

ARTICLE



The irrationality of folk metaethics

Ross Colebrook

Baruch College, City University of New York, New York, NY, USA

ABSTRACT

Many philosophers and psychologists have thought that people untutored in philosophy are moral realists. On this view, when people make moral judgments, they interpret their judgments as tracking universal, objective moral facts. But studies of folk metaethics have demonstrated that people have a mix of metaethical attitudes. Sometimes people think of their moral judgments as purely expressive, or as tracking subjective or relative moral facts, or perhaps no facts at all. This paper surveys the evidence for folk metaethical pluralism and argues for an explanation of this mix of folk metaethical attitudes: without philosophical education, these attitudes are typically caused by factors that are insensitive to their truth. Moreover, unless they can be justified by other means, metaethical attitudes with this etiology are, as a result, irrational, and ought not be used as evidence for or against moral realism.

ARTICLE HISTORY

Received 12 June 2018
Accepted 15 January 2021

KEYWORDS

Moral psychology;
metaethical pluralism; moral
realism; folk metaethical;
Irrationalism; Folk
Metaethics

Introduction

Your sense of right and wrong defines your identity, drives your emotions, and guides your actions. Moral judgments are among the strongest motivators in our personal, social, and political lives. But despite the centrality of moral judgments, we seldom stop to reflect on what our moral judgments are, what facts they are supposed to be tracking, and what would make them true or false. In fact, most of us have no settled, standing opinions about what we're doing when we make moral judgments. Though this may appear to be a truism, significant work in philosophy and psychology has disputed it.

In this paper I will argue that folk metaethical attitudes are as the truism suggests: unsettled, uncertain, and unprincipled. I will focus on a feature of folk metaethics that has most vexed philosophers and moral psychologists: whether ordinary people are moral realists. In §1 I will specify what metaethical attitudes are and lay out the major positions in the recent debate over the proper conception of folk metaethics. Then in §2 I will explore

recent research on folk metaethics and argue that this work supports *Folk Metaethical Pluralism*: people typically have a mix of realist and antirealist attitudes. In §3 I bolster the case for *Folk Metaethical Pluralism* by providing an explanation of it: *Folk Metaethical Irrationalism*. This hypothesis claims that folk metaethical attitudes are typically caused by factors that are insensitive to their truth. Finally, in §4, I will explore some of the implications of this hypothesis for metaethics.

§1 What are metaethical attitudes?

There are a variety of second-order attitudes we can take up toward our own moral judgments. One such family of attitudes – realist attitudes – can be summed up with at least four distinct theses about what moral judgments are supposed to be tracking and whether that tracking is successful. Broadly speaking, if moral judgments are tracking anything at all, they are tracking “moral facts,” like the wrongness of an act of murder, or the goodness of charity.¹ This family of realist attitudes corresponds with the following broad metaethical claims:

- (A) *Factualism*: some moral facts exist.
- (B) *Objectivism*: the existence and nature of some moral facts are independent of our attitudes.
- (C) *Universalism*: some of the same moral facts obtain in every culture.
- (D) *Cognitivism*: some moral judgments are truth-apt.

Though all four of these metaethical claims are often held by moral realists, I do not intend to imply that every one of them is necessary for a person to be a moral realist. Various philosophical accounts construe combinations of (A) through (D) as being definitive of moral realism (e.g. Boyd, 1988; Brink, 1989; Devitt, 2002; Enoch, 2011; Huemer, 2005; Jackson, 1998; Prinz, 2007; Railton, 1986; Sayre-McCord, 1988; Shafer-Landau, 2003; Smith, 1994). For our purposes, “realist attitudes” are metacognitive attitudes of types (A) through (D) that have first-order moral judgments as their semantic targets.

Similarly, antirealist attitudes are metacognitive attitudes that have first-order moral judgments as their semantic targets, but different content. This family of attitudes that can be summed up with at least four distinct metaethical claims. Once again, these claims are about what moral judgments are thought to be tracking and whether that tracking is successful:

- (A) *Non-Factualism*: no moral facts exist.
- (B) *Subjectivism*: the existence and nature of moral facts always depend on our attitudes.
- (C) *Relativism*: none of the same moral facts obtain in every culture.

(D) *Non-Cognitivism*: no moral judgments are truth-apt.

Note that each of these metaethical claims can, theoretically, be held in combination with any other, except non-factualism, which can only be held in combination with (D) or (H) (e.g. the combination of (D) and (E) represents moral error theory) (e.g. Joyce, 2001; Mackie, 1990). Some authors combine subjectivism and relativism, but there are reasons to think these notions are dissociable, as some philosophers have advanced subjectivist-universalist views (e.g. Korsgaard & O'Neill, 1996) and what might be described as objectivist-relativist views (e.g. Wong, 2006).²

To have a metaethical attitude, it is not necessary that a person conceive of her first-order moral judgments in the philosophical terms above. In fact, as I will cover later, one may have metaethical attitudes without ever consciously thinking about them. The above claims represent philosophical positions, but it is likely that the logical structure that best represents actual metaethical attitudes will vary. These attitudes may have a wide scope (the same attitude may have many, or even all first-order moral judgments as its semantic target), or they may have a narrow scope (an attitude may have a single first-order moral judgment as its semantic target). For instance:

(A_w) *Factualism_w*: Moral judgments M_{1-N} track facts F_{1-N}.

(A_n) *Factualism_n*: Moral judgment M tracks fact F.

Or alternatively,

(E_w) *Non-Factualism_w*: Moral judgments M_{1-N} do not track any facts F_{1-N}.

(E_n) *Non-Factualism_n*: Moral judgment M does not track fact F.

Until recently, these metaethical attitudes were not investigated on their own. Instead, philosophers and psychologists often considered realist attitudes to be conceptually necessary properties of first-order moral judgments (Hollos et al., 1986; Lawrence Kohlberg, 1981; Nucci, 2001; Nucci & Turiel, 1978, 1993; Nucci et al., 1983; Smetana, 2006; Smetana & Braeges, 1990; E. Turiel, 1979; Elliot Turiel, 1983).³ However, more recent work has discarded this assumption, and has gone on to examine folk metaethical attitudes directly.

The animating theoretical question in this debate is whether philosophically naïve people typically have realist attitudes by default, or if they typically have both realist and antirealist attitudes. As such, the two major positions in this debate might be called the following:

Folk Metaethical Realism: people typically have realist attitudes toward their own moral judgments.

Folk Metaethical Pluralism: people typically have a mix of realist and antirealist attitudes toward their own moral judgments.⁴

This paper will argue against the former and will aim to both support and explain the latter. This will require a detailed examination of many of the studies that bear on folk metaethics, explored in §2, and explaining *Folk Metaethical Pluralism* will involve a new thesis, argued for in §3:

Folk Metaethical Irrationalism: metaethical attitudes (both realist and antirealist) are typically caused by processes that are insensitive to the truth of these attitudes.

As I will argue, variations in metaethical attitudes are correlated with several factors that are rationally disconnected from the potential truth or falsity of these attitudes. Many of these factors plausibly cause metaethical attitudes, too, and this helps to explain why folk metaethics is a mix of realist and antirealist attitudes. Moreover, unless they can be justified by other means, metaethical attitudes produced this way are, as a result, irrational, and ought not be used as evidence for or against moral realism.

§2 Evidence for folk metaethical pluralism

In this section I will survey the evidence in favor of *Folk Metaethical Pluralism*, and against *Folk Metaethical Realism*. In short, in the central studies that examine folk metaethics, substantial minorities (and in some contexts, majorities) of participants express antirealist attitudes of various kinds. One possible response would be to try to explain these findings away; perhaps people are expressing attitudes that they do not really have or would give up after some reflection. The alternative, which I think is a better account of the evidence, is that the mix of attitudes that these studies have found is in fact the default folk metaethical view.

§2.1 Moss (2017), Pözlner (2018), and Sinclair (2012)

There are two major hurdles to overcome before beginning this survey of studies on folk metaethics, both epistemic and metaphysical. Starting with the epistemic, Moss (2017) and Pözlner (2018a) point out that studies that assess folk metaethics have not developed valid measures of various folk metaethical attitudes, either because participants might be misunderstanding questions aimed at eliciting these attitudes, or because the questions that researchers ask often elicit the wrong attitudes. This is a serious objection to the way folk metaethics has been investigated, and it casts doubt on whether we can draw any inferences at all about folk metaethical attitudes. However, these studies often plausibly capture some metaethical attitudes, even though they seldom reach the fineness of grain necessary to distinguish participants that hold specific variations of attitudes captured by (A) through (H) above.⁵ And in many of these cases, despite the fact that these studies are flawed, the best explanation of participant responses is

that they are being driven by antirealist attitudes, as captured by (E) through (H).

Some philosophers go beyond these methodological worries, and doubt that ordinary people have metaethical attitudes at all. Sinclair (2012), for instance, argues that metaethical attitudes do not show up in moral phenomenology. Even if people tend to see the world as “morally mandating” certain actions, this does not necessarily mean that people see these mandates as universal or objective (Sinclair, 2012, p. 167). According to Sinclair, “[t]he categories that define metaethical theories are to a large extent philosopher’s constructs rather than tacit categories of everyday moralisers” (p. 168). He argues that the kinds of semantic claims that realists make are too complex for us to posit non-philosophers having them (p. 168).

Though it is obvious that people untutored in philosophy do not have thoroughly conceived views on moral semantics or metaphysics, this argument conflates two separate issues in the study of folk metaethics. There are indeed a variety of complex philosophical interpretations of realism and its implications. But the folk do not need to be aware of these interpretations to be tacitly assuming them. Consider that non-physicists are plausibly realists about physical objects, even if they do not understand the semantics that allow them to talk about those objects, or the metaphysics that ground those objects. After all, people will attempt to resolve disagreements about medium-sized objects by adverting to apparently objective features of those objects, rather than accepting faultless disagreement. If two carpenters disagree about the length of a table, they measure it. Moreover, unless one is willing to endorse controversial commitments about cognitive phenomenology, it is not obvious why metaethical attitudes should be expected to have a distinctive phenomenology in the first place. Rather, we ought to posit realist attitudes when people have the disposition to act, reason, and speak as though their judgments are true of objective, universal facts. The same ought to hold, *mutatis mutandis*, for antirealist attitudes (see also Zijlstra, 2019).⁶

Having flagged these methodological issues, I will begin the survey of studies on folk metaethics. Many of the earlier studies on folk metaethics appeared to show that people are realists by default. But when we look more closely, we will see that these studies often contain distinct minorities that express antirealist attitudes or fail to give participants who have antirealist attitudes a chance to express them.

§2.2 Nichols and Folds-Bennett (2003)

I’ll begin with Nichols and Folds-Bennett (2003), which purports to show that children are moral objectivists. This study compares children’s treatment of paradigmatically response-dependent properties (i.e. properties

that depend on our attitudes for their existence and nature), with children's treatment of moral properties (e.g. "good" or "bad"). Children of ages four through six years old were given six different questions, two of which involved a moral property ("good"), two of which involved an esthetic property ("beautiful"), and two of which involved properties often regarded as response-dependent ("yummy" and "fun") (Nichols & Folds-Bennett, 2003, pp. 27–28). The children were asked whether a property applied to something (e.g. whether grapes were yummy). Then, to measure whether the children treated the property as depending on their preferences or not, they were asked another question: if someone disagreed with them, would the property apply "for some" or "for real." For example, the prompt for the response-dependent property "yummy" was the following: "You know, I think grapes are yummy too. Some people don't like grapes. They don't think grapes are yummy. Would you say that grapes are yummy for some people or that they're yummy for real?" (p. 27).

Nichols & Folds-Bennett then scored children's responses, with participants scoring 0 for answering "for some" and 1 for answering "for real." Adding responses for each domain together gave a total "preference-independence" score for each domain between 0 and 2. They then compared responses between domains, and found that both moral and esthetic judgments were treated as significantly more preference-independent than the response-dependent judgments (Nichols & Folds-Bennett, 2003, p. 27). The same results were found when the study was repeated with negatively valenced properties (e.g. "bad" and "icky") (p. 27–30).

On first inspection, this study appears to support *Folk Metaethical Realism*. After all, the children in the experiment seem to be treating moral properties as though they do not depend on our preferences, which is an expression of objectivism. There are several reasons why we should not draw this conclusion, though. First, though most children categorized moral properties as preference-independent, this was not the case for all of them. Of the nineteen children in the study, three categorized the moral properties as unambiguously preference-dependent, and four answered in an ambiguous way (judging one moral property as preference-dependent and another moral property as preference-independent). Overall, children expressed a pattern of objectivist responses about moral properties, and subjectivist responses about putatively response-dependent properties. However, though there were few antirealist responses in this study, we cannot draw the conclusion that children have realist attitudes by default while a minority of apparently subjectivist answers persist unexplained. Moreover, because realist and antirealist attitudes come in many flavors, other types of antirealist attitudes might have been driving the children's responses. Answering that a given property applies "for some" versus "for real" does not adequately

distinguish between relativism and subjectivism. As a result, for example, a child with implicit objectivist-relativist attitudes might see no difference between answering that a monkey helping another monkey is good “for some” and good “for real” (Nichols & Folds-Bennett, 2003, p. 26). On such a view, it is true that one monkey helping another is good “for some” (that fact may only obtain in our culture) but also good “for real” (that fact obtains in our culture, regardless of what our culture thinks). Because this study does not distinguish between these different forms of realism and antirealism, there are permutations of attitudes it may be failing to measure.

§2.3 Wainryb et al. (2004) and Nichols (2004)

Wainryb et al. (2004) also tested children’s metaethical attitudes, this time by presenting participants with disagreements about moral judgments, taste judgments, factual judgments, and ambiguous factual judgments (p. 691). After being given a disagreement between two characters, children aged five, seven, and nine were asked, among other questions, “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, p. 691). They then scored responses as either 1 (indicating non-relativism) or 2 (indicating relativism) (p. 692). They found that nearly all participants made non-relative judgments about both moral judgments and factual judgments.

This study also seems to be strong evidence in favor of *Folk Moral Realism*. But like Nichols and Folds-Bennett (2003) it is not clear that it can validly detect antirealist attitudes. Though it may appear as though participants choosing the first disjunct are expressing realist attitudes, we cannot infer this. The prompts for participants do not specify Sophie and Sarah’s cultural backgrounds, so participants with tacit relativist attitudes might have judged that only one of them could be right – assuming Sophie and Sarah come from the same culture – or that both could be right, assuming they each come from different cultures (see also Pölzler, 2017, pp. 464–466, 2018, p. 647).

Another early study that bears on folk metaethics is Nichols (2004). Nichols provided participants with four scenarios where two people (John and Fred) from different cultures have a disagreement, involving a moral scenario, a factual scenario, a disgust scenario, and a conventional scenario.⁷ Participants were given the option of responding to disagreements in one of three ways:

- It is okay to hit people just because you feel like it, so John is right and Fred is wrong.

- It is not okay to hit people just because you feel like it, so Fred is right and John is wrong.
- There is no fact of the matter about unqualified claims like “It’s okay to hit people just because you feel like it.” Different cultures believe different things, and it is not absolutely true or false that it’s okay to hit people just because you feel like it (Nichols, 2004, pp. 9–10).

The same type of response options were given for the disagreement about the other scenarios. Nichols found that six of the participants (13%) gave the last, non-objectivist answer to the factual statement, and, to control for global subjectivism, excluded them from analysis. Even after removing these participants from consideration, Nichols found that seventeen participants (42%) gave non-objectivist responses to the moral scenario (Nichols, 2004, p. 10). In a subsequent experiment, a significant portion of participants once again expressed non-objectivist responses (excluding the twelve global subjectivists, 26% of participants chose a subjectivist answer to the ethical question) (p. 18). This experiment added another scenario where a human and an alien disagree about the permissibility of torturing puppies, including the same response options. Despite the emotionally charged nature of the scenario, participants were more likely to choose the non-objectivist answer (38%) (p. 18).

This study also has significant limitations. As Beebe (2015) points out, the first and second responses are more plausibly construed as first-order judgments, and the third possible response forces participants to suspend judgment if they want to say there is no fact of the matter about the action being right or wrong. This prevents participants from expressing many antirealist attitudes. For instance, some participants may have wanted to both express a first-order moral judgment and a subjectivist metaethical judgment, but would not have been able to agree to the third option (Beebe, 2015, pp. 17–18; Pölzler, 2018a, pp. 657–658). Moreover, as Pölzler (2017) argues, this experimental design excludes participants who may have wanted to express non-factualist or non-cognitivist attitudes. Non-factualists would not be able to endorse the first two choices, and could not accept the third choice, since it might be taken to imply that there *are* facts about right and wrong actions, assuming they are properly relativized, and this is a judgment that non-factualists would reject. Non-cognitivists might have been able to accept either of the first answers, as long as they are interpreted “it is/is not okay to hit people just because you feel like it” as an expression of feeling, or an imperative, but this would be an expression of an antirealist attitude, and the experimental design treats such a response as a realist attitude (Pölzler, 2017, p. 467).

These methodological problems might lead one to conclude that Nichols and Folds-Bennett (2003), Wainryb et al. (2004), and Nichols (2004) can’t

give us any evidence that might tell between *Folk Metaethical Pluralism* and *Folk Metaethical Realism*. But this is not the case. The general problem in each of these studies is that they prevent participants from expressing antirealist attitudes. After all, a hypothetical participant who has all the realist attitudes discussed in §1 would have no problem expressing their attitudes in each of these studies. Such a participant could easily agree, for example, with the first or second option in Nichols (2004). Suppose the hypothetical participant chose the second option in the experiment (“it is not okay to hit people just because you feel like it”). Even if we suppose that she interprets this wording as expressing a first-order judgment, the second option also includes a conclusion: “so Fred is right and John is wrong” (Nichols, 2004, pp. 9–10). A realist could easily use this to express her realist attitudes (both objectivist and universalist).

In each of these studies, realists are not prevented from expressing their attitudes, but various forms of antirealist attitudes cannot be expressed, given the wording of the questions. What this shows, then, is that the criticism of these studies shouldn’t be taken to imply that they show nothing about folk metaethics. At least two of these studies, Nichols and Folds-Bennett (2003) and Nichols (2004), do show that a substantial minority of participants will express some form of antirealism about some moral judgments, even when the experimental design prevents other participants from expressing other types of antirealist attitudes.

§2.4 Goodwin and Darley (2008, 2012)

Yet more evidence in favor of *Folk Metaethical Pluralism* comes from another study that purports to support *Folk Metaethical Realism* – Goodwin and Darley (2008). Goodwin and Darley (2008) gave participants a series of factual, ethical, and conventional statements.⁸ In the first phase of the experiment, the researchers asked whether the participants agreed with the statement (on a six-point scale), and asked them to judge whether the statement the statement was (1) true, (2) false, or (3) an opinion or attitude (p. 1344). Researchers then selected five of the statements (two ethical, one conventional, one taste, and one factual) for a second phase of the experiment. In the second phase, participants were told that no statement on the list produced 100% agreement and they were asked to interpret their apparent disagreement with other participants in the study. They were given four options: (1) the other party is surely mistaken, (2) it is possible that neither you nor the other person is mistaken, (3) it could be that you are mistaken, and the other person is correct, and (4) other (p. 1344).

Some of the responses in this experiment seem to be evidence for metaethical pluralism. Notably, participants by and large refused to say that their judgments on the ethical statements were true in cases where

the statement was controversial, despite agreeing with those statements. So, for example, only between 2% and 8% of participants rated statements on abortion, euthanasia, and stem cell research to be true, despite agreeing with these statements. Instead, participants preferred to describe their judgments as opinions. For less controversial items, participants were more comfortable describing their own judgments as true, (for example, 61% of participants rated a statement condemning robbery as true).

This pattern of responses is best explained by metaethical pluralism. After all, this is a case where people's metaethical intuitions appear to vary dramatically, and many participants are comfortable describing their judgments as mere opinions. Yet Goodwin and Darley (2008) do not draw that conclusion. In fact, their conclusion is that ethical statements were treated by participants as more objective than both conventional and preference statements, and only treated less objectively than factual statements (p. 1348). This is because Goodwin & Darley operationalize a participants' objectivity about a judgment with a scale that combines results from the first phase of the experiment with the results from the second phase. If participants view their own judgments as true, and view those who disagree with them as surely mistaken, this counts as a 'fully objective' answer, and is given a score of '3' on the objectivity scale. Participants who interpret their judgment as a matter of opinion, and then say that it is possible that neither they nor the other person is mistaken are treated as 'least objective,' and this counts as a score of 1 on the objectivity scale. This leaves two more possibilities (shown in Table 1 below).⁹ When scores are combined in this way, Goodwin & Darley get the result that judgments about ethical statements ($M = 2.56$) rank somewhat below factual statements in objectivity ($M = 2.91$), and somewhat above judgments about conventional statements ($M = 2.00$) and statements of preference ($M = 1.34$) (Goodwin & Darley, 2008, p. 1348).

The major problem with this scale derives from its treatment of the two 'intermediately objective' categories. If a participant gives answers which fall into either of these categories, her level of objectivity for a given statement is

Table 1. Operationalization of objectivist attitudes in experiment 1, Goodwin and Darley (2008). Subscripts added.

Participants' Metaethical Attitude	Response in Phase 1	Response in Phase 2	Objectivity Scale Score
Fully Objective	Statement is true.	The other party (who disagrees) is surely mistaken.	3
Intermediately Objective ₁	Statement is true.	It is possible that neither you nor the other person is mistaken.	2
Intermediately Objective ₂	Statement is an opinion or attitude.	The other party is surely mistaken	2
Least Objective	Statement is an opinion or attitude.	It is possible that neither you nor the other person is mistaken.	1

treated the same, whether she gave one or the other combination of answers. Yet these two positions are in fact quite distinct – they are not properly treated as positions between full realism and full antirealism. Of course, Goodwin & Darley are aware of the fact that these positions are conceptually separable: they point out that *intermediately objective*₁ is consistent with philosophical subjectivism, and that *intermediately objective*₂ might involve an “oscillation (or possibly confusion) regarding a statement’s objectivity” (p. 1345). But much more can be said about these categories. *Intermediately objective*₁ would be endorsed by any participant who holds a subjectivist position. Yet Goodwin & Darley’s methodology does not lead to the conclusion that one might expect: that participants whose answers fall into this category have expressed an antirealist attitude, just as much as participants who fall into the *least objective* category. Moreover, Goodwin & Darley leave out an alternative interpretation of *intermediately objective*₂ which the participants may have had in mind when making their judgments: they were not confused about the judgment’s objectivity, but rather saw little difference between “true” (for statements they agreed with) and “opinion.” Goodwin & Darley clearly intended the word to imply that there was no right or wrong answer to the question, but this is not necessarily how participants saw it. After all, using the word “opinion” to signify a belief that might be true or false is a common usage (see also Beebe, 2015, p. 13). If this is how participants interpreted the word, then participants rated as *intermediately objective*₂ might have intended to express realist attitudes. The experiment thus fails to adequately operationalize objectivity about the statements in question, since it treats as identical two positions that have very different cognitive and metaethical significance (see also Moss, 2017, pp. 188–191; Pözlner, 2017, pp. 461–463, 2018, pp. 454–455).¹⁰

If we were to treat *intermediately objective*₁ as equivalent to *least objective* and *intermediately objective*₂ as *fully objective*, this leads to a somewhat different interpretation of Goodwin & Darley’s data.¹¹ For this first experiment, if we interpret the data on a binary, rather than three-pointed scale, 67% of responses to ethical questions count as *fully objective*, and 22% are *least objective*, with the remainder excluded due to answering ‘other’ in the second phase of the experiment. By comparison, for the factual statement, 92% of responses count as fully objective, and 6% count as *least objective*.¹² This makes the difference between participant responses for moral and factual statements look much more pronounced.

The second experiment in Goodwin and Darley (2008) is much the same as the first, with one notable exception: in the first phase, rather than asking whether a statement is true, false, or an opinion, they asked participants whether there can be a correct answer as to whether the statement is true (pp. 1351–2). This second experiment is an improvement over the first, as it excludes the possibility that participants interpreted “opinion” as an

Table 2. Operationalization of objectivist attitudes in experiment 2, Goodwin and Darley (2008). Subscripts added.

Participants' Metaethical Attitude	Response in Phase 1	Response in Phase 2	Objectivity Scale Score
Fully Objective	Yes, there can be a correct answer to whether this statement is true.	The other party (who disagrees) is surely mistaken.	3
Intermediately Objective ₃	Yes, there can be a correct answer to whether this statement is true.	It is possible that neither you nor the other person is mistaken.	2
Intermediately Objective ₄	No, there cannot be a correct answer to whether this statement is true.	The other party is surely mistaken	2
Least Objective	No, there cannot be a correct answer to whether this statement is true.	It is possible that neither you nor the other person is mistaken.	1

expression of a belief. Again, though, we have reason to doubt that *intermediately objective*₃ and *intermediately objective*₄ constitute theoretically grounded intermediate positions. *Intermediately objective*₃ could be accepted by any subjectivist, since, for a subjectivist, a statement is true “for her” despite the fact that, speaking objectively, neither she nor a disagreeing party need be wrong. And a non-cognitivist could accept *intermediately objective*₄. On such a view, moral statements are not true or false (because moral judgments are purely emotive, expressive, or prescriptive), yet a disagreeing party could be “mistaken” in the second phase, insofar as her moral judgment emotes, expresses, or prescribes differently. These non-cognitivist options have to be considered because there is evidence that some people have non-cognitivist interpretations of their own moral judgments (Beebe, 2015, pp. 25–28; Moss, 2017, pp. 188–191; Pözlner, 2017, pp. 461–463, 2018a, pp. 454–455; Wright, 2018, pp. 135–141), (see Table 2).

When we reconceptualize the data along the lines just suggested, we again get a different picture of folk metaethics from Goodwin & Darley’s data. For ethical statements, 70% of responses fell into the *fully objective* category (Goodwin & Darley, 2008, p. 1353). If each of the intermediately objective categories in this experiment is treated as *least objective*, 25% of participants would be classified as *least objective*, and another 5% cannot be classified. By comparison, for factual statements, 71% of responses were *fully objective* and 23% were *least objective*. The most striking fact about this result (compared with the first experiment) is that participants expressed far less objectivism about factual statements. This can be explained by the fact that the researchers (to test the effect of controversial statements) substituted a factual question about geography with a factual question about evolution for some respondents in the second phase of the experiment. In the first experiment 92% received a statement claiming that Boston is further north than Los Angeles. By contrast, in the second experiment, only 20% received this question, and 67% received the statement that “Homo sapiens evolved from more primitive primate species” (Goodwin & Darley, 2008, p. 1368). Despite agreeing to the statement, those who received this more

controversial statement expressed significantly fewer objectivist judgments.¹³ (I will return to the effects of controversy on metaethical attitudes in the next section).

In a later study, Goodwin & Darley employ somewhat different methodology (Goodwin & Darley, 2012). Rather than measuring objectivity by combining responses from two categorical variables (as in experiment 1 and 2), these experiments measure objectivity by combining two interval variables (both on a scale of 1–6). The first variable measures the extent to which participants believe there is a correct answer about whether a statement is true. Goodwin & Darley then tell participants that another participant has disagreed with them, and the second variable measures the extent to which participants think a disagreeing other needs to be mistaken (1 = neither of us need to be mistaken, 6 = the other person is clearly mistaken) (Goodwin & Darley, 2012, p. 251). They then average these two measures to derive a combined six-point objectivity scale. They again find that factual statements are treated as more objective ($M = 5.23$) than moral statements ($M = 4.25$), but more objective than conventional statements ($M = 3.16$) and matters of taste ($M = 1.60$) (p. 252).

There are a few issues with this latter study that prevent it from supporting *Folk Metaethical Realism*. First, it is unclear whether treating objectivity as an interval rather than categorical variable is theoretically warranted. Unlike attitude strength or certainty, which are not plausibly binary mental states, metaethical attitudes may in fact be binary. That is, we may either have them or not have them, not have them to some extent or other. Second, in terms of measurement, treating objectivity as an interval value introduces the possibility that many participants gave anti-realist answers, but that these answers are being obscured by the mean measurement of objectivity.

To briefly reiterate, Goodwin & Darley's hypothesis is that, overall, ethical statements are treated as more objective than conventional and preference statements, but somewhat less objective than factual statements. Supporting this hypothesis is complicated by their operationalization of metaethical attitudes, which unjustifiably treats some anti-realist positions as intermediate between objectivism and subjectivism. Yet even if an improved methodology could reveal that, on average, people treat moral statements as more objective than conventional or taste statements, this would not be enough to support *Folk Metaethical Realism*. The best defense of *Folk Metaethical Realism* would require explaining the cases that deviate from this generalization (i.e. the relativist attitudes participants often express, even if those same participants usually have realist attitudes toward moral judgments in general). Future research on folk metaethics must devise measures that help us explain these antirealist responses.

§2.5 Wright et al. (2013), Wright et al. (2014) & Wright (2018)

One possible explanation for the prevalence of antirealist attitudes might be that participants don't interpret some issue-items as moral in the first place, and so do not have realist metaethical attitudes toward them. Using methods similar to Goodwin and Darley (2008), Wright et al. (2013) set out to test whether participants would express more objectivist judgments if given the chance to classify issue items for themselves. Participants were asked to classify each issue item as either (1) a moral issue, (2) a social convention/norm, (3) a personal preference, or (4) a scientific fact, and the study employed the same two-phase question setup as Goodwin and Darley (2008), combining the scores from two measures into a "grounding" score that ranged from 0 (nonobjective) to 2 (objective) (Wright et al., 2013, p. 340).

The results, however, reinforce *Folk Metaethical Pluralism*. Wright et al. found that participants had more antirealist attitudes toward issue-items they rated as personal ($M = .32$), and realist attitudes toward items they judged scientific ($M = 1.82$). Social and moral issue items fell in between these items, and could not be statistically differentiated ($M = 1.32$ vs $M = 1.26$ respectively) (Wright et al., 2013, pp. 341–342). And though they confirmed that participants' judgments about which issue-items counted as moral (as opposed to personal, social, or scientific) were much the same as Goodwin and Darley (2008) supposed, there was more disagreement than might have been expected, a priori. A substantial portion of participants judged self-classified moral issues in an antirealist way (e.g. assisted suicide was judged moral by 61% of participants, yet 60% of participants scored 0 on the combined measure of objectivity) (Wright et al., 2013, p. 340).

In order to preclude the possibility that participants were confused, and issuing antirealist judgments that they *thought* were realist judgments, Wright et al. (2013) conducted a second study that asked participants about their justifications for judgments about various issue-items, then coded such responses as relative (antirealist) or objective (realist). They found significant confusion: many participants provided realist justifications for antirealist attitudes, and vice-versa. Even when taking participants' justifications as measures of their realist (or antirealist) attitudes instead of their pattern of responses to potential disagreement, a significant portion of participants still expressed antirealist attitudes about various issue-items (Wright et al., 2013, pp. 343–353).

Wright et al. (2014) expands on these results, finding that the "only domain with significant grounding variation was the moral domain: of the thirteen issues classified as moral, seven were (dominantly) given objective groundings, one was mixed, four were nonobjective, and one was split

between all three” (p. 36).¹⁴ The same study also tested many of the issue items that Goodwin & Darley defined as moral, and, again, found major variation in participants’ metaethical judgments. Some issue-items, such as cheating on an exam, were interpreted in an objectivist way by most participants (75%), whereas other items – particularly controversial ones – like abortion, were interpreted in highly non-objectivist ways, with only 14% interpreting the issue in an objectivist way (Wright et al., 2014, p. 40).

Wright (2018) explores new ways of demonstrating *Folk Metaethical Pluralism*, with experiments aimed at eliciting many types of antirealist attitudes, including subjectivism, relativism, and non-cognitivism. Again, participants were asked to identify the issues that they considered to be moral. Next, participants were asked to imagine someone who sincerely made a statement reflecting the opposing position. They were then asked to report what they believed about this person’s actions, and given three response options (in relation to a given action x):

- (1) It would be morally (un)acceptable for that person to x. The person would be correct because the rightness/wrongness of x is determined individually, by each person’s beliefs, attitudes, or feelings about the act of x-ing or type of action that x-ing is.
- (2) It may or may not be morally (un)acceptable for that person to x. Whether the person was correct would depend on the community in which that person lives. The rightness/wrongness of x is determined by a community’s collective beliefs, attitudes, or feelings about the act of x-ing or type of action that x-ing is.
- (3) It would not be morally (un)acceptable for that person to x. The person would be mistaken (as would anyone else who made this claim). The rightness/wrongness of x is determined by the action or type of action it is, independently of the person’s or his/her community’s beliefs, attitudes, or feelings about it. That is, there is something about x-ing or the type of action x-ing is that makes it right/wrong regardless of what that person or other people think or feel about it (Wright, 2018, pp. 125–126).

The first of these two responses were coded as subjective, and the third response objective (Wright, 2018, pp. 125–126). Coded this way, 78% of the participants gave objective responses to some of the moral statements and subjective responses to some of the others, with only 19% being consistent objectivists, and 3% being consistent subjectivists (p. 126).

In a follow-up experiment, Wright (2018) gave participants a short description of the difference between relative terms (that either apply or do not apply to their object depending on the context in which the term was uttered – like the term “tall”) and non-relative terms (that either apply or do not apply to their object independent of context – like the term “triangle”)

(p. 127). The study then asked participants to imagine a disagreement between two people over whether an issue-relevant action committed by a third party was correct. In assessing this hypothetical disagreement, participants were asked to agree with one of the following options:

- (1) ONLY ONE of these statements would be correct – either it is true that it was wrong for that person to x or it is true that it was not wrong for that person to x, regardless of who is making the statement or the contexts in which it is being made. Both statements cannot be correct.
- (2) BOTH of these statements could be correct – whether it is true that it was wrong for that person to x or not wrong for that person to x depends on who is making the statement and/or the context in which it is made (Wright, 2018, p. 128).

Wright coded response (1) as non-relative and (2) as relative. Here, again, participants gave a mix of realist and antirealist answers. Only 8% were consistent in interpreting the issues they identified as moral in a non-relative way, 7% consistently interpreted these issues as relative, and 85% were inconsistent in how they interpreted them, with some judged to be relative and others non-relative.

In a significant advance over previous research, Wright (2018) reports two more studies that address the question of folk non-cognitivism. Wright distinguishes between semantic non-factualism (the denial that moral statements express propositions or have truth conditions) and psychological non-cognitivism (the denial that the mental states that moral statements are conventionally intended to convey are beliefs) (p. 131; see also Joyce, 2016). Before each study, participants were presented with an explanation of the difference between factual and non-factual statements (and non-cognitive mental states and cognitive mental states), and then tested on their understanding of the distinctions (pp. 131–2 and pp. 135–7). After eliminating those participants that did not understand the distinction, participants were again asked to classify the same issue-items as a moral issue, a social convention/norm, or a personal preference, and asked whether they would classify each of these items as truth-apt or non-truth-apt in the first experiment, and whether the items expressed beliefs or positive/negative feelings, or both. Wright found, again, that the majority of respondents were pluralists: 78% expressed some mix of non-factualist and factualist judgments, and 75% expressed a mix of non-cognitivist and cognitivist judgments (p. 135 and p. 140).

§2.6 Sarkissian et al. (2011)

One final study that supports *Folk Metaethical Pluralism* is Sarkissian et al. (2011). This study showed that when people evaluate situations in their own

culture, they tend to express more realist attitudes, and when they evaluate situations outside their own culture, they tend to express more antirealist attitudes. Participants were provided with two vignettes that included an agent performing an action and two evaluators that disagreed about whether the action is permissible. Participants in the same-culture condition were told that both the evaluators were Americans like themselves. Those in the other-culture condition were told that one of the evaluators was from a very different culture and a different way of life. Finally, those in the extraterrestrial condition were told that one of the disagreeing evaluators was from an extraterrestrial culture (p. 486–7). In the first moral vignette, the behavior was the following: “Horace finds his youngest child extremely unattractive and therefore kills him.” The second was similar: “Dylan buys an expensive new knife and tests its sharpness by randomly stabbing a passerby on the street” (Sarkissian et al., 2011, p. 487). Participants were then told that one of the evaluators thought the action was permissible, and the other thought it was impermissible. Finally, they were told:

Given that these individuals [i.e. the evaluators] have different judgments about this case, we would like to know whether you think at least one of them must be wrong, or whether you think both of them could actually be correct. In other words, to what extent would you agree or disagree with the following statement concerning such a case: *Since your classmate and Sam have different judgments about this case, at least one of them must be wrong* (2011, p. 487, emphasis original).

Participants then indicated their level of agreement to the statement on a seven-point scale.¹⁵ The instructions were the same for participants in the other-culture and extraterrestrial condition, except for the description of one of the evaluators. In the other-culture condition, rather than describing the both the evaluators as Americans much like themselves, one of the evaluators was described as a member of an isolated Amazonian warrior tribe called the Mamilons. In the extraterrestrial condition, one of the disagreeing evaluators was an extraterrestrial from a race of beings called the Pentars, who are not interested in friendship or love, and only care about increasing the total number of pentagons in the universe (Sarkissian et al., 2011, pp. 489–490).

The results of this study provide yet more evidence for *Folk Metaethical Pluralism*. Participants in the same-culture condition largely expressed objectivist judgments ($M = 5.4$, $SD = 2.15$). Those in the other-culture condition centered on the mid-point of the scale ($M = 4.4$, $SD = 2.05$). And participants in the extraterrestrial condition were the least objectivist ($M = 3.2$, $SD = 2.28$). The difference in responses between the conditions were significant ($P < .05$ between the same-culture and other-culture conditions and $P < .01$ between other-culture and extraterrestrial conditions). Later experiments reported in the same study demonstrated that these

results could be replicated cross-culturally, when all three conditions were presented to participants at once, and when the actor in the vignette (e.g. Horace) was presented as a member of the participants' own culture (pp. 490–4). Like Nichols (2004), a study was performed that demonstrated that participants were not simply expressing global subjectivism (pp. 496–7). Sarkissian et al. conclude that “people can have different reactions to judgments about the very same act, performed by the very same agent, so long as we vary the identity of the judge” (Sarkissian et al., 2011, p. 494). This would be hard to explain unless participants were expressing subjectivist attitudes (the judge's attitudes constitute the rightness or wrongness of the act) or relativist attitudes (the judge may be adverting to rules established in his culture).

§2.7 Folk metaethical pluralism is the most plausible

As is becoming clear, studies on metaethical attitudes have produced many reasons to doubt *Folk Metaethical Realism*, and many reasons to accept *Folk Metaethical Pluralism*. In short, these studies demonstrate that ordinary people express various antirealist attitudes across many different conditions. This finding is quite robust, so researchers that wish to defend *Folk Metaethical Realism* must show how realist attitudes can be the default folk metaethical attitudes, even though people very often deviate from this default. This would require explaining why these deviations – produced by various extraneous factors – occur. But as I will argue in the next section, there is good reason to think that these factors are not extraneous at all. Instead, I will propose that they constitute the source of most metaethical attitudes, and this accounts for the fact that folk metaethical attitudes are so variable.

§3 Folk metaethical irrationalism

If people have a mix of metaethical attitudes rather than realist attitudes, it is worth asking why this is. In this section I will argue for *Folk Metaethical Irrationalism*, a thesis that constitutes a broad explanation of, and supports, *Folk Metaethical Pluralism*. Restated, the thesis is the following:

Folk Metaethical Irrationalism: metaethical attitudes (both realist and antirealist) are typically caused by processes that are insensitive to the truth of these attitudes.

If this thesis is true, it ought to be unsurprising that people have a mix of metaethical attitudes, rather than realist attitudes by default; they are not produced by a rational inferential process that we can uncover and explain. Before going into the factors that seem produce metaethical attitudes, though, it is worth considering what rational standard they fail to live up

to. Consider the following two first-order judgments, one moral and another non-moral, both drawn from Beebe and Sackris (2016):

Global Warming: global warming is due primarily to human activity (for example, the burning of fossil fuels).

Anti-Racism: treating someone poorly on the basis of their race is morally wrong.

Setting aside the question of whether we believe these judgments, should we be realists or antirealists about them? Presumably the correct metacognitive attitude about *Global Warming* is realism. Even people who disagree with the clear majority of scientists on the question of anthropogenic climate change still ought to agree that the facts that would make *Global Warming* true or false are objective, universal facts about the temperature of the Earth over time.¹⁶ Furthermore, like any attitude, the justification for realism about *Global Warming* depends on the process that produced it. Realist attitudes about global warming are rational insofar as they are caused by reliable evidence of, and/or valid inferences about, the mind-independent and universal facts that constitute global warming as a phenomenon. Conversely, antirealist attitudes about other phenomena are rational insofar as they are caused by reliable evidence of, and/or valid inferences about the mind-dependent or relative facts that constitute those phenomena, or (in the case of rational non-factualist attitudes) the perceptions of or valid inferences about the lack of such facts. Various theories in epistemology may differ as to what more is required for these metacognitive attitudes to be rational, but few would dispute these basic requirements.

In their study of folk metaethical attitudes, Beebe and Sackris (2016) found that fewer than half of their participants expressed objectivist (and thus realist) attitudes about *Global Warming* (p. 924). This is a surprising result, given how obvious it is that realist attitudes are true of the judgment. Clearly, the process that produced these antirealist attitudes is insensitive to their truth, and thus, these attitudes are (absent other justification) irrational. If we want to explain why people have metacognitive attitudes, and whether those attitudes are rational, we should look for their causal or inferential sources, and investigate whether those sources are sensitive to the truth of those attitudes. This is the case whether we're talking about metacognitive attitudes about factual judgments like *Global Warming* or metaethical attitudes about moral judgments like *Anti-Racism*.

One might think that *Anti-Racism* is different from *Global Warming* in a way that prevents us from claiming that metaethical attitudes about it are irrational. Namely, whereas most philosophers and psychologists agree that realism is the correct and justified metacognitive attitude to have about *Global Warming*, and this might allow us to call antirealist attitudes about it irrational, there is no such agreement about the metaethical status of *Anti-Racism*. This is true: we cannot assume that antirealist attitudes about moral

issues are irrational, since this is a matter of substantive philosophical dispute. But philosophers and psychologists need not agree about which metacognitive attitude is correct and justified in order to determine which *sources* of metacognitive attitudes are *not* justified. It may be the case that metacognitive attitudes in general are caused by factors that are insensitive to their truth. In what follows, though, I will focus only on the factors that have been demonstrated to correlate with, and thus might be causing, metaethical attitudes.

53.1 *The challenges of incoherentism and pragmatic function*

Before I go on to explore these factors, though, it is worth situating *Folk Metaethical Irrationalism* in the context of other attempts at explaining *Folk Metaethical Pluralism* and other cognate claims. For instance, Loeb (2008a) argues that the meaning of central moral terms (such as ‘good’ and ‘bad’) is incoherent, and that these terms cannot be made coherent without changing the subject. Though this view appears similar to the one I am advancing in this section, it is important to note that *Folk Metaethical Irrationalism* is a thesis about folk *metaethical* attitudes (metacognitive states that have first-order moral judgments as their semantic targets), whereas Loeb’s thesis is about the semantics of first-order moral judgments.¹⁷

One potential explanation of *Folk Metaethical Pluralism* comes from Wright (2018), who argues that such pluralism serves “a pragmatic function – namely, that it promotes civility and aids in the individual and collective navigation of normative space within a morally imperfect world” (p. 119). Wright’s explanation is not just that antirealist attitudes help people tolerate each other’s differences, but also that realist attitudes allow people to remove an issue “from the realm of legitimate personal/social negotiation” (Wright et al., 2014, p. 31). This thesis might imply that it is practically rational for people to have variable and even contradictory metaethical attitudes. This may be the case even if these attitudes are not theoretically rational, since social utility of such attitudes can be obtained irrespective their truth.

This explanation is surely on the right track, but we have reason to suspect that it is too narrow. This is because even if these attitudes usually serve this pragmatic function, the factors that have been shown to correlate with metaethical attitudes are not always social. These attitudes may allow people to conceptualize their disagreements with others in socially useful ways, but they are not always deployed this way. Furthermore, as we will explore below, many other factors may affect metaethical attitudes, including personality, developmental factors, and features of the first-order judgments they target. These factors plausibly cause people to have realist

attitudes even though it would be more socially useful to have antirealist attitudes, and vice versa.

As we've seen, then, researchers have already begun to explore some of the explanations for *Folk Metaethical Pluralism*. The goal in what follows will be to provide an overarching explanation of why people have the metaethical attitudes they do, and to explore some of the ways that this explanation can open more avenues of research in folk metaethics. Though more work will be needed to isolate the extent to which each of these factors causally influence metaethical attitudes, there are several studies that demonstrate interesting correlations between folk metaethical attitudes and various factors. These correlations are detailed in Table 3. These factors can be grouped into at least four categories: judgment factors, developmental factors, personality factors, and social factors. In each of these cases, we have reason to doubt that the factor that may well be driving metaethical attitudes is rationally connected with the truth of those attitudes.

§3.2 Judgment factors

I will start with judgment factors, which are features of first-order moral judgments that correlate well with either realist or antirealist attitudes. The most widely documented type of judgment factor is the perceived consensus about a first-order judgment. For example, Goodwin and Darley (2008) showed that only 24% of participants were willing to say there was a correct

Table 3. Factors that correlate with metaethical judgments and the direction of correlation.

	Factor	Direction of Correlation with Metaethical Attitudes
Judgment Factors	Perceived consensus about a judgment.	More realist (Beebe, 2014; Beebe et al., 2015; Beebe & Sackris, 2016; Goodwin & Darley, 2008, Goodwin & Darley, 2010, Goodwin & Darley, 2012; Heiphetz & Young, 2017; Wright et al., 2013, Wright et al., 2014)
	judgment strength.	More realist (Beebe, 2014; Beebe et al., 2015; Goodwin & Darley, 2010, Goodwin & Darley, 2012; Heiphetz & Young, 2017); though see Wright et al., 2013.
	Non-harmful first-order judgments.	More antirealist (Feltz & Cokely, 2008, p. 1773; Goodwin & Darley, 2010, p. 172).
	Negative judgment valence.	More realist (Beebe, 2014, pp. 180–182; Goodwin & Darley, 2010, p. 172, Goodwin & Darley, 2012, p. 254).
	Judgments of imperfect duties.	More antirealist (Beebe, 2014)
Development Factors	Desire to punish.	More realist (Rose & Nichols, 2019)
	More years spent in college.	More antirealist (Nichols, 2004)
	Age.	Varies (Beebe et al., 2015; Beebe & Sackris, 2016; Heiphetz & Young, 2017; Nichols & Folds-Bennett, 2003; Schmidt et al., 2017; Wainryb et al., 2004).
Personality Factors	Tendency toward disjunctive thinking.	More antirealist (Goodwin & Darley, 2010).
	Openness to experience.	More antirealist (Feltz & Cokely, 2008).
Social Factors	Deriving morality from religion.	More realist (Goodwin & Darley, 2008, Goodwin & Darley, 2010, Goodwin & Darley, 2012; Sarkissian & Phelan, 2019)

answer about whether “scientific research on embryonic human stem cells that are the product of in vitro fertilization is morally permissible” (p. 1362). This was the case despite the statement enjoying relatively high levels of agreement ($M = 4.84$ on a scale of 1–6). Other, less controversial moral statements enjoyed both high mean agreement and high percentages of participants willing to say there was a correct answer as to whether the statements were true (p. 1362). More tellingly, Goodwin and Darley (2012) explicitly tested the hypothesis that differences in perceived consensus predict objectivism (p. 252). To measure perceived consensus, participants were asked to estimate the percentage of US citizens that agreed with any given statement. This perceived consensus was highly correlated with participants’ strength of agreement. Even combining strength of agreement and perceived consensus, though, perceived consensus uniquely (if marginally) predicted participants’ level of objectivity (p. 252). These findings are by no means the only ones supporting the claim that perceived consensus affects expressions of folk metaethical attitudes (Beebe et al., 2015; Beebe & Sackris, 2016; Goodwin & Darley, 2010; Heiphetz & Young, 2017; Wright et al., 2014).

Though perceived consensus and metaethical attitudes are related, researchers have been cagey as to the causal relationship between the two phenomena. One possibility, proposed by Goodwin and Darley (2012), is that perceptions of high consensus cause realist attitudes, and perceptions low consensus cause antirealist attitudes (p. 251–2). An alternative explanation, proposed by Wright et al. (2014), is that realist attitudes “[involve] – and possibly even [generate] – the expectation of stronger and more widespread social consensus” (p. 55), and that conversely, antirealist attitudes could generate expectation of weaker or less widespread social consensus.

In assessing whether perceptions of consensus might affect the rationality of metaethical attitudes, the putative direction of causation matters. If Wright et al.’s interpretation is correct, this does not contradict *Folk Metaethical Irrationalism*, since the thesis concerns the rationality of the causes of folk metaethical attitudes, rather than the rationality of their effects. But if the direction of causation goes the other way, as Goodwin & Darley suggest, we ought to examine whether perceptions of consensus are good evidence for realist attitudes.

Prima facie, it is hard to see how realist attitudes could rationally follow from perceptions of high consensus about a putative fact, or how antirealist attitudes could rationally follow from perceptions of low consensus about it. Setting aside the accuracy of these perceptions, inferences in either direction are certainly not valid. This is especially clear in the case of metacognitive attitudes about scientific facts. To return to *Global Warming* from above, even if everyone agreed that global warming is not a real phenomenon, this would not mean that it is objectively or universally true that global warming

is not occurring. Moreover, if half the population believed that global warming is not a real phenomenon and the other half disagreed, this would not mean that the phenomenon of global warming is not a fact, nor that it is only subjectively or relatively true, nor that people fail to express truth-evaluable claims when they make statements about it. Arguably the same is true for moral facts, whatever they may be. For reasons like this, many philosophers are skeptical that disagreement (and conversely, consensus) can be taken to have metaethical implications (see e.g. Brink, 1984; Doris & Plakias, 2008; Dreier, 2009; Enoch, 2009; Fitzpatrick, 2014). Some philosophers argue that even disagreement in ideal conditions does not have metaethical implications (Bloomfield, 2001).

Some philosophers are more optimistic, however, and argue that it is sometimes rational to use consensus as evidence for metaethical claims. In a recent paper, Ayars and Nichols (2020) argue that perceived consensus causes people to adopt universalist attitudes, and that consensus could provide evidence that a claim is universally true. Ayars & Nichols argue that when comparing theories, fitting the data is not the only important consideration; very flexible theories are sometimes worse than inflexible theories, depending on the type of data they need to fit. Relativist theories about a given set of facts tend to be more flexible than universalist theories about those same facts, because, as compared with universalist theories that posit a single fact to explain many phenomena, relativist theories posit multiple facts to explain those same phenomena, depending on context or culture. In contexts of low consensus, this means there are more varied phenomena to explain (i.e. things appear different ways to different people), and relativist theories fare better. This is because they can accommodate all the data (i.e. everyone is correct), and universalists can only accommodate some of the data (i.e. some people are right, and others wrong). But in situations of high consensus, universalist theories fare better. In these contexts, relativist theories seem to fit the data *too well*, while universalist theories seem to capture something important that relativists theories miss. As Ayars and Nichols argue, “It will often be more plausible to count a small minority as mistaken about a universal truth rather than correct about a relative one” (Ayars & Nichols, 2020, p. 73). Applying this idea in the metaethical context, perceiving high consensus about a judgment may be reliable evidence that the judgment is universally true. If this is correct, perceptions of consensus are sensitive to the truth of some metaethical attitudes, and this would constitute an exception to *Folk Metaethical Irrationalism*.

However, we should be wary of conflating consensus, which may have an evidentiary role, with *perceived* consensus. As Ayars & Nichols note, they are not arguing that people “make these inferences in a way that precisely tracks the probabilities” (p. 14). And the evidentiary value of perceptions of

consensus can vary dramatically depending on how accurate those perceptions are. Perceptions of consensus can be well-grounded (when they are based on thorough, scientific surveys of representative samples), or they can be very poorly grounded (when they are based on a person's one-off, subjective estimation of how many people (dis)agree with her). And though no study has yet examined the question, it seems reasonable to conjecture that folk perceptions of consensus are in this latter category, rather than the former. To the extent that perceptions of consensus are causing metaethical attitudes, we should expect that this factor is typically distortive for folk attitudes. Whether this is the case for expert attitudes is another question for future research.

Other judgment factors that correlate with variance in metaethical attitudes also appear quite rationally problematic. Strength of agreement with a judgment is highly correlated with objectivism (J. Beebe et al., 2015; J. R. Beebe, 2014; Heiphetz & Young, 2017, cf. Wright et al., 2013). This pattern repeats for other judgments as well (factual, conventional, and preference judgments). When a participant feels strongly about a judgment of nearly any kind, she is more likely to have a realist attitude about it, and the same is true, *mutatis mutandis*, for antirealist attitudes. This is rationally problematic because, intuitively, the truth of a metacognitive attitude about a first-order judgment has nothing to do with one's feelings about that judgment. If someone thinks that Jimi Hendrix's *Purple Haze* is an excellent song, and feels very strongly about that judgment, her feeling this way certainly does not imply that it is an objective, universal fact that *Purple Haze* is an excellent song. Conversely, if she has a passing thought that a tree is an elm, and has no strong feelings about that judgment at all, this would not imply that the fact that the tree is an elm is only subjective or relative. The strength of a judgment is simply rationally disconnected from the truth of metaethical attitudes about that judgment. If judgment strength is causing metaethical attitudes, this factor is also distortive.

Yet another judgment factor that influences metaethical judgments is the harmfulness of the action being judged. Though I will consider personality factors below, Feltz and Cokely (2008) found an interaction between extroversion, harm judgments, and levels of objectivism: extroversion predicted the extent to which participants would be objectivists about non-harmful moral scenarios (Feltz & Cokely, 2008, p. 1773). This effect is less than certain, however, because Goodwin and Darley (2010) found that when controlling for strength of agreement, participants did not treat harmful actions as more objective than non-harmful (symbolic) actions (p. 172).

More broadly, there is good evidence that negatively valenced judgments are associated with greater objectivism, and positively valenced judgments

are associated with greater subjectivism. Goodwin & Darley found that judgments that involve wrongness are associated with greater levels of objectivity than judgments of goodness (Goodwin & Darley, 2010, p. 172). This was confirmed by later experiments as well (Beebe, 2014, pp. 180–182; Goodwin & Darley, 2012, p. 254). Like other judgment factors, the valence of a judgment has nothing to do with the truth of a metaethical attitude about that judgment. If someone judges that rocky road ice cream tastes great, this does not mean she should judge that it tastes great for everyone and would taste great regardless of anyone's opinion. And if someone cringes when his landlord slips the rent bill under his door, this feeling does not imply that the rent coming due is purely subjective and relative.

Another judgment factor that seems rationally disconnected from the truth of metaethical attitudes is the desire to punish a wrongdoer. Rose and Nichols (2019) report multiple studies in which participants' reported desire to punish correlates with universalist attitudes. For example, in one study, they chose a moral issue-item from Goodwin and Darley (2008) for which participants tended to express more universalist attitudes (conscious discrimination on the basis of race), and compared it with a moral issue-item from the same study, for which participants tended to express more relativist attitudes (assisting in the death of a terminally ill friend who wanted to die). Participants were given either of these issue items and asked how much a person committing these wrongs should be punished, using a scale of 1 (not at all) to 7 (very much) (Rose & Nichols, 2019, pp. 62–63). As in previous experiments, participants were asked whether people who disagree about each of these issue-items could both be correct, or if only one had to be correct (p. 63). As in Goodwin and Darley (2008), the discrimination issue-item received higher universalist ratings ($M = 3.79$, $SD = 1.52$) as compared with the euthanasia issue-item ($M = 3.02$, $SD = 1.49$) (Rose & Nichols, 2019, p. 63). But more importantly, participants reported a much higher desire to punish for the discrimination issue-item ($M = 5.12$, $SD = 1.53$) as compared with the euthanasia issue-item ($M = 1.94$, $SD = 1.41$).

It is hard to read this difference as anything other than a distortive influence on metaethical attitudes: there is simply no reason why desiring to punish someone for an action would affect whether that action is universally right or wrong. Indeed, Rose & Nichols draw much the same lesson in this case as I have drawn for the other judgment factors discussed above: “the motivation to punish is irrelevant to the truth of moral universalism” (p. 70).

One final judgment factor that appears to influence metaethical attitudes is the type of duty that the judgment involves. Beebe (2015) showed that participants view judgments about imperfect duties as less objective than judgments about perfect duties (p. 23). Again, given that the metaphysical

status of a duty is completely separable from the type of duty it is (perfect or imperfect), this factor also seems distortive.

§3.3 Developmental factors

In addition to judgment factors, there are several developmental factors that correlate with changes in metaethical attitudes. Most work to date has focused on age. Broadly speaking, younger children are more likely to express realist attitudes, while adolescents, college students, and young adults tend to express more antirealist attitudes, and older people tend to revert to realist attitudes (Beebe et al., 2015; Beebe & Sackris, 2016; Heiphetz & Young, 2017; Nichols & Folds-Bennett, 2003; Schmidt et al., 2017; Wainryb et al., 2004). The relationship between age and metaethical attitudes promises to be quite complex and is likely a combination of social and developmental factors. One likely explanation of this effect is that college students and young adults often find themselves in a new environment, having new experiences, meeting new types of people, and are exposed to ways of thinking that compete with the value systems they were raised with (Beebe & Sackris, 2016).¹⁸ This would mirror similar findings for first-order moral judgments (L. Kohlberg & Kramer, 1969). Another possibility that is that college-aged people are developmentally more open to new experiences, which is a personality factor that also correlates with antirealist attitudes (Nettle, 2009).

What matters for our purposes is whether these development factors are part of a rational, truth-tracking process. If it is true that changes in metaethical attitudes that correlate with age are driven by exposure to more (or less) moral disagreement at various stages of development, this means that changes in perceptions of consensus are likely to be driving these changes in metaethical attitude. And as we considered above, there are good reasons to think that folk perceptions of consensus are not usually a rational basis on which to form metaethical attitudes.

§3.4 Personality factors

Metaethical attitudes are also influenced by personality factors. Feltz and Cokely (2008) conducted a series of experiments which showed that these attitudes are “associated with stable individual differences such as personality traits and reflective cognitive styles” (pp. 1771). They conducted experiments similar to those found in Nichols (2004), but added a short version of the Big Five personality test (Gosling et al., 2003). Participants were then asked to adjudicate a disagreement between two people over moral and non-moral issues. The study showed that 69% of participants gave the subjectivist answer, and only 31% gave the objectivist answer (Feltz

& Cokely, 2008, p. 1772). In addition, participants who were high in openness to experience were significantly more likely to give the subjectivist answer (Feltz & Cokely, 2008, p. 1772). Moreover, there is some evidence that one's thinking style has an influence on objectivism. Goodwin and Darley (2010) demonstrated the effects of one thinking style in particular, disjunctive thinking, which is "the tendency to actively unpack alternative possibilities when reasoning" (p. 176). Participants who engaged in a more disjunctive thinking style were more likely to express subjectivist attitudes in response to ethical statements.

If any of these personality factors are causing changes in metaethical attitudes, it is hard to see how these could improve the epistemic warrant of these attitudes. As in the case of age, this is because the most plausible means by which these personality measures increase antirealist attitudes is by increasing exposure to disagreement, which thereby decreases perceptions of consensus. Goodwin & Darley argue that disjunctive thinking ought to increase subjectivist responses because "[s]uch individuals should possess a dispositional tendency to think disjunctively, and thus, in a moral context, they should be more inclined to actively consider the reasons why another person might disagree with their ethical beliefs" (Goodwin & Darley, 2010, p. 175). Similarly, Feltz & Cokely say that "[t]hose who are highly open to experience might be more likely to think that morals predominate in one's society are mistaken or otherwise flexible, and hence would be more open to the possibility that there is no single, correct ethical answer" (p. 1772). As I covered above, we should not take folk perceptions of high consensus about a judgment as evidence that a judgment ought to be treated in a realist way, and the same, *mutatis mutandis*, for low perceptions of consensus and antirealist attitudes. Unless we can find evidence that, independent of perceptions of consensus, these personality factors make one more sensitive to the potential truth of metacognitive attitudes, we ought to see these factors as distortive as well.

§3.5 Social factors

Social factors that affect metaethical attitudes are currently one of the least developed in the literature, but there is some evidence that religious grounding increases realist attitudes. Goodwin & Darley gave participants a chance to provide reasons that support their moral beliefs, including a religious grounding (one's moral beliefs are "ordained by a supreme being"), a consensus grounding ("every good person on earth, regardless of culture, holds these beliefs"), an instrumental grounding ("a society could not survive without its citizens holding these beliefs) and an intrinsic rightness grounding (the truth of one's moral beliefs "is self-evident") (Goodwin & Darley, 2008, p. 1363). Goodwin & Darley reason that "[g]rounding one's

ethics in religious belief is the most obvious way that one could be an objectivist about ethics” (Goodwin & Darley, 2008, p. 1354). When comparing participants who gave unreligious grounding with those who gave a religious grounding, the latter expressed significantly more objectivist attitudes. They also asked participants the question, “According to you, is it possible for there to be right and wrong acts, without the existence of God?” (p. 1354). Those that gave a “no” answer to this question expressed remarkably objectivist attitudes ($M = 2.95$).¹⁹

Other studies have also investigated the relationship between metaethical attitudes and religious belief. Yilmaz and Bahçekapili (2015) show that subjectivist attitudes are negatively correlated with multiple measures of religious belief, and that it is possible to increase expression of objectivist attitudes by priming participants with religious words in a word-scramble (pp. 96–7). The direction of effect may go the other way, as well, since Yilmaz & Bahçekapili show that participants express less confidence in the existence of God when they have read a persuasive essay about the subjectivity of morality (pp. 97–99). Even more recently, Sarkissian and Phelan (2019) present evidence that belief in God’s punishing characteristics is associated with expression of objectivist attitudes (pp. 2–3). In another experiment reported in the same study, believers in Abrahamic religions were more likely to express objectivist attitudes when thinking of God (pp. 3–5).

It may seem intuitive that highly religious people have realist metaethical attitudes. Goodwin & Darley say that “one way that the question of how there could be objective moral facts might be answered is to view ethical statements as having a religious foundation – for instance, they are the word of God” (p. 1341). Similarly, Sarkissian and Phelan (2019) consider it plausible “for someone who believed in God to believe that some moral dictates stem from Divine Command, are objectively true, and apply to all people” (p. 2). When examined from an epistemic standpoint, however, this is quite a puzzling idea. Despite first appearances, divine command theory (a position that makes right and wrong metaphysically dependent on the dictates of God) is a subjectivist theory, not an objectivist theory (see e.g. Brink, 2007; Huemer, 2005, pp. 54–55). It is thus strange that highly religious people who subscribe to divine command theory ought to have objectivist metaethical attitudes.²⁰ Of course, this is only a bizarre result if it we think that highly objectivist divine command theorists are supporting their objectivist attitudes through a rational process, whereby these metaethical attitudes can be inferred from their belief in divine command theory. The alternative, I submit, is more likely: people’s religious views may not rationally imply objectivist metaethical attitudes, but religious instruction often includes the claim that the possibility of objective morality depends on God’s dictates, and this causes believers to adopt these attitudes.

Seen from this perspective, we might think that grounding one's morality in religion, like the other influences we have covered, is distortive – it is a factor that causes metaethical attitudes without sensitivity to their truth.

§3.6 *These factors render folk metaethics irrational*

Having surveyed some of the factors that seem to have an influence on metaethical judgments, it will be useful to take stock. As I have argued, many of the factors discussed above plausibly cause people to adopt more realist or antirealist attitudes. What's more, these influences are distortive. I contend that this fact gives us some explanatory purchase on *Folk Metaethical Pluralism*. People have a mix of metaethical attitudes because various features of their first-order judgments, personality, development, and social situation cause them to adopt these attitudes without sensitivity to their truth. In fact, this explanation is not limited to cases where people have a mix of metaethical attitudes: in the few cases where folk metaethical attitudes are more uniformly realist (e.g. in the case of religious grounding) these attitudes are also plausibly caused by factors that are not sensitive to their truth. What this means is that, unless some alternative justification is discovered, folk metaethical attitudes are irrational.

§4 *Implications for metaethics*

Perhaps philosophers should be unsurprised by the disordered state of folk metaethics. One of the major purposes of philosophy is to learn how to reason better, and to use that reasoning to replace unjustified thinking with careful, rational thought. If there were ever a subject about which we should expect plenty of unjustified, uncaredful thinking, it is metaethics. After all, metaethics concerns abstract philosophical questions about what we are doing when we make moral judgments, what facts our judgments might be tracking, and how we are justified in making those judgments – these are esoteric subjects! Without philosophical training, people have no principled reasons for being realists or antirealists, and so the process they use to derive metaethical attitudes is distorted by the factors discussed above. If this is correct, we might hope that when people get philosophical training that includes metaethics, the process they use to derive their metaethical attitudes will be more rational, and they will be able to see why these factors are irrelevant to the truth of realist or antirealist attitudes (in all their forms, as captured by (A) through (H) in §1).

Given that this is the state of folk metaethics, one may reasonably wonder what it implies for metaethics. Philosophers sometimes argue that moral realism best explains or vindicates the intuitive metaethical commitments of ordinary moral reasoners. As Brink (1989) argues, for instance, “We begin

as (tacit) cognitivists and realists about ethics . . . [w]e are led to some form of antirealism (if we are) only because we come to regard the moral realist's commitments as untenable" (p. 23).^{21,22} Suppose it is true that (1) folk metaethical attitudes are typically pluralistic, rather than realist, and (2) these attitudes are caused by factors that are insensitive to their truth. By virtue of (1) alone, moral realism no longer looks like the best explanation or vindication of folk attitudes (see also Hopster, 2019; Pözlner, 2018). To put it simply, if moral realism is true, and people have a mix of metaethical attitudes, they are incorrect whenever they have antirealist attitudes.

A more fitting explanation might be metaethical pluralism: some moral judgments have objective facts as their intended targets, while others have subjective facts as their targets, or (in the case of non-cognitivist or non-factualist moral attitudes, no semantic target at all). Pözlner and Wright (2019) suggest as much, claiming that something like *Folk Metaethical Pluralism* might "ground nontraditional metaethical positions, according to which different metaethical views are true for different parts of moral language and thought" (p. 11). I advance an argument for such a nontraditional position in Colebrook (2018). More strongly, the truth of *Folk Metaethical Pluralism* suggests yet another intriguing possibility: that the claims to vindication of "common sense" currently employed by moral realists can better be claimed by moral pluralists.

If *Folk Metaethical Irrationalism* is true, though, we should be cautious about such an inference. It is only a mark in favor of a metaethical theory that it vindicates common sense if we have some reason to think common sense is epistemically valuable in the first place. As I have argued above, this seems unlikely to be the case for most (if not all) metaethical attitudes. They are simply influenced by too many extraneous factors to enjoy any special epistemic status. If, however, further research reveals a subset of metaethical attitudes that are caused by factors that we have good reason to believe *are* sensitive to the truth of metaethical realism, pluralism, or antirealism, such "vindication" or "explanation" arguments for any of these metaethical positions would become viable again.

Some philosophers suggest that the irrational source of folk metaethical attitudes suggests a debunking argument against moral realism itself (e.g. Ruse, 1998). In my view, this inference would not be warranted. In principle, any of these metaethical views could be true, even if its being true does not vindicate or explain anyone's attitudes. Even if it is true that realist attitudes are typically caused by factors that are not sensitive to their truth, and as a result, these attitudes are irrational, this does not imply they always are irrational, or that they will always be irrational. Common sense is not the only source of evidence that philosophers appeal to in favor of moral realism. To draw an analogy, if psychological research on folk intuitions about physics revealed that these intuitions were typically irrational, this

would not even be presumptive evidence in favor of antirealism about physics, though it would undermine arguments for realism about physics that rest on the epistemic warrant of those intuitions. (Luckily, physicists have more to rely on than folk physical intuitions). The same ought to apply to moral realism in its various dimensions.

Conclusion

Though philosophers have often held that ordinary people are moral realists, we now have good reason to doubt that this is the case. At the very least, a substantial minority of people have antirealist attitudes about many of their moral judgments. In the absence of defusing explanations that show why these people abnormally express antirealist judgments, the fact of robust and persistent metaethical diversity ought to be enough to undermine *Folk Metaethical Realism*. If *Folk Metaethical Pluralism* is correct, we are left with a question as to why people typically have this mix of attitudes. *Folk Metaethical Irrationalism* provides the answer. Metaethical attitudes are typically caused by factors that are insensitive to the potential truth of realism or antirealism, and as a result, they are irrational.

Notes

1. Some work that argues in favor of *Folk Metaethical Pluralism* defines moral objectivism, and thus, moral realism, as a claim about the truth of moral sentences, rather than about the existence of moral facts (e.g. Pölzler, 2017). There are reasons to prefer the latter, as construing the issue as one about truth, rather than facts, makes realism about a given subject matter dependent on the proper theory of truth (Devitt, 2002, 2010).
2. Whether Wong counts as an objectivist-relativist is somewhat ambiguous. In *Natural Moralities*, Wong advances naturalistic conception of morality that identifies various constraints on possibly adequate moralities but does not specify that a single moral code must be correct. Whether a given moral code is adequate or not depends on many features that are independent of the attitudes present in a given culture, though these constraints do depend on psychological and social facts about human beings. As a reviewer points out, Wong might plausibly be construed as an objectivist about some issues and a subjectivist about others.
3. See Quintelier and Fessler (2012) for a sustained argument about why researchers should not assume that realism is a conceptually necessary feature of first-order moral judgments.
4. As a reviewer points out, a third option may also be available. In a new paper, Pölzler and Wright (2019) argue for “Anti-Realist Pluralism.” According to this view, ordinary people have antirealist attitudes about the majority of their moral beliefs.
5. I thank an anonymous reviewer for this point.
6. An anonymous reviewer suggests that the fact that these attitudes are often unconscious and implicit might explain why these attitudes are often caused by factors that are insensitive to their truth. I think this is on the right track, and we may be able to

say a little more. In general, it is true that unconscious thoughts are seldom critiqued unless they are brought to consciousness. Unless one takes part in one of these experiments, or takes a course in philosophy, these attitudes will seldom have to endure criticism, either from oneself or from others. Moreover, when our attitudes are produced by an unconscious process, the rationality of those attitudes will depend, *inter alia*, on the reliability of the process that produced them. And such processes might be reliable in some contexts but not others.

7. In the moral scenario, the disagreement is about whether it is okay to hit other people just because you feel like it. In the factual scenario, the disagreement is about whether the Earth is flat. The disgust scenario involves a disagreement about whether it is okay to drink your own vomit if you have microwaved it first.
8. For example, their factual statements included “Boston (Massachusetts) is further north than Los Angeles (California).” Ethical statements included “Before the 3rd month of pregnancy, abortion for any reason (of the mother’s) is morally permissible. Their conventional statements included “Calling teachers by their first name, without being given permission to do so, in a school that calls them “Mr.” or “Mrs.” is wrong behavior.”
9. This excludes the very small number of participants who judged that a statement is true, and then judged that it is possible that they themselves were mistaken, rather than the disagreeing other.
10. Goodwin & Darley’s third reported experiment employs a similar methodology, although it changes the second phase of the experiment to employ a six-point scale that ranges from “neither of us need be mistaken” to “the other person is clearly mistaken” (Goodwin & Darley, 2008, p. 1356).
11. Goodwin and Darley (2008) does not include descriptive data about how many participants fell into each *intermediately objective* category for factual statements. Dr. Goodwin kindly shared this data with me in personal correspondence.
12. In the second phase, 20% of participants chose “it could be that you are mistaken, and the other person is correct.” Just one participant chose this answer in response to the two ethical statements, and Goodwin & Darley treated this response as fully objective (Goodwin & Darley, 2008, p. 1348). The percentage mentioned here for factual statements – 92% – assumes that we ought to treat this pattern of responses for factual statements the same way that we treat this pattern of responses for ethical statements.
13. Goodwin and Darley (2008) report that participants who received the more controversial statement about evolution were marginally less objectivist ($M = 2.64$ vs $M = 2.87$, $t(31.57, \text{unequal variances}) = 1.99$, $p < .06$.) (p. 1352). When we analyze the data with the new conceptualization proposed above (where 1 = *fully objective* and 0 = *least objective*), this pattern repeats, and is somewhat more robust ($M = .93$ vs $M = .72$, $t(43.01, \text{unequal variances}) = 2.28$, $p < .03$).
14. “Grounding” in these studies refers to how participants responded to the first and second phase of the experiment, with either objectivist or subjectivist answers on a scale of 0 to 2.
15. Instructions were the same for the other moral vignette.
16. As it happens, 66% of Americans believe global warming is being caused by human activity (Saad, 2019, p. 6).
17. For a discussion of the potential strengths and weakness of this incoherence argument, see Gill (2008, 2009), Sayre-McCord (2008), and Loeb (2008b).

18. Beebe & Sackris also investigate whether age effects on metaethical attitudes could be a cohort effect – that one generation may be more or less realist than another. However, they conclude that this is not the case (Beebe & Sackris, 2016, p. 921).
19. Bear in mind that, as mentioned previously, there are methodological problems with this measure of objectivity. Taking seriously the critiques found in Pözlzer (2018a), later research addressing the relationship between a religious grounding and metaethical attitudes ought to change the operationalization of objectivism.
20. It is not, one should note, very surprising that people who base their metaethical attitudes on divine command theory ought to be universalists, cognitivists, and factualists. After all, divine command theory is a subjectivist, universalist, cognitivist, factualist theory.
21. Shafer-Landau (2003) makes a similar argument for moral cognitivism.
22. See Rose and Nichols (2019) for a specific version of this objection and Sinclair (2012) for a more general critical overview of these arguments.

Acknowledgments

For their generous help in the improvement of this paper, I thank two anonymous reviewers, as well as Catherine Wilson, Geoff Goodwin, Hagop Sarkissian, Jesse Prinz, and my wife, Kathleen Rader.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Notes on contributor

Ross Colebrook is an Adjunct Assistant Professor at Baruch College, City University of New York, working on issues at the intersection of metaethics and moral psychology.

References

- Ayars, A., & Nichols, S. (2020). Rational learners and metaethics: Universalism, relativism, and evidence from consensus. *Mind & Language*, 35(1), 67–89. <https://doi.org/10.1111/mila.12232>
- Beebe, J Sarkissian H & Wright J. C. (Eds.). (2014). How different kinds of disagreement impact folk metaethical judgments. In, *Advances in experimental moral psychology* (pp. 167–187). Bloomsury Academic
- Beebe, J. (2015). The empirical study of folk metaethics. *Etyka*, 50, 11–28. <https://doi.org/10.14394/etyka.486>
- Beebe, J., Qiaoan, R., Wysocki, T., & Endara, M. A. (2015). Moral objectivism in cross-cultural perspective. *Journal of Cognition and Culture*, 15(3–4), 386–401. <https://doi.org/10.1163/15685373-12342157>
- Beebe, J., & Sackris, D. (2016). Moral objectivism across the lifespan. *Philosophical Psychology*, 29(6), 912–929. <https://doi.org/10.1080/09515089.2016.1174843>
- Bloomfield, P. (2001). *Moral reality*. Oxford University Press.

- Boyd, R. (1988). How to be a moral realist. In G. Sayre-McCord (Ed.), *Essays on Moral Realism* (pp. 181–228). Cornell University Press.
- Brink, D. (1984). Moral realism and the sceptical arguments from disagreement and queerness. *Australasian Journal of Philosophy*, 62(2), 111–125. <https://doi.org/10.1080/00048408412341311>
- Brink, D. (1989). *Moral realism and the foundations of ethics*. Cambridge University Press.
- Brink, D. (2007). The autonomy of ethics. In M. Martin (Ed.), *The Cambridge companion to atheism* (pp. 149–165). Cambridge University Press.
- Colebrook, R. T. (2018). *Toward a science of morals* [PhD diss.]. CUNY Academic Works. https://academicworks.cuny.edu/gc_etds/2648
- Devitt, M. (2002). Moral realism: A naturalistic perspective. *Croatian Journal of Philosophy*, 2(4), 1–15.
- Devitt, M. (2010). *Putting metaphysics first: Essays on metaphysics and epistemology*. Oxford University Press.
- Doris, J. M., & Plakias, A. (2008). How to argue about disagreement: Evaluative diversity and moral realism. In W. Sinnott-Armstrong (Ed.), *Moral psychology* (Vol. 2, pp. 303–331). MIT Press.
- Dreier, J. (2009). Relativism (and expressivism) and the problem of disagreement. *Philosophical Perspectives*, 23(1), 79–110. <https://doi.org/10.1111/j.1520-8583.2009.00162.x>
- Enoch, D. (2009). How is moral disagreement a problem for realism? *The Journal of Ethics*, 13(1), 15–50. <https://doi.org/10.1007/s10892-008-9041-z>
- Enoch, D. (2011). *Taking morality seriously: A defense of robust realism*. Oxford University Press.
- Feltz, A., & Cokely, E. T. (2008). The fragmented folk: More evidence of stable individual differences in moral judgments and folk intuitions. In B. C. Love, K. McRae, & V. M. Sloutsky, *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 1771–1776). Cognitive Science Society.
- Fitzpatrick, S. (2014). Moral realism, moral disagreement, and moral psychology. *Philosophical Papers*, 43(2), 161–190. <https://doi.org/10.1080/05568641.2014.932953>
- Gill, M. B. (2008). Metaethical variability, incoherence, and error. In W. Sinnott-Armstrong (Ed.), *Moral psychology* (Vol. 2, pp. 387–401). Bradford Books.
- Gill, M. B. (2009). Indeterminacy and variability in meta-ethics. *Philosophical Studies*, 145(2), 215–234. <https://doi.org/10.1007/s11098-008-9220-6>
- Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106(3), 1339–1366. <https://doi.org/10.1016/j.cognition.2007.06.007>
- Goodwin, G. P., & Darley, J. M. (2010). The perceived objectivity of ethical beliefs: Psychological findings and implications for public policy. *Review of Philosophy and Psychology*, 1(2), 161–188. <https://doi.org/10.1007/s13164-009-0013-4>
- Goodwin, G. P., & Darley, J. M. (2012). Why are some moral beliefs perceived to be more objective than others? *Journal of Experimental Social Psychology*, 48(1), 250–256. <https://doi.org/10.1016/j.jesp.2011.08.006>
- Gosling, S. D., Rentfrow, P. J., & Swann, W. B. (2003). A very brief measure of the big-five personality domains. *Journal of Research in Personality*, 37(6), 504–528. [https://doi.org/10.1016/S0092-6566\(03\)00046-1](https://doi.org/10.1016/S0092-6566(03)00046-1)
- Heiphetz, L., & Young, L. L. (2017). Can only one person be right? The development of objectivism and social preferences regarding widely shared and controversial moral beliefs. *Cognition*, 167, 78–90. <https://doi.org/10.1016/j.cognition.2016.05.014>

- Hollos, M., Leis, P. E., & Turiel, E. (1986). Social reasoning in Ijo children and adolescents in Nigerian communities. *Journal of Cross-Cultural Psychology*, 17(3), 352–374. <https://doi.org/10.1177/0022002186017003007>
- Hopster, J. (2019). The meta-ethical significance of experiments about folk moral objectivism. *Philosophical Psychology*, 32(6), 831–852. <https://doi.org/10.1080/09515089.2019.1627304>
- Huemer, M. (2005). *Ethical intuitionism*. Palgrave Macmillan.
- Jackson, F. (1998). *From metaphysics to ethics: A defence of conceptual analysis*. Oxford University Press.
- Joyce, R. (2001). *The myth of morality*. Cambridge University Press.
- Joyce, R. (2016). Moral anti-realism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016). Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2016/entries/moral-anti-realism/>
- Kohlberg, L. (1981). *Essays on Moral Development: The psychology of moral development*. Harper & Row.
- Kohlberg, L., & Kramer, R. (1969). Continuities and discontinuities in childhood and adult moral development. *Human Development*, 12(2), 3–120. <https://doi.org/10.1159/000270857>
- Korsgaard, C. M., & O'Neill, O. (1996). *The sources of normativity*. Cambridge University Press.
- Loeb, D. (2008a). Moral incoherentism: How to pull a metaphysical rabbit out of a semantic hat. In W. Sinnott-Armstrong (Ed.), *Moral psychology* (Vol. 2, pp. 355–385). Bradford Books.
- Loeb, D. (2008b). Reply to Gill and Sayre-McCord. In W. Sinnott-Armstrong (Ed.), *Moral psychology* (Vol. 2, pp. 413–421). Bradford Books.
- Mackie, J. (1990). *Ethics: Inventing right and wrong*. Penguin. Original work published 1977.
- Moss, D. (2017). Experimental philosophy, folk metaethics and qualitative methods. *Teorema: Revista Internacional de Filosofía*, 36(3), 185–203.
- Nettle, D. (2009). *Personality: What makes you the way you are* (1st ed.). Oxford University Press.
- Nichols, S. (2004). After objectivity: An empirical study of moral judgment. *Philosophical Psychology*, 17(1), 3–26. <https://doi.org/10.1080/0951508042000202354>
- Nichols, S., & Folds-Bennett, T. (2003). Are children moral objectivists? Children's judgments about moral and response-dependent properties. *Cognition*, 90(2), B23–B32. [https://doi.org/10.1016/S0010-0277\(03\)00160-4](https://doi.org/10.1016/S0010-0277(03)00160-4)
- Nucci, L. P. (2001). *Education in the moral domain*. Cambridge University Press.
- Nucci, L. P., & Turiel, E. (1978). Social interactions and the development of social concepts in preschool children. *Child Development*, 49(2), 400. <https://doi.org/10.2307/1128704>
- Nucci, L. P., & 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. <https://doi.org/10.2307/1131547>
- Nucci, L. P., Turiel, E., & Encarnacion-Gawrych, G. (1983). Children's social interactions and social concepts. *Journal of Cross-Cultural Psychology*, 14(4), 469–487. <https://doi.org/10.1177/0022002183014004006>
- Pözlner, T. (2017). Revisiting folk moral realism. *Review of Philosophy and Psychology*, 8(2), 455–476. <https://doi.org/10.1007/s13164-016-0300-9>
- Pözlner, T. (2018). *Moral reality and the empirical sciences* (1st ed.). Routledge.
- Pözlner, T. (2018a). How to measure moral realism. *Review of Philosophy and Psychology*, 9(3), 647–670. <https://doi.org/10.1007/s13164-018-0401-8>

- Pölzler, T., & Wright, J. C. (2019). Empirical research on folk moral objectivism. *Philosophy Compass*, 14(5), e12589. <https://doi.org/10.1111/phc3.12589>
- Prinz, J. (2007). *The emotional construction of morals*. Oxford University Press.
- Quintelier, K. J. P., & Fessler, D. M. T. (2012). Varying versions of moral relativism: The philosophy and psychology of normative relativism. *Biology & Philosophy*, 27(1), 95–113. <https://doi.org/10.1007/s10539-011-9270-6>
- Railton, P. (1986). Moral realism. *The Philosophical Review*, 95(2), 163–207. <https://doi.org/10.2307/2185589>
- Rose, D., & Nichols, S. (2019). From punishment to universalism. *Mind & Language*, 34(1), 59–72. <https://doi.org/10.1111/mila.12191>
- Ruse, M. (1998). *Taking Darwin seriously: A naturalistic approach to philosophy* (2nd ed.). Prometheus.
- Saad, L. (2019, March 25). Americans as concerned as ever about global warming. Gallup News Service. <https://news.gallup.com/poll/248027/americans-concerned-ever-global-warming.aspx>
- Sarkissian, H., Park, J., Tien, D., Wright, J. C., & Knobe, J. (2011). Folk moral relativism. *Mind & Language*, 26(4), 482–505. <https://doi.org/10.1111/j.1468-0017.2011.01428.x>
- Sarkissian, H., & Phelan, M. (2019). Moral objectivism and a punishing god. *Journal of Experimental Social Psychology*, 80, 1–7. <https://doi.org/10.1016/j.jesp.2018.08.012>
- Sayre-McCord, G. (2008). Moral semantics and empirical enquiry. In W. Sinnott-Armstrong (Ed.), *Moral psychology* (Vol. 2, pp. 403–412). Bradford Books.
- Sayre-McCord, G. (1988). *Essays on moral realism*. Cornell University Press.
- Schmidt, M. F. H., Gonzalez-Cabrera, I., & Tomasello, M. (2017). Children's developing metaethical judgments. *Journal of Experimental Child Psychology*, 164, 163–177. <https://doi.org/10.1016/j.jecp.2017.07.008>
- Shafer-Landau, R. (2003). *Moral realism: A defence*. Clarendon Press.
- Sinclair, N. (2012). Moral realism, face-values and presumptions. *Analytic Philosophy*, 53(2), 158–179. <https://doi.org/10.1111/j.2153-960X.2012.00558.x>
- Smetana, J. G. (2006). Social-cognitive domain theory: Consistencies and variations in children's moral and social judgments. In M. Killen & J. G. Smetana (Eds.), *Handbook of moral development* (pp. 119–153). Lawrence Erlbaum Associates.
- Smetana, J. G., & Braeges, J. L. (1990). The development of toddlers' moral and conventional judgments. *Merrill-Palmer Quarterly*, 36(3), 329–346.
- Smith, M. (1994). *The moral problem* (1st ed.). Blackwell.
- Turiel, E. (1979). Social convention and morality: Two distinct conceptual and developmental systems. *Nebraska Symposium on Motivation*, 25, 77–116.
- Turiel, E. (1983). *The development of social knowledge: Morality and convention*. Cambridge University Press.
- Wainryb, C., Shaw, L. A., Langley, M., Cottam, K., & Lewis, R. (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. <https://doi.org/10.1111/j.1467-8624.2004.00701.x>
- Wong, D. B. (2006). *Natural moralities: A defense of pluralistic relativism*. Oxford University Press.
- Wright, J. C., McWhite, C. B., & Grandjean, P. T. (2014). The cognitive mechanisms of intolerance. In J. Knobe, T. Lombrozo, & S. Nichols (Eds.), *Oxford studies in experimental philosophy* (Vol. 1, pp. 28–61). Oxford University Press.
- Wright, J. C. (2018). The fact and function of meta-ethical pluralism. In T. Lombrozo, J. Knobe, & S. Nichols (Eds.), *Oxford studies in experimental philosophy* (Vol. 2, pp. 119–150). Oxford University Press.

- Wright, J. C., McWhite, C. B., & Grandjean, P. T. (2013). The meta-ethical grounding of our moral beliefs: Evidence for meta-ethical pluralism. *Philosophical Psychology*, 26(3), 336–361. <https://doi.org/10.1080/09515089.2011.633751>
- Yilmaz, O., & Bahçekapili, H. G. (2015). Without god, everything is permitted? The reciprocal influence of religious and meta-ethical beliefs. *Journal of Experimental Social Psychology*, 58, 95–100. <https://doi.org/10.1016/j.jesp.2015.01.003>
- Zijlstra, L. (2019). Are people implicitly moral objectivists? [*Unpublished manuscript*].