judgments. Social punishment, as a more overt form of blame, begins to establish behavioral markers of underlying moral judgments.

We also explore whether attitudes about rights predict different kinds of judgments. The situations that we describe often bring into conflict an individual right and a kind or generous thing that could be done for others. For example, someone has the right to keep the seat they paid for on the airplane, but it is a kind thing to offer one's seat when asked to switch. Because of this, we thought that attitudes towards *rights*, instead of duties, might usefully predict moral judgments. This would partially vindicate the rule-based conception of morality, as one might hold that there are generalized functions from rights to prohibited behaviors (even if there are fewer—if any—such functions from rights to *prescribed* behaviors). Moreover, we found this pattern of reasoning in the open responses for Experiment 1. All of this suggests that personal

sense of rights might provide some details about one range of inputs to which one kind of moral judgment is sensitive

Participants

309 participants (165 women and 144 men, mean age = 39.89, SD =13.76), recruited through Prolific Academic, took part in the study in exchange for 60 pence. The average completion time was of 4.87 minutes and none of the participants had taken part in the previous experiments.

Materials and procedure

We used the same materials and procedures from Experiment 4 with two changes. First, people had to judge whether the response merited social sanction or recognition by using a slider with an underlying scale going from 0 to 100 (social sanction = 0, neither = 50, social recognition=100). Second, people provided information about their attitudes towards rights by answering these three questions with a 7-point Likert scale (1= Strongly disagree, 7 = Strongly agree): 1) Everyone has a right to do anything they want on their property (as long as they are not hurting anyone else); 2) Everyone has the right to bodily autonomy; 3) Everyone has a right to make use of their money (or goods) as they see fit.

Results

As in Experiments 2 and 4, we fitted a series of random effect models which resulted in a best fitting model that includes the interactive effect of judgment type, omission, erogation condition and scenario. Crucially, including terms for attitudes about rights did not significantly improve the fit of this model.

Table 5. Summary of Models fitted for Experiment 5

	score
Omission	-8.084*** (-10.949, -5.220)
Judg TypeOBLIGATORY	-22.945*** (-26.742, -19.148)
Judg TypePERMISSIBLE	19.742*** (15.945, 23.539)
Judg TypeSANCTION	-2.658 (-6.455, 1.139)
ErogationSUPER EROGATORY	44.904*** (42.214, 47.593)
ScenarioMowing	-5.245*** (-8.252, -2.238)
ScenarioNewlyweds	-5.136*** (-7.038, -3.234)
ScenarioRaffle	3.685** (0.678, 6.691)
condiomi:Judg TypeOBLIGATORY	15.276*** (9.898, 20.655)
condiomi:Judg TypePERMISSIBLE	9.128*** (3.749, 14.507)
condiomi:Judg TypeSANCTION	7.976*** (2.597, 13.355)
Judg TypeOBLIGATORY:ErogationSUPER EROGATORY	-39.761*** (-45.131, -34.391)
Judg TypePERMISSIBLE:ErogationSUPER EROGATORY	-22.181*** (-27.551, -16.811)
Judg TypeSANCTION:ErogationSUPER EROGATORY	-14.171*** (-19.541, -8.801)
condiomi:Judg TypeOBLIGATORY:ErogationSUPER EROGATORY	-3.294 (-10.901, 4.313)
condiomi:Judg TypePERMISSIBLE:ErogationSUPER EROGATORY	-5.618 (-13.225, 1.989)
condiomi:Judg TypeSANCTION:ErogationSUPER EROGATORY	-3.650

Constant	(-11.257, 3.956) 45.964*** (43.028, 48.900)
N	4944
Log Likelihood	-22788.000
AIC	45617.000
BIC	45747.000

*** $p < .01$; ** $p < .05$; * $p < .1$

The pattern found in Experiments 2 and 4 was replicated here (see Figure 8). Supererogatory behaviors were judged as better and more permissible ($M = 85.16$, $SD = 19.31$, $n = 618$ and $M = 84.47$, $SD = 20.18$, $n = 618$) than suberogatory behaviors ($M = 40.26$, $SD = 23.99$, $n = 618$ and $M = 64.55$, $SD = 30.09$, $n = 618$) Welch's t -test Good: $t(1180) = -44.9$, $p < .001$, $M_{diff} = -44.4$, 95% CI [-47.3, -42.5] and Welch's t -test Permissible: $t(1079) = -13.7$, $p < .001$, $M_{diff} = -19.9$, 95% CI [-22.8, -17.1], regardless of being actions or omissions. Supererogatory behaviors were again judged to be less optional ($M = 28.43$, $SD = 30.61$, $n = 618$) than the corresponding suberogatory responses ($M = 24.93$, $SD = 26.34$, $n = 618$), Welch's t -test: $t(1207) = -2.15$, $p = .03$, $M_{diff} = -3.50$, 95% CI [-6.68, -0.31]. As in Experiment 4 this difference only emerged when an additional measure was included (blameworthiness in Exp 4 and social sanction in this Exp) suggesting the presence of a joint evaluation effect (Hsee, Blount, Loewenstein, & Bazerman, 1999).

Supererogatory behaviors were judged to merit more social recognition than suberogatory responses ($M = 70.49$, $SD = 25.41$, $n = 618$ vs $M = 41.59$, $SD = 20.57$, $n = 618$, Welch's t -test: $t(1183) = -22.0$, $p < .001$, $M_{diff} = -28.9$, 95% CI [-31.5, -26.3]). Crucially, both means for social recognition and social sanction are significantly different from the indifference point (50 in our

scale) (Supererogatory Welch's t -test: $t(617) = 20.1, p < .001, M = 70.5, 95\% \text{ CI } [68.5, 72.5]$ and suberogatory Welch's t -test: $t(617) = -10.2, p < .001, M = 41.5, 95\% \text{ CI } [40.0, 43.2]$. Mirroring the asymmetric pattern observed for blameworthiness ratings, supererogatory responses deserve more recognition than sanction is deserved by suberogatory responses. Also, as in the previous studies there are important differences between scenarios, even when the overall pattern is consistent. For example, not giving up the raffle prize does not merit neither recognition nor sanction ($M = 50.82, SD = 22.06, n = 155$) while mowing your lawn early morning is clearly disapproved ($M = 34.23, SD = 19.76, n = 155$).

Figure 8. Scores by type of judgment and condition. Error bars represent 95% confidence interval.

To explore the association between judgment scores and attitudes towards rights we first observed the strength of the association between attitudes towards rights. Correlations are medium to small in size (see Table 6) and there is low internal consistency (Cronbach's alpha = 0.50). We then calculated the correlations between judgment scores and these attitudes towards rights (Table 6). Similar to what happened with questions about duties, the correlations are low and only significant for two associations: (1) bodily autonomy and permissibility, and; (2) obligatoriness and right to property. Only when considering particular scenarios, moderately stronger correlations start to emerge (for example, permissibility ratings in the *Kidney* scenario are significantly correlated with attitudes on bodily autonomy $r(309) = .12, p < 0.01$)

Table 6. Bivariate correlations between attitudes towards rights and judgment scores

	Permissible	Good	Obligatory	Sanction	Right to Money	Bodily autonomy
Permissible						
Good	0.54***					
Obligatory	-0.24**	-0.04***				
Sanction	0.24***	0.40***	0.14***			
Right to money	0.03	0.04	-0.05	0.01		
Bodily autonomy	-0.16**	0.05	-0.14	0.00	0.16***	
Right to property	0.05	0.02	0.06*	0.01	0.44***	0.17***

** $p < .001$; * $p < .01$; * $p < .05$

Discussion

We again find the same dissociations as in Experiments 2 and 4. Including the category of sanctions, however, brings out something interesting. There is a much closer relationship between goodness and sanction than there is between goodness and blame in suberogatory cases. This suggests that people might be thinking of badness in terms of how others will react, with judgments of blame tracking something else. The relationship does not hold in cases of

supererogatory behavior, where judgments of goodness were statistically distinct from judgments of sanction (positive recognition, in this case). This suggests that people think suberogatory actors should be sanctioned, though supererogatory actors need not necessarily be positively recognized. This, however, might be a function of the cases we used rather than marking out an intrinsic difference between commonsense thinking about the sub- and supererogatory.

We also identified an interesting relationship between sanction and goodness. Actions rated as bad received similar ratings of sanction. In the previous experiment, we saw that ratings of badness do not align with ratings of permissibility. Permissibility ratings are more closely aligned with ratings of blame. This suggests that the permissibility of suberogatory behaviors is related to blameworthiness, whereas the badness of suberogatory behaviors is related to public sanction. However, there is an asymmetry between suberogatory and supererogatory behaviors. We discuss this patter more in the General Discussion.

We did not find any interesting correlations between beliefs about rights and different kinds of moral judgment. This was similar to our findings about individual sense of duty in Experiment 2. This shows that there is no generalized function from facts about individual rights to how one should evaluate behavior in morally charged situations. This provides further evidence for the claim that commonsense thinking about morality does not reflect a structured system of rules.

Experiment 6

So far we have shown that judgments of permissibility, goodness, and obligation track distinct features of a situation. Evaluations of supererogatory and suberogatory behaviors brings these distinctions to light. One question that remains is what factors drive these different

741 judgments. The previous experiments offered some suggestions. In Experiment 1, we found some
742 people using the language of rights to offer justifications for suberogatory behavior. Additionally,
743 people used virtue terminology and aretaic categories in free descriptions of the scenarios.

744 Experiments 2, 4 and 5 suggest, however, that people are not explicitly taking into account their
745 attitudes towards duties (Exps. 2 and 4) or rights (Exp 5). This suggests two things. Judgments of
746 goodness might be sensitive to what people think a decent or virtuous person would do in a
747 situation, while judgments of permissibility track what some individual has the right to do in a
748 given situation. If these suggestions are correct, it would explain why judgments of permissibility
749 and goodness often coincide: what is decent or virtuous often overlaps with what one has a right
750 to do. However, it also explains why the two concepts dissociate in cases of suberogatory
751 behavior. Suberogatory situations arise when someone has a right to do something that a virtuous
752 or decent person would not do. If this is the case, it is the absence of duties and the corresponding
753 presence of rights that are characteristic of suberogatory situations. These characteristics are,
754 however, not framed as rules, but are instead low-level features of the situation.

755 In this experiment, we explore the justifications people offer for their judgments of
756 permissibility and goodness to see whether these different factors explain the distinction between
757 these two kinds of judgment. We also included different scenarios that are more properly moral
758 situations. In the previous experiments, we used scenarios that might be construed as
759 conventional (e.g. lawn mowing or seat ownership) without being unmistakably moral. In this
760 experiment, we aimed to use prototypical moral situations.

761 **Method**

762 **Participants**

763 We recruited 316 participants (160 women and 153 men, mean age = 33.30, SD =10.81)
764 through Prolific Academic to participate in the study in exchange for 60 pence. The average
765 completion time was 10 minutes and none of the participants had taken part in the previous
766 experiments.

767 **Materials and procedure**

768 We constructed two new scenarios and created a suberogatory and supererogatory version
769 of each. The two scenarios (adapted from Thomson, 1971 and Nozick, 1974, respectively) are
770 described below [suberogatory version in brackets]:

771 **Violinist** Alex is driving home from work on the highway when she gets into an accident that knocks her
772 unconscious. When she wakes up, she finds herself in a hospital bed. She's also connected to another individual
773 through a series of wires and tubes. A doctor enters the room and explains to Alex that she is fine, but the
774 individual she's connected to suffered some severe damage to internal organs. Alex has the right blood type to
775 help, and—since she was unconscious—the doctor decided to connect Alex to keep the other individual alive for
776 the time being. The doctor explains that Alex can unplug herself if she chooses, but the individual will most likely
777 die. The individual will recover from these injuries in about a month (give or take a few days), after which time Alex
778 can unplug herself and leave. After a few hours of pondering what to do, Alex decides stay plugged in for the
779 month [*to unplug herself*].

780 **Well** Jones finds a large freshwater source on his property, so he digs a well as a way of claiming the
781 water. A few weeks later, the town where he lives begins experiencing a drought, which was completely
782 unpredictable. Town representatives visit Jones to ask whether they can use his water to alleviate some of the
783 drought. Without Jones' help, the town will likely run out of water in a few days. If Jones donates some of his

water, however, he might experience the effects of the drought in the unlikely event that the drought prolongs for too long. After considering what to do, Jones decides to offer his water [*declines to offer his water*]

To test variation against a known benchmark, we also included the *Newlyweds* scenario (See Exp. 5). Participants read the informed consent and then rated each one of the three vignettes. The order of the vignettes was randomized across participants. Each participant was randomly assigned to see either the supererogatory or suberogatory condition of each scenario and the rest of the procedure was the same as in Experiment 4. The only difference is that, after providing judgment for the dimensions requested, participants were asked to explain their ratings in their own words (cf. Christensen et al., 2014).

Results

Overall, numerical ratings closely followed the pattern observed in prior Experiments (See Figure 9). Suberogatory behaviors are rated as worse ($M = 34.19$, $SD = 23.71$, $n = 468$ vs $M = 86.37$, $SD = 19.80$, $n = 480$), more blameworthy ($M = 64.13$, $SD = 22.46$, $n = 468$ vs $M = 14.92$, $SD = 19.12$, $n = 480$), and less permissible ($M = 28.17$, $SD = 32.31$, $n = 468$ vs $M = 83.15$, $SD = 22.27$, $n = 480$) than supererogatory responses (all significant differences). Unlike Experiment 5, there is no significant difference in obligatoriness judged (Suberogatory $M = 28.17$, $SD = 32.31$, $n = 468$ vs $M = 26.64$, $SD = 31.75$, $n = 480$). As before, we fitted random effect models, where the best model includes the interaction of judgment type, scenario and erogation condition (results can be seen in Table 6).

Figure 9. Mean scores by type of judgment Experiment 4. Error bars 95% represent confidence intervals

There is also a significant between-scenario variation. For example, only in the *Well* scenario people judged the suberogatory behavior, not giving water, as significantly less optional than the supererogatory behavior ($M = 31.26$, $SD = 31.75$, $n = 160$ vs $M = 38.81$, $SD = 34.63$, $n = 156$, Welch's t -test Good: $t(310) = 2.02$, $p < .05$, $M_{diff} = 7.55$, 95% CI [0.19, 14.9]). In this scenario also occurred the most extreme values for suberogatory responses.

Table 6. Summary of Models fitted for Experiment 6

	score
Judg TypeGOOD	-18.64*** (-24.32, -12.96)
Judg TypeOBLIGATORY	-43.06*** (-48.74, -37.37)

Judg TypePERMISSIBLE	14.01*** (8.33, 19.69)
Erogation: SUPER EROGATORY	-35.48*** (-41.08, -29.89)
ScenarioViolinist	3.61 (-2.02, 9.24)
ScenarioWell	13.13*** (7.48, 18.79)
Judg Type:GOOD:ErogationSUPER EROGATORY	77.22*** (69.31, 85.13)
Judg Type:OBLIGATORY:ErogationSUPER EROGATORY	36.02*** (28.11, 43.93)
Judg Type:PERMISSIBLE:ErogationSUPER EROGATORY	50.07*** (42.16, 57.9)
Judg Type:GOOD:ScenarioViolinist	-7.79* (-15.75, 0.16)
Judg Type:OBLIGATORY:ScenarioViolinist	10.88*** (2.92, 18.84)
Judg Type:PERMISSIBLE:ScenarioViolinist	-10.40** (-18.37, -2.45)
Judg Type:GOOD:ScenarioWell	-25.91*** (-33.91, -17.91)
Judg Type:OBLIGATORY:ScenarioWell	10.22** (2.22, 18.22)
Judg Type:PERMISSIBLE:ScenarioWell	-29.89*** (-37.88, -21.89)
Erogation:SUPER EROGATORY:ScenarioViolinist	-17.16*** (-25.07, -9.25)
Erogation:SUPER EROGATORY:ScenarioWell	-24.187*** (-32.096, -16.277)
Judg Type:GOOD:ErogationSUPER EROGATORY:ScenarioViolinist	27.74*** (16.56, 38.93)
Judg Type:OBLIGATORY:ErogationSUPER EROGATORY:ScenarioViolinist	19.657*** (8.47, 30.84)
Judg Type:PERMISSIBLE:ErogationSUPER EROGATORY:ScenarioViolinist	12.94** (1.75, 24.12)
Judg Type:GOOD:ErogationSUPER EROGATORY:ScenarioWell	44.97***

Judg Type:OBLIGATORY:ErogationSUPER EROGATORY:ScenarioWell	(33.78, 56.15) 16.09***
Judg Type:PERMISSIBLE:ErogationSUPER EROGATORY:ScenarioWell	(4.91, 27.28) 39.84***
Constant	(28.65, 51.03) 58.52***
N	(54.50, 62.54) 3792
Log Likelihood	-17589.00
AIC	35231.00
BIC	35393.00

***p < .01; **p < .05; *p < .1

Qualitative Data

For the qualitative analysis of the open responses, we used two coders. Both were blind to the individual ratings associated with each open response and one coder was completely blind to the objective of the study. Coders used ten predefined categories to sort responses. We report agreement between coders (Cohens Kappa) and results from fitting loglinear models on the classification frequencies.

There were 4003 unique written explanations of ratings (1004 for goodness ratings, 999 for obligatory ratings, 997 for permissibility ratings and 1003 for blame/praise ratings). Coders were given ten categories. *Character* when the response included some mention of personality or character traits; *Rights (presence or absence)* when responses appealed to what the protagonist can morally do, is allowed to do, or its explicit negation; *Duties (presence or absence)* when responses appeal to what the protagonist is morally obliged to do, or its explicit negation; *Neither* when participants appealed to norms that could not be considered rights nor duties, etc; *Outcomes* when the justification was based on the consequences for the “victims”, negative or positive;

Values actions performed had a clear valence; *Justification* when appealing to reasons not based on consequences; *Other* if the response did not fall in the previous categories but was common enough to merit its own. Coders were given an example of each one of the categories. Raters completed coding independently. They were also told that they could classify any given response in more than one category.

For example, one response to the *Well* vignette reads: “He didn’t have to donate the water, and even if he didn’t that wouldn’t be necessarily a bad thing.” This was rated as indicating ‘Absence of Duties’ and ‘Values’. Another response (to the *Newlyweds* vignette) reads: “No one should have to switch a seat unless they want to. They bought the seat they were in and would have been justified in staying in that seat.” This was rated as indicating ‘Justification’ and ‘Presence of Rights’.

Raw agreement between coders and Cohen’s Kappa for each category are presented in Table 7. The category *Neither* is not presented because there was virtually no agreement in its use. Most of the unique responses were classified into at least 2 categories (58%). Half of the categories show a weak reliability between coders (outcomes, values, justification and absence of rights) while the other half shows an overall good agreement (presence of duties, character, presence of rights and absence of duties).

Table 7. Cohen's kappa agreement between coders for the eight categories. Columns L and U are the Lower and Upper limits of Cohen’s confidence interval.

	Cohen	Inf	Sup	% Classified as	% Agreement
Abs duties	0.81	0.79	0.84	11.02	95.68
Pres rights	0.75	0.73	0.77	21.06	90.50

Character	0.69	0.65	0.72	8.61	94.18
Pres duties	0.63	0.58	0.68	5.65	95.25
Outcomes	0.39	0.35	0.43	11.86	85.08
Values	0.38	0.35	0.42	20.28	77.18
Justification	0.37	0.34	0.41	13.7	82.35
Abs rights	0.34	0.27	0.41	3.13	95.15

To determine the association between justification category and the variables manipulated in this experiment, we submitted the classification frequencies to a loglinear analysis with a saturated model including erogation condition (sub and supererogatory), judgment dimension (goodness, permissibility, obligatoriness, blameworthiness) and justification classification category (presence of duties, character, presence of rights and absence of duties). The three-way loglinear model produced a model that retained all interactive effects (likelihood ratio $\chi^2(0) = 0$, $p = 1$), indicating a significant interaction between the variables fitted to the model ($\chi^2(12) = 102.53$, $p < 0.001$). Judgments of permissibility and obligatoriness were mainly sensitive to presence of rights and these explanations appealing to the presence of rights were more prevalent in suberogatory cases. Judgments of permissibility were also justified by appealing mostly to the absence of duties and the presence of rights (>40%) but in this case there is no asymmetry between sub and supererogatory response. Judgments of goodness were mostly justified by appealing to character and again the suberogatory responses were justified by relying on presence of rights. Additionally, there are more character justifications for supererogatory responses based on character than for the corresponding suberogatory responses.

The qualitative analysis of the open explanations suggests that people do systematically call to mind considerations on duties and rights when making judgments about permissibility and obligatoriness. While this is not surprising, the interesting thing here is that these considerations

are also significant for judgments of blameworthiness and goodness. Duties and rights showed up, not as blank generalizations – as explored in Experiments 2 and 5 – but as situated appeals to situational factors.

Discussion

As before, we replicated the basic pattern found in our other experiments despite using two entirely different situations that are more clearly moral situations. This provides strong evidence for our claim that differences in evaluative categories track differences in kinds of judgments.

We also used open responses to allow participants to supply their own reasoning for making different moral judgments. Our results are striking. People appear to be sensitive to many kinds of information differentially when making moral judgments. Judgments of praise and blame are sensitive to character considerations. Judgments of goodness seem primarily sensitive to character considerations, whereas judgments of badness seem sensitive to both character considerations and the presence of rights. This, however, might be a function of the situations we provided. When people make judgments of badness, they were often pointing to the fact that the bad behavior was fully within the rights of the individual, even though the behavior was indicative of vicious character. The fact that people mention both points to the relevance of both kinds of information to making a judgment of badness. Judgments of permissibility appear to track the presence of rights and the absence of duties, whereas judgments of obligatoriness track the presence of rights and the presence or absence of duties. We offer an interpretation of this in the general discussion.

Finally, we should note that the categories used to sort open responses were generated *a priori*. Bottom-up, data-driven methods might reveal a more perspicuous classificatory scheme that might generate surprising results about the kinds of information relevant to making different kinds of moral judgment.

General Discussion

In this paper, our goal was to develop the idea that morality is not exhausted by a system of rules and explore some of its implication for our understanding of the psychology of moral judgment. To do this, we've made a case for moving beyond the use of dilemmas in the study of moral judgment and propose a wider array of measures and situations. In this General Discussion, we summarize how our findings apply to the measurement of moral judgment, the cases used to elicit judgments, and the rule-based view of morality.

Measures and Categories of Moral Judgment

We found that judgments across different moral categories are dissociable. In particular, judgments of goodness and permissibility come apart depending on the situation presented. People form judgments about situations that conform to the category of the suberogatory (bad, but permissible). This challenges the idea that different measures of moral judgment tap into the same underlying reasoning process. The resulting picture of moral judgment is that people are variably sensitive to different kinds of considerations when making different kinds of judgments (see Barbosa and Jiménez-Leal, 2017). While this presents a richer picture of moral judgment, it also opens up new questions. When do people prefer different kinds of information in making

904 moral judgments? What sorts of considerations predominantly drive different kinds of
905 judgments?

906 Our results provide some initial suggestion for how to answer these questions. In many of
907 our experiments, ratings of obligation remained consistently low. This indicates that people
908 considered both suberogatory and supererogatory behaviors as optional rather than obligatory.
909 The qualitative data from Experiment 6 suggests that ratings of obligation are predominantly
910 tracking the presence of duties or explicit rules that prohibit or prescribe conduct. Because
911 morally charged situations are not composed of conflicts between rules, there is an absence of
912 duties that defines how one ought to behave. Hence, ratings of obligation are low.

913 Results from Experiment 5 indicate that ratings of badness align with ratings of sanction.
914 That is, the degree to which one rates a suberogatory action as bad is related to amount of
915 sanction one deserves in light of doing something suberogatory. This is importantly different
916 from the results of Experiment 4, which indicated that ratings of permissibility are aligned with
917 ratings of blame.

918 These results suggest that folk psychological categories of blame and sanction might
919 dissociate. Because these categories differentially associate with others, it might be the case that
920 these judgments track different features of a situation. Understanding this difference might
921 provide further insight into what judgments of permissibility and goodness indicate about folk
922 psychological evaluation. This is an important result to investigate in future work. Moral
923 psychologists and philosophers have assumed that these constructs significantly overlap and,
924 accordingly, have used them interchangeably (Bendor & Swistak, 2001; Deutsch & Gerard,
925 1955; Cialdini, Reno, & Kallgren, 1990; Scanlon, 2008; McKenna, 2013; Bennett, 2012). Future

work should investigate further whether and under what circumstances there are reliable dissociations between these categories.

Notably, some have mentioned the need for using new measures in studying moral judgment (see Uhlmann, Pizarro, & Diermeier, 2015), arguing that folk psychological categories of judgment are fundamentally directed at personal evaluation rather than behavior evaluation and routinely employs aretaic rather than deontic concepts. Our results suggest that this is partially true. People do show an interest in personal evaluation using aretaic concepts. However, we also find that behavior evaluation and deontic concepts also play a significant role. This suggests that social cognition is sensitive to both act-based and person-based considerations, as well as a wide conceptual repertoire in normative evaluation.

Moral Situations

Using a variety of measures is not sufficient to bring out the underlying variability of moral judgment. In Experiment 3 (on sacrificial dilemmas), we did not find similar dissociations between judgment categories as we did in our other experiments. This was the case even when controlling for underlying ethical tendencies (i.e., low or high trait utilitarianism). Even among low trait utilitarians, different kinds of judgments never significantly differed from each other. This suggests that dilemmas themselves are qualitatively distinct from the morally charged situations we used in our experiments.

As discussed earlier, assumptions about methodology and dilemmas mutually reinforce each other. We can now clarify this point further. The use of dilemmas as stimuli functions as a demand effect that gives the appearance of commonsense moral thinking exhibiting a rule-based, hierarchical structure. In providing scenarios that explicitly bring different sets of rules into

conflict, researchers have set people up to exhibit judgments that appear to preferentially select some rules over others. This, in turn, blurs the distinctions between different judgment categories. Once we free moral judgment of the constraints of dilemmas, the variegated measures are able to capture the variability of judgment. That variability is only possible when using situations that do not bring different sets of rules into conflict.

Rule-based Morality

There is something important about presenting people with situations that extend beyond dilemmas beyond addressing concerns of ecological validity. Even the most realistic dilemma scenarios still present only a narrow slice of the moral life (Kahane, 2015). Dilemmas arise from conflicts between different sets of rules. We have seen that not all moral situations reflect conflicts between rules. There are many moral situations that cannot be navigated by finding and applying the right rules.

Similarly, moral reasoning does not rely on applying abstract principles. Some of our results align with this. We failed to find correlations between abstract claims about the nature of rights and duties and people's judgments about particular situations. This reflects a tendency of people to identify features of situations that guide judgment without having explicit representations of abstract principles play a causal role in forming moral judgments (Graham et al., 2013). The open responses we collected in Experiment 6 suggest that there are no rules that structure even part of commonsense moral thinking. People display a willingness to recognize various exempting conditions, exceptions, or room to maneuver. Rather than rules, people seem to have a series of highly particularized heuristics that guide normative evaluation.

For all we have said, it may turn out that precise algorithms adequately characterize the computations performed in forming moral judgments. Alternatively, moral judgment might be supported by an architecture that functions as a model-based system without having to explicitly represent rules (see Cushman, 2015; Brownstein, 2018). Our interest here, however, does not lie with the computations or cognitive architecture behind moral judgment, but with the content of morality. Our claim is that most research on moral judgment has, thus far, assumed that the content of moral judgment can be measured in terms of its alignment with some structured system of rules. But fixating on rules provides only a partial window into the moral life, one that fails to reveal the complexities and subtleties of commonsense moral judgment. When moral judgment is removed from the narrow frame of sacrificial dilemmas, the appearance of rules in moral judgment evaporates.

Finally, our argument here does not assume that if morality has a rule-based structure, then commonsense moral judgment ought to exhibit perfectly coherent and systematic principles. Even Sidgwick (1981) admits that the maxims of commonsense morality are “somewhat vague generalities” (p. 342). Still, on his view and on the view of those who follow him, commonsense morality can be refined through rigorous theorizing to *approximate* the structure of an ideal normative theory. It’s made of the right sort of stuff, as it were, to function as a rule-based system.

The judgments that we have identified in these experiments, however, imply that this is an incomplete view of the moral framework reflected in moral cognition. Consider that people regularly judge that some behavior is bad but permissible across a range of different scenarios. If there are rules constitutive of commonsense morality, what must they be like such that they allow for such judgments? How can we make sense of such judgments within a rule-governed system?

992 Recall that one key feature of the rule-based system of morality is the inter-definability of moral
993 concepts. Whatever is permissible is in line with the rules of morality; whatever is bad goes
994 against those rules.

995 One option is to say that separate domains of rules govern separate judgments. But that
996 requires giving up on the notion of inter-definability that is central to the rule-based system of
997 morality. Another option is to say that one kind of judgment is properly moral and the other is
998 not. However, there appears to be no principled way of stating that either judgments of
999 permissibility or goodness are properly moral while excluding the other. A third option is to say
1000 that there *are* systematic principles underlying these judgments, but people do not understand
1001 what these are. Hence, they are making a mistake in dissociating permissibility and goodness.
1002 But this begins to look like the imposition of researcher assumptions rather than an investigation
1003 into commonsense moral thinking. A final possibility is that morality is self-contradictory.

1004 At this point, one begins to wonder if the substantive assumption about morality being a
1005 system of rules is worth the cost. We suggest that it is not, and that it is time to consider what
1006 morality beyond the rules would look like.

References

- Anderson, A.R. (1958). A reduction of deontic logic to alethic modal logic. *Mind*, 67, 100-103.
- Andrade, G. (2019). "Medical ethics and the trolley Problem," *Journal of Medical Ethics and History of Medicine*, 12:3.
- Archer, A. (2018). Supererogation. *Philosophy Compass*, 13(3), 1–9. doi:[10.1111/phc3.12476](https://doi.org/10.1111/phc3.12476)
- Barbosa, S., & Jiménez-Leal, W. (2017). It's not right but it's permitted: Wording effects in moral judgement. *Judgment and Decision Making*, 12(3), 308.
- Bartels, D. M., & Pizarro, D. A. (2011). The mismeasure of morals: Antisocial personality traits predict utilitarian responses to moral dilemmas. *Cognition*, 121(1), 154–161.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting linear mixed-effects models using lme4. *ArXiv Preprint ArXiv:1406.5823*.
- Bauman, C. W., McGraw, A. P., Bartels, D. M., & Warren, C. (2014). Revisiting external validity: Concerns about trolley problems and other sacrificial dilemmas in moral psychology. *Social and Personality Psychology Compass*, 8(9), 536–554.
- Białek, M. & De Neys, W. (2017). Dual processes and moral conflict: Evidence for deontological reasoners' intuitive utilitarian sensitivity. *Judgment and Decision Making*, 12(2), 148–167.
- Bennett, C. 2012. The Expressive Function of Blame. In J. Coates and N. Tognazzini (Eds.) *Blame: Its Nature and Norms* (Oxford: Oxford University Press).
- Björklund, F. (2003). Differences in the justification of choices in moral dilemmas: Effects of gender, time pressure and dilemma seriousness. *Scandinavian Journal of Psychology*, 44(5), 459–466.
- Christensen, J. F., Flexas, A., Calabrese, M., Gut, N. K., & Gomila, A. (2014). Moral judgment reloaded: A moral dilemma validation study. *Frontiers in Psychology*, 5(JUL), 1–18. <https://doi.org/10.3389/fpsyg.2014.00607>
- Christensen, J. F., & Gomila, A. (2012). Moral dilemmas in cognitive neuroscience of moral decision-making: A principled review. *Neuroscience & Biobehavioral Reviews*, 36(4), 1249–1264.
- Conway, P., & Gawronski, B. (2013). Deontological and utilitarian inclinations in moral decision making: a process dissociation approach. *Journal of Personality and Social Psychology*, 104(2), 216.
- Cushman, F. (2008). Crime and punishment: Distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition*, 108(2), 353–380.
- Cushman, F., Young, L., & Greene, J. D. (2010). Our multi-system moral psychology: Towards a consensus view. *The Oxford Handbook of Moral Psychology*, 47–71.

- 1042 Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in
1043 moral judgment: Testing three principles of harm. *Psychological Science*, 17(12), 1082–
1044 1089.
- 1045 Dahl, F. A., & Ofstedal, G. (2019). Trolley Dilemmas Fail to Predict Ethical Judgment in a
1046 Hypothetical Vaccination Context. *Journal of Empirical Research on Human Research*
1047 *Ethics*, 14(1), 23–32.
- 1048 Darley, J. M., & Shultz, T. R. (1990). Moral rules: Their content and acquisition. *Annual Review*
1049 *of Psychology*, 41(1), 525–556.
- 1050 Djeriouat, H., & Trémolière, B. (2014). The Dark Triad of personality and utilitarian moral
1051 judgment: The mediating role of Honesty/Humility and Harm/Care. *Personality and*
1052 *Individual Differences*, 67, 11–16.
- 1053 Driver, J. (1992). The suberogatory. *Australasian Journal of Philosophy*, 70(3), 286–295.
- 1054 Everett, J. A. C., & Kahane, G. (2020). Switching Tracks? Towards a Multidimensional Model of
1055 Utilitarian Psychology. *Trends in Cognitive Sciences*.
- 1056 Fauconnet, P. (2013). *La responsabilité. Étude de sociologie*. Presses Électroniques de France.
- 1057 Gilligan, C. (1977). In a different voice: Women's conceptions of self and of morality. *Harvard*
1058 *Educational Review*, 47(4), 481–517.
- 1059 Gilligan, C. (1982). *In a different voice: Psychological theory and women's development*.
1060 Harvard University Press.
- 1061 Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S. P., & Ditto, P. H. (2013). Moral
1062 foundations theory: The pragmatic validity of moral pluralism. In *Advances in experimental*
1063 *social psychology* (Vol. 47, pp. 55–130). Elsevier.
- 1064 Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An
1065 fMRI investigation of emotional engagement in moral judgment. *Science*, 293(5537), 2105–
1066 2108. <https://doi.org/10.1126/science.1062872>
- 1067 Haidt, J. (2001). The emotional dog and its rational tail: a social intuitionist approach to moral
1068 judgment. *Psychological Review*, 108(4), 814.
- 1069 Hofmann, W., Wisneski, D. C., Brandt, M. J., & Skitka, L. J. (2014). Morality in everyday life.
1070 *Science*, 345(6202), 1340–1343.
- 1071 Holyoak, K. J., & Powell, D. (2016). Deontological coherence: A framework for commonsense
1072 moral reasoning. *Psychological Bulletin*, 142(11), 1179.
- 1073 Hornik, K., Zeileis, A., & Meyer, D. (2006). The strucplot framework: visualizing multi-way
1074 contingency tables with vcd. *Journal of Statistical Software*, 17(3), 1–48.
- 1075 Hsee, C. K., Loewenstein, G. F., Blount, S., & Bazerman, M. H. (1999). Preference reversals
1076 between joint and separate evaluations of options: a review and theoretical analysis.
1077 *Psychological Bulletin*, 125(5), 576.

- 1078 Hurd, H. M. (1998). Duties beyond the call of duty. *JRE*, 6, 3.
- 1079 Kahane, G. (2015). Sidetracked by trolleys: Why sacrificial moral dilemmas tell us little (or
1080 nothing) about utilitarian judgment. *Social Neuroscience*, 10(5), 551–560.
1081 <https://doi.org/10.1080/17470919.2015.1023400>
- 1082 Kahane, G., Everett, J. A. C., Earp, B. D., Caviola, L., Faber, N. S., Crockett, M. J., & Savulescu,
1083 J. (2018). Beyond sacrificial harm: A two-dimensional model of utilitarian psychology.
1084 *Psychological Review*, 125(2), 131.
- 1085 Koenigs, M., Kruepke, M., Zeier, J., & Newman, J. P. (2012). Utilitarian moral judgment in
1086 psychopathy. *Social Cognitive and Affective Neuroscience*, 7(6), 708–714.
- 1087 Kohlberg, L., & Hersh, R. H. (1977). Moral development: A review of the theory. *Theory into*
1088 *Practice*, 16(2), 53–59.
- 1089 Kohlberg, L., Levine, C., & Hower, A. (1983). *Moral stages: A current formulation and a*
1090 *response to critics*.
- 1091 Lee, J. J., & Gino, F. (2015). Poker-faced morality: Concealing emotions leads to utilitarian
1092 decision making. *Organizational Behavior and Human Decision Processes*, 126, 49–64.
- 1093 McKenna, M. 2013. *Conversation and Responsibility* (Oxford: Oxford University Press).
- 1094 Mikhail, J. (2011). *Elements of moral cognition: Rawls' linguistic analogy and the cognitive*
1095 *science of moral and legal judgment*. Cambridge University Press.
- 1096 Moll, J., & de Oliveira-Souza, R. (2007). Moral judgments, emotions and the utilitarian brain.
1097 *Trends in Cognitive Sciences*, 11(8), 319–321.
- 1098 Monin, B., Pizarro, D. A., & Beer, J. S. (2007). Deciding versus reacting: Conceptions of moral
1099 judgment and the reason-affect debate. *Review of General Psychology*, 11(2), 99–111.
- 1100 Monroe, A. E., Dillon, K. D., Guglielmo, S., & Baumeister, R. F. (2018). It's not what you do,
1101 but what everyone else does: On the role of descriptive norms and subjectivism in moral
1102 judgment. *Journal of Experimental Social Psychology*, 77(March), 1–10.
1103 <https://doi.org/10.1016/j.jesp.2018.03.010>
- 1104 Nosek, B. A., Beck, E. D., Campbell, L., Flake, J. K., Hardwicke, T. E., Mellor, D. T., van't
1105 Veer, A. E., & Vazire, S. (2019). Preregistration is hard, and worthwhile. *Trends in*
1106 *Cognitive Sciences*, 23(10), 815–818.
- 1107 Nozick, R. (1974). *Anarchy, state, and utopia* (Vol. 5038). New York: Basic Books.
- 1108 O'Hara, R. E., Sinnott-Armstrong, W., & Sinnott-Armstrong, N. A. (2010). Wording effects in
1109 moral judgments. *Judgment and Decision Making*.
- 1110 Piaget, J. (2013). *The moral judgment of the child*. Routledge.
- 1111 Pizarro, D., Uhlmann, E., & Salovey, P. (2003). Asymmetry in judgments of moral blame and
1112 praise: The role of perceived metadesires. *Psychological Science*, 14(3), 267–272.

- 1113 Rest, J. R. (1992). *Development in judging moral issues*. U of Minnesota Press.
- 1114 Rosen, B. (1980). Moral dilemmas and their treatment. *Moral Development, Moral Education,*
1115 *and Kohlberg. B. Munsey (Ed).*(1980), 232–263.
- 1116 Scanlon, T.M. 2008. *Moral Dimensions* (Cambridge, MA: Harvard University Press).
- 1117 Sidgwick, H. (2019). *The methods of ethics*. Good Press.
- 1118 R Core Team. (2019). *R: A language and environment for statistical computing*.
- 1119 Thomson, J. J. (1976). A defense of abortion. In *Biomedical ethics and the law* (pp. 39–54).
1120 Springer.
- 1121 Uhlmann, E. L., Pizarro, D. A., & Diermeier, D. (2015). A person-centered approach to moral
1122 judgment. *Perspectives on Psychological Science*, 10(1), 72–81.
- 1123 Young, I. M., (1985). Impartiality and the civic public: Some implications of feminist critiques of
1124 moral and political theory. *Praxis International*, 5(4), 381–401.
- 1125