

Anti-Realist Pluralism: A New Approach to Folk Metaethics

ABSTRACT: Many metaethicists agree that as ordinary people experience morality as a realm of objective truths, we have *prima facie* to believe that it actually is such a realm. Recently, worries have been raised about the validity of the extant psychological research on this argument's empirical hypothesis. Our aim is to advance this research, taking these worries into account. First, we propose a new experimental design for measuring folk intuitions about moral objectivity that may serve as a model for future studies. Then we report and discuss the results of a survey that was based on this design. In our study, most of our participants denied the existence of objective moral truths. In particular, many of them had the intuition that whether moral sentences are true depends both on their own moral beliefs and on the dominant moral beliefs within their culture ("anti-realist pluralism"). This finding suggests that the realist presumptive argument must be rejected and that instead anti-realism may have a presumption in its favor.

KEY WORDS: moral realism; moral psychology; metaethics; experimental philosophy; moral objectivity; metaethical pluralism

Introduction

Are there things that are morally right or wrong, good or bad, virtuous or vicious, and so on independently of what anybody thinks about them? Moral realists believe that such objective moral truths exist (e.g., Brink 1989; Huemer 2005; Shafer-Landau 2003). Anti-realists, in contrast, deny them. As they see it, moral sentences are either not truth-apt at all (non-cognitivism; Ayer 1952; Blackburn 2000); are all false (error theory; Joyce 2001; Mackie 2011); or are, if true, only subjectively true (subjectivism; Harman 1996; Prinz 2007).¹

Several metaethical arguments appeal to what ordinary people think about the existence of objective moral truths. Most important among them is the widely held "presumptive argument" (e.g., Brink 1989; Enoch 2017a, 2017b; Dancy 1986; Huemer 2005; McNaughton 1988). According to this argument, it seems to ordinary people that morality is a realm of ob-

¹ Some contemporary non-cognitivists argue that moral sentences are truth-apt in a deflationary sense (e.g., Blackburn 2000: 79). We will come back to this issue in Sec. 1.2 and 2.4 below. Moreover, error theory has often been restricted to particular kinds of moral sentences (such as positive ones) (e.g., Joyce 2007; Pigden 2007: 451), and has sometimes been stated in terms of these sentences being "untrue" rather than false (Joyce 2001: 6-9, 2007). For a plausible defense of my above formulation see Loeb 2007a.

jective truths; therefore we have a *prima facie* reason to believe that it actually is such a realm.² McNaughton, for example, writes:

The realist insists on an obvious, but crucial, methodological point: there is a presumption that things are the way we experience them as being [...]. Moral value is presented to us as something independent of our beliefs or feelings about it; something which may require careful thought or attention to be discovered. There is a presumption, therefore, that there is a moral reality to which we can be genuinely sensitive. (McNaughton 1988: 40)

Despite the presumptive argument's importance, philosophers have so far not provided any serious empirical evidence for or against its underlying hypothesis that ordinary people experience morality as a realm of objective truths. At first sight recent research in psychology promises to remedy for this shortcoming.

In the last 15 years an increasing number of psychologists have begun to study the content, correlates, and causes of folk intuitions about the existence of objective moral truths. Their results suggest that rather than being realists, ordinary people intuitively tend towards "metaethical pluralism" (Wright et al. 2013; Wright 2018): they regard moral realism as true with regard to some moral sentences or circumstances and anti-realism as true with regard to other moral sentences or circumstances (e.g., Beebe 2014; Beebe and Sackris 2016; Sarkissian et al. 2011; Wright et al. 2013, 2014; Wright 2018). In light of this finding it is tempting to regard the presumptive argument as empirically implausible (and likewise any other argument that presupposes dominantly realist intuitions). However, such a conclusion may be premature.

A number of serious doubts have been raised about the validity of the existing psychological studies on folk moral realism. Most importantly, it has been argued that many of these studies failed to fully and exclusively measure what they set out to measure. They have not adequately methodologically distinguished between moral realism and anti-realism; have conflated these positions with distinct metaethical positions; have elicited irrelevant intuitions (such as first-order moral intuitions); and so on (see Beebe 2015; Pölzler 2018, forthcoming a). This means that there are important improvements to be made or follow-up studies to be

² This argument has been put forward in several different versions. For example, it has been claimed that realism has a presumption in its favor because our realist intuitions are best explained by realism (e.g., Brink 1989; Enoch 2017b), because these intuitions are best justified or vindicated by realism (e.g., Brink 1989; Shafer-Landau 2003), or because we have a *prima facie* reason to believe the content of any of our intuitions (e.g., Dancy 1986; Huemer 2005; McNaughton 1988). In what follows we will not be concerned with the argument's philosophical merits but exclusively focus on its empirical premise.

conducted before psychological findings can be appealed to as evidence for or against philosophers' hypotheses about folk moral objectivity.

Moreover, previous research on folk moral realism has also mainly attempted to identify subjects' immediate or gut intuitions about the existence of objective moral truths. But several proponents of the presumptive argument have suggested that what grounds a *prima facie* reason in favor of moral realism are actually our more reflective intuitions; that is, how we would answer questions about moral objectivity after having given them some thought. Brink, for example, explicitly points out that he is not concerned with “unreflective and untutored metaphysical or metaethical views” (1989: 25). In addition, reflective intuitions also seem to be more philosophically relevant when it comes to drawing potential semantic conclusions from research on folk moral realism (see Hannon 2018; Kauppinen 2007; Pölzler forthcoming b).³

The aim of this paper is to advance psychological research of folk moral realism, and in a way that is more likely to allow for philosophical conclusions. We begin by introducing previous research and explain why it has lacked in validity (Section 1). Then we propose a new experimental approach designed to overcome these worries as far as methodologically feasible, which may serve as a model for future studies (Section 2). Finally, and most importantly, we report and discuss the results of a psychological study that was based on this new design. In this study, most of our participants neither favored realism nor pluralism, but rather denied the existence of objective moral truths *tout court*. In particular, many of them had the intuition that whether moral sentences are true depends both on their own moral beliefs and on the dominant moral beliefs within their culture (“anti-realist pluralism”) (Sections 3 and 4). This suggests that realists' presumptive argument must indeed be rejected and that instead anti-realism may have a presumption in its favor.

1 Previous Research

In the last 15 years psychologists have conducted a number of studies on folk moral objectivity, typically labelled as being about realism (this is the label we will use), objectivism, subjectivism or relativism (e.g., Beebe 2014, forthcoming; Beebe and Sackris forthcoming; Beebe et al. 2015; Cova and Ravat 2008; Fisher et al. forthcoming; Goodwin and Darley 2008, 2010,

³ Kauppinen of course denies that such conclusions can be drawn, as in his view, reflective intuitions are not accessible to quantitative research. But studies like ours (and many others) suggest that at least some reflection can in fact be primed or be tested for.

2012; Heiphetz and Young forthcoming; Nichols 2004; Nichols and Folds-Bennett 2003; Sarkissian et al. 2011; Wainryb et al. 2004; Wright 2018; Wright et al. 2013, 2014; Young and Durwin 2013). In this section we introduce the methods of these studies and then recapitulate several worries about them.

1.1 Methods

The terms “moral truth” and “moral objectivity” can be understood in different ways (see, e.g., Lynch 2013; Raz 2001). In order for psychological research on folk moral realism to possibly bear on metaethics it must assume the very same understanding as this area of philosophical inquiry. So what do metaethicists mean when they discuss moral truth and moral objectivity? Here we assume a correspondence theory of moral truth and a mind-independence conception of moral objectivity. That is, we take the moral realism/anti-realism debate to be about whether moral sentences are true in that they match (objective) moral facts; and about whether these facts are objective in that they would obtain even if observers had different or no mental states towards them (see, e.g., Huemer 2005; Joyce 2007).

Suppose a study asked: “Do you believe that there are objective moral truths?” It is unclear whether participants would understand this question according to the above conception of objective moral truths. Psychologists have therefore attempted to measure realist and anti-realist intuitions in indirect ways. They have asked participants simpler and more concrete questions whose answers, in their view, logically entail realist or anti-realist positions. Two kinds of questions have been particularly prominent. They constitute what we will henceforth refer to as “disagreement tasks” and “truth-aptness tasks”.

Disagreement tasks represent the most common way of studying folk moral realism. In their study on children between the age of five and nine Wainryb et al. (2004), for example, asked participants to imagine two persons who morally disagree with each other. Then participants were asked to interpret this disagreement.

Sarah believes that it’s okay to hit and kick other children, and Sophie believes that it’s wrong to hit and kick other children. [...] Do you think that only one belief [what Sophie believes] is right, or do you think that both beliefs [what both Sophie and Sarah believe] are right? (Wainryb et al. 2004: 691)

According to Wainryb et al., “only one belief is right” responses entail realist intuitions, and “both beliefs are right” responses entail anti-realist intuitions (2004: 692).⁴

Truth-aptness tasks have been introduced by Goodwin and Darley. In their influential 2008 study participants were presented a moral sentence (e.g., “Anonymously donating a significant proportion of one’s income to charity is a morally good action”; 2008: 1361-1362). Then they were asked, among others, whether they take this sentence to be “true”, “false“ or “an opinion or attitude“:

How would you regard the previous statement? Circle the number. (1) True statement. (2) False statement. (3) An opinion or attitude. (Goodwin und Darley 2008: 1344)

In Goodwin and Darley’s interpretation, both “true” and “false” responses entail realist intuitions, and “opinion or attitude” responses entail intuitions in favor of subjectivism (see 2008: 1345).⁵

1.2 Worries

One minor problem with research such as Wainryb et al.’s and Goodwin and Darley’s is that it has lacked in external validity. Many studies have mainly surveyed US students (e.g., Goodwin and Darley 2008; Wright et al. 2013)⁶, have mainly used particular kinds of item statements (e.g., such as statements that involve “thin” moral concepts and action types) (e.g., Beebe and Sackris 2016; Goodwin and Darley 2008; Nichols 2004)⁷, or have involved unrealistic or humorous experimental stimuli (e.g., Sarkissian et al. 2011).

In addition, this research is also subject to a more fundamental worry about its construct validity. Researchers have typically assumed adequate definitions of moral realism and anti-realism; that is, they have understood these positions as affirming or denying the existence of objective moral truths in the sense explained above (e.g., Goodwin and Darley 2008; Nichols

⁴ Wainryb et al. used the labels “objectivism” and “relativism” to refer to claims that affirm/deny objective moral truths.

⁵ Goodwin and Darley used the labels “objectivism” and “subjectivism” to refer to claims that affirm/deny objective moral truths.

⁶ Studies involving subjects from different cultures include Beebe et al. (2015) and Sarkissian et al. (2011). By investigated subjects between the ages of 12 and 88 Beebe and Sackris (2016) also attempted to shed light on the question of how intuitions about the existence of objective moral truths vary across the lifespan.

⁷ Some researchers have successfully increased the variety of their item statements in some respects. Goodwin and Darley (2012), for example, included several statements about positive moral actions (in addition to the more usual statements about moral transgressions).

2004). But these definitions have then not been properly operationalized. This means that their research likely failed to fully and exclusively measure participants' intuitions about the existence of objective moral truths.

Elsewhere one of us has identified ten common sources of this (potential) lack in construct validity ([references removed for review process]). Many of them straightforwardly apply to the measures by Wainryb et al. and Goodwin and Darley introduced above.⁸

(W1) *Failure to Cover Variants of Anti-Realism*

Most studies' answer options have failed to cover main variants of anti-realism (i.e., subjectivism, error theory or non-cognitivism). Participants who felt drawn towards these variants hence had to respond in ways that did not reflect their actual intuitions — which may have been ways that were classified as realist.

(W2) *Failure to Cover Variants of Subjectivism*

Many studies also have not sufficiently considered variants of subjectivism that may plausibly be held by ordinary people (i.e., divine command theory, cultural relativism and individual subjectivism). Again, this may have systematically distorted their results.

(W3) *Conflation with Cognitivism/Non-Cognitivism Debate*

With regard to truth-aptness tasks, researchers have assumed that by measuring intuitions about the truth-aptness of moral sentences (cognitivism/non-cognitivism debate) they can illuminate intuitions about moral realism in a comprehensive sense. But anti-realism is compatible with both moral sentences being truth-apt and not being truth-apt.

(W4) *Conflation with Universalism/Relativism Debate*

The realism/anti-realism debate has occasionally also been conflated with the debate about the scope of moral sentences; that is, about how widely these sentences apply (universalism/relativism debate).

⁸ In the research mentioned above, one of us has argued that in addition to the studies by Goodwin and Darley and Wainryb et al., one or more of the following worries also apply to Beebe and Sackris 2016; Nichols 2004; Nichols and Folds-Bennet 2003; and Wright 2018. Moreover, the worries apply to any research that has adopted (part of) these studies' methodology (such as Beebe et al. 2015; Goodwin and Darley 2012; Wright et al. 2013, 2014).

(W5) *Elicitation of Irrelevant First-Order Moral Intuitions*

Even though moral realism and anti-realism are claims in metaethics, and hence descriptive, some measures have prompted first-order moral intuitions. These intuitions may partly explain participants' answers.

(W6) *Elicitation of Irrelevant First-Order Moral Intuitions*

Participants' responses may also have been unduly systematically influenced by epistemic intuitions; that is, intuitions about whether persons know, are justified in believing or are certain about moral propositions.

(W7) *Ex ante Classification of Statements*

So far most studies' item statements have been classified as moral or non-moral by the researchers themselves. Participants' responses therefore may not only be explained by their realist or anti-realist intuitions but also by their not having regarded some statement as moral at all.

(W8) *Misleading or Biasing Instructions*

Some studies have involved misleading or biasing instructions of the realism/anti-realism debate or their truth-aptness or disagreement tasks.

(W9) *Unsubstantiated Assumption about Moral Truth*

So far no evidence has been provided for the claim that participants understand studies' underlying concepts of truth, rightness or correctness in a correspondence theoretic sense.

(W10) *Unsubstantiated Identification of Responses with Intuitions*

It is also possible that participants have systematically conflated the studies' questions with independent questions (such as about the amount of societal consensus about moral sentences) or that they do not have any determinate intuitions about the existence of objective moral truths at all.

While some of the existing research (e.g., Sarkissian et al. 2011; Wright 2018) has managed to escape some of the above worries, none of the studies conducted thus far emerge completely unscathed. In light of these worries it seems reasonable to conclude that the extant psychological research, while uncovering many interesting aspects of people's underlying moral

psychology, has thus far failed to yield data that can conclusively weigh in as evidence about the content of folk metaethical intuitions. More careful work needs to be done before intuitions can be classified according to metaethical categories and philosophical conclusions can possibly be drawn.

2 New Experimental Design

To improve upon the situation described in the previous section, we developed a new experimental design for measuring folk intuitions about moral objectivity. This design purports to overcome W1 to W10 as far as methodologically feasible and may serve as a model for future research. It involves three main parts: (1) instructions, (2) abstract tasks, and (3) concrete tasks. In what follows we will explain each of these parts and then provide some general discussion.

2.1 Instructions

Our design starts with an explanation that purports to prevent participants from misreading its main tasks as asking for first-order moral intuitions (W6, W10), and to promote reflection (so as to make subjects' responses more philosophically relevant). To begin with, we explain the difference between what we call "normative" and "meta-ethical" sentences about morality:

Normative sentences about morality express moral judgments. In uttering these sentences we evaluate something morally; we indicate that we regard something as morally right or wrong, good or bad, virtuous or vicious, and so on.

[...]

Meta-ethical sentences about morality do not express moral judgments. In uttering them we remain evaluatively neutral. Instead, we are making claims about the nature of morality itself.

This explanation is supplemented by the least controversial and least biasing examples of normative and metaethical sentences that we could think of (W8).⁹ We also explicitly note that the classification of sentences as normative or metaethical is independent from the extent

⁹ Our example of normative sentences about morality were: "It is wrong to break promises," "The US has a duty to reduce its greenhouse gas emissions," "Hitler was morally depraved," and "Parents should be willing to make sacrifices for their children." Our example of metaethical sentences were: "No person has any moral knowledge," "There are facts about what is morally right or wrong, good or bad, virtuous or vicious, and so on," "Sentences about what is the case never entail sentences about what morally ought to be case," "Judgments such as "'Hitler was morally depraved' are neither true nor false."

to which one agrees to them. Finally, participants are informed that the present study concerns *metaethical* sentences and that they are therefore asked to abstract from their first-order moral intuitions.

Given that we are interested in your intuitions about meta-ethical sentences, we ask you to “bracket” your views about the normative sentences that we will present you (to ignore these intuitions or put them to the side). For the purposes of this study it does not matter whether, for example, you judge that breaking promises is wrong, that the US has a duty to reduce their greenhouse gas emissions, and so on.

Following these explanations, we test and improve participants’ understanding of the normative/metaethical distinction by two comprehension checks: (1) a theoretical question about what they have just read (right answer: “Normative sentences express moral judgments and meta-ethical sentences make claims about the nature of morality itself”), and (2) an exercise that asks them to classify several sentences as either normative or metaethical (e.g., “Eating meat is morally wrong” and “Gaining moral knowledge requires careful rational reflection”). Participants who do not fully succeed on either check are presented the above explanations once again and have to complete the relevant check/s one more time.

2.2 Abstract Measures

Previous research on folk moral realism has exclusively measured responses to concrete moral sentences.¹⁰ This experimental design, in contrast, also involves abstract tasks (formulated as being about morality in general). Our motivation for including such tasks is threefold. First, abstract responses are less likely to be influenced by first-order moral intuitions. Second, folk intuitions about moral responsibility, knowledge and other non-metaethical philosophical issues have been found to vary across the abstract/concrete distinction (e.g., Nichols and Knobe 2008; Sinnott-Armstrong 2008). And third, no single measure of folk metaethical intuitions will ever be perfect. So the more comprehensive and varied one’s measures, the more robust and complete (and thus valid) the picture one can gain from them.

The five abstract tasks that we developed are (1) a theory task, (2) a metaphor task, (3) a comparison task, (4) a truth-aptness task, and (5) a disagreement task.

¹⁰ To our knowledge, there is only one exception: a study by Ziljstra (forthcoming) that is currently under review.

2.2.1 Theory Task

The theory task begins with an explanation of the three main questions that determine one's position in the moral realism/anti-realism debate: Do moral sentences intend to state moral facts? If yes, do these facts exist? And if yes, are they independent from what anybody thinks about them? Then participants are asked to choose their preferred answer to these questions, with their options corresponding to all main variants of moral realism and anti-realism (W1), including what we take to be the most plausibly held variants of subjectivism (divine command theory, cultural relativism, and individual subjectivism) (W2).¹¹

- When a person says that something is morally right or wrong, good or bad, etc. she intends to state a fact. Such facts exist – and they are independent from what anybody thinks about them. For example, an action that is morally wrong is wrong no matter what anyone thinks. So it would still be wrong even if you yourself, or the majority of the members of your culture, thought that it is not morally wrong. [REALISM]
- When a person says that something is morally right or wrong, good or bad, etc. she intends to state a fact. Such facts exist – and they depend on God's will. For example, an action is only morally wrong if God forbids us to perform the action. If God did not forbid us performing the action, it would not be wrong. [DIVINE COMMAND THEORY]
- When a person says that something is morally right or wrong, good or bad, etc. she intends to state a fact. Such facts exist – and they depend on what the majority of the members of her culture think about them. For example, an action is only morally wrong if the majority of the members of your culture believe that it is wrong. If the majority of the members of her culture did not believe the action to be wrong, it would not be wrong. [CULTURAL RELATIVISM]
- When a person says that something is morally right or wrong, good or bad, etc. she intends to state a fact. Such facts exist – and they depend on what individuals think about them. For example, an action is only morally wrong if you yourself believe that it is morally wrong. If you did not believe the action to be wrong, it would not be wrong. [INDIVIDUAL SUBJECTIVISM]
- When a person says that something is morally right or wrong, good or bad, etc. she intends to state a fact. Such facts do not exist. Thus, it is never the case that something is morally right or wrong, good or bad, etc. No such moral statement can be true. [ERROR THEORY]
- When a person says that something is morally right or wrong, good or bad, etc. she does not intend to state a fact. Instead, she intends to communicate/express her feelings, emotions, intentions or attitudes about it. For example, by saying that something is wrong, you only express feelings of disapproval towards it (and that is the only thing you are doing). Moreover, there are no facts about what is morally right or wrong, good or bad, etc. [NON-COGNITIVISM]

¹¹ Needless to say, subjectivism can be held in many other forms as well. Testing intuitions about all of these versions at once does not seem methodologically feasible. As pointed out above, this study therefore only addresses three versions of subjectivism that can reasonably be expected to be held by a noteworthy proportion of the population.

2.2.2 Metaphor Task

Moral realism and anti-realism are claims about moral facts. Our second abstract task asks participants to choose between a number of metaphors for these facts which correspond to the main variants of realism and anti-realism (W1, W2).¹²

This task is about moral facts. Moral facts are facts about what is morally right or wrong, good or bad, virtuous or vicious, and so on. For example, it could be a moral fact that it is (or is not) wrong to break promises, or that the US has (or does not have) a duty to reduce their greenhouse gas emissions. Below moral facts are explained in terms of several metaphors. Which of these metaphors seems most appropriate to you?

- Moral facts are “discovered”. They can be discovered in the same way in which we discover other facts about the objective world. [REALISM]
- Moral facts are “divine commandments”. They are introduced and determined by God. [DIVINE COMMAND THEORY]
- Moral facts are “cultural inventions”. They are introduced and determined by cultures. [CULTURAL RELATIVISM]
- Moral facts are “individual inventions”. They are introduced and determined by individuals. [INDIVIDUAL SUBJECTIVISM]
- Moral facts are “illusions”. While it may seem to us that they exist they actually do not exist at all. [ERROR THEORY or NON-COGNITIVISM]

2.2.3 Comparison Task

Next, participants are asked to which non-moral domain morality is most similar. The answer options are again formulated in such a way that they reflect the main variants of moral realism and anti-realism (W1, W2).

Below morality is being compared to various types of matters. Please indicate which comparison seems most appropriate to you.

- Morality is akin to science or mathematics. There are objective facts about what is right or wrong (facts that are independent from what anybody thinks about them). We cannot change these facts, we have to discover them. [REALISM]
- Morality is akin to religion. What is morally right or wrong is determined by what God wants us to do. Individuals cannot, by themselves, change the moral facts. [DIVINE COMMAND THEORY]

¹² It may be objected that the formulation of this task biased subjects in favor of thinking that moral facts exist. However, while this may be true to some extent, note that the task also involved an answer option according to which moral facts are “illusions”; and that even if we exclude the metaphor task from our analysis this does not change our overall results, as reported in Sec. 3.3.

- Morality is akin to social conventions. In each culture different things can be morally right or wrong. Cultures determine the moral facts. Individuals within cultures cannot, by themselves, change those facts. [CULTURAL RELATIVISM]
- Morality is akin to personal taste or preferences. For each person different things can be morally right or wrong. The individual determines the moral facts. [INDIVIDUAL SUBJECTIVISM]
- Morality is akin to superstition. It is based on a fundamental error. It assumes things exist (namely facts about rightness and wrongness) that do not actually exist. [ERROR THEORY]
- Morality is akin to exclamations (such as “Yeah!” or “That sucks!”). We use terms such as “right” and “wrong” to express our feelings, emotions, intentions or attitudes, but that is all. There are no moral facts. [NON-COGNITIVISM]

2.2.4 Disagreement Task

Our fourth task is an improved version of traditional disagreement tasks. It presents participants a case of moral disagreement about an unspecified situation and then asks them to interpret this case in ways that entail variants of moral realism and anti-realism (W1).

Consider the following scenario. Two people from the same culture are evaluating the exact same situation and utter conflicting moral sentences about it. One person says that what happened is morally bad (wrong, vicious, etc.).¹³ The other person says that what happened is not morally bad (wrong, vicious, etc.). Which interpretation of this disagreement seems most appropriate to you?

- One of these two people is right and the other one is wrong (Please note that this could be the case for several reasons: for example, because the truth of moral sentences is objective, or because it is determined by the dominant moral beliefs in their culture, or because it is determined by the commandments of God). [REALISM, CULTURAL RELATIVISM or DIVINE COMMAND THEORY]
- Both people are right (because the truth of moral sentences is determined by the moral beliefs of individuals). [INDIVIDUAL SUBJECTIVISM]
- Both people are wrong (because although moral sentences intend to state moral truths, there are no such truths). [ERROR THEORY]
- Neither person is right or wrong (because moral sentences do not intend to state moral truths, and are therefore neither true nor false). [NON-COGNITIVISM]¹⁴

Participants who respond that one of the persons is right and the other one is wrong are presented one or two otherwise identical disagreement scenarios to discern whether they favor

¹³ Half of the participants received a vignette in which the two people disagree about whether a situation is “good (right, virtuous, etc.)” (rather than “bad (wrong, vicious, etc.)”).

¹⁴ In a pilot study we found that if answer options are exclusively stated in terms of persons being right or wrong then participants often ignore or misunderstand their implications for the moral realism/anti-realism debate. It is therefore advisable to explicitly note these implications (as we did in parentheses).

realism or non-individualistic variants of subjectivism (W2). First they are presented a follow-up scenario in which the disagreeing parties are members of different cultures, with each of these parties' moral judgements conforming to the dominant view within their culture. Participants who respond that under these circumstances both people are right are classified as cultural relativists. Those who again choose "One of the persons is right and the other one is wrong" are presented a third scenario in which the disagreeing parties are subject to different commands by God, with each of their moral judgements corresponding to these commands. Regarding this scenario, "One of the persons is right and the other one is wrong" answer are classified as realist and "Both people are right" answers as indicating divine command theory.

2.2.5 Truth-Aptness Task

Finally, our abstract part also involves an improved version of traditional truth-aptness tasks. Participants receive an explanation of what it is for a moral sentence to be or not be truth-apt.¹⁵ For example, they are told that "[t]ruth-apt sentences express beliefs about facts", and that even when we do not know whether a moral sentence is true or false it can still be truth-apt. Their understanding of this explanation is tested and improved by comprehension checks analogous to those reported in Sec. 2.1. Then participants are presented the following task (W3):

Think about moral sentences (sentences that express that something is morally good or bad, right or wrong, virtuous or vicious, and so on). Are these sentences truth-apt or not truth-apt?

- Yes, moral sentences are "truth-apt" -- that is, they intend to express how things are; what is the case (either with regard to the objective world or with regard to what particular individuals, cultures, etc. think about morality). Thus, these sentences are either true or false. [COGNITIVISM: REALISM, DIVINE COMMAND THEORY, CULTURAL RELATIVISM, INDIVIDUAL SUBJECTIVISM or ERROR THEORY]
- No, moral sentences are not "truth-apt" -- that is, they do not intend to express beliefs about objective or subjective facts, but rather only express feelings, emotions, intentions or attitudes. Thus, these sentences are neither true nor false. [NON-COGNITIVISM]

2.3 Concrete Measures

The third main part of our experimental design consists of concrete tasks. These tasks are essential to measuring people's intuitions about moral realism because they allow conclusions about the extent to which their metaethical intuitions vary with moral sentences or circum-

¹⁵ Participants who receive our concrete measures first are provided with this explanation prior to the concrete truth-aptness task.

stances. One problem with some of the previous concrete tasks has been that participants were presented purportedly moral sentences that they themselves might not have regarded as moral (W7). We therefore explain the distinction between sentences that make dominantly moral versus sentences that make dominantly non-moral evaluations (in most general and non-biasing terms).¹⁶ Then we ask participants to rate the following sentences as belonging to either of these categories.

- A country with the death penalty is morally worse than a country without.
- Before the third month of pregnancy, abortion for any reason is morally impermissible.
- Men who violently physically punish their children are cruel.
- John cheating (committing adultery) on his wife Elizabeth for no other reason than boredom with his marriage is morally permissible.
- Consciously discriminating against someone, by not hiring them for a job they are clearly qualified for, just because of their race is not wrong.
- It was horrendous of Martin Shkreli to overprice drugs that he knew sick people really needed.
- A world in which wealth is distributed equally is more just than a world in which it is distributed unequally.
- Martin Luther King was a righteous man.
- It is good to do onto others as you would have them do onto you.
- It is morally reprehensible of the current US administration to limit or prohibit immigration.

In the subsequent two concrete tasks participants are only presented those sentences that they themselves rated as dominantly moral. Note also that our above selection is less homogenous than that of most previous research on folk moral realism. For example, it involves sentences about particular moral cases and moral principles, sentences that involve thick moral concepts, a negated sentence, highly counterintuitive sentences, and so on. This increases the generalizability of our findings to morality as a whole.

¹⁶ In particular, this explanation reads as follows: “The main point of some sentences is to make moral evaluations (i.e., evaluations about a moral matter, evaluating something as being morally right or wrong, good or bad). Here are some sentences that one might think belong to this category: It is wrong to break promises. The US has a duty to reduce its greenhouse gas emissions. Hitler was morally depraved. Parents should be willing to make sacrifices for their children. The main point of other sentences is to make other sorts of (non-moral) evaluations (i.e., evaluations that don't have anything to do with morality, evaluating something as being correct or incorrect). Here are some sentences that one might think belong to this category: You put your shoe on the wrong foot. That chocolate ice cream tastes good. It is illegal for you to park on campus without a permit. It is rude to talk with your mouth full.

2.3.1 Disagreement Task

In the disagreement task's concrete version participants are asked to interpret cases of disagreement about each of their self-classified moral sentences. Here is one example:

Men who violently physically punish their children are cruel.

Consider the following situation. Two people from the same culture discuss whether men who violently physically punish their children are cruel. One person says that men who violently physically punish their children are cruel. The other person says that it is not the case that men who violently physically punish their children are cruel. Which interpretation of this disagreement seems most appropriate to you?

- One of these two people is right and the other one is wrong (Please note that this could be the case for several reasons: for example, because the truth of the people's sentences is objective, or because it is determined by the dominant moral beliefs in their culture, or because it is determined by the commandments of God). [REALISM, DIVINE, CULTURAL RELATIVISM or DIVINE COMMAND THEORY]
- Both people are right (because the truth of their sentences is determined by the moral beliefs of individuals). [INDIVIDUAL SUBJECTIVISM]
- Both people are wrong (because although their sentences intend to state moral truths, there are no such truths). [ERROR THEORY]
- Neither person is right or wrong (because their sentences do not intend to state moral truths, and are therefore neither true nor false). [NON-COGNITIVISM]

Analogously to the abstract disagreement task, answering in the first way triggers one or two follow-up scenarios (cross-cultural and God-related) that allow to distinguish realist from cultural relativist and divine command theorist intuitions.

2.3.2 Truth-Aptness Task

In the concrete truth-aptness task participants are asked to classify all of the sentences that they rated as moral as either truth-apt or not-truth-apt (W3):

Consider the following sentences. Here we are not interested in whether you believe these sentences are actually true. Our focus is rather on their truth-aptness.

Are these sentences truth-apt (i.e., they intend to express what is the case about individuals' moral beliefs, culturally dominant moral beliefs or the objective world; and thus, they are either true or false) or not truth-apt (i.e., they intend only to express feelings, emotions, intentions or attitudes, and thus, they are neither true, nor false)?

truth-apt [COGNITIVISM: REALISM, DIVINE, CULTURAL RELATIVISM, DIVINE COMMAND THEORY, INDIVIDUAL SUBJECTIVISM or ERROR THEORY]

not truth-apt [NON-COGNITIVISM]

2.4 Discussion

In its above form our new experimental design alleviates most of the worries that have been raised about previous research on folk moral realism (W1-W8). The two most fundamental worries have not yet been (fully) addressed, however; namely that participants may not understand studies' underlying concepts of truth and rightness in a correspondence theoretic sense (W9), and that they systematically conflate their questions with independent questions or do not have any determinate intuitions about the existence of objective moral truths at all (W10).

In a pilot study we attempted to investigate whether moral truth is regarded as correspondence with facts, as coherence, or as the reaffirmation of first-order moral commitments. Participants' verbal explanations indicated serious misunderstandings about what we were asking. The issue seemed too abstract and too intricately related to the moral realism/anti-realism debate to be addressed *en passant*. In our design we therefore only prime participants to think in correspondence-theoretic ways. For example, with regard to our truth-aptness measures, we explain that to say of moral sentences that they are truth-apt is to say that "they intend to express how things are; what is the case" (see Sec. 2.2.5 and 2.3.2).

According to the second fundamental worry, participants may systematically conflate the issue of moral objectivity with independent issues or may not have any determinate intuitions about it at all. One way in which our design addresses this worry is by involving detailed instructions of the normative ethics/metaethics and truth-aptness/not truth-aptness distinctions, and by testing participants' understandings of these distinctions (see Sec. 2.1 and 2.2.5). In addition, we also introduced the following measures: open-ended questions that ask participants to explain some of their responses and response patterns (in particular, variation between the abstract and concrete disagreement and truth-aptness responses, and within the concrete disagreement and truth-aptness responses), completion times recordings for each task and the survey as a whole, additional comprehension checks, and attention checks.

3 Study

By applying the experimental design introduced in Sec. 2 we gathered (what we deem to be) more valid and philosophically relevant data about the folk's metaethical intuitions than previous research. In what follows we will describe the participants, methods and results of our study.

3.1 Participants

We surveyed 172 participants. Of that, 98 participants were from Amazon MTurk (49% male) and were paid \$7.25; and 74 participants were from the College of Charleston (26% male) and received two research credits.

Prior to analysis the quality of participants' responses was assessed according to the following criteria (in the order of their weighting): (1) performance in attention checks (failure in any check automatically eliminated participants from the study), (2) study completion time, both overall and for each individual task, (3) performance in comprehension checks, (4) minimal relevance of verbal explanations. On the basis of these criteria 55 participants (32%) were excluded from analysis. For example, we did not consider the responses of participants who finished the complete survey in less than 20 minutes (the average time was 45 minutes); who copied their verbal explanations from Wikipedia and other online sources; and who took the first option in all comprehension checks.

In the end we had 117 participants, 67 from MTurk and 50 from the College of Charleston. Of these participants 63% were female; their age varied between 18 to 64 years ($M = 29.6$); and they were 86% Caucasian, 4% African American, 4% Asian American, 3% Hispanic, and 3% other.

3.2 Methods

The survey was administered online through Qualtrics. Each participant was presented both the abstract measures (theory, metaphor, comparison, disagreement and truth-aptness tasks) and the concrete measures (disagreement and truth-aptness tasks). The order of these measures was counterbalanced, and many of the individual tasks, as well as most of the answer options, were randomly ordered. Responses to the tasks were coded as being indicative of realist or anti-realist intuitions according to the information added in square brackets in Sec. 2.

One potential objection against our methodology concerns our classification of divine command theory as a version of subjectivism, and hence anti-realism. According to divine command theory, a thing is morally right, good, virtuous, etc. if and only if it is commanded by God, and morally wrong, bad, vicious, etc. if and only if it is forbidden by God.¹⁷ At first

¹⁷ Often divine command theory is defended in a "restricted" sense: as the claim that only deontic moral properties (such as rightness) are dependent on the commands of God (e.g., Quinn 2000). Our study focused on the above comprehensive version (e.g., Evans 2013) because it seems more likely to be held by ordinary people.

sight this view seems to entail moral objectivity. However, there are good reasons to think otherwise. For moral facts to be objective means that these facts would obtain even if observers had different or no mental states towards them (Sec. 1.1). But on divine command theory moral facts do not meet this condition. They are contingent upon the will of God (Huemer 2005: 54-55). In fact, analogous to other versions of subjectivism, this mind-dependence has traditionally been the main objection against divine command theory (if moral truths depended on God's will then if God commanded us to murder, rape, steal, and so on these actions would be right — which is highly counterintuitive; see Plato [380 B.C.] 2010).¹⁸

Another important methodological question is which proportion of a participant's intuitions about concrete cases must favor realism or anti-realism in order for the participant to count as being drawn towards these views. Elsewhere one of us has argued that from a metaethical perspective there must not be any genuine intrapersonal variation at all ([references removed for review process]). That said, when a participant is asked to give many responses, as in our concrete tasks, the likelihood of performance errors increases. We thus stipulated that participants have “consistently” realist or anti-realist intuitions if 100% of their responses to the abstract tasks and at least 90% of their responses to the concrete tasks were in favor of only one of these views. This is the threshold above which the likelihood of realist or anti-realist intuitions having an effect on the participant's responses is significantly greater than the likelihood of there being no such effect, as calculated by a binomial test, based on the number of separate tasks involved in each set of measures.

Finally, most of our tasks involve five anti-realist but only one realist answer option. It may be objected that this unequal distribution biases our results in favor of anti-realism. We agree that this feature is potentially problematic, especially to the extent that participants answer without sufficient effort or comprehension. But it is also hard to avoid. Our study's most important tasks are the abstract and concrete disagreement tasks. Prior research has shown

¹⁸ We would like to supplement our above argument by three brief remarks. First, those who suggest classifying divine command theory as a version of realism might respond that objectivity only requires independence from the mental states of *human* observers. This response strikes us as ad hoc, in particular as this view's observer-independence gives rise to the same problem as the observer-independence of paradigmatic subjectivist views (see above). Second, our hunch is that the inclination to regard divine command theory as a version of realism derives from a conflation of the realism/anti-realism debate with the universalism/relativism debate. Divine command theory construes morality as universal in the sense that moral sentences apply to all human beings at all times and places (at least in our actual world). However, as one of us has shown in previous work [reference removed for review process], this issue of the scope of moral sentences is logically distinct from the question of their objectivity. And third, while the objection against divine command theory mentioned in the main text roughly originates in Plato, his “Euthyphro dilemma” was of course concerned with “piety” rather than with morality (“Is the pious loved by the gods because it is pious, or is it pious because it is loved by the gods?” ([380 B.C.] 2010: 10a).

that in order for such tasks to be valid researchers need to fully tease apart the logically possible answers to them (see W1 and W2). By doing so one automatically ends up with only one option that (potentially) entails realism, and a number of options that entail anti-realism. Moreover, to alleviate the problem of our unequally distributed answer options, we decreased insufficient effort and insufficient comprehension responding in various ways, so as to increase the likelihood that subjects who opted for anti-realist options really felt drawn towards these options (as explained in Sec. 3.1).

3.3 Results

3.3.1 Comprehension Checks

Our metaethics/normative ethics comprehension check involved a theoretical question and a classification exercise (see Sec. 2.1). 114 participants (97%) answered the theoretical question correctly on the first pass, and all three who got it wrong successfully corrected their answer on the second pass. The classification exercise showed a marked difference in performance between the MTurk and student participants. On average, 69% of the MTurkers classified all 14 sentences correctly at first pass, with the correction rate being 79%. In contrast, only 35% of the student participants initially classified all sentences correctly, with the correction rate being 72%.

Participants had to complete several additional comprehension and memory checks as well. 81%, 82%, 94%, and 71% correctly selected what the abstract theory, metaphor, comparison, and disagreement tasks were about; and our question of what had been presented in the concrete disagreement tasks received 83%, 75%, 75%, 72% and 77% correct responses. The MTurk participants again did somewhat better on most of these checks. This also holds for the extensive comprehension checks prior to each participant's first truth-aptness task. While both the MTurkers and students mostly correctly answered our theoretical question (69% and 70%), the MTurkers did better at our classification exercise (48% versus 20% correct classifications at first attempt) and were able to correct more of their misclassifications (86% versus 69%).

In total, the analyzed participants' performance in the comprehension checks was good; and where possible, mistakes were mostly corrected (indicating that participants learned from their mistakes).

3.3.2 Realist/Anti-Realist Responses

Collapsing across all of the opportunities participants had to give either a realist or an anti-realist response, and then averaging these together, we found that participants dominantly gave anti-realist responses. In response to the abstract measures the proportion of anti-realist responses was 87%, and in response to the concrete measures 92% (as opposed to 13% and 8% of realist responses) (see Tables 1 and 2). The order of the measures' presentation did not have any significant effect; that is, participants who received the abstract measures first responded in the same way as participants who received the concrete measures first.

	Theory	Comparison	Metaphor	Disagreement	Truth-Aptness	TOTAL (not incl. Truth-Aptness)
Realist	20%	10%	13%	8%	≤73%	13%
Anti-Realist	80%	90%	87%	92%	≥27%	87%

Table 1: Proportion of realist and anti-realist responses in abstract measures.

	Martin Luther King	Wealth Distribution	Cheating	Overpricing Drugs	Physical Punishment	Immigration	Abortion	Racial Discrimination	Death Penalty	Golden Rule	TOTAL (not incl. Truth-Aptness)
Realist	4% ≤65%	10% ≤85%	0% ≤78%	17% ≤63%	21% ≤75%	0% ≤80%	10% ≤84%	0% ≤82%	6% ≤84%	7% ≤62%	8%
Anti-Realist	96% ≥34%	90% ≥14%	100% ≥21%	83% ≥36%	79% ≥24%	100% ≥19%	90% ≥15%	100% ≥17%	94% ≥15%	93% ≥37%	92%

Table 2: Proportion of realist and anti-realist responses in concrete measures. Upper value = disagreement task, lower value = truth-aptness task.

Most anti-realist responses were anti-realist in a cognitivist sense. On average, 73% of the responses to the abstract and 76% of the responses to the concrete truth-aptness tasks were in favor of moral sentences being truth-apt. Almost all other tasks showed an even higher proportion of cognitivist vis-à-vis non-cognitivist anti-realism. In particular, participants' responses dominantly indicated intuitions in favor of cultural relativism (36% in the abstract and 24% in the concrete condition) and individual subjectivism (25% in the abstract and 42% in the concrete conditions). Divine command theorist and error theoretic responses were less widespread (see Tables 3 and 4).

	Theory	Comparison	Metaphor	Disagreement	TOTAL
Divine Command Theory	12%	11%	11%	6%	10%
Cultural Relativism	27%	46%	53%	19%	36%
Individual Subjectivism	7%	28%	19%	47%	25%
Error Theory	4%	3%	4% ¹⁹	5%	4%
Non-Cognitivism	30%	3%		23%	19%

Table 3: Breakdown of anti-realist responses in abstract measures (excluding truth-aptness task).

	Martin Luther King	Wealth Distribution	Cheating	Overpricing Drugs	Physical Punishment	Immigration	Abortion	Racial Discrimination	Death Penalty	Golden Rule	TOTAL
Divine C. Theory	5%	1%	0%	10%	10%	1%	5%	3%	3%	6%	4%
Cultural Rel.	15%	21%	43%	19%	27%	18%	24%	41%	14%	23%	24%
Individual Subj.	35%	40%	34%	43%	33%	48%	59%	35%	51%	41%	42%
Error Theory	10%	10%	11%	6%	6%	15%	3%	8%	11%	10%	9%
Non-Cognitivism	35%	28%	12%	22%	24%	18%	9%	14%	21%	20%	20%

Table 4: Breakdown of anti-realist responses in concrete measures (excluding truth-aptness tasks).

Even though subjects dominantly responded as anti-realists across all of our concrete disagreement tasks, there was nonetheless a significant spread between those tasks that received the least amount of anti-realist responses and those that received the most, $t(192) = 5.5$, $p < .001$. We also observed a significant difference in strength of agreement and perceived consensus between these two tasks: strength of agreement, $t(190) = 18.7$, $p < .001$, consensus, $t(190) = 11.5$, $p < .001$ (similar to much previous research, e.g., Beebe and Sackris 2016; Goodwin and Darley 2008, 2012).

We did not find a significant correlation between divine command theory and moral realism ($r = -.067$, $p = .47$). This suggests that these views are not only conceptually distinct but also represent distinct psychological constructs.²⁰ Moreover, even if divine command theory were classified as a version of realism our results would nevertheless strongly suggest that the folk are moral anti-realists (with 77% anti-realist responses in the abstract measures, and 88% anti-realist responses in the concrete measures).

¹⁹ Our metaphor task did not discriminate between error theory and non-cognitivism (see Sec. 2.2.2). This is because we were not able to come up with a metaphor that clearly only reflected one but not the other of these views.

²⁰ Based on different measures of moral realism, this distinctness has recently also been observed by Zijlstra forthcoming.

3.3.3 Inter- and Intrapersonal Variation

Above we defined a “consistent” participant as one who gives the same kind of responses in 100% of our abstract and at least 90% of our concrete tasks. None of our participants turned out consistently realist (with 4 participants giving consistently realist responses in the abstract measures). 72 participants (62%) gave consistently anti-realist responses. This percentage is higher than what we would expect if people were answering randomly (binomial test, $p = .016$). Of the 45 participants who did not give consistent responses across both kinds of measures, 25 gave consistent anti-realist responses to the concrete measures (but not the abstract), 9 gave consistent anti-realist responses to the abstract measures (but not the concrete), and only 11 did not give consistent anti-realist responses across any of the two kind of measures.

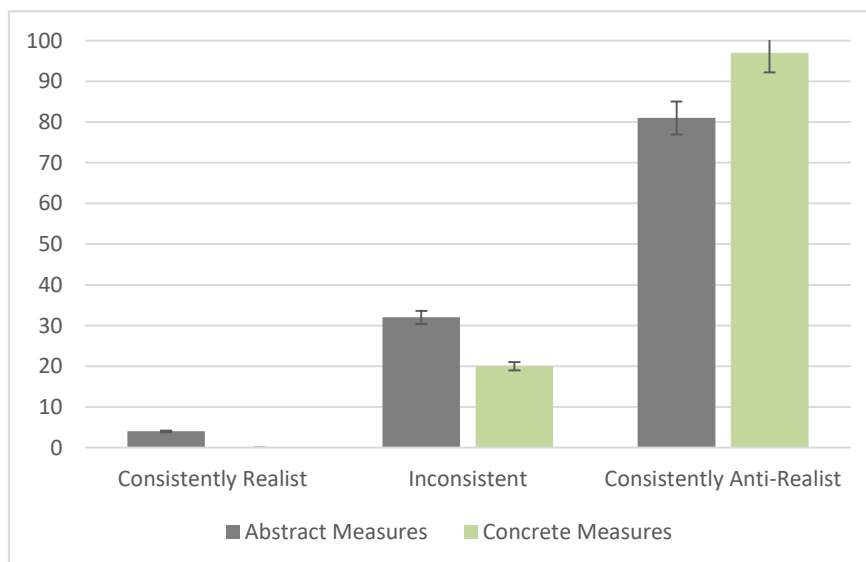


Figure 1: Number of participants who gave consistently realist, inconsistent and consistently anti-realist responses to the abstract and concrete measures.

Breaking participants’ anti-realist responses down into divine command theory, cultural relativism, individual subjectivism, error theory and non-cognitivism, only few consistently favored one of these views across or within the abstract and concrete measures. There was less variation among anti-realist responses in the abstract than in the concrete measures. Most variation occurred across the cultural relativism/individual subjectivism line. That is, many participants seemed to have the intuition that while moral truth is sometimes (mainly) determined by the dominant moral beliefs within cultures, it is other times (mainly) determined by the moral beliefs of individuals (see Tables 3 and 4).

There was a strong positive correlation between participants' realist abstract and concrete responses, $r(117) = .35, p < .001$ as well as a strong negative correlation between their concrete realist and their abstract anti-realist responses, $r(117) = -.35, p < .001$. This suggests some consistency between abstract and concrete responses. Nonetheless, we did find a significant difference across the abstract and concrete measures. In the abstract tasks the anti-realist response rate was significantly higher than in the concrete ones (93% vs. 87%, $t(116) = 2.5, p = .013$). However, this abstract/concrete variation may not be genuine. It may actually simply be explained by differing responses to the disagreement tasks on the one hand and the theory, comparison and metaphor tasks on the other hand (see Tables 1 and 2).

3.3.4 Demographics

In our concrete measures females provided less anti-realist responses than males, $t(114) = 2.3, p = .021$. The strength of participants' religious beliefs was negatively correlated with their anti-realist responses, $r(116) = -.32, p < .001$, as well as positively correlated with their realist responses (in the concrete condition), $r(116) = .18, p = .048$. These latter correlations suggest that divine command theory does not represent how religious participants tend to conceive of moral objectivity — a finding that is also supported by their verbal explanations (and see also Goodwin and Darley 2008: 1356-1357)

A previous study (Beebe and Sackris 2016) suggested that people's metaethical intuitions change across their lifespans. Our results support this hypothesis. However, they do so only partially, and suggest a different direction of change than this previous study. Across our abstract and concrete measures 18-29 year olds provided more realist responses than 30-51 years olds, $t(105) = 2.4, p = .018$. No differences between either of these age groups and 52-88 year olds was observed.

Finally, participants' political orientation was not significantly related to their realist or anti-realist responses in either the abstract or concrete conditions.

4 Discussion

As mentioned in the introduction to this article, previous research has converged on the hypothesis that ordinary people's metaethical intuitions are pluralist in the sense that they regard moral realism as true with regard to some moral sentences or circumstances and anti-realism with regard to other moral sentences or circumstances (e.g., Beebe 2014; Beebe and Sackris 2016; Sarkissian et al. 2011; Wright et al. 2013, 2014). Our results cast doubt on this hypothe-

sis. Most of our participants' intuitions did not vary across the realism/anti-realism distinction. They were rather consistently anti-realist.

Proponents of metaethical pluralism may respond that there are plausible explanations of this finding that are consistent with their view. First, our low proportion of realist responses at the concrete level may be partially a function of the particular item statements that we used in our study (especially the lack of more "extreme" statements such as about murder or rape).²¹ Second, by informing participants upfront about the survey's complexity and its attention and comprehension checks, as well as by excluding those who did badly on these checks, we may also have biased our sample.²² And third, our efforts to make participants engage in reflection (by presenting them with explanations and requiring them to complete comprehension checks) may have had an effect too.

All of these explanations strike us as initially plausible. Yet, there are good reasons to believe that the best account of our divergent results will also and mainly appeal to our methodological improvements. Most importantly, we observed less realist responses and hence less traditional pluralism because we accounted for non-individualistic versions of subjectivism (W2), and did not wrongly identify cognitivism with realism (W3).

First, many previous studies have exclusively used intra-cultural disagreement scenarios. Participants' "One person is right and the other one is wrong" responses to these scenarios were interpreted as being indicative of realism (e.g., Beebe 2014; Beebe and Sackris 2016; Goodwin and Darley 2008; Wainryb et al. 2004; Wright et al. 2013, 2014).²³ But as long as the parties of a moral disagreement are members of one and the same culture not only realists, but cultural relativists should be drawn to this response as well (because in their view moral sentences are made true by the dominant moral beliefs within cultures; and within a particular culture there can only be one dominant moral belief about a moral sentence). Previous disagreement tasks hence likely exaggerated the proportion of realist responses.²⁴

²¹ Though note that even the cruelty of violent physical punishment of children was considered to be a non-objective matter by 79% of our participants; and many previous studies did not involve (many) such extreme statements either.

²² Though note that previous studies only found a correlation between participants' disjunctive reasoning ability and their realist intuitions. Traits like need for cognition and reflective capacity did not correlate with these intuitions (Goodwin and Darley 2010)

²³ The particular formulation of this response varied across studies.

²⁴ This is supported by the fact that two studies that did include cross-cultural disagreement scenarios (Nichols 2004; Sarkissian et al. 2011) reported disproportionately strong tendencies towards anti-realism as well. Moreover, note that the above-mentioned intra-cultural disagreement scenarios also do not account for the possibility of participants being divine command theorists.

A second plausible explanation of our lower proportion of realist responses concerns our study's truth-aptness tasks. So far "truth-apt" responses have exclusively been interpreted as being indicative of moral realism (e.g., Goodwin and Darley 2008; Wright et al. 2013, 2014). But again, this interpretation is inadequate. Cognitivist variants of anti-realism entail the truth-aptness of moral sentences as well. Subjectivists believe that moral sentences are true or false depending on whether they correctly represent the subjective moral facts (such as facts about individual's or culturally dominant moral beliefs); and error theorists believe that all moral sentences are false. Our study accounted for this fact by classifying "truth-apt" responses as reflecting *either* realism *or* cognitivist anti-realism. Thus, it raised the bar for providing a realist response in this regard as well.

To test these potential explanations, we returned to our data and did the following: First, we reinterpreted "One person is right and the other person is wrong" responses to our intra-cultural disagreement task as realist (regardless of how participants responded to the disagreement tasks that followed). In the abstract measures, this significantly shifted participants' responses towards realism, from 8% to 35%, $t(117) = 6.6, p < .001$. In the concrete measures, it shifted them even more significantly towards realism, from 7% to 40%, $t(116) = 11.4, p < .001$. A similar finding emerged when we reinterpreted all truth-aptness responses as realist (if attributing truth-aptness) or anti-realist (if denying truth-aptness). In the abstract measure, this significantly shifted participants' responses towards realism, from 13% realist to 25% realist, $t(117) = 13.8, p < .001$. In the concrete measure, the proportion of realist responses even rose from only 7% to 42%, $t(116) = 23.9, p < .001$.

In our view, these considerations suggest that traditional metaethical pluralism might need to be replaced by what we will henceforth refer to as "anti-realist pluralism". At least on a more reflective level most people's metaethical intuitions may not vary across the realism/anti-realism distinction. They (relatively) consistently deny moral objectivity. But at the same time our results suggest that there is significant intrapersonal variation *within* anti-realism. Specifically, people may view the truth of moral sentences as depending on their own moral beliefs (individual subjectivism) and on the dominant moral beliefs within their culture (cultural relativism) (see Tables 3 and 4). Here are some representative explanations by our participants:

- The underlying theme between all of these answers [the participant's answers] is that there is not one universal code of morality. I think that senses of morality come from several sources such as culture/society, or individual convictions shaped by personal experience.

- I think morality is dependent on a person’s personal opinion and the culture they grew up in.
- Along with the idea of culture as a construct for what a persons morals may be with regards to certain topics it comes down to the individual and what they believe inside themselves is moral and immoral.

On the concrete level, this anti-realist pluralism might be understood in two distinct ways. First, ordinary people may have the intuition that some moral sentences are made true by their own moral beliefs, while other moral sentences are made true by the dominant moral beliefs within their culture (“either-or pluralist anti-realist”). Second, it may seem to them that the truth of every moral sentence is determined by both their own moral beliefs and the dominant moral beliefs within their culture, but to varying degrees (“conciliatory pluralist anti-realism”). The prevalence of these distinct intuitions is to be settled by future research that employs a more fine-grained methodology.

Our argument for anti-realist pluralism is limited in two ways that we have already touched upon. First, some of our participants may not have assumed a correspondence-theoretic understanding of moral truth. It would be inappropriate to attribute intuitions in favor of anti-realist pluralism (or any other realist or anti-realist view) to these participants (Sec. 2.4). Second, our results also may not be widely generalizable. Not only may they be contingent on our item statements, our participants scoring high in traits such as need for cognition or general intelligence, and the presence of prior reflection; people from non-Western cultures may think differently about metaethics as well (see Henrich et al. 2010).

That said, we are confident that we have gone quite some way in ruling out what is probably the most serious objection against research on folk moral realism, namely that participants systematically conflate the issue of moral objectivity with independent questions or may not have any determinate intuitions about it at all. These explanations strike us as implausible. We observed a higher proportion of anti-realist responses than we would expect if participants were answering randomly (Sec. 3.3.3). Some potentially distorting mental processes that might plausibly drive such a pattern can be (partly) ruled out. For example, participants’ responses likely are not attributable to emotional reactions (our design involved a number of abstract measures which half of the participants received first) or to their strength of agreement or perceptions of consensus (as they even responded in dominantly anti-realist ways to those moral sentences that they most strongly agreed to and took to be most widely accepted).

There are also positive reasons for believing that our survey responses reflect determinate intuitions about moral objectivity. First, all analyzed participants paid at least minimal atten-

tion throughout the whole survey (as evinced by the fact that they passed our attention checks). Second, most of them showed a good understanding of the survey's basic subject matter (as evinced by the fact that they performed well in our comprehension checks, see Sec. 3.3.1). Third, participants' verbal explanations mostly at least roughly corresponded to their responses (at least in the sense that they reflected the responses' basic realist or anti-realist direction). And fourth, having intuitions in favor of an individual subjectivism/cultural relativism brand of anti-realist pluralism makes sense from a psychological perspective. This view reflects and helps to navigate a central struggle in any person's life: the struggle between individuality (a desire to be different) and conformity (a desire to fit in).

Conclusion

In this paper we addressed the question of what ordinary people think about moral objectivity. Previous psychological research on this question has been subject to a number of methodological worries. Our aim thus was to advance this research, taking these worries into account, and specifically targeting more reflective intuitions. First, we proposed a new experimental design for measuring folk intuitions about moral objectivity that may serve as a model for future studies. Then we reported and discussed the results of a survey that was based on this design. We found that upon some reflection most of our participants denied the existence of objective moral truths. In particular, many of them had the intuition that whether moral sentences are true depends both on their own moral beliefs and on the dominant moral beliefs within their culture ("anti-realist pluralism").

What does this result imply for metaethics? We can think of several potential implications (such as for the question of what, if anything, moral sentences purport to refer to). Here, however, we again focus on the presumptive argument that we introduced at the beginning of this paper. Many philosophers accept that as ordinary people experience morality as a realm of objective truths, we have a *prima facie* reason to believe that it actually is such a realm (e.g., Dancy 1986; Enoch 2017a, 2017b; Huemer 2005; McNaughton 1988). This argument can be understood in various ways (see Loeb 2007; Pölzler forthcoming b). At least on the face of it, our results cast doubt on many of its versions. It just does not seem true that ordinary people (at least in the US) experience morality as a realm of objective truths — especially if they have given the issue some thought. In fact, if there were a philosophically valid way

of inferring a presumption from metaethical intuitions then it would rather be appropriate to grant such a presumption to anti-realism.²⁵

The failure of realists' traditional presumptive argument is significant. So far even most anti-realists have accepted their presumptive disadvantage (e.g., Blackburn 2006: 153; Brink 1989: 24, 36; Mackie 2011: 35). This has strongly influenced the debate's dialectic. While anti-realists have proposed numerous arguments against the existence of objective moral truths, realists have primarily only attempted to refute these arguments (e.g., Huemer 2005; Shafer-Landau 2006). Our study suggests that the time has come to rethink the rules of this game. Metaethicists may have to take realism and anti-realism to at least start on an equal footing. Those who advocate moral objectivity may owe us convincing positive arguments in its favour; and those who deny it may be in a better position than has long been acknowledged.

²⁵ For objections against such inferences see again Loeb 2007; Pölzler forthcoming b.

References

- Ayer, Alfred J. (1952): *Language, Truth and Logic*. New York: Dover.
- Beebe, James R. (2014): How Different Kinds of Disagreement Impact Folk Metaethical Judgments. In: Wright, Jennifer C; Sarkissian, Hagop (eds): *Advances in Experimental Moral Psychology: Affect, Character, and Commitments*. London: Bloomsbury, 167-178.
- Beebe, James R. (2015): The empirical study of folk metaethics. *Etyka* 15, 11-28.
- Beebe, James R.; Qiaoan, Runya; Wysocki, Tomasz; Endara, Miguel A. (2015): Moral objectivism in cross-cultural perspective. *Journal of Cognition and Culture* 15, 386-401.
- Beebe, James R.; Sackris, David (2016): Moral objectivism across the lifespan. *Philosophical Psychology* 29 (6): 919-29.
- Blackburn, Simon (2000): *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Oxford University Press.
- Blackburn, Simon (2006): Antirealist expressivism and quasi-realism. In: Copp, David (ed.): *The Oxford Handbook of Ethical Theory*. Oxford: Oxford University Press, 146-162.
- Brink, David O. (1989): *Moral Realism and the Foundations of Ethics*. New York: Cambridge University Press.
- Cova, Florian; Ravat, Jérôme (2008): Sens commun et objectivisme moral. Une introduction par l'exemple à la philosophie expérimentale. *KLESIS – Revue philosophique* 9, 180-202. English translation available at: https://www.academia.edu/1884720/Common_sense_and_moral_objectivism_Global_or_Local_A_Case_Study_of_Experimental_Philosophy.
- Dancy, Jonathan (1986): Two conceptions of moral realism. *Proceedings of the Aristotelian Society*, Supplemental Volume 60, 167-187.
- Enoch, David (2017a): Non-naturalistic realism in metaethics. In: McPherson, Tristram; Plunkett, David (eds.): *The Routledge Handbook of Metaethics*. New York: Routledge.
- Enoch, David (2017b): Why I am an objectivist about ethics (and why you are, too). In: Shaffer Landau, Russ (ed.): *The Ethical Life*. Oxford: Oxford University Press.
- Evans, Stephen C. (2013): *God and Moral Obligation*. Oxford: Oxford University Press.
- Fisher, Matthew; Knobe, Joshua; Strickland, Brent; Keil, Frank C. (forthcoming): The influence of social interaction on intuitions of objectivity and subjectivity. *Cognitive Science*.

- Goodwin, Geoffrey P.; Darley, John M. (2008): The psychology of meta-ethics: Exploring objectivism. *Cognition* 106 (3), 1339–1366.
- Goodwin, Geoffrey P.; Darley, John M. (2010): The perceived objectivity of ethical beliefs: Psychological findings and implications for public policy. *Review of Philosophy and Psychology* 1 (2), 161-188.
- Goodwin, Geoffrey P.; Darley, John M. (2012): Why are some moral beliefs perceived to be more objective than others? *Journal of Experimental Social Psychology* 48 (2012), 250-256.
- Hannon, Michael (2018): Intuitions, reflective judgments, and experimental philosophy. *Synthese* (online first).
- Harman, Gilbert (1996): Moral relativism: In: Harman, Gilbert; Thompson, Judith J. (eds.): *Moral Relativism and Moral Objectivity*. Cambridge, MA: Blackwell, 1-64.
- Heiphetz, Larisa; Young, Liane L. (forthcoming): Can only one person be right? The development of objectivism and social preferences regarding widely shared and controversial moral beliefs. *Cognition*.
- Henrich, Joseph; Heine, Steven J.; Norenzayan, Ara (2010): The weirdest people in the world. *Behavioral and Brain Sciences* 33 (2-3), 61-83.
- Huemer, Michael (2005): *Ethical Intuitionism*. Basingstoke: Palgrave MacMillan.
- Joyce, Richard (2001): *The Myth of Morality*. Cambridge: Cambridge University Press.
- Joyce, Richard (2007): Moral anti-realism. In: Zalta, Edward N. (ed.): *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/moral-anti-realism/>.
- Kauppinen, Antti (2007): The rise and fall of experimental philosophy. *Philosophical Explorations* 10 (2), 95-118.
- Loeb, Don (2007a): Moral statements for error theorists. In: *PEA Soup — A blog dedicated to philosophy, ethics, and academia*. <http://peasoup.typepad.com/peasoup/2007/09/moral-statement.html>.
- Loeb, Don (2007b): The argument from moral experience. *Ethical Theory and Moral Practice* 10 (5), 469–484.
- Lynch, Michael P. (2013): Truth in Ethics. In: LaFollette, Hugh (ed.): *International Encyclopedia of Ethics*. Oxford: Wiley-Blackwell.
- Mackie, John L. (2011): *Ethics: Inventing Right and Wrong*. London: Penguin.

- McNaughton, David (1988): *Moral Vision: An Introduction to Ethics*. Oxford: Blackwell.
- Nichols, Shaun (2004): After objectivity: An empirical study of moral judgment. *Philosophical Psychology* 17 (1), 3-26.
- Nichols, Shaun; Folds-Bennett, Trisha (2003): Are children moral objectivists? Children's judgments about moral and response-dependent properties. *Cognition* 90 (2), B23-32.
- Nichols, Shaun; Knobe, Joshua (2008): Moral responsibility and determinism: The cognitive science of folk intuitions. *Nous* 41, 663-685.
- Pigden, Charles R. (2007): Nihilism, Nietzsche and the doppelganger problem. *Ethical Theory and Moral Practice* 10 (5), 441-456.
- Plato ([380 B.C.] 2010): Euthyphro. In: Jowett, Benjamin (ed.): *Dialogues of Plato. Translated into English, with Analyses and Introduction. Vol. 2*. New York: Cambridge University Press, 293-318.
- Pölzler, Thomas (2018): *Moral Reality and the Empirical Sciences*. New York: Routledge.
- Pölzler, Thomas (forthcoming a): How to measure moral realism. *Review of Philosophy & Psychology*.
- Pölzler, Thomas (forthcoming b): Psychologische Forschung zum moralischen Realismus und ihre philosophischen Implikationen. In: Paulo, Norbert; Bublitz, Christoph J.: *Experimentelle Philosophie*. Berlin: Suhrkamp.
- Prinz, Jesse J. (2007): *The Emotional Construction of Morals*. Oxford: Oxford University Press.
- Quinn, Phillip L. (2000): Divine command theory. In: LaFollette Hugh (ed.): *The Blackwell Guide To Ethical Theory*. Oxford: Blackwell, 53-73.
- Raz, Joseph (2001): Notes on value and objectivity. In: Leiter, Brian (ed.): *Objectivity in Law and Morals*. Cambridge: Cambridge University Press.
- Sarkissian, Hagop; Tien, David; Wright, Jennifer; Knobe, Joshua (2011): Folk moral relativism. *Mind and Language* 26 (4), 482-505.
- Shafer-Landau, Russ (2003): *Moral Realism: A Defense*. Oxford: Oxford University Press.
- Shafer-Landau, Russ (2006): Ethics as philosophy: A defense of ethical nonnaturalism. In: Horgan, Terry; Timmons, Mark (eds.): *Metaethics After Moore*. New York: Oxford University Press, 209-232.

- Sinnott-Armstrong, Walter (2008): Abstract + concrete = paradox. In: Knobe, Joshua; Nichols, Shaun (eds.): *Experimental Philosophy*. New York: Oxford University Press, 209-230.
- Wainryb, Cecilia; Shaw, Leigh; Langley, Marcie; Cottam, Kim; Lewis Renee (2004): Children's thinking about diversity of belief in the early school years: judgments of relativism, tolerance, and disagreeing persons. *Child Development* 75 (3), 687-703.
- Wright, Jennifer C. (2018): The Fact and Function of Meta-Ethical Pluralism: Exploring the Evidence. In: Lombrozo, Tania; Knobe, Joshua; Nichols, Shaun (eds.): *Oxford Studies in Experimental Philosophy, Vol. 2*. Oxford: Oxford University Press. Cited according to https://www.academia.edu/12324262/The_Fact_and_Function_of_Meta-Ethical_Pluralism_Exploring_the_Evidence.
- Wright, Jennifer C.; Cullum, Jerry; Schwab, Nicholas (2008): The cognitive and affective dimensions of moral conviction: Implications for tolerance and interpersonal behaviors. *Personality and Social Psychology Bulletin* 34(11), 1461-1476.
- Wright, Jennifer C.; Grandjean, Piper T.; McWhite, Cullen B. (2013): The meta-ethical grounding of our moral beliefs: Evidence for meta-ethical pluralism. *Philosophical Psychology* 26 (3), 336-361.
- Wright, Jennifer C.; McWhite, Cullen B.; Grandjean, Piper T. (2014): The cognitive mechanisms of intolerance: Do our meta-ethical commitments matter? In: Knobe, Joshua; Lombrozo, Tania; Nichols, Shaun (eds.): *Oxford Studies in Experimental Philosophy, Vol. 1*. Oxford: Oxford University Press, 28-61.
- Young, Liane; Durwin, A. J. (2013): Moral realism as moral motivation: The impact of meta-ethics on everyday decision-making. *Journal of Experimental Social Psychology* 49 (2), 302-306.
- Ziljstra, Lieuwe (forthcoming): Folk moral objectivism and its *measurement*. *Manuscript under review*.