

Folk Platitudes as the Explananda of Philosophical Metaethics: Are They Accurate? And Do They Help or Hinder Inquiry?

Hagop Sarkissian¹

Received: 30 June 2017 / Revised: 6 July 2017 / Accepted: 7 July 2017 /
Published online: 24 August 2017
© ICPR 2017

Abstract The field of metaethics, the branch of moral philosophy that examines the nature and status of morality, is rich in theoretical diversity. Nonetheless, a majority of professional philosophers embrace a subset of theories that affirm the existence of objective moral facts. I suggest that this may be related to the very method that philosophers use to construct metaethical theories. This method involves analyzing how ordinary people think and argue about morality. Analysis of ordinary moral discourse is meant to reveal common platitudes (or truisms) about the nature of morality itself, including the platitude that morality trades in objective moral facts. But do philosophers investigate ordinary moral discourse in any systematic way? How do they arrive at such platitudes? On what grounds are they justified? In this paper, I critically examine these questions and argue that a) any such platitudes need to be investigated systematically through empirical research and b) philosophers ought to be engaged in this research themselves.

Keywords Metaethics · Experimental philosophy · Folk morality · Folk metaethics · Moral objectivism · Moral platitudes

The field of metaethics, the branch of moral philosophy that examines the nature and status of morality, is rich in theoretical diversity (e.g., Miller 2003). There are philosophers who maintain that morality is fictional, relative, constructed, expressive, subjective, quasi-real, real, just one big error, and many other variations besides. And yet, the majority of professional philosophers adhere to one theoretical position—namely, some form of moral realism. To cite a relevant data point: a 2009 survey conducted by PhilPapers of over 3000 professional philosophers and

✉ Hagop Sarkissian
hagop.sarkissian@baruch.cuny.edu

¹ Baruch College & Graduate Center, The City University of New York, New York, NY, USA

graduate students found that more than half (56%) embraced moral realism (Bourget and Chalmers 2014). Why?

This is a question I cannot fully answer in this paper. However, one reason for realism's popularity stems from the very method that philosophers use to construct metaethical theories. This method involves analyzing how ordinary people conceive of or argue about morality. Analysis of such ordinary usage is meant to reveal common *platitudes* (or truisms) about the nature of morality itself. One central platitude embraced by philosophers maintains that ordinary moral discourse is objectivist in its character, that people tend to believe there are moral facts, and that moral disagreement is concerned with such objective moral facts. The precise characterization of this platitude varies across accounts, but the basic idea is often explained in terms of how ordinary people react to cases of moral disagreement: if ordinary people are not committed to the existence of such real facts about morality, the argument goes, it is difficult to understand why they engage in moral debate at all, and why they proffer arguments and attempt to reason through moral issues.

But how do philosophers arrive at this platitude about ordinary folk morality? On what grounds do they justify claims about ordinary moral discourse and practice? Are the emerging platitudes the proper explananda for a metaethical theory? In what follows, I will discuss this method used in analytic metaethics, and scrutinize the claim that folk moral objectivism emerges as a platitude requiring explanation. A key feature of this method is to collect a number of purported commonplace notions about morality, along with some uses of moral language, so as to uncover conceptual truths that any metaethical theory must accommodate. However, as I'll argue, while this method claims to examine ordinary usage, it fails to do so, thus risking enshrining *philosophers'* beliefs about morality as *ordinary* beliefs. I'll then turn to attempts by experimental philosophers to redress this issue by undertaking the piecemeal and messy work of investigating, in systematic ways, the nature of folk morality. The emerging picture is that ordinary folk morality is far less uniform (and far more interesting) than philosophers have traditionally maintained. By expanding the explananda of philosophical metaethics, experimental philosophers are poised to contribute to it in fundamental ways.

What are Platitudes?

A widespread commitment in analytic metaethics involves studying platitudes (or commonplace beliefs) about morality. These platitudes can consist of explicit, well-worn truisms (e.g., “if an action is morally wrong, one ought not do it”) or implicit commitments that emerge from an analysis of the ways in which ordinary people are disposed to judge, act upon, and evaluate moral concepts. The assumption is that ordinary folk have competence with moral concepts (a kind of “know-how”), so understanding morality requires examining ordinary folk usage. According to an influential account by Michael Smith,

in acquiring a concept C we come to acquire a whole set of inferential and judgemental dispositions connecting facts expressed in terms of the concept C

with facts of other kinds. A statement of all of these various dispositions constitutes a set of platitudes surrounding C. And an analysis of a concept is then best thought of as an attempt to articulate all and only these platitudes. (Smith 1994, pp. 37–38)

Platitudes about morality, then, are commonplace truths that emerge from an analysis of the numerous facts surrounding the ordinary uses of moral concepts.

Smith goes on to claim that there are two platitudes about morality that metaethicists must, above all others, accommodate in their theories. These platitudes “are manifest in ordinary moral practice as it is engaged in by ordinary folk,” and the “philosopher’s task,” we are told, “is to make sense of a practice having these features” (Smith 1994, p. 5). One platitude maintains that when persons make moral judgments they are motivated to act on them.¹ The other, and the one I’ll focus on, maintains that persons engaging in moral discourse assume that moral facts exist—that there are objectively correct answers to moral questions.

How does Smith arrive at these particular platitudes? In spite of the central role they play in his theoretical account, there is little to directly answer this question. This is not to say that Smith fails to provide well-argued considerations in favor of these platitudes. However, their “ordinariness” is largely established in passing and through stipulation. For example, we are simply told that

it is a distinctive feature of engaging in moral practice that the participants are concerned to get the answers to moral questions right. And this concern itself seems to force certain meta-ethical conclusions... [We] seem to think moral questions have correct answers; that the correct answers are made correct by objective moral facts; that moral facts are wholly determined by circumstances; and that, by engaging in moral conversation and argument, we can discover what these objective moral facts determined by the circumstances are. (see, e.g., Smith 1994, pp. 5–6)

The use of the pronoun “we” is meant to signify that these notions are *ordinary*, in the sense of being widespread and representative of how most users of moral concepts think (or would think, if only they engaged in such thoughts). However, apart from this, there is little else speaking to the actual process of settling on the veracity of the platitudes themselves.

Smith is by no means alone in thinking of metaethics as the study of platitudes that emerge from an examination of ordinary moral practice. It is, in fact, a default method in the field. Consider, for example, how Terence Cuneo discusses the explananda of metaethics.

When we talk of a commonsensical conception of a putative moral feature, we mean to pick out a way of thinking about that feature that incorporates various platitudes concerning the content and objectivity of that feature that are sufficiently entrenched in the beliefs and practices of ordinary folk that moral realists and antirealists ordinarily attempt to accommodate them in their respective theories. (Cuneo 2007, pp. 36–39)

¹ For critical discussion of this platitude, see Gill (2009).

Here, what is commonsensical is characterized as incorporating platitudes entrenched among ordinary folk usage. And, as evidence of their ordinariness, Cuneo appeals to the fact that various professional philosophers agree that they are ordinary (and Cuneo is certainly correct about that).

Cuneo also has two platitudes that play a central role in his account. These concern the *content* and the *authority* of morality (Cuneo 2007, pp. 36–39). Regarding the content, he claims that (a) only actions that promote or frustrate human flourishing merit the terms “moral” and “immoral” (respectively), and (b) there are limits to what counts as authentic human flourishing. Regarding the authority, he claims that (a) moral facts provide reasons for action, and (b) they are *authoritative*, inescapably governing a person’s conduct. Once again, readers are provided with no further data to support the claim to ordinariness; instead, it seems taken for granted that readers (invariably, other professional philosophers) will find them unobjectionable. And they routinely do.

Examples like these abound in the literature (e.g., Darwall 1998, p. 25; Jackson 2000, p. 137; Joyce 2011, p. 529). And some of the platitudes may turn out to be true. But, as I’ll argue, whether or not they are true cannot be established from philosophers’ armchairs.

The Trouble with Platitudes

How do philosophers verify or establish that they’ve articulated the correct platitudes? How do they come to know these ordinary truths? As argued above, there is little here beyond stipulations that philosophers themselves find compelling. As Michael Gill notes, systematic, empirical investigation of ordinary folk usage by philosophers is, historically, exceedingly rare. Philosophers have “typically presented some examples of ordinary discourse. But they didn’t gather data in any kind of comprehensive and systematic way. A handful of illustrative cases were taken to be a sufficient starting point” (2009, p. 217). Of course, it can be maintained that philosophers are ordinary folk too, differing from others only in the rigor and systematicity with which they approach the moral domain (Smith 1994, p. 1). But just how ordinary are philosophers? And how can readers evaluate claims to ordinariness?

Cuneo is admirably forthright about this issue. He acknowledges that talk of what’s *ordinary* or *commonsensical* “is bound to arouse suspicion” (Cuneo 2007, p. 32).

What, after all, could we mean by the claim that there is an ordinary or commonsensical conception of a kind? And what could be meant by the claim that there is an ordinary or commonsensical conception of specifically moral features?... (Cuneo 2007, p. 32).

Suppose some philosophers reject the platitudes that others find compelling? Suppose some characterize ordinary discourse in ways that diverge from their peers? What to do?

Philosophers rarely tackle this question head on, as there seems to be no good way, given the methods they embrace, of resolving such cases of disagreement. As Richard Joyce writes, philosophers “have no settled views on how to adjudicate or even conceptualize such a debate. If one person asserts that something is a non-negotiable feature of some concept and another person denies this, where should they take their dispute?” (Joyce 2006, p. 142). We are left with little more than assertion and counter-assertion, supported by overall considerations that seem compelling to some but not to others.²

Gill is, once again, instructive. He notes that proponents on opposite sides of such debates have tried to show that their views make better sense of folk usage *as a whole*, even while admitting to the existence of troubling counterexamples. Any such counterexamples can, after all, be explained away in some form or other. Besides, they can uncover troubling counterexamples for their opponents. But their opponents will then simply explain these counterexamples away in turn. The dialectic has been such that “each side contended that while their position had some awkward elements, the opposing position had *even more* awkward elements, and that this was a powerful reason for thinking that their own position was true” (2009, pp. 217–218).

For his own part, Cuneo’s responds to imagined skeptics who question the veracity (or ordinariness) of his proffered platitudes by claiming that the platitudes themselves need not be widespread after all in order to serve as the proper explananda of a metaethical theory. They need only be representative of a plausible subpopulation of persons.

But suppose that I am wrong about [these platitudes]. Then what I say can be easily recast as saying that our ‘commonsensical’ conception of moral features is constituted by a certain stock of platitudes that are sufficiently entrenched in the beliefs and practices of *some subset of the ordinary folk*, and that philosophers have typically found *these* platitudes *plausible enough* to want to try to preserve them in their moral theories. (Cuneo 2007, pp. 32–33, emphasis added)

Of course, this retreat to a subset of the ordinary folk only raises further questions. What justifies relying on this subpopulation? How is it picked out, and why should it be considered representative? Is talk of what’s commonsensical or ordinary just a way for philosophers to claim broader legitimacy for their idiosyncratic notions? After all, it seems as though the relevant criterion for evaluating talk of what is ordinary is agreement by some other professional

² A parallel worry is raised in the *Zhuangzi*, a compendium of early Daoist thought, concerning the adjudication of competing claims among philosophers about the substance of morality. “Once you and I have started arguing... Is one of us right and the other one wrong? Or are both of us right and both of us wrong?... Whom shall we get to set us right? Shall we get someone who agrees with you to set us right? But if they already agree with you how can they set us right? Shall we get someone who agrees with me to set us right? But if they already agree with me, how can they set us right? Shall we get someone who disagrees with both of us to set us right? But if they already disagree with both of us, how can they set us right? Shall we get someone who agrees with both of us to set us right? But if they already agree with both of us, how can they set us right? If you and I and they all can’t understand each other, should we wait for someone else?” (Ivanhoe and Van Norden 2005, p. 223).

philosophers.³ If it turns out that philosophers have uncovered the wrong platitudes, the right thing to do (we are told) is not to *revise* them but instead to *ignore ordinary usage* and press ahead with the platitudes that are of interest to philosophers; after all, they are representative of *some group or other*, and at any rate *philosophers* (ordinary folk!) find them plausible enough.

In recent years, though, there's been a growing recognition that claims about folk morality ought to be addressed not from one's armchair but rather through systematic empirical investigation. Gill, as might be expected, calls for such empirical research (2009, p. 232). But so do others. Richard Joyce, for example, agrees that there is a pressing need for such research, as the resulting data would ultimately trump anything that philosophers might aver from the armchair. Thus,

what needs to be examined in detail are the conventions surrounding moral discourse in our natural language. Quiet and careful introspection is not a reliable source of knowledge on such matters; we really need to observe language being used across a wide range of everyday settings. (Joyce 2006, p. 138)

There are many ways to approach the question, and I don't want to deny that sitting in one's armchair and speculating may be among them (such activity often casts up promising ideas)—but *the ultimate arbiter must be the body of a posteriori data issuing from such disciplines as psychology, anthropology, and experimental economics*. (Ibid, pp. 143–144, emphasis added)

Research on folk morality by social and behavioral scientists is indeed of great value when pursuing questions about the nature and status of morality. But it is also true that metaethicists have interests, commitments, and hypotheses that may not be obvious targets of empirical investigation by researchers in these other disciplines. So philosophers ought to be involved in generating the hypotheses and collecting and analyzing the data.

Experimental Metaethics

Consider, again, the central claim (noted above) that ordinary moral discourse and practice seems committed to objectivism of some sort. On first pass, this seems like a plausible claim. For example, in ordinary discourse, assertions about rightness or wrongness are not normally qualified in any way. People do not ordinarily say that “racial discrimination is wrong insofar as one values human dignity or equality.” Instead, the language is baldly absolutist: “racial discrimination is wrong!” People also engage in protracted debates about moral issues and seem interested in arguing for the correctness of their views. What does this tell us about the nature of folk morality? The dominant interpretation is that these aspects of folk morality (and related ones) reveal a folk commitment to the existence of objective moral facts.

³ Compare Doris and Stich (2005, pp. 124–125) on claims about conceptual truths concerning motivational internalism.

Only such a commitment, the argument goes, can explain these and related features of ordinary discourse. Is this interpretation correct?

Some experimental forays by philosophers and psychologists alike have begun to explore whether or not this is the case, and many have followed procedures and suggestions found in the philosophical literature outlined above.⁴ Consider, for example, the following passage from Smith.

... it is a platitude that our moral judgements at least purport to be objective.... Thus if A says 'It is right to ϕ in circumstances C' and B says 'It is not right to ϕ in circumstances C' then we take it that A and B disagree; that at most one of their judgements is true... (Smith 1994, 86)

One way to test this platitude is to investigate how ordinary folk react to just such cases of disagreement. Do they agree that at most one of the two judgments can be true? Or do they disagree, and find the conflicting judgments equally acceptable?

This basic paradigm is readily amenable to empirical investigation, and has been used in several different permutations across a number of studies. The results suggest a complicated picture that should militate against any attempt to depict folk morality as having some definite, invariant, objectivist shape. For even while some early studies by Nichols (2004) and Goodwin and Darley (2008), for example, showed that a majority of individuals tend to think moral disagreements admit of only a single correct answer, these studies routinely found sizable minorities of participants (approximately 30–40%) who disagree, and think the disagreeing judgments could both be correct. Goodin and Darley also uncovered wide disparities in people's tendencies to embrace objectivism across specific moral issues. When participants were presented with cases of disagreement on the permissibility of discrimination based on race, cheating on a lifeguard qualification exam, opening gunfire on a crowded street, or robbing a bank to pay for an expensive holiday, majorities (54–68%) said one of the disagreeing parties had to be wrong. Yet when it came to other topics such as the use of stem cells in research, the permissibility of abortion during the first trimester, or the permissibility of euthanasia for the terminally ill, very few seemed to think so (only 8–10%).

Some researchers found these results baffling. One attempt to understand these large differences *within* the moral domain focused on experimental design (Wright et al. 2013). The studies above adopt a *third-person* approach to design, where the experimenters themselves decide what constitutes moral disagreements (as opposed to, say, factual or aesthetic disagreements) and present these disagreements to participants to adjudicate. But, of course, some participants might not agree with the experimenters' classifications. Take, for example, donating to charity. Some might think this a paradigmatically moral act, whereas others might think it a matter of personal choice, or even a self-interested act (if motivated by considerations of self-esteem or tax breaks). If that's so, these individuals might judge that there is no objectively correct answer about whether it's right to donate to charity because they don't take this to be a *moral* act at all.

⁴ For a more detailed overview of the existing literature, see Sarkissian (2016).

To overcome this potential limitation, Jennifer Wright et al. (2013) used a *first-person* approach and asked participants to classify the topics of disagreement themselves.⁵ They were given four choices for classification: factual, moral, conventional, or taste. Their hypothesis was that participants would show much higher and more consistent levels of objectivism for the items they themselves deemed to be moral. This prediction failed. The researchers found no difference between average levels of objectivism for the moral items as compared to the conventional items, with considerably variability within the moral domain. The evidence, they argue, supports pluralism about the moral domain among ordinary folk; for some topics they deem moral, they believe there are objectively correct answers. For others, they demur.

What explains these differences in objectivity judgments *within* the moral domain? Some have found evidence that folk use their judgments about *consensus* on the particular moral issue as a basis for their judgments about their objectivity (Ayars and Nichols in prep; Beebe 2014; Goodwin and Darley 2012; Nichols forthcoming), and this might explain the differences in objectivity scores noted above; after all, topics such as stem cell research, abortion, and euthanasia are paradigmatic cases of moral *controversies* (so much so that they are mainstays of applied ethics courses). And studies on children show that this tendency to treat consensus as a proxy for objectivity emerges early on in development, persisting through adulthood (Heiphetz and Young 2016). Of course, if such an account is correct, philosophers would be mistaken to infer that a commitment to objectivism underlies ordinary moral practice. Other studies point to individual differences to explain why individuals diverge in their judgments about moral objectivity, including aspects of their personality (Feltz and Cokely 2008; Goodwin and Darley 2010), or age (Beebe and Sackris 2016).

Of course, as noted in the previous section, not all philosophers think the objectivity platitude is central to ordinary discourse. Consider, for example, David Wong's (2006) characterization of "moral ambivalence," a feeling that arises when one disagrees with someone on a moral issue, even though one finds this person to be knowledgeable, reasonable, and morally competent. In such cases, when we fail to persuade others whom we find neither obstinate nor uninformed, when we gain new insight into an issue about which we previously felt certainty, Wong claims that it is normal to step back and question the superiority or unique correctness of one's own judgment.

Understanding other moral codes and the ways of life in which they are embedded is not to see them as alien and incomprehensible but in some respects familiar and in other respects constituting a challenge to our own codes and ways of life. Since we ourselves are complex and ambivalent moral beings, we are able to see that at least some other codes and ways of life may just as reasonably be adopted by decent and informed human beings as our own. (Wong 2006, 20)

⁵ The 'third-person' approach is dominant within experimental psychology, though there are strengths and weaknesses to both approaches. See Meindl and Graham (2014) for helpful discussion.

If this characterization of ordinary practice has merit, then we ought to be able to see it revealed in experimental settings.

Indeed, there are now a number of studies that support this phenomenon of moral ambivalence. Sarkissian et al. (2011), for example, note that most cases of disagreement portrayed in the experimental studies discussed above depict the two disagreeing parties as either explicitly or implicitly *peers* from the same (and sometimes very local) cultural group, and argue that it would be natural for folk to think there could be correct answers to moral disagreements relative to such specific contexts. However, folk might react differently to cases of disagreement between individuals of differing cultural backgrounds and value orientations, revealing an ambivalence not unlike the kind described by Wong above. In their studies, they used standard cases of disagreement between two people described as peers, but also included cases in which the disagreement was between someone described as a peer and a person of a foreign culture with different value orientations—either a member of an isolated tribe in the Amazonian rainforest or an extraterrestrial being with a different psychological profile. While in the first type of case a majority of participants thought there could only be a single correct answer, they tended to drop this attitude as the disagreeing individuals were described as dissimilar and foreign. If folk were indeed objectivists in the face of disagreement, it would not matter what values or cultural perspective individuals brought to bear on the topic of discussion; there would still be only one correct answer. Thus, Sarkissian et al. conclude that ordinary folk might tacitly believe moral judgments are qualified in certain ways—for example, being relative to a cultural framework. They found this basic effect across a number of experimental variations, and the finding has been replicated in other research, together with an account of the semantics of disagreement at play (Khoo and Knobe 2016).

Indeed, this same basic phenomenon has been uncovered using a completely different experimental paradigm. Fisher et al. (2016) changed the context of disagreement by prompting participants to adopt particular *mindsets* before engaging in debate themselves. Some participants were encouraged to take an “argue-to-win” mindset (promoting a goal of scoring points over opponents), whereas others were encouraged to take an “argue-to-learn” mindset (promoting discovery of new information). Across a series of experiments, they found consistent evidence that when prompted to learn and discover new information, participants were more likely to think that a moral issue admitted of more than one correct answer. By contrast, those in the competitive condition did not differ from those in a control condition (where no particular mindset was suggested) in affirming a more objective viewpoint.

More work needs to be done to fine-tune this picture. But if these data are right, they suggest that philosophers may take either combative instances of moral discourse, or moral discourse within cultures, to be paradigmatic. If this is so, it is troubling for at least two reasons. First, and most obviously, philosophers’ accounts of ordinary usage would fail to capture how folk think about the nature of morality when they are engaged in moral discussion with those they find reasonable, knowledgeable, or competent, even while occupying very different value orientations. Yet many may find such forms of discussion to be the ideal form of moral

dialogue. Second, and more generally, this research shows folk morality to be far more complex and nuanced than philosophers' accounts tend to portray it.

Conclusion

Metaethicists have been committed to working out the platitudes lying behind ordinary moral discourse and practice, and have tended to assume that there is something determinate and invariant there to study (Gill 2009), something readily amenable to theorizing and capable of being summarized in a small number of truisms. However, if anything like the picture just sketched is true, this basic method stands to be revised. Experimental philosophers point to a way forward by carrying out the kind of messy, systematic research about ordinary moral usage that seems necessary to come to a more accurate and detailed account of the explananda of metaethical theories.

Indeed, one effect of having canvassed the experimental literature may be to come to a more vivid awareness of the insular way that analytic metaethics can proceed. The standard method deployed by philosophers risks mistaking their own, professionally sanctioned platitudes *qua* philosophers as those of ordinary folk, when they might instead be peculiar or idiosyncratic. The worry is all the more pressing when one considers that most philosophers in the Anglo-American world are a self-selected group of white males who read and respond to one another's work on the topic, with comparatively little interest in what others in disciplines such as anthropology or psychology might have to say about it, let alone what ordinary users of moral language think.

This does not, of course, necessarily impugn philosophers' beliefs about ordinary platitudes; philosophers *are* products of the broader cultures in which they develop, and it would be surprising if their notions were *wildly* off the mark. Nonetheless, we might wonder what effects professionalization and clustering around a few core claims might do for the nature of metaethical inquiry. At the very least, we should examine whether much of ordinary moral life is simply left out of the picture, and whether philosophers analyzing folk morality from the armchair might instead be generating idiosyncratic platitudes that seem compelling to them but are of questionable relevance to actual moral discourse and practice.

References

- Ayars, A., & Nichols, S. (in prep). Rational learners and metaethics: Universalism, relativism, and evidence from consensus.
- Beebe, J. R. (2014). How different kinds of disagreement impact folk metaethical judgments. In H. Sarkissian & J. C. Wright (Eds.), *Advances in experimental moral psychology* (pp. 167–187). London: Bloomsbury Publishing.
- Beebe, J. R., & Sackris, D. (2016). Moral objectivism across the lifespan. *Philosophical Psychology*, 29(6), 912–929.

- Bourget, D., & Chalmers, D. J. (2014). What do philosophers believe? *Philosophical Studies*, 170(3), 465–500.
- Cuneo, T. (2007). *The normative web: An argument for moral realism*. New York: Clarendon Press.
- Darwall, S. L. (1998). *Philosophical ethics*. Boulder: Westview Press.
- Doris, J., & Stich, S. (2005). As a matter of fact: Empirical perspectives on ethics. In F. Jackson & M. Smith (Eds.), *The oxford handbook of contemporary philosophy*. New York: 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 (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. (2016). The influence of social interaction on intuitions of objectivity and subjectivity. *Cognitive Science*. doi:10.1111/cogs.12380.
- Gill, M. B. (2009). Indeterminacy and variability in meta-ethics. *Philosophical Studies*, 145(2), 215–234.
- Goodwin, G. P., & Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106(3), 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(2), 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(1), 250–256.
- Heiphetz, L., & Young, L. L. (2016). Can only one person be right? The development of objectivism and social preferences regarding widely shared and controversial moral beliefs. *Cognition*. doi:10.1016/j.cognition.2016.05.014.
- Ivanhoe, P. J., & Van Norden, B. W. (2005). *Readings in classical chinese philosophy* (2nd ed.). Indianapolis: Hackett Publishing Company.
- Jackson, F. (2000). *From metaphysics to ethics: A defence of conceptual analysis*. New York: Oxford University Press.
- Joyce, R. (2006). Metaethics and the empirical sciences. *Philosophical Explorations: An International Journal for the Philosophy of Mind and Action*, 9(1), 133–148.
- Joyce, R. (2011). The error in “The error in the error theory”. *Australasian Journal of Philosophy*, 89(3), 519–534.
- Khoo, J., & Knobe, J. (2016). Moral disagreement and moral semantics. *Noûs*. doi:10.1111/nous.12151.
- Meindl, P., & Graham, J. (2014). Know thy participant: The trouble with nomothetic assumptions in moral psychology. In H. Sarkissian & J. C. Wright (Eds.), *Advances in experimental moral psychology* (pp. 233–252). London: Bloomsbury Publishing.
- Miller, A. (2003). *An introduction to contemporary metaethics*. Cambridge: Polity Press.
- Nichols, S. (2004). After objectivity: An empirical study of moral judgment. *Philosophical Psychology*, 17(1), 3–26.
- Nichols, S. (forthcoming). Debunking and vindicating in moral psychology. In A. Goldman & B. McLaughlin (Eds.), *Metaphysics and cognitive science*. Oxford: Oxford University Press.
- 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., Park, J., Tien, D., Wright, J. C., & Knobe, J. (2011). Folk moral relativism. *Mind and Language*, 26(4), 482–505.
- Smith, M. (1994). *The moral problem*. Oxford: Blackwell Publishing.
- Wong, D. B. (2006). *Natural moralities: A defense of pluralistic relativism*. New York: Oxford University Press.
- Wright, J. C., Grandjean, P. T., & McWhite, C. B. (2013). The meta-ethical grounding of our moral beliefs: Evidence for meta-ethical pluralism. *Philosophical Psychology*, 26(3), 336–361.