The meta-ethical significance of experiments about folk moral objectivism

Jeroen Hopster

Published online: 21 Jun 2019.

Submit your article to this journal ↗

Article views: 746

View related articles ↗

View Crossmark data ↗

Routledge
Taylor & Francis Group

ARTICLE

🔓 OPEN ACCESS ✓ Check for updates

The meta-ethical significance of experiments about folk moral objectivism

Jeroen Hopster

Department of Philosophy, Utrecht University, The Netherlands

ABSTRACT

The meta-ethical commitments of folk respondents – specifically their commitment to the objectivity of moral claims – have recently become subject to empirical scrutiny. Experimental findings suggest that people are meta-ethical pluralists: There is both inter- and intrapersonal variation with regard to people's objectivist commitments. What meta-ethical implications, if any, do these findings have? I point out that current research does not directly address traditional meta-ethical questions: The methods used and distinctions drawn by experimenters do not perfectly match those of meta-ethicists. However, I go on to argue that, in spite of this mismatch, the research findings should be of interest to moral philosophers, including meta-ethicists. Not only do these findings extend the field of moral psychology with new data and hypotheses, but they also provide tentative evidence that touches on the adequacy of theses in moral semantics and moral metaphysics. Specifically, they put pressure on arguments in support of moral realism.

1. Introduction

Over the last decade, several philosophers and psychologists have begun to test non-philosophers' commitments concerning the objectivity of moral judgments in an experimentally rigorous way (e.g., Beebe, 2014; Fisher, Knobe, Strickland, & Keil, 2017; Goodwin & Darley, 2008, 2012; Khoo & Knobe, 2018; Nichols, 2004; Sarkissian, Parks, Tien, Wright, & Knobe., 2011; Theriault, Waytz, Heiphetz, & Young, 2017; Wright, Grandjean, & McWhite, 2013; Wright, McWhite, & Grandjean, 2014). Empirical research on people's meta-ethical commitments is not entirely new; the moral–conventional distinction has long been a topic of experimental research (e.g., Turiel, 1983), and specific kinds of moral violations, such as the violation of harm norms, have also been studied experimentally (e.g., Haidt, Koller, & Dias, 1993). However, only in recent years have experimenters begun to inquire systematically into how lay individuals think

---

about issues that have traditionally been focal points of meta-ethical theorizing – for example, the nature of moral motivation, moral disagreement, moral truth, and moral objectivity.

In this article, I focus on one specific topic of this new wave of research – folk moral objectivism – and assess its meta-ethical implications. These implications are not straightforward. Critics of experimental philosophy (e.g., Bengson, 2013; Kauppinen, 2007; Ludwig, 2010) even question whether experimental research can touch upon traditional philosophical questions at all. As these critics point out, the survey-methodology typically employed by experimenters is ill-suited to scrutinize people's reflective commitments. By contrast, when meta-ethicists make a claim about lay people's objectivist commitments, their claim typically does concern people's reflective commitments. Hence, there seems to be a mismatch between the *explanandum* of experimenters and that of meta-ethicists. Moreover, this mismatch is aggrandized by the fact that the distinctions drawn by experimenters on folk moral objectivism do not always perfectly match those drawn by meta-ethicists.

I submit that these worries should be taken seriously, and I call for great caution in extrapolating the results of experimental findings to a meta-ethical context. However, I will argue, nonetheless, that there may be fertile interactions between philosophical meta-ethics and psychological studies of people's meta-ethical commitments. The primary upshot of the new wave of experiments is to extend the field of moral psychology with a new set of questions, for instance, with those regarding the psycho-social strategies underlying people's meta-ethical commitments. The theories thereby advanced – such as Wright et al.'s (2013, 2014) theory of 'folk pluralism' – should not be regarded as alternatives to traditional meta-ethical theories: They shed light on the *psychology* of moral judgment, whereas meta-ethicists have traditionally been interested in the *concept* of moral judgment. In this respect, the new findings and hypotheses are complementary to traditional meta-ethical work, but, additionally, some experimental findings provide indirect evidence that touches upon issues in moral semantics and moral metaphysics.

The article proceeds as follows. In Section 2, I summarize the new wave of experimental research on folk moral objectivism and outline some of its most well-established findings. In Section 3, I make some critical remarks about the methodology used by several experimenters and explain why their findings do not directly address traditional meta-ethical concerns, but I proceed to argue that, nonetheless, experimental research does have the potential to shed light on some traditional meta-ethical theses. In Section 4, I discuss the *concept* of moral judgment, and I argue that, although the conceptual implications of experimental findings can only be indirect, these findings create argumentative pressures in extant debates, for

example, in the debate about the tenability of error theory. In Section 5, I turn to the relevance of experimental findings for moral metaphysics, and I argue that the findings give reason to question 'presumptive arguments' in support of moral realism. In Section 6, I highlight recent hypotheses about the psycho-social function of meta-ethical judgments, which build on the experiments discussed. These hypotheses are more naturally understood as belonging to the field of moral psychology than to that of meta-ethics; nonetheless, I argue that they may be taken up as a further resource to argue against the 'presumptive support' for moral realism. I conclude that although research on folk moral objectivism is still in its early stages, and various methodological improvements are called for, this research is beginning to fill an empirical hiatus and has the potential to become an important resource in meta-ethical theorizing.

2. Experimental findings: An overview

In this section, I highlight some of the main findings of the recent wave of experiments on folk objectivism, assuming the validity of the underlying research methodology (I discuss limitations of this methodology in Section 3). These experiments have been conducted across different cultures and with different age-groups, using different research methods and asking respondents a variety of questions. The most common research method has been to present respondents with surveys designed to test people's objectivist commitments. For instance, respondents might be asked whether or not they regard the statement "Consciously discriminating against someone on the basis of race is unacceptable" as a moral statement, whether they think that the statement is true, false, or a matter of opinion, and whether they think that, in cases of disagreement, one of the disagreeing parties would surely be mistaken or both parties would be equally justified in holding their view (e.g., Goodwin & Darley, 2008; Wright et al., 2013). Experimenters commonly operationalize 'objectivity' as the combined commitment to moral statements being either true or false, and to moral disagreement being indicative of mistake. For example, if respondents judge that the claim "Consciously discriminating against someone on the basis of race is unacceptable" is true and that in cases of disagreement, one of the disagreeing parties will surely be mistaken, then they will be classified as objectivists. If they think that this claim is a matter of opinion and that in cases of disagreement, both parties might be equally correct, then they will be classified as non-objectivists.

With concern to the moral domain in general, it has been found that people generally regard moral judgments as much more objective than judgments of taste, somewhat more objective than judgments of convention, and slightly less objective than judgments of fact (Goodwin & Darley, 2008; though see Theriault et al., 2017 for a somewhat different result).

However, there is substantial variation with regard to the objectivity attributed to specific moral judgments, which depends *inter alia* on the topic under consideration (Wright et al., 2013). Hence, rather than morality being regarded as uniformly objective, research suggests that people perceive moral objectivity as more of a gliding scale: Some moral issues are regarded as being as objective as matters of fact, while others are regarded as being more akin to questions of taste.

The variability of test-subjects' objectivist endorsements has been one of the most notable and consistent findings of the experimental research. By now, the fact that people are pluralists with regard to their commitment to moral objectivity has been well established. Indeed, a good case can be made that the results of earlier studies which suggested that people are objectivists should be re-interpreted as supporting the view that people are meta-ethical pluralists (Pölzler, 2017a). This pluralism manifests itself in several dimensions:

- *Interpersonal variation*: There exists variation in the objectivist leanings of different people. This variation has been found to correlate with differences in personality traits (Feltz & Cokely 2008) as well as with different stages in a person's lifespan (Beebe & Sackris, 2016): People tend to become gradually more objectivist over time but increasingly relativist during their early twenties. Goodwin and Darley (2010) found that a religious grounding of ethics is another predictor of an objectivist meta-ethical stance, whereas a better capacity for 'disjunctive thinking' correlates with diminished moral objectivism.[1]

- *Intrapersonal variation*: There also exists intrapersonal variation in the objectivity ascribed to moral judgments along different dimensions. One of these is the aforementioned variability in terms of topic. Depending on the moral issue under consideration, moral judgments may be regarded as very objective or much less so (Wright et al., 2013). Typically, moral issues that are regarded as subjective tend to be issues that are politically controversial – for example, abortion and assisted suicide, among American respondents. Some experimenters have found that perceived social consensus is an important predictor of this variability (Beebe, 2014; Goodwin & Darley, 2012): The less social consensus there is on a given moral issue, the less objective people regard the issue to be. Moreover, already at an early age, the idea emerges that disagreements about widely shared moral beliefs have only one right answer, whereas disagreements about controversial moral beliefs do not (Heiphetz & Young, 2017).

The level of objectivity attributed to moral judgments also varies with contextual and situational factors. Sarkissian et al. (2011) found that if two disagreeing parties are depicted as belonging to the same cultural group, then, in situations of disagreement, respondents typically judge that one of the disagreeing parties must be mistaken. However, when the disagreement involves individuals from different cultural or evolutionary backgrounds, then the answers to moral disputes are regarded as more relativist, and respondents are inclined to judge that both parties in the disagreement can be correct. Hence, standards of moral correctness are taken to be at least somewhat relative to a specific culture: Objectivist intuitions vary as a function of cultural distance, with increased distance leading to decreased attributions of objectivity. Using a similar research setup, Khoo and Knobe (2018) found that there are moral-conflict cases in which people are inclined to say both that two speakers disagree and that it is not the case that at least one of them must be saying something wrong. Hence, it seems that people allow for the possibility that there are moral disagreements in which the claims of the disagreeing parties do not exclude each other.

Goodwin and Darley (2012) found that greater moral objectivity is associated with 'closed' rather than 'open' responses in the face of moral disagreement. Following up on this finding, Fisher et al. (2017) found that, at least with regard to controversial issues, people's mode of social inter-action influences their construal of moral truth. When people adopt a cooperative 'argue-to-learn' mindset, they typically think that controver-sial moral issues admit of more than one correct answer. When people adopt a competitive 'argue-to-win' mindset, they typically take on a more objectivist stance.

Lastly, there have been some findings suggesting that people's meta-ethical commitments affect their moral behavior and everyday moral decision-making. For instance, Wright et al. (2014) found that subjects stated that they are less willing to interact with and help people who disagree with them about moral issues, especially if these issues are taken to be objective. The authors suggested that regarding a moral issue as objective is related to lower tolerance for disagreement, whereas non-objectivism leaves an opening for diversity and dialogue.

These findings shed light on a new area of moral psychology – the psychology of meta-ethics. They may provide insight about the psycholo-gical and social functions that attributing a certain meta-ethical status to one's judgments might serve (see Section 6). However, none of the experi-ments reviewed here engage directly with traditional meta-ethical objects of study, such as the concept of moral judgment or the metaphysics of moral objectivity. As a result, and as I will emphasize in the next section, attempts to draw direct meta-ethical conclusions based on current experi-mental findings are often problematic.

3. Meta-ethical implications? Pitfalls and promises

The aforementioned findings and hypotheses are steadily beginning to fill an empirical hiatus in discussions about moral objectivity, but their meta-ethical promise should not be overstated, as recent commentators have pointed out (Beebe 2015; Pölzler, 2018 chapter, p. 3). In fact, some philosophers maintain that the ambitions of experimentalists to shed light on traditional philosophical questions is misguided in general (e.g., Bengson, 2013; Kauppinen, 2007; Ludwig, 2010). Building on these criticisms, I will highlight some of the pitfalls of experiments on folk moral objectivism, and I will explain why current studies do not lend themselves to the drawing of direct meta-ethical conclusions.[2] These limitations need to be taken seriously. Nonetheless, I will conclude this section by arguing that they do not fully undermine the promise of current experimental findings to serve as a resource in the meta-ethical debate.

One potential pitfall of many of the studies mentioned in Section 2 is their reliance on a survey-methodology. Several critics of experimental philosophy have underlined the deficiencies of surveys as a means for investigating people's philosophical stances. When philosophers make an assertion about the contents of people's intuitions, or about people's conceptual commitments, they typically mean to make a claim about these intuitions or commitments under idealized conditions. Performance errors, semantic incompetence, and other philosophically irrelevant factors may detract from these idealized conditions in survey-based research. Indeed, some philosophers maintain that the kinds of intuitions that philosophers are interested in can only become apparent in Socratic dialogue (Kauppinen, 2007). The survey-methodology used by experimental philosophers, by contrast, typically serves to probe respondents' immediate responses, often prompted by means of a forced choice paradigm. Such responses are typically quick guesses, hypotheses, and emotional reactions (Bengson, 2013), which clearly differ from the "intellectual intuitions" that philosophers are interested in (Ludwig, 2010).

This criticism also applies to research on folk moral objectivism, which tends to rely heavily on a survey methodology. Surveys seem to capture people's *first-off* meta-ethical responses. As such, they differ from meta-ethicists' traditional object of analysis, which is people's *endorsed* commitments. Therefore, the experimental findings cannot directly shed light on traditional meta-ethical questions.

Adding to this mismatch is the fact that the distinctions experimenters draw do not always capture existing meta-ethical distinctions. To give just one example, consider experimenters' operationalization of 'moral objectivity' as the dual commitment to a moral statement being either true or false and moral disagreement being indicative of a mistake (e.g., Goodwin

& Darley, 2008). On a standard meta-ethical classification, the former part of the experimentalist's operationalization is compatible with various meta-ethical views, including moral realism, subjectivism, and constructivism. Yet, many meta-ethicists regard only the first of these views as a form of objectivism, in virtue of the realist's commitment to moral truth being stance-independent (Shafer-Landau, 2003). Hence, what many experimentalists regard as being a condition that adds to a moral statement's objectivity, many meta-ethicists do not. This illustrates, once more, that experimentalists tend to study different phenomena than meta-ethicists have traditionally been interested in.

Insofar as they aspire to shed light on traditional meta-ethical theses about moral objectivity, then, we should conclude that the current experiments suffer from serious flaws. However, if we look at the aspirations of these experiments from another angle, as shedding light on a different set of phenomena altogether – namely, people's immediate responses rather than their reflective intuitions and 'moral objectivity' in the experimenter's psychological sense rather than the meta-ethicist's philosophical sense – then this mismatch does little to discredit the experiments. This, I submit, is the most charitable interpretation of extant experiments about folk objectivism: They shed light on a topic that lies in the vicinity of, but differs from, traditional meta-ethical concerns.

Apart from its mismatch with traditional meta-ethics, the methodology of experiments on folk moral objectivism has also been criticized on other grounds.[3] As Beebe (2015) points out, in several existing studies, the questions posed by experimenters can be read in multiple ways, and there is no guarantee that the answers respondents give are tracking the meta-ethically relevant distinctions. More generally, it may be questioned whether current studies really succeed in pulling apart task demands, experimenter effects, and people's meta-ethical commitments. Additionally, many current studies suffer from the "stimuli as fixed effects fallacy" (Judd, Westfall, & Kenny, 2012): They generalize from findings about specific examples to conclusions about a domain in general, even though there is no statistical basis for this generalization.

All of these methodological flaws may be overcome, at least in principle (see, e.g., Theriault et al., 2017), but they do detract from the general reliability of existing results. In what follows, I will proceed on the assumption that, upon overcoming these methodological flaws, the main research findings, as outlined in Section 2, will remain intact. It should be kept in mind, however, that future research may prove this assumption wrong.

In closing this section, let's start to consider whether, and how, current findings may still be meta-ethically relevant, in spite of the criticisms outlined above. While these findings do not have direct meta-ethical implications, they do lie in the vicinity of traditional meta-ethical

concerns, and I will argue that the relations between them are worth exploring – both with regard to traditional meta-ethical questions (see also Section 4 and Section 5) and with regard to the prospect of new philosophical questions that psychological research might give rise to (see also Section 6).

To start with the latter, one of the upshots of the experiments is that they highlight previously unrecognized patterns of stability in folk meta-ethical judgments. These patterns might be relevant for meta-ethical practice – for instance, by highlighting how the choice of examples can influence meta-ethical analysis (see Section 4). Additionally, these patterns may themselves be proper targets of meta-ethical analysis, and they may be connected with theoretical insights about the psychological, social, and evolutionary functions of folk meta-ethical judgments. While some of these insights go beyond the strictly meta-ethical domain, they can be connected with theses that have traditionally been advanced in a meta-ethical context. Consider, for instance, Ruse's (1995 [1986]) argument for a version of error theory, which rests on an empirical hypothesis about the biological function of objectivist moral beliefs. While the philosophical tenability of Ruse's argument is contested (see, e.g., Joyce, 2016), Ruse's empirical hypothesis might be of meta-ethical interest nonetheless. Research on folk moral objectivism has recently contributed to an elaborate novel hypothesis in the same evolutionary ballpark (see Stanford, 2018), which raises questions that are of interest to moral psychologists and, possibly, to meta-ethicists as well. For instance, as I will argue in Section 6, Stanford's suggestion that the objectivist features of our moral phenomenology can be evolutionarily explained can be taken up in an argument against moral realism.

Secondly, with regard to their relevance for traditional meta-ethical questions, and in defense of the survey method used by experimenters, it should be noted that although people's prompted responses about moral objectivity may come apart from their considered intuitions, it is by no means guaranteed that they will. In fact, there is some evidence that the patterns in philosophical responses unveiled by experimental research on first-off intuitions do not disappear after greater reflection (Weinberg et al., 2012). Further experimental research is needed here; awaiting such research, we can only guess how people's immediate responses relate to their reflective intuitions. However, given the pervasiveness of rationalization in moral reasoning, it seems reasonable to guess that findings about people's first-off moral responses provide at least tentative evidence about the contents of their reflective intuitions.

Thirdly, in response to the criticism that experimenters often do not frame moral objectivism in the same way as meta-ethicists – particularly, realists – tend to do, it should be pointed out that there is, in fact,

substantial variation in the way that meta-ethicists conceive of moral objectivity. To give just one example of this diversity, the aforementioned realist account of moral objectivity in terms of stance-independence has been challenged – partly on experimental grounds – by an antirealist account which frames moral objectivity in terms of stance-invariance (Hopster, 2017). The fact that different meta-ethicists adhere to different accounts of moral objectivity should give some pause to critics who argue that experimentalists fail to capture a meta-ethically relevant phenomenon since they do not capture moral objectivity in the realist's sense. As Beebe (2015, p. 27) rightly points out, "the investigation of folk views about the objectivity of morality should be broadened to encompass more than just the issue of moral realism." Perhaps, depending on the context, moral judgments may purport to be objective in a variety of ways (e.g., in being stance-independent, stance-invariant, authority-independent, and so on). At the very least, the meta-ethical primacy of a realist conception of moral objectivity is not sufficiently well established to regard research about alternative conceptions as meta-ethically irrelevant.

Lastly, the fact that there exists substantial diversity in the views of professional meta-ethicists, all of whom have dedicated much time to reflecting on the topic of moral objectivity, may be taken as suggestive – if anecdotal – evidence in support of the hypothesis that the pluralism of folk respondents does not merely reflect their incompetence but constitutes a real feature of people's meta-ethical intuitions.

4. The concept of moral judgment

Suppose that current experimental findings (Section 2) are indeed along the right lines, and the meta-ethical status that people attach to moral judgments varies along multiple dimensions. How might this be relevant for traditional meta-ethical analysis?

First, the findings are in tension with traditional philosophical approaches to meta-ethics, many of which do not allow for much flexibility or hybridization. As Gill (2009) argues, much of 20[th]-century meta-ethics proceeded on the assumption that ordinary moral discourse is uniform and determinate. Classical meta-ethical questions – for example, whether there is a necessary connection between moral judgment and moral motivation (Gill, 2009) and whether moral language is cognitivist or non-cognitivist (Loeb, 2008) – were assumed to have invariant and determinate answers. Experimental findings provide preliminary evidence that this assumption is mistaken: There is inter- and intrapersonal variation regarding the meta-ethical status of moral claims, which problematizes any view that assumes that this status is invariant (Beebe, 2014; Khoo & Knobe, 2018).

Secondly, the findings give pause to meta-ethicists who defend a thesis on the basis of a narrow set of examples. One important lesson for meta-ethicists to draw from the experimental findings is that their choice of examples might heavily influence the resulting analysis (Sarkissian 2017, p. 574). Experiments show that situational and contextual factors substantially impact respondent's objectivist proclivities, which should make meta-ethicists more cautious about regarding a few core examples as being representative of moral discourse as a whole. For instance, if a meta-ethicist's example of choice is a particularly uncontroversial case ("torturing innocent people is morally bad"), then people's meta-ethical intuitions will tend to be more objectivist than they are with disputed moral claims. Similarly, if an example is presented in a situation of intra-cultural conflict, then people's meta-ethical intuitions will be more objectivist than in a situation of inter-cultural dialogue. Here, meta-ethicists can clearly profit from experimental findings: By providing a better understanding of the general correlates of people's moral intuitions, these experiments can serve as a corrective to a biased choice of examples and provide a better understanding of the respects in which alleged paradigm cases of moral judgment are really representative of moral judgment across the board.

Thirdly, the findings may help to shed light on the contents of the concept of moral judgment. It seems plausible that, with regard to our concepts, there has to be a fairly tight connection between their actual contents and the contents that most people take them to have. If this connection were very loose, then it would seem that in analyzing the contents of our concepts, philosophers are simply changing the topic: They are no longer analyzing the concepts that people ordinarily use (. Loeb, 2008). During the latter half of the 20th century, meta-ethicists have often assumed that in making moral judgments, people are conceptually committed to a claim to objectivity. The finding of folk pluralism, however, suggests that this may not always be the case; at the very least, this shifts the burden of argument to the meta-ethicist who maintains that they are.

One possible response for meta-ethicists is to withdraw to the armchair and deny that experimental findings have any conceptual relevance at all. This position may, for example, be advanced by externalists with regard to the contents of moral concepts. According to content-externalists, features of the world function as reference-fixers for moral concepts. What lay people make of the concept of moral judgment need not carry much evidential weight: The true meaning of the concept may simply be opaque to them.

Most meta-ethicists, however, find it more plausible to pursue the route of content-internalism – to view moral concepts as mental constructs that are transparent to reflection (Laskowski & Finlay, 2017). For the content-internalists who claim that the concept of moral judgment commits us to

a claim of objectivity, experimental findings pose an argumentative challenge. Perhaps this challenge can be met, with the support of reasons similar to those mentioned in Section 3. For instance, content-internalists commonly assume that upon reflection, competent participants in moral discourse will tend to agree on whether a proposition (such as the proposition that moral judgments make a claim to objectivity) should be regarded as a platitude or not. By contrast, existing experiments scrutinize respondents' immediate commitments to moral objectivity, rather than their stance on reflection. (Additionally, they do not check for the conceptual competence of test subjects.) As a result, present experimental evidence may be too shaky to have a great impact on this debate. With improved experimental methodologies, however, this research could, potentially, have such an impact.

The outcome of this debate, in turn, is also relevant for the tenability of more general meta-ethical theories, such as John Mackie's (1977) error theory. According to Mackie, our concept of 'moral judgment' entails a reference to objectively prescriptive properties. Current experimental findings put some pressure on Mackie's empirical claim by suggesting that people's commitment to objectivism is both more moderate and more flexible than Mackie assumed. We should keep in mind that, as of yet, no research has been conducted that specifically questions whether people intend for the concept of 'moral judgment' to refer to objectively prescriptive properties (Fraser, 2014). Instead, people's objectivist commitments are tested using a somewhat different measure, which can only provide indirect evidence for or against Mackie's error theory. However, even if the current evidence is tentative, it provides a preliminary reason to think that Mackie's conceptual claim may not hold up – at least not for all moral judgments (Sarkissian, 2016, p. 222). In this indirect way, current experimental findings can be relevant for extant meta-ethical debates – by shifting the burden of argument. However, their capacity to do so often depends on additional meta-ethical assumptions, and it may, in part, require further empirical exploration (see Pölzler, 2017b for further discussion).

5. The metaphysics of moral judgment

In the previous section, I pointed out that with regard to our conceptual commitments, it seems plausible that there has to be a fairly tight connection between what these commitments are and what people take them to be. With regard to our metaphysical commitments, the same reasoning may not apply. The fact that many people once thought that witches exist, or, today, think that God exists, does little to enhance the probability that witches or God actually exist (Joyce, 2015). Similarly, even if it were

experimentally shown that many people believe in the existence of objectively prescriptive properties, this does not establish that such properties indeed exist. The folk might simply be mistaken.

We should, therefore, expect experimental findings about folk objectivism to be of little relevance to the metaphysics of moral judgment. Nonetheless, as I will argue in the present section, given the actual dialectic of the meta-ethical debate, such findings are not completely mute, so to speak. An influential type of argument which several meta-ethicists have advanced does purport to derive metaphysical conclusions on the basis of a claim about folk judgments: the so-called 'experiential' or 'presumptive' argument in support of moral realism (e.g., Brink, 1989; Shafer-Landau, 2003; for criticisms, see Kirchin, 2003; Loeb, 2007; Sinclair, 2012; Pölzler, 2017a). Since experiments about folk moral objectivism touch upon the adequacy of this claim about folk judgments, they also touch upon the adequacy of the presumptive argument in support of realism.

There are, in fact, various kinds of presumptive arguments playing a role in meta-ethics which engage with different aspects of moral discourse and practice and differ in how they claim these aspects to be meta-ethically relevant (e.g., Loeb, 2007).[4] Some presumptive arguments claim to best accommodate moral discourse in the sense of best *explaining* it; others claim to best accommodate moral discourse in the sense of best *justifying* it. In what follows, I shall be interested in the latter version of the argument, as used in support of moral realism. Following Sinclair (2012, pp. 158–159), this argument relies on three premises:

- *Premise 1*: A claim about the face-value commitments of moral discourse and practice.
- *Premise 2*: The claim that moral realism best justifies these face-value commitments.
- *Premise 3*: The claim that it is a *desideratum* of meta-ethical theories to justify the face-value commitments of moral discourse and practice.

The first claim is empirical, and it concerns findings of the sort that experiments about folk objectivism shed light upon. The second is a philosophical claim about the justificatory power of moral realism versus that of rivalling theories. The third claim concerns the meta-ethical relevance of folk intuitions about moral metaphysics. (It contradicts the suggestion with which this section began – that such intuitions are of limited meta-ethical relevance.)

The presumptive argument in support of moral realism is of particular importance to the meta-ethical debate, as it has an alleged potential to shape the philosophical dialectic. If the argument holds up, then the starting point of meta-ethical inquiry should be to vindicate some form

of moral realism. Indeed, this is precisely what David Enoch (2017), another defender of this type of argument, has suggested in recent work. In the remainder of this section, I shall critically discuss Enoch's argument on empirical grounds.

According to Enoch, moral discourse is implicitly committed to a realist – specifically, a non-naturalist – ontology. By way of illustration, he considers the following six considerations about the face-value commitments of moral discourse (Enoch, 2017, pp. 29–42):

(1) Moral language functions much like representational language does in other domains.
(2) Moral discourse seems to exhibit objective purport. For example, we are comfortable applying moral standards to other people without inquiring about their own moral inclinations, and we don't withdraw once we find out that their moral commitments differ from ours.
(3) We seem to endorse – pre-theoretically – counterfactuals that do not sit well with response-dependence. (If eating meat is morally wrong, then, presumably, it would be wrong even if no one had ever acknowledged its wrongness.)
(4) We think that, in cases of moral disagreement, only one of the disagreeing parties can be right.
(5) Paradigmatic moral facts and properties appear to be very different from paradigmatic natural ones.
(6) Moral evaluation is not just the evaluation of something-qua-something, but of goodness *simpliciter*.

As Enoch is aware, these considerations do not vindicate non-naturalist realism. Other meta-ethical theories might also be able to accommodate face-value commitments. Nonetheless, they do seem to give non-naturalist realism an advantage over rivalling theories, as it has a particularly good fit with the data. Hence, Enoch contends that, all other things being equal, face-value observations make non-naturalist realism "the view to beat" in meta-ethics.

5.1. Criticism of premise 1

Is Enoch's presumptive argument in support of moral realism tenable in the light of current evidence about folk moral objectivism? Let's assess how Enoch's six considerations hold up in the light of this research:

(1) Whether people think that moral language in general is representational is a question that has not yet been a specific topic of

experimental research. Indeed, some might want to argue (e.g., Sinclair, 2012, p. 168) that this is not a kind of question that is apt for experimental research. After all, that the function of moral language is representational is a sophisticated philosophical claim, and it is unlikely to be part of folk meta-ethical discourse. However, if this is correct, then Enoch's presumption begs the question: He simply presumes that people are moral realists in a technical sense.

Alternatively, it may be argued that the question of whether moral language is representational is implicit in moral discourse and can be experimentally tested. Several experiments have tested whether participants think that moral statements are true or false, or matters of opinion (e.g., Goodwin & Darley, 2008, 2012; Wright et al., 2013, 2014). The findings of these studies do not corroborate Enoch's contention. By contrast, they suggest that the face-value commitments of moral discourse are a thoroughly mixed bag, and they suggest that whether respondents think that moral claims are truth-apt or not is *inter alia* dependent on the topic under consideration. These findings suggest that the claim that moral language functions much like representational language is likely to hold up only under a specific choice of examples.

(2) Does moral discourse exhibit objective purport? Again, there is experimental evidence that the answer depends on contextual factors. Whether people are comfortable applying moral standards to other people depends on the cultural and evolutionary proximity of these people (Sarkissian et al., 2011). This suggests that, in some situations, moral discourse does exhibit objective purport, but, in other situations, it does not.

(3) What moral counterfactuals people are inclined to endorse has not been experimentally studied. It won't come as a surprise that realists like Enoch endorse counterfactuals that do not sit well with response-dependent views, but neither would it be a surprise if countervailing counterfactuals were endorsed by antirealists (e.g., Street, 2009). Awaiting further research, it seems premature (and question-begging) to attach much argumentative weight to these suspicions.

(4) Enoch's claim about the semantics of disagreement has not been borne out by Khoo and Knobe's (2016) experimental research. In fact, their findings suggest that, contrary to what was previously thought, many people think that there can be moral disagreements in which none of the disagreeing parties are mistaken. As a result, if we want to accommodate the folk's face-value commitments, there is a pressure to justify the intuition that non-exclusionary moral disagreement is possible.

(5) Whether paradigmatic moral facts appear to be different from paradigmatic natural facts may depend on what one regards as paradigm cases. Therefore, without further argument, not much argumentative weight can be attached to single examples.

(6) That moral evaluation is generally regarded as the evaluation of goodness *simpliciter* is not brought out by experimental findings. In fact, current research suggests that various contextual features are relevant to moral evaluation (e.g., Fischer et al. 2017; Sarkissian et al., 2011), which puts some pressure on Enoch's claim.

5.2. Criticism of premise 2

The face-value commitments of moral discourse can only provide support for moral realism if realists are best able to account for these findings. If rivalling theories are able to account for them equally well, then the presumptive support for moral realism is undermined. Traditionally, the purported objectivism of moral discourse has been cited as the most obvious face-value commitment in support of realism. However, antirealist proposals justifying this purported objectivism have also been advanced (Hopster, 2017). Therefore, further argument is needed to get the presumptive argument in support of realism started.

Importantly, realism actually would have a clear advantage over rivalling theories if the distinctive metaphysical commitment of realism – the existence of mind-independent properties – were presupposed by ordinary moral discourse and practice. However, thus far, experimental findings have provided no indication that the face-value commitments of moral discourse and practice support a distinctively realist metaphysics. Indeed, as Kirchin (2003, pp. 249–51) argues, it is implausible that the notion of mind-independent moral facts is part of our moral experience. Similarly, as Sinclair (2012, p. 168) observes:

> Everyday moralisers seldom show appreciation of metaethical issues and categories, let alone appreciation of a particular metaethical theory ... Every day moralisers hold that some things are right, others wrong, that some moral judgments are true, that there are correct answers to moral questions, that moral disagreement is possible, and so on. But these are not categories that define realism. What defines realism is the view that moral judgments have a characteristic linguistic function, express states of mind with a characteristic representational function and (therefore) that their truth consists in correspondence between the representational content of such states and the moral way of the world.

On this more technical characterization, we could properly distinguish realist commitments from the commitments of rivalling views, but it is implausible that realism in this more technical sense is presupposed by moral discourse.

5.3. Criticism of premise 3

Even if our moral discourse seemed to be realist, it would remain an open question whether our discourse is justified or not.[5] As I said, people's metaphysical intuitions may be in error. Indeed, it is not obvious why we should think that common intuitions about the (non-)existence of mind-independent moral properties are on track. An argument is required to explain why folk commitments constitute any argumentative weight – why meta-ethicists should accommodate face-value intuitions in the sense of justifying them. After all, these are merely pre-theoretical intuitions concerning a subject matter for which the importance of pre-theoretical intuitions is not evident. Plausibly, it doesn't pose much of an intellectual cost to rid ourselves of such intuitions and replace them with more reflective understanding of moral metaphysics.

Consider a weaker version of Premise 3: Conformity to our metaphysical intuitions provides some inductive support for a meta-ethical view. This weaker premise may be defensible: Just like the plausibility of a scientific claim might be somewhat enhanced if it conforms to our scientific intuitions, a metaphysical claim might derive some support from conforming to intuition. The extent to which it does, however, crucially depends on the issue under consideration. For instance, ordinary people's intuitions about the behavior of particles at the micro-level should carry extremely little weight for scientists studying micro-level phenomena: After all, there is little reason to think that such intuitions are truth-apt. How about metaphysical intuitions concerning the reality of mind-independent moral properties? Perhaps such intuitions might carry more weight, for instance, if the properties in question are directly perceptible. However, to assume that they are is a meta-ethical assumption which lies at the heart of the debate between realists and antirealists. It would be question-begging to rely on this assumption in defense of a presumptive argument for moral realism. If people's metaphysical intuitions are products of philosophical argument, then the meta-ethical *desideratum* is not to capture these intuitions; rather, it is to evaluate the cogency of the argument that generates them.

6. The psychology of meta-ethics

The findings highlighted in Section 2 have generated new questions about the cognitive processes associated with meta-ethical judgments, as well as the evolutionary and psycho-social functions that meta-ethical commitments serve. These questions fall in the dominions of what we might call *the psychology of meta-ethics*. In this section, I highlight some of the hypotheses that this new area of research has given rise to. They give an

impression of the various avenues of theorizing triggered by research on folk moral objectivism – research avenues that go beyond the traditional subject-matter of meta-ethicists but are nonetheless likely to attract philosophical interest.

One hypothesis, formulated by Wright et al. (2013), (2014)), addresses the widespread intrapersonal variation regarding the objectivity of moral judgments. Wright et al. hypothesize that judgments about moral objectivity regulate how tolerant individuals and communities are to divergent moral judgments. If people think that a judgment is not to be tolerated, they adopt an objectivist stance; if they are willing to enter into discussion about the judgment, they ground the judgment subjectively. This tendency is particularly clear with regard to moral issues – and much less so, for instance, with issues that are regarded as matters of convention or taste. As the authors explain, the capacity to view moral issues pluralistically modulates

> (…) the level of permissible choice and dialogue about moral issues, both within and between sociocultural groups. Viewing a moral issue as objectively grounded removes it from the realm of legitimate personal/social negotiation (i.e., individual and/or social attempts to condone it will be deemed unacceptable, and censorship/prohibition will be supported). Viewing a moral issue as non-objectively grounded, on the other hand, allows people to acknowledge its moral significance (i.e., that it is not simply a personal matter), while at the same time maintaining room for choice, dialogue, and debate (…) (Wright et al., 2014, p. 31).

This hypothesis gains further support from Fischer et al.'s (2017) finding that a competitive mind-set invites an objectivist construal of moral truth, whereas a cooperative mind-set invites a subjectivist construal of moral truth. In a competitive setting, moral issues are taken out of the realm of open discussion, people's attitudes become less flexible, and disagreeing parties are judged to be increasingly immoral. In a cooperative setting, people are more willing to consider different moral viewpoints.

Naturally, meta-ethical judgments may serve more than one psychosocial function. Consider the relation mentioned earlier between increased social consensus concerning a moral truth and increased objectivism concerning that truth – a relation that is, in fact, causal (Beebe, 2014; Goodwin & Darley, 2012). Perhaps, apart from signaling that a judgment is not to be tolerated, judging a moral issue to be objective also serves as a means to signal allegiance to a certain in-group.

The former signaling hypothesis merits special interest, as it ties in with recent hypotheses (DeScioli & Kurzban, 2013; Stanford, 2018) from evolutionary psychology. If an agent signals that a certain moral judgment is not to be tolerated, this might, additionally, be taken as a public signal that she is a reliable cooperation partner with regard to that issue. After all, the agent thereby signals a willingness to further a given moral goal and not to cooperate with agents who do not will it. This signal is at the basis of

Stanford's (2018) recent hypothesis about the evolutionary origins of the phenomenological perception of morality's objectivity. Stanford argues, in part on the basis of the research outlined in Section 2, that the 'externalizing' of moral demands has played a key role in the evolution of human cooperation. Following Joyce (2006), he observes that moral demands have a distinct phenomenology: They often come across as if they are imposed on us externally, and it seems that we have to comply with them regardless of our subjective preferences and desires. How should we account for this feature of our moral experience? Stanford argues that the externalizing of moral motivation was favored in ancestral environments because it allowed prosocial, altruistic, and cooperative agents to quickly and efficiently correlate their interactions, paving the way for a more cooperative form of social life:

> Experiencing moral demands and obligations as externally imposed simultaneously on both ourselves and others ensures that if I myself come to be motivated to conform to a particular norm or standard of behavior that I experience as distinctively moral in character, I *automatically* demand that others conform to it as well, judging them to be less attractive potential partners in social interaction generally if they do not ... From an evolutionary point of view, our characteristic externalization of moral motivation thus represents a mechanism for establishing and maintaining *correlated interaction under plasticity*. (Stanford, 2018, pp. 8–9)

This hypothesis belongs to the realm of evolutionary moral psychology. Like the hypotheses about the psycho-social function of meta-ethical judgments mentioned above, it has no direct implications for the conceptual and metaphysical issues that tend to be at the heart of meta-ethical debate. Stanford (2018, p. 45) himself considered the suggestion that his account might constitute an error theory of morality but concluded that it does not; he contended that people's ordinary views about moral objectivity are not articulate enough to be proven false.

Nonetheless, Stanford's hypothesis, if correct, does add to the psychological challenge to moral realism that I outlined in the previous section. As noted, several meta-ethicists, such as Enoch (2017), support their case for moral realism with appeals to intuition, phenomenology, and the seemingly objectivist features of moral discourse. Allegedly, these features are best explained on the assumption that moral realism is true. Indeed, this suggestion is what fuels the "presumptive support" for realism. Suppose that Enoch's phenomenological observation is apt and that it generalizes: People generally perceive of moral demands as highly objective, as if they were imposed on them externally. Stanford's evolutionary account suggests that the seeming objectivity which is characteristic of the phenomenology of moral discourse can be explained in purely evolutionary terms. While in and of itself this does not discredit moral realism, it does undercut any presumed meta-ethical "plausibility points" that realist

theories might have been thought to have, in virtue of the fact that they vindicate morality's seeming objectivity. Indeed, this seeming objectivity may be little more than the culmination of various cognitive mechanisms, biases, and illusions placed on us by evolution.

As this example suggests, there may be interesting ways in which psychological findings can touch upon traditional meta-ethical theses. From a philosophical point of view, one might expect these realms to be clearly distinct, at least where the tenability of largely metaphysical doctrines – such as moral realism – is concerned. However, since realists commonly defend their views by appealing to empirical or phenomenological claims, there turns out to be some overlap between them. To further explore this area of overlap – where it lies and whether experimental work might touch upon philosophical arguments – is among the tasks for future work in the psychology of meta-ethics.

7. Conclusion

I have shown that the recent wave of experiments about folk moral objectivism has opened up various avenues of research in moral psychology, leading to novel hypotheses about the psychological function of our meta-ethical commitments. Moreover, I have argued that, notwithstanding the mismatch between the experimenter's and the meta-ethicist's typical objects of study, this research holds promise as an important resource of empirical insight that touches on traditional meta-ethical concerns. Whether this promise can be fully delivered is, as of yet, an open question – one which depends *inter alia* on the success of experimenters in correcting for existing methodological flaws. However, its current limitations notwithstanding, this research is beginning to fill an empirical hiatus, and it has already delivered some suggestive findings which may be indirectly relevant for analyses in moral semantics and moral metaphysics. Specifically, the findings provide resources to challenge meta-ethical arguments that rely heavily on the alleged objectivist commitments of ordinary moral discourse, such as Mackie's argument for error theory and Enoch's presumptive argument in support of moral realism. Also, they serve as a reminder to philosophers that concerns about moral objectivity are more encompassing than concerns about moral realism. There may be other notions of objectivity at play in moral discourse, which merit the attention of both psychologists and meta-ethicists.

Notes

1. A disjunctive reasoning ability denotes the capacity to actively unpack alternative possibilities when reasoning, which is tested, for instance, by presenting respondents with the "five blocks" task: There are five blocks in a stack, where the second one

from the top is green and the fourth is not green. Is a green block definitely on top of a non-green block? Participants have to choose between three response options: yes, no, cannot tell. The intuitive response "cannot tell" is incorrect, as becomes clear after explicitly considering what follows if the third block is green and if the third block is not green.

2. As Hannon (2018) rightly stresses, the philosophical significance of findings from experimental philosophy should be evaluated on a case-to-case basis. While the general observations made in this section apply to some studies on folk moral objectivism, it should be kept in mind that they do not apply to all of them; some existing studies are more careful in drawing meta-ethical distinctions than others.

3. I thank the reviewers for pressing me to point out these methodological pitfalls.

4. I thank one of the reviewers for pressing me to point this out.

5. As Sinclair (2012, p. 163) observes, it is plausible that large parts of moral discourse and practice are pragmatically justified. That is, we have good reason to go on engaging in this practice, as it allows us to relate to the world and to others in worthwhile ways. Moreover, the fact that practically all known human societies have developed some sort of moral practice suggests that this practice has clear functional benefits. However, none of this gives us reason to think that our moral practice is also *epistemically* justified – that it allows us to relate to the world and each other in a veridical way.

Disclosure statement

Funding

Notes on contributor

Jeroen Hopster is a postdoctoral researcher, currently working on the project 'Ethics of Evolution and Technology'.

References

Beebe, J. R. (2014). How different kinds of disagreement impact folk metaethical judgments. In J. C. Wright & H. Sarkissian (Eds.), *Advances in experimental moral psychology: Affect, character, and commitments* (pp. 167–187). London: Continuum.

Beebe, J. R. (2015). The empirical study of folk metaethics. *Etyka*, *15*, 11–28.

Beebe, J. R., & Sackris, D. (2016). Moral objectivism across the lifespan. *Philosophical Psychology*, *29*, 912–929.

Bengson, J. (2013). Experimental attacks on intuitions and answers. *Philosophy and Phenomenological Research*, *86*, 495–532.

Brink, D. (1989). *Moral realism and the foundations of ethics*. Cambridge: Cambridge University Press.

DeScioli, P., & Kurzban, R. (2013). A solution to the mysteries of morality. *Psychological Bulletin*, 139, 477–496.

Enoch, D. (2017). Non-naturalistic realism in meta-ethics. In T. McPherson & D. Plunkett (Eds.), *The routledge handbook of metaethics*. New York: Routledge, 29–42.

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 (Eds.), *Proceedings of the 30th annual conference of the cognitive science society* (pp. 1771–1776). Austin: Cognitive Science Society.

Fisher, M., Knobe, J., Strickland, B., & Keil, F. C. (2017). The influence of social interaction on intuitions of objectivity and subjectivity. *Cognitive Science*, 41, 1119–1134.

Fraser, B. (2014). Moral error theories and folk metaethics. *Philosophical Psychology*, 27, 789–806.

Gill, M. B. (2009). Indeterminacy and variability in meta-ethics. *Philosophical Studies*, 145, 215–234.

Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106, 1339–1366.

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, 161–188.

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, 250–256.

Haidt, J., Koller, S. H., & Dias, M. G. (1993). Affect, culture and morality, or is it wrong to eat your dog? *Attitudes and Social Cognition*, 65, 613–628.

Hannon, M. (2018). Intuitions, reflective judgments, and experimental philosophy. *Synthese*, 195, 4147–4168. (online first).

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.

Hopster, J. K. G. (2017). Two accounts of moral objectivity: From attitude-independence to standpoint-invariance. *Ethical Theory and Moral Practice*, 20, 763–780.

Joyce, R. (2006). *The evolution of morality*. Cambridge (MA): MIT Press.

Joyce, R. 2015. Moral anti-realism versus realism: Intuitions. *Stanford Encyclopedia of Philosophy*: https://plato.stanford.edu/entries/moral-anti-realism/moral-realism-intuitive.html

Joyce, R. (2016). *Essays in moral skepticism*. Oxford: Oxford University Press.

Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103, 54–69.

Kauppinen, A. (2007). The rise and fall of experimental philosophy. *Philosophical Explorations*, 10, 95–118.

Khoo, J, & Knobe, J. (2018). Moral disagreement and moral semantics. *Noûs*, 52, 109–143. doi:10.1111/nous.2018.52.issue-1

Kirchin, S. (2003). Ethical phenomenology and metaethics. *Ethical Theory and Moral Practice*, 6, 241–264.

Knobe, J. (2016). Experimental philosophy is cognitive science. In, *In j* (pp. 37–52). Sytsma & W. Buckwalter (Eds.), A companion to experimental philosophy, New York: Wiley-Blackwell.

Laskowski, N., & Finlay, S. (2017). Conceptual analysis in metaethics. In T. McPherson & D. Plunkett (Eds.), *Routledge handbook of metaethics*. New York: Routledge, 536–551.

Loeb, D. (2007). The argument from moral experience. *Ethical Theory and Moral Practice*, *10*, 469–484.

Loeb, D. (2008). Moral incoherentism: How to pull a metaphysical rabbit out of a semantic hat. In W. Sinnott-Armstrong (Ed.), *Moral psychology: The cognitive science of morality* (Vol. 2, pp. 355–385). Cambridge: MIT Press.

Ludwig, K. (2010). Intuitions and relativity. *Philosophical Psychology*, *23*, 427–445.

Mackie, J. L. (1977). *Ethics: Inventing right and wrong*. Harmondsworth: Penguin Books.

Nichols, S. (2004). After objectivity: An empirical study of moral judgment. *Philosophical Psychology*, *17*, 5–28.

Pölzler, T. (2017a). Revisiting folk moral realism. *Review of Philosophy and Psychology*, *8*, 455–476.

Pölzler, T. (2017b). Can the empirical sciences contribute to the realism/antirealism debate?. *Synthese*, *195*, 4907–4930.

Pölzler, T. (2018). *Moral reality and the empirical sciences*. New York: Routledge.

Ruse, M. (1995 [1986]). Evolutionary ethics: A phoenix arisen. In P. Thompson (Ed.), *Issues in evolutionary ethics* (pp. 225–247). SUNY Press: Buffalo, NY.

Sarkassian. (2017). Folk platitudes as the explananda of philosophical metaethics: Are they accurate? And do they help or hinder inquiry? *J Indian Council of Philosophical Research*, *34*, 565–575.

Sarkissian, H. (2016). Aspects of folk morality: Objectivism and relativism. In W. Buckwalter & J. Sytsma (Eds.), *A companion to experimental philosophy* (pp. 212–224). London: Blackwell.

Sarkissian, H., Parks, J., Tien, D., Wright, J. C., & Knobe., J. (2011). Folk moral relativism. *Mind & Language*, *26*, 428–505.

Shafer-Landau, R. (2003). *Moral realism*. Oxford: Oxford University Press.

Sinclair, N. (2012). Moral realism, face-values and presumptions. *Analytic Philosophy*, *53*, 158–179.

Stanford, P. K. (2018). The difference between ice cream and nazis: Moral externalization and the evolution of human cooperation. *Behavioral and Brain Sciences*, 1–57.

Street, S. (2009). In defense of future tuesday indifference: Ideally coherent eccentrics and the contingency of what matters. *Philosophical Issues*, *19*, 273–298.

Theriault, J., Waytz, A., Heiphetz, L., & Young, L. (2017). Examining overlap in behavioral and neural representations of morals, facts, and preferences. *Journal of Experimental Psychology: General*, *146*, 1586–1605.

Turiel, E. (1983). *The development of social knowledge: Morality and convention*. Cambridge, MA: MIT Press.

Weinberg, J. M., Alexander, J., Gonnerman, C., & Reuter, S. (2012). Restrictionism and reflection: Challenge deflected, or simply redirected? *The Monist*, *95*, 200–222.

Wright, J. C., McWhite, C. B., & Grandjean, P. T. (2014). The Cognitive Mechanisms of Intolerance: Do our Meta-Ethical Commitments Matter?. In T. Lombrozo, J. Knobe, & S. Nichols (Eds.), *Oxford studies in experimental philosophy* (Vol. 1). Oxford: Oxford UP.

Wright, J. C., Grandjean, P. T., & McWhite, C. B. (2013). The meta-ethical grounding of our moral beliefs: Evidence for metaethical pluralism. *Philosophical Psychology*, *26*, 336–361.