1 Immutable morality: Even God could not change some moral facts

2 Madeline G. Reinecke

3 Yale University

4 Zachary Horne

5 Arizona State University

6 Author Note

7 This research was presented at the 2018 Annual Meeting of the Cognitive Science
8 Society, the 2016 Meeting of the Society for Philosophy and Psychology, and the 2016
9 Meeting of the Buffalo Experimental Philosophy Conference.

Abstract

10

11 The idea that morality depends on God is a widely held belief. This belief entails that the
12 moral “facts” could be otherwise because, in principle, God could change them. Yet, some
13 moral states-of-affairs seem so foundational (e.g., that it is immoral to kill someone just for
14 pleasure) that it is hard to imagine how they could be otherwise. In two experiments, we
15 investigated people’s intuitions about the immutability of moral facts. Participants judged
16 whether it was even possible, or possible for God, to change moral, logical, and physical
17 facts about the world. In both experiments, people judged that altering some moral facts
18 was impossible—not even God could turn morally wrong acts into morally right acts.
19 Strikingly, people thought that God could make physically impossible and logically
20 impossible events occur. These results demonstrate the strength of people’s metaethical
21 commitments, shedding light on the nature of morality and its centrality to thinking and
22 reasoning.

23

Keywords: morality, possibility, religion, philosophy

Immutable morality: Even God could not change some moral facts

Introduction

Morality is contentious. People disagree about many, if not most, ethical issues (e.g., Skitka, Bauman, & Sargis, 2005). Yet, some moral beliefs appear universal. For example, few people think that it is ever permissible to murder an innocent person just for fun; this kind of behavior must be morally wrong. What does this observation tell us about the strength of people's metaethical commitments? The spirit of this question dates back over 2000 years to a Platonic dialogue between Euthyphro and Socrates. There, Euthyphro famously asserted that what is "pious" is the same as "that which is loved by the gods" (Geach, 1966). As Socrates demonstrates, this entails that if God saw fit, God could change the moral facts—obvious moral facts need not be true. However, Socrates argues that this belief has seemingly untenable implications: God could change the facts in unthinkable ways. Even the most heinous transgressions could become morally virtuous just in case God decided those transgressions were morally virtuous. Here, we examine this dilemma empirically, investigating the extent to which people think certain moral facts are objective and immutable, unchangeable even by God.

Let us introduce some terminology before proceeding. We denote a fact as *weakly objective* just in case (i.e., iff) that fact is true, but could be otherwise. For example, the statement "Dogs are pets" is weakly objective, or contingently true, because we can imagine a world in which this statement is false. In contrast, we denote a fact as *strongly objective* just in case that fact is true and could not be otherwise. For example, mathematicians and philosophers typically think of statements like, "A triangle is a three-sided figure" or "Dogs are dogs" as strongly objective, or necessarily true (Benacerraf, 1973; Garson, 2018). Does this way of thinking about objectivity extend to morality?

The existing literature suggests that people may take an objectivist stance toward some moral statements, but it does not definitively speak to whether people tend to be strongly or weakly objectivist. Research on the psychology of metaethics has revealed that

moral statements are understood to be more objective than statements concerning matters of taste and nearly as objective as statements concerning scientific facts (e.g., Goodwin & Darley, 2008; Beebe & Sakris, 2016; but see Theriault, Waytz, Heiphetz, & Young, 2017). Consistent with the hypothesis that people may think some moral facts are objective, anthropological evidence indicates that some moral beliefs consistently emerge across cultures (e.g., Mikhail, 2011): People value protecting children from harm, maintaining friendships, and treating others fairly (although what constitutes “harm,” “friendship,” and “fairness” are culturally dependent; Fiske & Rai, 2014; Haidt, 2012; Miller, 2006). Some of these moral attitudes may even emerge early in development. Preverbal infants appear to have a rudimentary awareness of prosociality (Cowell & Decety, 2015; Hamlin, Wynn, Bloom, & Mahajan, 2011), indicating that some core moral knowledge may be innate rather than furnished by socialization alone (e.g., Hamlin, Wynn, & Bloom, 2007; Sloane, Baillargeon, & Premack, 2012; Tasimi & Wynn, 2016). These findings motivate the prediction that people may be predisposed to believe that at least some aspects of morality are objective, either in its weaker or stronger form.

Cross-cultural and developmental research indicates that people likely have some objectivist commitments, but the strength of these objectivist commitments has not yet been established. Evidence from research on the psychology of morality and religion builds on these findings, but more directly predicts that people are *only* weakly objectivist about morality. First, it is commonly believed that morality and religion are inextricably linked (e.g., Norenzayan & Shariff, 2008; Norenzayan, Shariff, Gervais, Willard, McNamara, et al., 2016). It is practically a platitude to say that “without God, everything is permitted” (Dostoevsky, 1991; Sinnott-Armstrong, 2009). Second, national polling indicates that people think that being a morally good person requires believing in God (Pew Research Center, 2017; Sinnott-Armstrong, 2009). Even atheists associate moral goodness with religiosity (Gervais et al., 2017). Third, many people believe that God is all-powerful, meaning that God “has the power to direct/change everything” (Pew Research, 2018).

These data suggest that people believe that (1) morality depends on God, and that (2) God has the power to change the moral facts. To deny either of these claims would conflict with strongly held religious and social convictions, formed by the association between morality and God. Taken together, these findings predict that people are “weakly” objectivist. People believe that moral facts are only contingently true, because God could change them.

To summarize, recent psychological research raises an untested question about the nature and strength of people’s metaethical commitments: Do people consider foundational moral facts weakly or strongly objective? Research from social, developmental, and cognitive psychology indicates that (1) people understand moral statements to be nearly as objective as scientific claims, (2) certain moral considerations are nearly omnipresent across cultures, and (3) some prosocial attitudes are early emerging. These results suggest that people are objectivists of some form. However, the nature of people’s religious convictions, and the strength of the association between religion and morality (Gervais et al., 2017), seems to dictate that people are only weakly objectivists. We tested this prediction in two preregistered experiments.

The Present Experiments

Over the last ten years, researchers have begun to investigate how people reason about the objectivity of morality (Goodwin & Darley, 2008; Theriault, Waytz, Heiphetz, & Young, 2017; Wright, Grandjean, & McWhite, 2013), which may shed light on the roots of human prosociality (Young & Durwin, 2013). Despite the theoretical and historical importance of research on the psychology of metaethics, it remains unknown whether people believe that fundamental aspects of morality could radically change. Furthermore, it remains unclear how beliefs about the mutability of morality may compare to intuitions concerning the immutability of societal conventions, logic, and the natural world. To address these open questions, we pitted people’s intuitions about obvious moral facts against the belief that morality depends on God, providing a strong test of people’s

objectivist commitments. Specifically, we examined whether people thought it was possible (or possible for God) to turn morally wrong acts into morally right acts, and societal conventions into morally wrong acts. In light of prior research indicating the centrality of religion in people’s moral lives (e.g., Fiske & Rai, 2014) and the common belief that God is all-powerful (Pew Research Center, 2018), we hypothesized that people would judge that God could make even foundational moral facts otherwise. We anticipated that people would preserve their religious and social commitments over their metaethical commitments.

Nevertheless, if people are reluctant to accept the implication that God could make some moral transgressions become morally virtuous, then this would suggest that people do have strong metaethical commitments after all. Not even an all-powerful agent, like God, could make some moral facts otherwise. This finding would be consistent with a stronger interpretation of the existing cross-cultural and developmental data. Early emerging, fundamental intuitions about moral principles may lead people to think that some moral facts must be true.

Experiment 1

Method

Preregistration. We preregistered the data collection plan and analytic syntax for this study. The raw data and code used to generate our analyses and figures are available on the Open Science Framework: <https://osf.io/jq9gz/>.

Participants. We determined our sample size based on a power analysis that indicated that approximately 350 participants would be needed in order to detect an Odds-Ratio of 1.33 with 80% power for a within-subjects condition effect. This effect size is equivalent to a Cohen’s d of .15, which is the modal effect size in social psychology (see Figure 1 in Richard, Bond, & Stokes-Zoota, 2003).

Participants in Experiment 1 were 347 Amazon Mechanical Turk workers (51% female, $M_{age} = 36$ years) who were paid \$0.60 for participating in the study. Twenty-six

participants were excluded for failing to correctly answer questions checking their attention. These exclusion rules were decided according to the data collection plan outlined in our registration. After excluding these participants, a sensitivity analysis indicated that the smallest effect we could detect with 80% power was a within-subjects Cohen's d of .157. This research was approved by the Arizona State University Institutional Review Board (IRB: STUDY00007785).

Materials and Procedure. In Experiment 1, we tested whether people think that some moral facts could be otherwise. Participants read ten moral propositions, judging for each whether that proposition could be changed. Half of the moral propositions concerned states-of-affairs where morally wrong acts become morally right acts (henceforth, denoted “morally wrong” items) and half of the moral propositions concerned states-of-affairs where societal conventions become morally wrong acts (denoted “societal convention” items). Participants also made judgments about propositions that did not concern morality (discussed in more detail below; see *Logical and physical propositions*). We presented participants with each proposition with one of two stems in a within-subjects design: “Is it possible that [proposition]?” (Bare possibility condition) or “Is it possible for God to make it that [proposition]?” (God possibility condition). For example, participants were asked, “Is it possible for it to be morally right to kill someone just for fun?” or “Is it possible for God to make it morally right to kill someone just for fun?” The presentation of these items was randomized and their assignment to a stem was counterbalanced.

Instructions. As discussed, participants made judgments about whether it was possible (Bare possibility) or possible for God (God possibility) to make a moral state-of-affairs otherwise. It is likely that most participants have never been asked questions of this kind and thus may find the question odd. With this in mind, we provided participants with instructions about how to understand what they were being asked. We stated that they should judge only whether an event *could* be the case, not whether it *would* be the case:

157 In this study, we will ask you to answer two kinds of questions:

158 First, we will ask you whether certain moral, physical, and logical events are
159 even possible.

160 Second, we will ask you whether it is even possible for God to bring about
161 certain moral, physical, and logical events.

162 We then provided further clarification about the questions:

163 We are **not** asking you to judge whether these events **would** actually occur.

164 We are **not** asking you whether God **would** actually make these events occur.

165 We are asking whether these events **could** occur, or whether God **could** make
166 them occur.

167 We took special care to emphasize that participants should not interpret our questions as
168 investigating what God would do. These directions were included because some people
169 may be unwilling to say that a morally perfect God would, rather than could, turn morally
170 wrong acts into morally right acts. Anticipating that some participants would be agnostics
171 or atheists, we also directed participants to suppose that God existed for the purposes of
172 our survey.

173 For the questions concerning God, we ask that you please suppose that God
174 exists whether you are religious or not. We understand that you might disagree
175 with this supposition, but it is just for the purposes of answering questions in
176 the survey. Please do your best to keep this in mind as you complete the study
177 even if you disagree with this supposition.

178 After finishing the study, we collected open-ended feedback from the participants. Their
179 responses indicated that they had no trouble understanding and adhering to this
180 supposition, even when they disagreed with it.

Logical and physical propositions. We compared participants' beliefs about the mutability of moral propositions with their beliefs about the mutability of physical and logical propositions. Physical and logical propositions, as we are defining them, concerned changes to the laws of nature or changes to the laws of logic (e.g., "Is it possible for God to make the moon shine brighter than the sun?" [God possibility condition, physical proposition]; "Is it possible for there to be a perfectly round square?" [Bare possibility condition, logical proposition]; see Table 3 of the Supplemental Online Materials). Including these additional propositions allowed us to determine whether people think morality is alterable relative to other kinds of states-of-affairs. More specifically, it allowed us to test (1) whether God—who people conceive of as an all-powerful agent—can make some propositions otherwise (effect of condition), (2) whether some propositions are more alterable than others (effect of proposition type), and (3) whether the effect of God on possibility judgments, so to speak, varies as a function of the proposition participants are considering (condition $\times$ proposition interaction). In particular, we were interested in the interaction between condition and proposition, as this would indicate whether people believe that God could alter moral facts more than logical facts. Many logical propositions (e.g., mathematical facts) are often considered to be paradigmatic examples of necessarily, objectively true propositions (see Benacerraf, 1973). Thus, comparing moral and logical propositions provided a strong test of people's objectivist metaethical commitments.

Religiosity Measures. We anticipated that people's responses in the God possibility condition would be affected by their religious beliefs. To test whether participants' religious views affected their responses, we collected two indicators of religiosity: First, we asked participants whether they believed in God (Yes or No). Second, we included a shortened version of a reliable measure of religiosity (Duke University Religiosity Index, or "DUREL"; Koenig & Büssing, 2010). This allowed us to detect whether differences in possibility judgments might be driven by engagement in religious practice. Both measures of religiosity were completed after the main task (see Table 4 of

SOM for scale items).

Results

Preregistered Analyses. As discussed, anthropological and psychological data suggest that people have some objectivist metaethical commitments. Common religious and social convictions, however, support the prediction that these commitments will be weakly objectivist. These bodies of research motivated three predictions: First, we predicted that participants would think that God could make false propositions true. Second, we predicted that, regardless of condition (i.e., Bare possibility vs. God possibility), people would judge that some propositions were more mutable than other propositions. Specifically, we predicted that people would think physical and moral propositions were more mutable than propositions about logical facts. Third, we predicted that these factors would interact: We anticipated that the effect of God would be uniquely strong for moral propositions, consistent with thinking that God is all-powerful, and that morality depends on God. This is to say that people would think God could change the moral facts more than the logical facts (see Benacerraf, 1973). This pattern of results would suggest that people are not strong moral objectivists; the moral facts are seen as only contingently true.

To test our predictions, we performed Bayesian logistic mixed-effects modeling using the R package `brms` (Bürkner, 2017). We regressed possibility judgments on Condition (Reference = Bare Condition), proposition type (Reference = Logical propositions), and the interaction between these predictors, along with a group-level effect that allowed the intercept of each subject to vary. For brevity, we specify the model in `brms` syntax:

```
Possibility Judgments ~ Condition*Proposition + (1|Subject)
```

Bayesian analyses formulate model parameters as probability distributions wherein the joint probability distribution of the data, y , and a given parameter θ is computed via the prior probability of θ and the probability of y given θ .

$$p(y, \theta) = p(y \mid \theta) \times p(\theta)$$

We can derive this result from Bayes Rule, which is used to calculate the posterior probability of θ given y :

$$p(\theta | y) \propto p(y | \theta) \times p(\theta) = p(y, \theta)$$

To model the joint probability distribution of participants' responses, we specified the following priors over the possible effects each parameter could have on the response variable. These priors are recommended because they regularize parameter estimates (e.g., Bürkner, 2017; McElreath, 2016).

$$\alpha \sim \text{Student-}t(3, 0, 10)$$

$$\text{All } \beta \sim \mathcal{N}(0, 1)$$

$$\sigma \sim \text{Half Student-}t(3, 0, 10)$$

Table 1

Bayesian Mixed-Effects Logistic Regression Predicting Possibility Judgments (Higher = More Possible) on the Basis of Condition (0 = Bare, 1 = God), Proposition (0 = Logical, Moral, Physical, Conventional), and the interaction of these predictors.

<i>Parameter</i>	<i>Estimate</i>	<i>Error</i>	<i>Lower</i>	<i>Upper</i>
Intercept	-1.86	0.11	-2.08	-1.64
Condition	1.84	0.11	1.64	2.05
Moral	0.05	0.12	-0.19	0.30
Conventional	1.27	0.12	1.06	1.50
Physical	0.07	0.12	-0.18	0.31
Condition × Moral	-0.81	0.16	-1.12	-0.49
Condition × Conventional	-0.93	0.15	-1.24	-0.63
Condition × Physical	0.82	0.16	0.50	1.13

Note: The Bare possibility condition and logical proposition type were the reference groups in this analysis. Consequently, “main effects” of condition and proposition are estimates of a given parameter in the Bare-Logical cell of the 2×4 design. Group-level effects are omitted for brevity. These estimates can be found on the Open Science Framework project page.

This model revealed that condition and proposition both predicted participants' possibility judgments (see Figure 1). People thought God could make some events obtain that they otherwise judged impossible (as indicated by judgments in the Bare possibility condition), and some propositions were judged to be more mutable than others. Furthermore, we

Figure 1. A marginal effects plot of possibility judgments by type of proposition and condition in Experiment 1. Error bars represent 95% credible intervals (see Cumming & Finch, 2005 for interpreting within-subjects 95% CIs). The figure indicates that the effect of God on possibility judgments is weaker for moral propositions (both societal conventions and morally wrong acts) than for logical or physical propositions.

observed an interaction between these predictors (see Table 1). However, contrary to our predictions, people were *more likely* to say that God could change the logical and physical facts than the moral facts. That is, the effect of God on logical and physical propositions was larger than the effect of God on moral propositions. In contrast, regardless of condition, people believed that societal conventions were more mutable than “morally wrong” propositions. This finding is consistent with the hypothesis that norms acquired through socialization (e.g., societal conventions) may be more easily viewed as alterable compared to norms that may not be the product of socialization alone (e.g., foundational moral norms). We return to this possibility in the general discussion.

The results of Experiment 1 provide preliminary support for the hypothesis that people are strong objectivists about some moral facts. In the Bare possibility condition, we found that morally wrong propositions were judged to be as immutable as both physical and logical propositions. Strikingly, in the God condition, participants judged that moral facts were less mutable than even mathematical or physical facts.

Exploratory Analyses. We predicted that people would judge that God could make things possible that they would otherwise have thought impossible, that some propositions would be judged less alterable than others, and that these effects would interact. We observed all three of these effects. However, contrary to our hypothesis, people viewed moral propositions as less alterable than logical or physical facts—evidence that people perceive some moral facts as strongly objective.

Religiosity is likely to impact people’s beliefs about morality, possibility, and the scope of God’s omnipotence. For instance, it is possible that people who are more religious will attribute more power to God. To examine the possible interactions between religiosity and the effects observed, we conducted a series of exploratory analyses. Following the recommendations of McElreath (2016, p. 165 – 203), we fit every possible model containing condition, proposition, religiosity, and the interactions between these predictors. We put regularizing priors on the population-level effects of the model to reduce the likelihood of overfitting. We then assessed the predicted out-of-sample performance of each of these models based on its leave-one-out cross-validation information criterion (loo-ic)¹, which is calculated from the posterior samples of a model using Pareto-smoothed importance sampling (Vehtari, Gelman, & Gabry, 2017). This procedure revealed that the most

¹ Cross-validation is a statistical technique for assessing how well a model fits unobserved data. Cross-validation can proceed in several ways (see Arlot & Celisse, 2010 for a review) but broadly all cross-validation techniques involve (1) training a model on a dataset and (2) testing how well the model fits data the model has not “seen.” As noted, the analyses we report below relied on assessing information criterion scores of an approximation of leave-one-out cross-validation (loo-ic). Loo-ic has been shown to asymptotically match the Watanabe-Akaike information criterion (i.e., WAIC; Watanabe, 2010), a common information criterion used to assess the fit of Bayesian mixed-effects models. However, we assessed model performance based on loo-ic rather than WAIC in light of recent work demonstrating it more accurately predicts out-of-sample performance in many real-world cases (Vehtari, Gelman, & Gabry, 2017).

279 predictive model includes the interaction between condition and proposition, along with the
 280 interaction between religiosity and proposition. The inclusion of the interaction between
 281 religiosity and proposition substantially improved the predicted out-of-sample performance
 282 of our model ($\text{loo-ic} = 7099.0$) compared to a less flexible model which only contained the
 283 main effects and interaction of condition and proposition ($\text{loo-ic} = 7198.1$; difference in
 284 expected-lpd, $\Delta\text{elpd} = -49.5$, $SE = 9.8$; see Table 5 in the SOM for complete details). For
 285 brevity, we focus our discussion on interpreting this model, which is displayed in Figure 2.

Figure 2. Marginal effects spaghetti plots of possibility judgments by religious beliefs and proposition in Experiment 1. Each line represents a random draw from the posterior distribution.

286 Figure 2 highlights two consistent relationships between religiosity and possibility
 287 judgments: First, regardless of condition, religious participants are more willing to say the
 288 physical and logical facts could be otherwise (i.e., the red and blue regression lines have
 289 positive slopes in both the left and right panels of Figure 2). Second, regardless of
 290 condition, religious participants are slightly less likely to say that morally wrong acts could
 291 be turned into morally right acts (i.e., the green regression lines have negative slopes in
 292 both the left and right panels of Figure 2). These results suggest that even though people's
 293 religious beliefs impact possibility judgments—to our surprise—this pattern seems

unaffected by condition (i.e., God vs. Bare possibility). Neither a model that included the interaction between religiosity with condition ($\Delta\text{elpd} = -1.1$, $SE = 2.3$) nor a model testing whether religiosity moderated the relationship between proposition type and condition (i.e., the three-way interaction between condition, proposition type, and religiosity; $\Delta\text{elpd} = .7$, $SE = 2.1$) performed better than a simpler model, containing only the interaction between proposition and condition and religiosity and proposition.

To summarize, Experiment 1 sheds new light on people’s beliefs about the necessity of some moral facts. We found that people think moral propositions are less alterable than logical or physical facts—even God could not turn some morally wrong acts into morally right acts. Thus, Experiment 1 provides evidence for the claim that people may be strong objectivists about at least some foundational moral facts. We sought to confirm these results in Experiment 2 by conducting a direct replication.

Experiment 2

Method

Participants. Participants in Experiment 2 were 350 Amazon Mechanical Turk workers (50% female, $M_{age} = 37$ years old) who were paid \$0.60 for participating in the study. Our sample size was determined via an identical power analysis to Experiment 1. We excluded 56 participants for failing to correctly answer questions checking their attention: <https://osf.io/8hpm9/>. After excluding these participants, a sensitivity analysis indicated that our design could detect a within-subjects Cohen’s d of .165 with 80% power. This research was approved by the Arizona State University Institutional Review Board (IRB: STUDY00007785).

Procedure. The materials, procedures, sample size, and analytic strategy were identical to Experiment 1. After we completed data collection in Experiment 2, we collected additional norming data, which we detail in the Exploratory Analyses section below. We again set the same regularizing priors for all population-level effects in our models.

320 **Results**

Figure 3. A marginal effects plot of possibility judgments by type of proposition and condition in Experiment 2. Error bars represent 95% Credible Intervals. As in Experiment 1, the figure in the left panel indicates that the effect of God on possibility judgments is weakened for moral propositions (both for societal conventions and morally wrong acts) relative to logical and physical propositions. A more conservative model (right panel) provides stronger evidence for the moral $\times$ condition interaction.

321 Experiment 2 directly replicated the findings of Experiment 1. We found that moral
 322 propositions were viewed as less mutable than either logical or physical propositions, as
 323 indicated by the weaker effect of God on possibility judgments for morally wrong
 324 propositions (Figure 3; also see Table 6 in the SOM). As in Experiment 1, we also found
 325 that participants who were more religious thought logical and physical facts were more
 326 mutable. However, there was increased variability in the relationship between religiosity
 327 and possibility judgments for moral propositions. In Experiment 2, religious participants
 328 were as likely as non-religious participants to think that God could turn morally wrong
 329 acts into morally right acts, and slightly more likely to think that God could turn societal
 330 conventions into morally wrong acts (i.e., the green lines are flat rather than negatively

Figure 4. Marginal effects spaghetti plots of possibility judgments by religious beliefs and proposition in Experiment 2. Each line represents a random draw from the posterior distribution.

sloped; see Figure 4). We did not observe this pattern in Experiment 1. Consequently, the relationship between religiosity and beliefs about the mutability of moral propositions in particular should be interpreted with some caution.

Exploratory Analyses

Alternative models. We directly replicated the central findings from Experiment 1. However, one limitation of the present analyses is that they do not model the maximal group-level (i.e., “random”) effect structure dictated by our design (Barr, Levy, Scheepers, & Tily, 2013). Specifically, they do not model heterogeneity across participants in the interaction between proposition and condition. Consequently, this model does not represent the most conservative model possible for testing our hypotheses. To further interrogate the effect of interest, we reran our primary analysis, but allowed for the slope of proposition, condition, and the interaction between these treatments to vary by participant. Because participants may differ in how they interpret the main task, this model better reflects the underlying data-generating model (McElreath, 2016).

Possibility Judgments $\sim$ Condition*Proposition + (1 +
Condition*Proposition|Subject)

A marginal effects plot of this analysis is shown in the right panel of Figure 3. Directly modeling the maximal group-level structure suggests two primary conclusions. First, the uncertainty around every estimated cell mean is larger, which is typical of models that account for variability in the effect of a treatment (again see Barr et al., 2013). Second, and importantly, the effect of God on moral propositions—that is, the main effect of interest—is *smaller*, whereas the effect of God on logical and physical propositions does not materially differ. This analysis provides converging evidence for our primary finding. People view fundamental moral facts as less mutable than either logical or physical facts (see Table 7 in the SOM).

Another concern about the preceding studies is that our materials may have unforeseen item-level differences that could account for the effects we observed. These item-level differences could produce patterns of judgments that we mistook as effects of proposition type, for example, when in fact the difference was simply due to variability in individual propositions (however, see Figure 6 in the SOM). For instance, it is possible that people are more *confident* in the truth of some propositions compared to others (e.g., being more confident that killing for fun is wrong than being confident that people breathe CO₂). These differences might make people more likely to think that it is possible to change the truth value of a proposition, not because of how people conceive of possibility for moral and physical propositions, but merely because of their knowledge about the truth of the propositions in question.

To test this possibility, we obtained norming data for all of the items used in Experiments 1 and 2. We then reran the analyses from Experiment 2 using these norming ratings as predictors in a series of exploratory analyses relying using the same procedure described in Experiment 1 (again see McElreath, 2016 for recommendations for conducting exploratory data analysis). We describe this norming study and our analyses below.

Norming Study Analyses. We collected additional norming data from 148 participants ($M_{age} = 36$, 45% female) who were tasked with providing three norming ratings about all of the items used in Experiments 1 and 2. Specifically, we collected data measuring (1) how confident participants were that the propositions we used in the study were true (e.g., *Confidence*: “How confident are you that people breathe CO₂, not Freon gas?”), (2) the perceived change that resulted from the truth-value of a proposition being changed (e.g., *Perceived Change*: “Assume that fire requires oxygen to burn. Now try to imagine that our world changes in a single way: It is now the case that fire does not require oxygen to burn. How much do you think that the world has meaningfully changed?”), and (3) people’s beliefs about the “severity” of moral and societal conventions (e.g., *Severity*: “How severe of a moral violation is drinking orange juice for breakfast?”). Participants made their ratings using a ten-point Likert scale (see Table 8 in the SOM for complete details). Confidence and severity items were provided first in a counterbalanced order, followed by perceived change items. Participants performed the perceived change task last because these questions required participants to assume that certain propositions were true, which may have affected their confidence and severity evaluations. We included several additional items in the norming study to confirm participants were performing the task as instructed. For example, we confirmed that participants were not always selecting the endpoint of the scale by including several propositions that participants should be less confident about (e.g., “How confident are you that dark matter composes most of the known universe?”). Check items were created for all four types of propositions and all of these check items behaved as expected. The results from this norming study are summarized in Table 2.

The results of the norming study provide several further insights about the connection between possibility judgments, proposition type, and condition. First, we observed that participants were extremely confident in the truth of the propositions used in both experiments. For instance, when participants were asked “How confident are you that

Table 2

Median item-level norming ratings and possibility judgment averages from Experiment 2.

<i>Item</i>	<i>Confidence</i>	<i>Perceived Change</i>	<i>Severity</i>	<i>Bare</i>	<i>God</i>	<i>Difference</i>
<i>Conventional</i>						
Driving	9	6	1	.49	.54	.05
Juice	10	6	1	.41	.47	.05
Showering	10	7	1	.41	.52	.11
Spoon	10	6	1	.36	.55	.18
Socks	10	5	1	.30	.51	.21
<i>Moral</i>						
Lying	9	8	7	.42	.47	.05
Rape	10	10	10	.18	.28	.10
Child Abuse	10	10	10	.15	.30	.15
Steal from poor	10	9	10	.30	.45	.15
Kill for fun	10	10	10	.15	.31	.16
<i>Physical</i>						
Steam	10	8	—	.23	.52	.19
Moon	10	8	—	.31	.68	.37
CO ₂	10	9	—	.14	.58	.44
Fire	10	8	—	.19	.64	.45
Animal	10	9	—	.15	.64	.49
<i>Logical</i>						
Line	7	7	—	.54	.66	.12
Color	8	7	—	.34	.55	.21
Add	10	8	—	.07	.30	.23
Square	10	7	—	.14	.38	.24
House	10	7	—	.26	.51	.25
Ounce	10	8	—	.14	.45	.31
Triangle	10	7	—	.04	.36	.32

Note: We computed medians for confidence, perceived change, and severity ratings because of the extreme left skew of the data. Additional analyses using means rather than medians yielded identical conclusions. The difference column indicates the mean difference between God and Bare conditions. Complete item descriptions can be found in Table 3 of the SOM.

people breathe CO₂, not Freon gas?”, the median response for this item was 10 out of 10.

Second, we observed that participants assigned different weights to the amount of change a given proposition would cause if it were no longer true. For example, making an animal run faster than the speed of light was perceived to result in a large change to the world (9 out of 10). In contrast, people thought it would be a less substantial change if it became morally wrong to wear socks to bed. This implies that intuitions about possibility and

perceived change can diverge: Although people tended to believe that God *could* change physical facts, for example, they also believed that changing these facts would lead to large changes in the world (see Figure 6 in the SOM). Third, as we would expect, societal conventions were rated to be considerably less severe than fundamental moral propositions.

We also explored whether model fit improved after including confidence and perceived change ratings in our regression model. First, confidence and perceived change ratings for a given proposition were, in fact, predictive of participants' judgments about whether it was possible to change the truth value of a given proposition. These predictors improved model fit over an intercept-only model, $\text{loo-ic} = 7060.6$, $\Delta\text{elpd} = -124.0$, $SE = 15.7$. Importantly, even after accounting for item-level variability in confidence and perceived change, adding proposition, condition, and their interaction to the model substantially improved model fit, $\text{loo-ic} = 6324.6$, $\Delta\text{elpd} = -368.0$, $SE = 26.4$ (see Table 9 in the SOM). This finding suggests that the interaction we observed between condition and proposition is robust.²

Discussion

Over 2000 years ago, Euthyphro claimed that God determines the moral facts. We predicted that this belief would reflect people's weak objectivist commitments—the moral facts could be otherwise. Contrary to this prediction, people seem to believe that some foundational moral propositions must be true: Regardless of religiosity, participants thought that even God could not make these moral facts otherwise. We found that people were more likely to judge that God could make logical contradictions true, or physical impossibilities occur, than that God could change foundational moral facts.

Many ethicists have argued that we can come to know what is morally right in the same way that we can come to know facts about mathematics and logic (see Clarke-Doane,

² As before, we explored the entire model space by interacting norming ratings with proposition and condition. This set of exploratory analyses indicated that the addition of the interaction between condition and perceived change, and condition and confidence, improved the predictive performance of the model. Namely, this model revealed that the effect of confidence on possibility judgments was weaker in the God condition compared to the Bare condition, $\text{loo-ic} = 6296.2$, $\Delta\text{elpd} = -14.2$, $SE = 5.2$; see Figure 6 in the SOM. Interacting proposition type with norming ratings, in contrast, did not meaningfully improve out-of-sample model performance, $\Delta\text{elpd} = -1.9$, $SE = 2.0$.

2014 for a review). This argument is known as the Ethics-Mathematics Analogy. The motivations for this argument are many, but, crudely put, their aim is to demonstrate that objective moral facts could be uncovered through reasoning alone, much in the way objective mathematical facts have been discovered (e.g., Locke, 1841). A further implication of this analogy, or so it is argued, is that some moral facts are necessarily true—much like logic or mathematical facts, some aspects of morality are not contingent. Our findings appear to turn this inference on its head. People believe that some moral facts are more likely to be necessarily true than are logical facts (again, see Clarke-Doane, 2014). Not even an all-powerful entity, like God, could change them. This is the first empirical evidence demonstrating that people believe that certain foundational moral facts are more objective—in the sense that they could not be otherwise—than either physical or logical propositions.

The present research is a step towards understanding how beliefs about the mutability of morality compare to beliefs about the mutability of logical and physical states-of-affairs. This work highlights several directions for future research. One open question concerns the boundary conditions of the finding that some fundamental moral facts are judged to be less mutable than logical or physical facts. Although people appear to consider some foundational moral propositions as necessarily true, it is important to exercise some care in interpreting the scope of this claim. We selected particular moral facts *because* we considered them to be some of the most fundamental and uncontroversially true. This provided a strong test of people's metaethical commitments. One possibility is that other, less obvious moral facts (e.g., sacrificing an innocent person to save a group of people, euthanasizing someone to prevent pain, and so on) would be perceived as more mutable than the moral facts examined here. We find some preliminary evidence to support this possibility: People do not seem to believe that all changes to the structure of morality are impossible. For example, in both experiments, people were more likely to think that societal conventions could become morally wrong than to think that

logical states-of-affairs could change. One interpretation of this finding is that some aspects of morality are perceived as being more mutable than others. Future work should investigate the dimensions along which foundational and non-foundational moral propositions differ. Even acknowledging these potential boundary conditions, however, we note that the moral propositions that participants considered in the present experiments are, without question, central to ethical theory, law, and social relationships. Consequently, understanding that these fundamental moral propositions are perceived as immutable provides new insights into the relationship between moral and modal cognition.

Our findings also lay the groundwork for further investigation into the question of whether conceiving of some moral facts as immutable reflects a cultural (i.e., Western) tendency rather than revealing something about the strength of people's moral commitments. The idea that God is perfectly good is widespread in Western religious traditions, but may be less common cross-culturally (see Wainwright, 2012 for a review). Our findings may owe to the fact that Westerners, specifically, are willing to narrow the scope of God's power to preserve God's omnibenevolence, because moral perfection may be more central to their concept of God than omnipotence is. This argument would suggest that our findings are the consequence of a cultural mechanism (e.g., socialization in Western society), rather than revealing something about the nature of people's moral convictions.

Two points are worth considering with respect to this cultural account: First, if our findings only reflected Western socialization, then religiosity should have reliably interacted with the effects of condition and proposition type. It seems reasonable to assume that participants who are more religious would internalize the maxims typical of Western religion to a greater extent than would less religious participants. If so, then religious participants should have more readily sought to preserve, so to speak, God's omnibenevolence; religiosity should have tracked with the tendency to deny that God could change the truth value of fundamental moral propositions. Yet, we observed no consistent effect of religiosity compatible with this explanation. Religious and non-religious

participants alike denied that God could change at least some foundational moral propositions. Furthermore, even if theists and non-theists internalized Western religious commitments to the same degree, these groups would presumably still differ in evaluating God's ability to affect fundamental moral states-of-affairs. Theists may preserve God's perfect moral goodness (i.e., by denying that God *would* ever change morality) over preserving God's powers, consistent with our present data. Non-theists, however, have no reason to preserve God's moral goodness over any other attribute. Because we found no relationship between religiosity and reasoning about the mutability of morality, we think it is implausible that exposure to Western religious traditions alone accounts for the present effects. Still, future work could more directly test this hypothesis by conducting research with non-Western participants.

Second, there may be some concern about the validity of our religiosity measures. One might believe that religious socialization is still likely to play a central role in our findings, but we failed to accurately measure people's religious commitments in a way that would allow us to detect this effect. Although it is possible that other religiosity measures would yield a different pattern of results, our primary measure of religiosity, the Duke University Religiosity Index, has been validated and widely employed as a measure of religious practice (Koenig & Büssing, 2010). In fact, we found that including this religiosity measure improved model fit in both Experiments 1 and 2, and we consistently observed that religiosity predicted possibility judgments about logical and physical propositions. This is to say that, although we did not observe a relationship between religiosity and beliefs concerning the malleability of morality, we did observe that religious participants were more likely to judge that God could change the logical and physical propositions. This suggests that religious socialization may affect some aspects of modal cognition, but this does not seem to extend to the moral domain. Independent of cultural norms and religious socialization, people may still believe that at least some moral states-of-affairs could not be otherwise.

In contrast to this cultural account, we propose that our findings reflect a general tendency for people's normative commitments to affect non-normative cognition (e.g., beliefs about possibility). In the last decade, researchers have revealed that morality plays a more central role in non-normative cognition than was once assumed (see Knobe, 2010 for a review). Moral attitudes can affect reasoning across a host of purportedly non-moral domains, including judgments of causality (Hitchcock & Knobe, 2009; Knobe & Fraser, 2008), intentionality (Cushman & Mele, 2008; Knobe, 2003), and the self (Strohming & Nichols, 2014; Strohming, Knobe, & Newman, 2017). Indeed, normative considerations have been shown to influence how people reason about what is possible (Phillips & Cushman, 2017; Phillips & Knobe, 2018) and prevent them from imagining that certain states-of-affairs could occur (e.g., Gendler, 2000; Kim, Kneer, & Stuart, in press). Our findings are consistent with this emerging research, highlighting the strong connection between morality and cognition (in this case, modal cognition). This connection may be even more deeply rooted than previously considered: People's moral commitments may be stronger even than their commitments to the truth and necessity of logical or physical facts. In part, this may be because it is unclear to laypeople what the world would be like if fire could burn without oxygen, or if $2 + 2$ could equal 5. Yet, it is both horrifying and salient what the world would be like if abusing a child, stealing from the poor, or killing for fun were morally right acts.

The differential implications of these imagined states-of-affairs may also coincide with the early emergence of basic moral attitudes, such as those regarding fairness and prosociality (e.g., Hamlin et al., 2007; Sloane et al., 2012; Tasimi & Wynn, 2016). To speculate, some moral intuitions may be so deeply embedded and reinforced in everyday social interactions that people are incapable of imagining that they could be otherwise (Gendler, 2000; Kim, Kneer, & Stuart, in press). Recent research suggests that this may be the case: Even preschoolers seem to link immorality with impossibility (Shtulman & Phillips, 2018). These children may also think that morality could not be otherwise. As

our data suggest, this may even be independent of children's early religious socialization (e.g., Nucci & Turiel, 1993; Nucci, 2001). Additional research—across cultures and across development—will be necessary to fully understand the nature and scope of people's objectivist commitments. It may be that the influence of normativity on non-normative cognition is innate and universal (Knobe & Burra, 2006; Leslie, Knobe, & Cohen, 2006). The present data provide some support for this possibility.

Two experiments reveal that people think that some foundational moral facts are immutable. Surprisingly, people appear to think that some aspects of morality are less alterable than even facts related to domains like logic or physics. These findings shed new light on the strength of people's metaethical commitments and, more generally, the centrality of morality to thinking and reasoning.

References

- Arlot, S., & Celisse, A. (2010). A survey of cross validation procedures for model selection. *Statistics Surveys*, 4, 40-79. doi: 10.1214/09-SS054
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3). doi: 10.1016/j.jml.2012.11.001
- Beebe, J. R., & Sackris, D. (2016). Moral Objectivism Across the Lifespan. *Philosophical Psychology*, 29(6), 912-929. doi: 10.1080/09515089.2016.1174843
- Benacerraf, P. (1973). Mathematical Truth. *The Journal of Philosophy*, 70(19), 661-679. doi: 10.2307/2025075
- Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1), 1-28. doi: 10.18637/jss.v080.i01
- Clarke-Doane, J. (2014). Moral Epistemology: The Mathematics Analogy. *Noûs*, 48(2), 238-255. doi: 10.1111/j.1468-0068.2012.00875.x
- Cowell, J. M., & Decety, J. (2015). Precursors to morality in development as a complex interplay between neural, socioenvironmental, and behavioral facets. *Proceedings of the National Academy of Sciences*, 112(41), 12657-12662. doi: 10.1073/pnas.1508832112
- Cumming, G., & Finch, S. (2005). Inference by eye: confidence intervals and how to read pictures of data. *American Psychologist*, 60, 170.
- Cushman, F., & Mele, A. (2008). Intentional action: Two-and-a-half folk concepts? In *Experimental Philosophy* (pp. 171-188). New York, NY: Oxford University Press.
- Dostoevsky, F. (1991). *The Brothers Karamazov: A novel in four parts with epilogue* (R. Pevear & L. Volokhonsky, Trans.). New York, NY: New York: Vintage Books.
- Fiske, A. P., & Rai, T. S. (2014). *Virtuous Violence: Hurting and Killing to Create, Sustain, End, and Honor Social Relationships*. Cambridge, England: Cambridge University Press. doi: 10.1017/CBO9781316104668

- Garson, J. (2018). *Modal logic*. Retrieved from
<https://plato.stanford.edu/entries/logic-modal/>
- Geach, P. T. (1966). Plato's Euthyphro: An Analysis and Commentary. *The Monist*, 50(3), 369–382. doi: 10.5840/monist196650327
- Gendler, T. S. (2000). The Puzzle of Imaginative Resistance. *The Journal of Philosophy*, 97(2), 55–81. doi: 10.2307/2678446
- Gervais, W. M., Xygalatas, D., McKay, R., Van Elk, M., Buchtel, E. E., Aveyard, M., . . . Klocova, E. (2017). The Cultural Evolution of Prosocial Religions. *Nature Human Behavior*, 1(8), 0151.
- Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106(3), 1339–1366. doi: 10.1016/j.cognition.2007.06.007
- Haidt, J. (2012). *The Righteous Mind: Why Good People are Divided by Politics and Religion*. New York, NY: Pantheon/Random House.
- Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, 450(7169), 557–559. doi: 10.1038/nature06288
- Hamlin, J. K., Wynn, K., Bloom, P., & Mahajan, N. (2011). How infants and toddlers react to antisocial others. *Proceedings of the National Academy of Sciences*, 108(50), 19931–19936. doi: 10.1073/pnas.1110306108
- Hitchcock, C., & Knobe, J. (2009). Cause and Norm. *Journal of Philosophy*, 106(11), 587–612. doi: 10.5840/jphil20091061128
- Kim, H., Kneer, M., & Stuart, M. T. (in press). The Content-Dependence of Imaginative Resistance. In F. Cova & S. Réhault (Eds.), *Advances in Experimental Philosophy of Aesthetics*. Bloomsbury.
- Knobe, J. (2003). Intentional Action and Side Effects in Ordinary Language. *Analysis*, 190–194.
- Knobe, J. (2010). Person as scientist, person as moralist. *The Behavioral and Brain Sciences*, 33(4). doi: 10.1017/S0140525X10000907

- Knobe, J., & Fraser, B. (2008). Causal Judgment and Moral Judgment: Two Experiments. In W. Sinnott-Armstrong (Ed.), *Moral Psychology*. MIT Press.
- Koenig, H. G., & Büssing, A. (2010). The Duke University Religion Index (DUREL): A Five-Item Measure for Use in Epidemiological Studies. *Religions*, 1(1), 78–85. doi: 10.3390/rel1010078
- Leslie, A. M., Knobe, J., & Cohen, A. (2006). Acting intentionally and the side-effect effect. *Psychological Science*, 17(5), 421–427. doi: 10.1111/j.1467-9280.2006.01722.x
- Locke, J. (1841). *An Essay Concerning Human Understanding*. New York, NY: Oxford University Press.
- McElreath, R. (2016). *Statistical Rethinking*. Boca Raton, FL: CRC Press.
- Mikhail, J. (2011). *Elements of moral cognition: Rawls' linguistic analogy and the cognitive science of moral and legal judgment*. New York, NY: Cambridge University Press. doi: 10.1017/CBO9780511780578
- Miller, J. G. (2006). Insights into moral development from cultural psychology. In *Handbook of moral development* (pp. 375–398). Mahwah, NJ: Lawrence Erlbaum Associates Publishers.
- Norenzayan, A., & Shariff, A. F. (2008). The Origin and Evolution of Religious Prosociality. *Science*, 322(5898), 58–62.
- Norenzayan, A., Shariff, A. F., Gervais, W. M., Willard, A. K., McNamara, R. R., Slingerland, E., & Henrich, J. (2016). The Cultural Evolution of Prosocial Religions. *Behavioral and Brain Sciences*, 39, 1–65. doi: 10.1017/S0140525X14001356
- Nucci, L., & Turiel, E. (1993). God's word, religious rules, and their relation to Christian and Jewish children's concepts of morality. *Child Development*, 64(5), 1475–1491. doi: 10.2307/1131547
- Nucci, L. P. (2001). *Education in the moral domain*. Cambridge, England: Cambridge University Press.
- Phillips, J., & Cushman, F. (2017). Morality constrains the default representation of what

is possible. *Proceedings of the National Academy of Sciences*, 114(18), 4649–4654.

doi: 10.1073/pnas.1110306108

Phillips, J., & Knobe, J. (2018). The psychological representation of modality. *Mind & Language*, 33, 65–95. doi: 10.1111/mila.12165

Pew Research Center. (2017). *A growing share of americans say it's not necessary to believe in god to be moral*. Retrieved from

<https://www.pewresearch.org/fact-tank/2017/10/16/a-growing-share-of-americans-say-its-not-necessary-to-believe-in-god-to-be-moral/>

Pew Research Center. (2018). *Americans' beliefs about the nature of god*. Retrieved from

<https://www.pewforum.org/2018/04/25/when-americans-say-they-believe-in-god-what-do-they-mean/>

Richard, F. D., Bond Jr., C. F., & Stokes-Zoota, J. J. (2003). One Hundred Years of Social Psychology. *Review of General Psychology*, 7(4), 331 - 363.

Shtulman, A., & Phillips, J. (2018). Differentiating “could” from “should”: Developmental changes in modal cognition. *Journal of Experimental Child Psychology*, 165, 161–182.

doi: 10.1016/j.jecp.2017.05.012

Sinnott-Armstrong, W. (2009). *Morality Without God?* New York, NY: Oxford University Press.

Skitka, L. J., Bauman, C. W., & Sargis, E. G. (2005). Moral Conviction: Another Contributor to Attitude Strength or Something More? *Journal of Personality and Social Psychology*, 88(6), 895–917. doi: 10.1037/0022-3514.88.6.895

Sloane, S., Baillargeon, R., & Premack, D. (2012). Do infants have a sense of fairness? *Psychological Science*, 23(2), 196–204. doi: 10.1177/0956797611422072

Strohming, N., Knobe, J., & Newman, G. (2017). The True Self: A Psychological Concept Distinct From the Self. *Perspectives on Psychological Science*, 12(4), 551–560. doi: 10.1177/1745691616689495

Strohming, N., & Nichols, S. (2014). The essential moral self. *Cognition*, 131(1),

- 159–171. doi: 10.1016/j.cognition.2013.12.005
- Tasimi, A., & Wynn, K. (2016). Costly rejection of wrongdoers by infants and children. *Cognition*, 151, 76–79. doi: 10.1016/j.cognition.2016.03.004
- Theriault, J., Waytz, A., Heiphetz, L., & Young, L. (2017). Examining overlap in behavioral and neural representations of morals, facts, and preferences. *Journal of Experimental Psychology: General*, 146(11), 1586–1605. doi: 10.1037/xge0000350
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5), 1413–1432. doi: 10.1007/s11222-016-9696-4
- Wainwright, W. (2012). Concepts of God. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Watanabe, S. (2010). Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11, 3571–3594.
- Wright, J., Grandjean, P., & McWhite, C. B. (2013). The meta-ethical grounding of our moral beliefs: Evidence for meta-ethical pluralism. *Philosophical Psychology*, 26(3), 336–361. doi: 10.1080/09515089.2011.633751
- Young, L., & Durwin, A. J. (2013). Moral realism as moral motivation: The impact of meta-ethics on everyday decision-making. *Journal of Experimental Social Psychology*, 49, 302–306. doi: 10.1016/j.jesp.2012.11.013