



Third factor explanations and disagreement in metaethics

Michael Klenk¹ 

Received: 27 November 2017 / Accepted: 7 July 2018 / Published online: 12 July 2018
© The Author(s) 2018

Abstract

Several moral objectivists try to explain the reliability of moral beliefs by appealing to a third factor, a substantive moral claim that explains, first, why we have the moral beliefs that we have and, second, why these beliefs are true. Folke Tersman has recently suggested that moral disagreement constrains the epistemic legitimacy of third-factor explanations. Apart from constraining third-factor explanations, Tersman's challenge could support the view that the epistemic significance of debunking explanations depends on the epistemic significance of disagreement. This paper aims to show that disagreement does not constrain the epistemic legitimacy of third-factor explanations in metaethics, and it suggests a way forward in addressing the view that debunking depends on disagreement. First, Tersman's constraints are impossible to violate, given the assumption that the justification relation exhibits monotonicity. Second, some disagreements are irrelevant, given that they cannot be about beliefs whose reliability the objectivist seeks to defend. Third, actual disagreement about moral beliefs is implausible, given recent ethnographic findings. In light of this discussion, the paper shows that the prospects of the disagreement view depend on which moral beliefs objectivists need to defend and the criteria we use to assess epistemically relevant moral disagreement.

Keywords Evolutionary debunking arguments · Third-factor explanations · Disagreement · Higher-order evidence

I am grateful for Folke Tersman's, Herman Philipse's, and Hein Duijf's constructive comments on earlier versions of this paper. Many thanks also to the audience at the St. Andrews CEPPA Graduate Conference 2018 and to two anonymous reviewers.

✉ Michael Klenk
m.b.o.t.klenk@uu.nl

¹ Department of Philosophy and Religious Studies, Utrecht University, Jankserkhof 13a, 3512BL Utrecht, The Netherlands

1 Introduction

The Darwinist view of morality claims that the human propensity to make certain moral evaluations, such as the widespread judgement that parents have most reason to support their children, can be explained in evolutionary terms. *Evolutionary debunkers of morality* have taken up the Darwinist's claim to draw metaethical conclusions. Most prominently, some debunkers argue that the Darwinist view of morality implies that *moral objectivism*, the view that there are mind-independent, non-natural moral truths that we can know about, would commit us to a form of moral scepticism. Even assuming that our objectivist moral beliefs are prima facie justified, none of our objectivist moral beliefs would remain justified after we had learned about the Darwinist view of morality (Street 2006, 2016; Joyce 2006, 2013, 2016).¹

A commonly cited construal of the debunker's challenge is what I call the *reliability view*: the challenge posed by the Darwinist view of morality is *explaining the reliability* of objective moral beliefs. It is often assumed, although this assumption is not without problems, that 'explaining the reliability' of some set of target beliefs at least requires one to explain why those beliefs are true more often than chance would predict (Street 2016: p. 305; cf. Schechter 2018).² If an adequate explanation proves to be impossible in principle, the challenge goes, then any prima facie justification that our objective moral beliefs might have is undercut.³ Several objectivists regard this challenge to be their most arduous test (Shafer-Landau 2012; Enoch 2010; Wielenberg 2014). In effect, the focal point of the current debunking debate, viewed from the perspective of the reliability view, is whether objectivists can adequately explain the reliability of objectivist moral beliefs (Vavova 2015: p. 111). It is widely assumed that so-called *third-factor explanations* are the most promising candidate explanations available to moral objectivists (Enoch 2010; Behrens 2013). Third-factor accounts appeal to "bridge principles" that "posit a relation between the facts in virtue of which our moral beliefs are true and the (non-moral) facts to which the evolutionary account

¹ The structure and conclusion of both Street's and Joyce's argument are subject to much debate; cf. Bogardus (2016) and Berker (2014). For example, Street's (2006) main aim is to show that moral objectivism should be rejected. As such, her argument may be better interpreted as suggesting that we have no reason to think that any objective moral belief is justified in light of a Darwinist view of morality. This claim, in turn, is plausibly seen as an intermediary step towards drawing the conclusion that all our objectivist moral beliefs are undercut in light of the Darwinist view of morality. I would like to thank an anonymous reviewer for prompting me to clarify this point. In what follows, I will assume that there is a legitimate way for a Darwinist view of morality to undercut all objectivist moral beliefs (though see Klenk 2017b for a critical discussion).

² The correct interpretation of 'explaining the reliability' is an interesting issue in itself that I cannot fully address in this paper. Common alternative interpretations to the one introduced above invoke modal conditions such as sensitivity or safety, while Tersman (2017) proposes the view that a belief is reliable to the extent that possessing it gives us reason to think its content is true. The correct interpretation of 'explaining the reliability' is an issue that can be set aside in this paper insofar as it may affect whether third-factor accounts offer any help against debunking arguments in the first place, while I focus on whether disagreement constrains third-factor accounts. Since Tersman regards third-factor accounts to be potentially viable responses to explaining reliability in his sense, too, the minimal characterisation invoked above should suffice for present purposes. Thanks are due to an anonymous reviewer for raising this point.

³ See Pollock and Cruz (1999: p. 195ff) for a discussion of undercutting defeat.

attributes them” (Tersman 2017: p. 765).⁴ If third-factor explanations work, then moral objectivists can pass their most arduous test.

This paper addresses a recent innovative proposal by Folke Tersman about how to determine the epistemic legitimacy of third-factor explanations (Tersman 2017). Tersman argues that the epistemic significance of disagreement constrains third-factor explanations. More precisely, objectivists rely on a substantive moral claim to get third-factor explanations off the ground, and radical moral disagreement might undermine the objectivist’s *prima facie* justification for maintaining the substantive moral claim. A disagreement is ‘radical’ in Tersman’s sense if it cannot be explained away as a cognitive shortcoming (e.g. fallacious reasoning, lack of imagination) or a translation error (Tersman 2017: p. 769).⁵ Tersman’s proposal threatens objectivists with a second-order problem: there might be no legitimate moral belief to rely on to get a third-factor explanation started. If objectivists cannot meet Tersman’s constraints, they lose their most promising answer to the evolutionary debunking challenge, viewed from the perspective of the reliability view. Moreover, as I demonstrate below, Tersman’s proposal might lend support to an alternative to the reliability view, according to which the success of evolutionary debunking arguments depends on the epistemic significance of disagreement.

The paper aims to show that constraints that have to do with disagreement do not pose a problem for objectivists for three reasons: given plausible assumptions about belief formation, one of Tersman’s constraints is *impossible* to violate. Another constraint proves *irrelevant* for the objectivist’s cause, and even if both constraints were acceptable, there is good empirical reason to conclude that moral disagreement that would violate Tersman’s constraints would be *implausible*. The paper thereby vindicates the legitimacy of third-factor accounts as explanations in metaethics as far as worries about disagreement are concerned. Moreover, by relating Tersman’s challenge to the broader view that debunking depends on disagreement, it gives some indication about the strength and the prospects of the view that debunking depends on disagreement.

Section 3 introduces third-factor explanations in relevant detail, and Sect. 4 reconstructs Tersman’s account. Section 5 makes a new connection between Tersman’s challenge and how it might lend support to the aforementioned alternative to the reliability view. Section 6 contains my argument against Tersman’s account, and in Sect. 7, I consider a reply on behalf of Tersman that puts pressure on my objection by appealing to the epistemic significance of higher-order evidence. I show that the rejoinder fails and conclude that worries about disagreement do not constrain third-factor accounts. In Sect. 7, I evaluate the relevance of my findings regarding the view that disagreement is at the heart of evolutionary debunking arguments. Rejecting Tersman’s constraints on third-factor accounts means bad news for proponents of the disagreement view, but, insofar as third-factor explanations are *prima facie* legitimate, the conclusion of this paper should be a boon for moral objectivists.

⁴ ‘Belief’ is ambiguous here: ‘true belief’ really means that the propositional content of a belief is true, whereas evolutionary explanations show why we have dispositions to hold true certain propositions.

⁵ Radical disagreement in Tersman’s sense is thus understood as rational and faultless; cf. Kölbel (2004).




Fig. 1 Structure of a third-factor explanation (adapted from Berker (2014: 230))

2 Third-factor explanations

Let M be the moral fact that you have conclusive reason to bring about a state of affairs ϕ . Let N be the non-moral features of ϕ . Suppose that evolutionary considerations explain why actions that brought about states of affairs with the non-moral property N were adaptive, that is, they increased the actor's relative fitness. A third-factor f then explains the relation between the moral fact, M , and the evolutionary explanation of our tendency to value acts with features N . f is a bridge principle of the form 'N is M' that links the natural state of affairs with a moral value. For example, Enoch's proposed third-factor is that "survival is at least somewhat good" (Enoch 2010).⁶ We can assume that organisms that tend to prefer things that aid survival prosper, whereas organisms that don't prefer things that aid survival fail in evolutionary terms. An evolutionary explanation along these lines at least partly explains why humans believe that survival is good. The bridge principle, then, explains why such beliefs are also true. Thus, as Fig. 1 illustrates, given an evolutionary explanation of why our ancestors favoured N , the third factor thereby also explains why M is instantiated:

If third-factor accounts work, then there is a "gap in the debunker's evolutionary argument against moral objectivism; a gap through which the non-sceptic might try to sneak out" (Tersman 2017: p. 767). Crucially, third-factor explanations rely on assumptions about the truth of the bridge principle. As we have seen, for example, Enoch assumes that "survival (N) is at least somewhat good (N)".⁷

Not surprisingly, the main discussion point about third-factor explanations themselves is whether assuming the truth of a substantive moral claim, that is, a normative claim rather than a claim about why humans adopt certain moral norms, in response to the debunking challenge is legitimate. Critics claim that it begs the question (Street 2016), but they are met with a *tu quoque* by their opponents: reliability in any domain of belief formation, understood as truth-conduciveness, can only be explained by assuming the truth of some beliefs of that domain (Berker 2014; Vavova 2014; Tersman 2017: p. 766). Hence, it would seem that debunkers rely on the very same assumptions that

⁶ See Behrends (2013) for a detailed discussion of Enoch's account. For related third-factor accounts see Brosnan (2011), Skarsaune (2011), Talbott (2015), Artiga (2015), Street (2008) and Copp (2008).

⁷ Note that proponents of third-factor accounts do *not* aim to *justify* moral norms but to *explain their reliability*. This is as it should be: the challenge raised by the evolutionary debunker of morality is that there is *no* explanation of the reliability of moral beliefs; proponents of third-factor accounts try to provide such an explanation.

they deem question-begging if made by objectivists. Tersman, for that matter, accepts the pro tanto legitimacy of third-factor accounts:

[T]he non-sceptic has a greater room for manoeuvring than one might initially think ... [A] non-sceptic can hope to accommodate the evolutionary account by invoking a “bridge principle”. (Tersman 2017: p. 767)

That is, there is a ‘bridge’ between the evolutionary explanation of why humans endorse certain moral norms and the truth of these norms as conceived by moral objectivists. To discuss Tersman’s challenge, I will assume for the sake of argument that the substantive moral assumption at the heart of a third-factor explanation is *prima facie* legitimate. For reasons of space, we must now sidestep the deep epistemological debate about the legitimacy of using beliefs produced by a faculty in explaining the reliability of that faculty.⁸

3 Tersman’s challenge: disagreement constrains third-factor explanations

Tersman’s main contention is that “not just any bridge principle [that] generates the conclusion that the target beliefs are reliable does the trick for the non-sceptic” (Tersman 2017: p. 767). Hence he proposes three constraints for evaluating the “plausibility” of possible bridge principles on behalf of non-sceptics, of which two constraints are relevant for present purposes.⁹

First, Tersman demands that *an acceptable bridge principle must be supported (i.e. we have reason to accept it)* by the “beliefs whose reliability is to be established [by the objectivist]” (Tersman 2017: p. 767). Objectivists need some (epistemic) justification to invoke a particular bridge principle, and the required justification must come from the beliefs whose reliability the objectivist wants to defend (Tersman 2017: p. 768). In other words, there must be a support relation between the contents of the beliefs the objectivist wants to defend and the bridge principle employed in the third-factor account. Moreover, argues Tersman, the content of the beliefs whose reliability is to be established must be “sufficiently varied and rich” and “cannot merely consist of uncontroversial platitudes” because otherwise the third-factor account would be “underdetermined” (Tersman 2017: p. 768). Thus, objectivists must feel the pull of two conflicting demands: on the one hand, endorsing a set of moral beliefs that are

⁸ As an anonymous reviewer notes, there may be an interesting relation between third-factor explanations and Clarke-Doane’s (2016) response to the debunking challenge, which focuses on the modal security of (basic) moral beliefs. The latter may be constrained by the extent of actual moral disagreement, too, as Clarke-Doane himself suggests (2016: 29). Addressing the relevance of Tersman’s constraints for Clarke-Doane’s account, however, would require a detailed comparison of third-factor explanations with Clarke-Doane’s response, which is beyond the scope of this paper. One apparent disanalogy, however, is that the former aims to show why a large swathe of our current moral beliefs is true, whereas the latter only aims to show that that our moral beliefs are modally secure, *assuming* that they are true.

⁹ Tersman also demands that the objectivist’s third-factor account must not be self-defeating. This constraint is entailed by the constraint that an acceptable bridge principle must be supported by the beliefs whose reliability is to be established insofar as the set of beliefs whose reliability is to be established can only provide sufficient support for the bridge principle if it is not in tension with the bridge principle.

sufficiently varied and rich raises the probability of radical disagreements about some of the members of that set. Alternatively, limiting the size of the set of moral beliefs whose reliability they want to establish may keep their contents uncontroversial, but that set might fail to provide the required justificatory support for the bridge principle.¹⁰

Second, Tersman requires that the bridge principle allows us to provide a theory of error in case there is disagreement between people who endorse moral beliefs that are explained by the third-factor account of choice (the disagreement has to be about *those* beliefs, of course). To illustrate Tersman's second demand, suppose that your belief that 'eating sweets is morally permissible' is influenced by selective pressures towards having a sweet tooth and it is true that 'sweet-tasting food is the best'. If I believe that 'eating sweets is morally impermissible', also because of evolutionary pressures, then we have to be able to explain what went wrong in my belief. If we cannot explain how I came to hold that belief as a cognitive shortcoming, a lack of imagination, or any other relevant failure, then we would have to conclude that we are in radical disagreement about whether sweet-tasting food is the best or not. This radical disagreement, Tersman contends, would then cast doubt on the truth of the chosen bridge principle 'sweet-tasting food is the best'. Most importantly, Tersman argues, both constraints combine as follows¹¹:

[A] third-factor account is plausible only if it generates the conclusion that there is a sufficiently varied and rich set of moral claims about which there is no radical disagreement. (Tersman 2017: p. 769)

Tersman thus concludes that

the plausibility of a defence of the reliability of our [objectivist] moral beliefs ... is going to depend on *which types of disagreements actually exist* [and] this is why appeals to disagreement might play a *crucial dialectical role* in the debunkers' strategy. (Tersman 2017: p. 769 emphasis added)

Note that Tersman relies on a *conciliatory* view of disagreement, according to which it is rational to at least reduce confidence in beliefs about which there is radical disagreement (e.g. Elga 2007). That view is controversial, but discussing it would lead us too far afield into the epistemology of disagreement, and thus I will assume it for the sake of argument.

Tersman's challenge is now in full view. Depending on how widespread radical moral disagreement is, the third-factor response to the debunker's challenge might yet turn out to be a dead-end street.

¹⁰ Note that Tersman does not claim that there *is* radical disagreement about any bridge principles. As such, he poses a challenge to rather than an argument against non-sceptics.

¹¹ By "generating the conclusion", Tersman means that there is a bridge principle of the form 'N is M' which is supported by the set of moral propositions that form the content of the beliefs whose reliability the objectivist wants to establish.

4 Implications beyond the reliability view

Tersman's argument has potential implications beyond the reliability view and the debate about third-factor accounts because it provides indirect support for the *disagreement view*. According to the disagreement view, there are consequences for the epistemic status of our moral beliefs caused by the Darwinist view of morality if and only if the Darwinist view of morality implies that there is counterfactual disagreement about human moral beliefs (Bergmann and Kain, eds. 2014; Bogardus 2016; White 2010; Mogensen 2016).

If the disagreement view were correct, the most commonly cited explanations of the epistemic significance of the Darwinist view of morality would be mistaken. Contrary to the conclusions of many discussions, the epistemic problems that may arise by uncovering the causal origins of our beliefs would not be due to epistemic insensitivity, a lack of safety, the presence of irrelevant influences, striking contingencies, accidentally true beliefs, or the worrying historical variability of moral beliefs.¹² Instead, the problem, if there is one, will turn out to be due to the epistemic significance of counterfactual disagreement, or so the proponents of the debunking–disagreement thesis claim.

Here is how Tersman's challenge may provide indirect support for the disagreement view. Suppose that the reliability view is correct, but third-factor replies work. In that case, undercutting defeat due to evolutionary explanations of morality could be prevented by a third-factor explanation, which could be taken as defeating the evolutionary defeater.¹³ However, if Tersman is correct and disagreement constrains or altogether invalidates third-factor explanations, then defeat due to the evolutionary debunking explanation cannot be prevented by third-factor accounts, and therefore our objective moral beliefs are undercut. Thus, disagreement is ultimately crucial for evaluating the epistemic significance of evolutionary debunking explanations, just as proponents of the disagreement view claim. Figure 2 depicts the relation between the disagreement view and Tersman's challenge. Arrows in the graph signify an 'attack' or defeat relation, which can be between arguments or theoretical positions (in circles) or between an argument or position and an attack relation.

To further illustrate the significance of Tersman's challenge, it should be noted that *any* implication that follows from a Darwinist view of morality that (a) threatens to undercut all objectivist moral beliefs and (b) may itself be defeated by a third-factor explanation could be put into the top-left ellipse in Fig. 2. Thus, notwithstanding the direct implications of the Darwinist view of morality (of which there might be several valid ones, of course), objectivists could be forced to a sceptical conclusion *only* because of the epistemic significance of disagreement. So, if Tersman's constraints successfully disallowed all third-factor explanations, the disagreement view would gain strong indirect support.

¹² For reasons of space, I cannot fully introduce these proposals here. For a good overview, see Wielenberg (2016).

¹³ See Pollock and Cruz (1999: p. 200ff) and Klenk (2017a) for a discussion of defeater–defeaters.

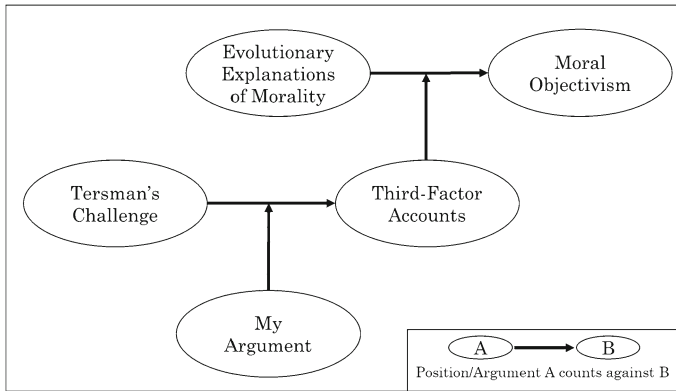


Fig. 2 Indirect support for the disagreement view

However, in the next section, I argue that the constraints proposed by Tersman do not constrain third-factor accounts. I then return to an assessment of the prospects of the disagreement view in Sect. 7.

5 Unconstrained by moral disagreement

My criticism has two prongs. First, Tersman’s constraints miss the mark against moral objectivists. I show that the first constraint is trivially satisfied on plausible assumptions about the set of beliefs whose reliability objectivists seek to defend and that the second constraint is irrelevant to objectivists’ cause. Second, recent ethnographic studies suggest that even if Tersman’s criteria were not trivially satisfied, there is a set of moral beliefs that is varied and rich and not subject to any disagreement, thus satisfying Tersman’s criteria.

5.1 Relevant moral disagreement is impossible

Tersman’s first constraint is plausible. However, it is unclear at a critical junction. The relevant sense of “support” (Tersman 2017: p. 768) between the beliefs whose reliability is to be established and the bridge principle needs to be clarified. However, no plausible understanding of support makes it the case that Tersman’s constraint could be violated. To support this claim, I will proceed in two steps. First, I consider the set of the contents of the beliefs whose reliability is to be established and how that set gets ‘populated’.¹⁴ Next, I consider whether there is a notion of ‘support’ between the contents of that set and the bridge principle such that there could be radical disagreement between believers who endorse propositions in the set of beliefs

¹⁴ I make a distinction between *the beliefs whose reliability is to be established by the objectivist* and *the set of the contents of the beliefs whose reliability is to be established by the objectivist* to make clear that there are justificatory relations between propositions, but that disagreement must be between believers who form beliefs about those propositions.

whose reliability is to be established (disagreement, that is, about the contents of that set).

The first step is to become clear about what the beliefs are whose reliability is to be established. Given that objectivists, by hypothesis, aim to defend the reliability of those beliefs, it stands to reason that they should aim to establish the reliability of only those beliefs that are worth keeping. Any viable epistemology will provide relevant constraints. For example, as an externalist process reliabilist, you would want to keep the beliefs that are formed by a reliable method. As an internalist evidentialist, you would want to keep the beliefs that, very roughly, are sufficiently supported by the evidence available to you or that *seem* to be reliable (to you). Moreover, whichever structure of justification you defend, be it a pyramidal foundationalist structure or a coherentist picture, the justified beliefs that are the beliefs whose reliability is to be established stand in justification-conferring relations to each other. Though it is logically possible that the set of the contents of the beliefs whose reliability is to be established may consist of (sets of) mutually independent propositions, objectivists would plausibly want to defend propositions that are interconnected through some normative theory. Thus, we can assume that the beliefs whose reliability is to be established will be *prima facie* justified and, given their content, stand in justification-conferring relations to each other. So much for the structural aims of ‘populating’ the set of beliefs whose reliability is to be established. Before proceeding to the question about support, consider a sceptical worry.

A moral sceptic might ask the following question: how can we select the beliefs whose reliability is to be established so that they stand in justification-conferring relations to each other? However, this is just the problem of determining which moral beliefs we can hold rationally or justifiably. There might be good reason to be sceptical about finding such an account. For example, we might doubt that there is a method through which we can determine what the moral beliefs are that are worth keeping—which ones are worth defending? Should we defend the reliability of the belief that the death penalty is morally impermissible or the reliability of the belief that the death penalty is morally permissible? These are difficult questions, of course, but they are questions that can be bracketed when discussing the debunking challenge, given that the debunking challenge is supposed to be distinct from the general sceptical challenge about having justified moral beliefs *at all*. We should, therefore, assume that once a set of beliefs whose reliability is to be established is populated, its members are *prima facie* justified in mutually reinforcing support relations.

Let us now turn to the second step of my first objection and consider different notions of support. The question is whether the beliefs whose reliability is to be established can *support* the bridge principle (satisfying Tersman’s first constraint) but allow the option of there being radical disagreement about these beliefs (violating Tersman’s second constraint), as is essential to Tersman’s argument.

Suppose the relevant support relation between the set of the contents of the beliefs whose reliability is to be established, or the set of target beliefs MT, and the bridge principle *f* is *deductive* inference, as depicted by the arrow in Fig. 3.

In the case depicted in Fig. 3, the contents of the beliefs whose reliability is to be established provide conclusive reason for accepting the bridge principle. Suppose that all propositions in the set of the contents of the beliefs whose reliability is to be




Fig. 3 A support relation from moral beliefs to a bridge principle

established (MT) are required to deduce the bridge principle f . In that case, there cannot be radical disagreement amongst people who endorse the beliefs whose reliability is to be established. The propositions would have to be inconsistent for radical disagreement to be possible (amongst the believers of those propositions), which is impossible given the assumption that the entire set of the contents of the beliefs whose reliability is to be established is part of the bridge principle's premises.

A problem seems to arise if there is another set of propositions that is mutually inconsistent with the set of the contents of the beliefs whose reliability is to be established. In that case, the propositions in both sets taken together would imply anything by material implication, including the bridge principle. Thus, though it is obvious that objectivists should only take consistent propositions to be part of the set of the contents of the beliefs whose reliability they seek to establish, as argued above, it seems problematic that there are many internally consistent but mutually inconsistent sets of (moral) propositions. The union of those sets of propositions would imply the bridge principle, but radical disagreement amongst people who endorse the premises of the bridge principle is certainly possible, contrary to what I have argued in the previous paragraph.¹⁵

The answer to this worry depends on whether the intersection of the deductive closures of the relevant sets is empty or not. If the intersection of the deductive closures of the relevant sets is non-empty, objectivists have an easy way out: they could take beliefs in the contents of the intersection as the beliefs whose reliability they seek to establish. For example, in light of Parfit's (2011) discussion, we might be 'climbing the same mountain from different sides' in normative ethics, and so it would be plausible that the intersection of the deductive closures of both sets is non-empty. In that case, the argument of the previous paragraph applies, and Tersman's first constraint is trivially satisfied. If the intersection is empty, however, then there are at least two mutually inconsistent sets of moral propositions. To illustrate, suppose that one set contains characteristically deontological claims and the other characteristically utilitarian claims. If we cannot legitimately choose one set over the other, based on normative theorising, then there will be radical disagreement about moral matters and a conciliatory view of disagreement will imply that no moral belief remains justified. Of course, such fundamental moral scepticism is a serious possibility, but one that can be set aside in the present context. That is because both proponents and opponents of moral objectivism commonly suppose that the evolutionary debunking challenge is an *additional* threat to moral objectivism (cf. Vavova 2015; Wielenberg 2016). If there were fundamental moral disagreement to the extent that we could not rationally choose between different and mutually inconsistent normative theories, then justi-

¹⁵ Many thanks to an anonymous reviewer for prompting me to address this objection.




Fig. 4 Three ways to support third-factor accounts

fication of moral beliefs seems forlorn from the start (as long as we assume moral objectivism). Worries about evolutionary debunking arguments would be superfluous (cf. Klenk 2017a). Thus, it is legitimate to assume that we can legitimately choose one set amongst the set of mutually inconsistent sets of moral propositions that together entail the bridge principle as the set of beliefs whose reliability is to be established. In that case, however, there cannot be radical disagreement about the propositions in that set, as argued above.

Tersman might object as follows: suppose that only *some* propositions in the set of the contents of the beliefs whose reliability is to be established, MT , are required to deduce the bridge principle. Let the propositions required to deduce the bridge principle from MT be the subset MT_f . Again, there cannot be radical disagreement about the members of MT_f because of the justificatory structure of MT : the members of MT stand in justificatory relations to each other, and they entail the bridge principle f . Hence, a subset of MT cannot fail to entail f .

It might seem possible that there is radical disagreement between believers who endorse the members of MT_f and those that endorse the members of MT_f^C (the complement of MT_f , i.e. the members of MT that are not members of MT_f). However, as Fig. 4 illustrates, MT_f^C is a proper subset of MT . If either of the above methods is used to determine the contents of MT , then MT_f^C can only contain members that are consistent with MT_f (Fig. 4).

Suppose instead that support is based on an inductive argument that goes from the set of the contents of the beliefs whose reliability is to be established to the bridge principle. The situation would be as in Fig. 3, but with an inconclusive support relation. If it is appropriate for the inductive base to confer justification on the bridge principle, then the inductive base must be consistent. So, as before, there could not be radical disagreement amongst believers who endorse the beliefs whose reliability is to be established. The inductive case can, of course, be defeated by bringing to light new information that was not but should have been part of the inductive base, and that would not serve as a base for the inductive case. However, that would mean that the set of beliefs whose reliability is to be established is incomplete. However, adding a member to the set of beliefs whose reliability is to be established would follow the same constraints as outlined above, in which case there could not be disagreement either.

Therefore, on any plausible understanding of ‘support’, if the set of the contents of the beliefs whose reliability is to be established supports a bridge principle, then radical disagreement about the beliefs whose reliability is to be established is impossible. There might be disagreement about beliefs other than those whose reliability is to be established, but it will not be disagreement that could make trouble for third-factor accounts. Let’s look at Tersman’s second constraint next.




Fig. 5 Beliefs that are causally explained by the third-factor (adapted from Berker (2014: 230))

5.2 Possible moral disagreement is irrelevant

Tersman's second constraint is that there must be no radical disagreement about the propositions that form the content of beliefs that are causally explained by the third factor. However, we will see that beliefs that are causally explained by the third factor are relevant for anti-sceptics only insofar as they are *also* members of the set of the contents of the beliefs whose reliability is to be established.

Tersman's constraint concerns the beliefs that share a common (ultimate) causal background factor. Let the contents of these beliefs form the set U , as depicted in Fig. 5. Tersman demands that there must be no radical disagreement about the propositions in U . Clearly, some propositions in U might be inconsistent with the propositions in the set of the contents of the moral beliefs whose reliability is to be established (MT). Thus, U and MT might diverge, as Tersman suggests. However, the intersection of both sets is the only aspect that should worry objectivists. Given that the intersection will be within the set of the contents of the beliefs whose reliability is to be established, there cannot be radical disagreement about the relevant area of the beliefs that are causally explained by the third factor (Fig. 5).

To support the claim that objectivists need only worry about propositions that are both in U and in MT , I show that beliefs caused by the third factor need not be amongst the beliefs whose reliability the objectivist needs to defend. Thus, disagreement about the beliefs that are causally explained by the third factor is possible, but it turns out to be irrelevant for our assessment of third-factor accounts.¹⁶ For example, suppose the causal factor in the third-factor account is 'aids survival'. It might be that some beliefs are ultimately causally explained by actions or events that have this property and nonetheless conflict. A good example is the killing of relatives, which might aid survival for those living in polar regions and thus lead to a more favourable opinion of that practice as compared to, say, the prevailing views in Western countries. The set of the contents of the beliefs whose reliability is to be established also has members that are *not* members of the set. For example, take a complex belief like <the right to

¹⁶ There are three further possibilities concerning the relation between the contents of the beliefs that are causally explained by the third factor and the set of the contents of the beliefs whose reliability is to be established that I do not discuss in greater detail because they are easily dealt with: (a) the former is a proper subset of the latter, (b) there are propositions that are in the former but not in the latter but consistent with the latter, and (c) there are propositions that are in the former but not in the latter and inconsistent with the latter. Given the argument in Sect. 5.1, there cannot be relevant disagreement in (a), and objectivists could just adopt the union of both sets in (b) and drop the intersection of inconsistencies in (c).

determine what happens to one's body trumps the concern for the life of a foetus>. The third factor identified by objectivists may not cause this belief, but it might nonetheless be a member of the beliefs whose reliability is to be established by objectivists. At the same time, the set of the contents of the beliefs that are causally explained by the third factor and the set of the contents of the beliefs whose reliability is to be established cannot be *disjunct* because that would violate the constraint that the bridge principle must be supported by the beliefs whose reliability is to be established.

It is possible that the contents of the beliefs that are causally explained by the third factor form a proper subset of the set of the contents of the beliefs whose reliability is to be established ($U \subseteq MT$). Thus, the relevant beliefs, about which there must not be radical disagreement, are the beliefs that are causally explained by the third factor and amongst the beliefs whose reliability is to be established. The considerations above have shown, however, that radical disagreement about those beliefs is impossible given plausible assumptions about the population of the set of beliefs whose reliability is to be established.

It might be objected that the set of beliefs whose reliability is to be established might be a proper subset of the set of beliefs that are (ultimately) causally explained by the third factor ($U \supseteq MT$). The beliefs about the contents of U might seem to be 'ascribed reliability' by the bridge principle insofar as objectivists have identified the relevant causal factor as conducive to reliability. In that case, some beliefs that are causally explained by the third factor might conflict with the beliefs whose reliability is to be established (and yet all seem supported by the bridge principle). That seems to create the predicament of a radical disagreement that cannot be easily resolved, for example if both your belief that 'killing your elderly relatives is permissible' and my belief that 'it is not the case that killing your elderly relatives is permissible' could be explained in reference to the fitness-enhancing effects of environmentally-sensitive family management. Here the *causal* explanation is in line with the third-factor account, but the resulting beliefs are not because we have two inconsistent sets of beliefs. That situation is possible if 'ascribing reliability'—or the support relation between the bridge principle and the beliefs whose reliability is to be established and the beliefs that are ultimately causally explained by the third factor—is understood as an explanation that shows that the evolutionary process that influenced our moral beliefs creates beliefs that are mostly, but not always, true. Hence, the process that leads to the beliefs whose reliability is to be established can be reliable even if it gives rise to some false beliefs.

However, even if the contents of the beliefs whose reliability is to be established form a proper subset of the contents of the beliefs that are ultimately causally explained by the third factor, there is no special problem in sorting out the possible disagreement. The answer is clear once we recognise that explaining the reliability of our beliefs by means of an evolutionary process is different from explaining the support that beliefs that are ultimately causally explained by the third factor gain from the bridge principle. Evolutionary considerations might explain why we tend to form the beliefs that are explained by the third factor. However, even if the contents of the evolutionarily explained beliefs conflict with the contents of the beliefs whose reliability is to be established, the bridge principle will either imply that the evolutionarily explained beliefs

are unsupported or provide reason to include the evolutionarily explained beliefs it supports amongst the beliefs whose reliability is to be established.

Therefore, there are no reasons to think that there could be radical disagreement about the beliefs that underwrite the objectivist's third-factor account. Thus, disagreement between thinkers who endorse beliefs that are causally explained by the third-factor account is possible, but only insofar as those beliefs are not amongst the beliefs whose reliability is to be established. In that case, however, the disagreement is irrelevant for the objectivist.

5.3 Actual moral disagreement is implausible

Even if Tersman's criteria were not trivially satisfied, recent ethnographic research provides strong support for the view that there is a rich and varied set of foundational moral beliefs that is not subject to radical disagreement. In other words, even if Tersman's conditions could work, it is *implausible* that there would be actual radical disagreement amongst the members of MT. So, objectivists would pass Tersman's test.

I should make clear that this section addresses the question of whether actual moral disagreement helps the debunker's case, assuming that Tersman's main point (that disagreement has an important role to play in debunking arguments) is correct. My arguments in Sects. 5.2 and 5.3 might fail, and disagreement could be relevant in the way Tersman describes and potentially negative for moral objectivists, but here I argue that disagreement does not have negative effects for moral objectivists. Of course, since I am making a claim about the actual extent of moral disagreement, this section is to some extent speculative: efforts are under way, but the empirical investigations required to establish the degree of actual moral disagreement are still in their infancy.

A widely endorsed hypothesis about the evolution of morality is that moral beliefs serve a cooperative function. This view is endorsed by authors across the metaethical spectrum (Gibbard 2003; Kitcher 2011; Wielenberg 2014; Street 2006; Joyce 2006). A recent defence of the hypothesis comes from Curry (2016). Using a game-theoretic approach, Curry predicts that moral beliefs will be concerned with four problem-centred domains: (1) the allocation of resources to kin; (2) coordination to mutual advantage; (3) exchange; and (4) conflict resolution (Curry 2016: p. 30). For reasons of space, the interesting details of these considerations cannot be recounted here. However, in a recent study, Curry, Mullins, and Whitehouse show that their predictions are born out in the ethnographic record of 60 societies (Curry et al. in press). Behaviour that fell into the four categories specified by Curry (2016) was regarded as morally good by all studied societies. The finding implies that there is tremendous agreement about several moral domains.

Of course, there might be cases of intractable, radical disagreement about moral matters too. For example, Doris and Plakias (2008) suggest, based on the work of Nisbett and Cohen (1996), that "North/South differences in attitudes toward violence and honour might well persist in ideal discursive conditions" (Doris and Plakias 2008: p. 331). However, as they themselves acknowledge, perhaps a tad too modestly, "one case is not an induction base" (ibid.). More importantly in the present context, the cases

of radical disagreement discussed in the literature (the locus classicus being Brandt's (1954) discussion of Hopi ethics) do not concern the domains of shared moral beliefs identified by Curry et al., and it is possible that apparent moral disagreements may depend on non-moral disagreements.

Therefore, leaning on the hypothesis of morality as stimulating cooperation, objectivists can argue that widely shared positive evaluations of kinship relations, cooperation, exchange, and conflict resolution provide an adequate basis for coming up with a third-factor account that meets Tersman's criteria.

6 Rejoinder: higher-order evidence

Disagreement does not seem to impugn the epistemic legitimacy of third-factor responses. However, Tersman's proposal seems suggestive of a wider problem that has to do with higher-order evidence. The thought goes as follows. The fact that there is actual disagreement about even deeply held moral beliefs, often even with those whom we consider to be our peers, provides evidence that our moral beliefs are not reliable. Debunking challenges provide us with higher-order evidence, not about particular moral beliefs, but rather about *our general ability to form reliable moral beliefs*, at least concerning the topics about which we disagree.

Perhaps interpreting the problem of disagreement as a higher-order evidence problem helps Tersman's account. According to Christensen (2010), a distinctive feature of sceptical higher-order evidence about p is that it requires the bracketing of one's evidence for p when assessing whether one should or should not endorse p . If that is correct, and if we can interpret the debunker's challenge as a higher-order evidence challenge, then the support from the beliefs whose reliability is to be established for the bridge principle must be bracketed when evaluating disagreement about the bridge principle. The question becomes how likely we would be to be correct in picking a bridge principle once we bracket the available evidence for the bridge principle (such as our moral intuitions, several 100 years of philosophising about the bridge principle, and so on). This response would avoid the problem of the critical notion of support that is at the heart of Tersman's original argument.

However, the appeal to higher-order evidence confuses two challenges, so it does not help Tersman's account. The debunker challenges moral objectivists to vindicate a third-factor account given the assumption that at least some moral beliefs have positive epistemic credentials. The higher-order evidence challenge to moral objectivism, in contrast, casts doubt precisely on the validity of that assumption.

Moreover, higher-order evidence challenges seem to break down in the limiting case where the reliability of beliefs about an entire domain is called into question. If higher-order evidence is applied to all beliefs of a type and there is no possibility of using other beliefs to vindicate the justification of that type, then higher-order evidence reduces to a general sceptical challenge. Tersman himself acknowledges that beliefs of the moral type are properly 'isolated'. Thus, given that the debunking challenge is supposed to be distinct from the general sceptical challenge, the higher-order-evidence interpretation cannot save Tersman's disagreement-based constraints on third-factor explanations.

7 Prospects for the disagreement view

Tersman's challenge might provide indirect support for the disagreement view, but his disagreement-based constraints have been shown to fail. A question about the broader prospects of the disagreement view remains: do evolutionary debunking arguments show that there is counterfactual disagreement that directly undercuts objective moral beliefs?

A full discussion of this question would require a separate paper. Here I have the modest aim of sketching two considerations that are crucial for assessing the disagreement view in full, on the one hand, and that are relevantly linked to the discussion of Tersman's challenge, on the other.¹⁷ Before introducing these points, let me illustrate the disagreement view in greater detail. One part of the disagreement view is negative and contains criticism of existing interpretations of the evolutionary debunking challenge to the effect that they all fail, based on arguments that need not concern us here (cf. Bogardus 2016). The positive idea is to link up Darwinist views of morality with the epistemic significance of disagreement: had humans taken a different evolutionary path, and given the influence of evolution on our moral beliefs, they might have formed moral beliefs that are incompatible with the moral beliefs that humans currently hold (cf. Bogardus 2016: p. 656; Mogensen 2016: p. 607). Learning about the Darwinist view of morality, the reasoning goes, thus informs us about a counterfactual disagreement and, assuming that counterfactual disagreement is undercutting, our moral beliefs will be undercut.

After setting aside the indirect support that the disagreement view might have gained from Tersman's challenge, the truth of the disagreement view depends on two very general criteria:

1. Darwinist views of morality must imply disagreement about the target beliefs.
2. Darwinist views of morality must imply disagreement that is defeating.

The level of generality of both criteria allows us to pinpoint two unobvious problems that will determine the success of the disagreement view.

First, what are the target beliefs of evolutionary debunking arguments? The evolutionary debunking arguments of Street and Joyce are aimed at the objective moral beliefs, but it is unclear how widely or narrowly we construe the set of moral beliefs that objectivists need to defend. Two key considerations are in tension with each other from the perspective of the objectivist, and the intuition that motivated Tersman's first constraint (about the contents of the set of beliefs whose reliability is to be established) resurfaces here, though it is applied to a different criterion.

On the one hand, for the disagreement view to be correct, there has to be an evolutionary story that makes it plausible that our counterfactual selves (whom we imagine having taken a different evolutionary path) could have formed beliefs whose content is inconsistent with the content of the target beliefs. The best available accounts of evolutionary psychology will be crucial here. For instance, it is plausible that our counterfactual selves have formed different beliefs about property rights in, say, resource-rich environments, but less plausible that they have formed different beliefs

¹⁷ See Klenk (2018: ch. 5) for a detailed discussion.

about whether a person ought to take the means to her ends or whether cooperation is good.¹⁸ It will thus be tempting for objectivists to restrict themselves to defending only very few basic moral beliefs. For example, beliefs about the goodness of cooperation are obvious candidates to be defended by objectivists because evolutionary game-theoretic considerations suggest that such beliefs will be stable across nearby evolutionary paths (cf. Bowles and Gintis 2011).

On the other hand, objectivists cannot restrict the set of beliefs they are trying to defend too much to cover only basic beliefs such as the belief that <you ought to take the means to your ends>. Rovane (2013: p. 207ff), for example, suggests that objectivists must vindicate the justification of a significant proportion of the current moral discourse for their view to be interesting. Thus, an adequate assessment of the disagreement view will call for a renewed focus on the question of which beliefs moral objectivists need to defend as well as an adequate approximation of the ‘stability’ of certain moral beliefs across different socio-environmental conditions. The anthropological research cited in Sect. 5.3 suggests starting points to address the latter question, but the former question has hitherto played only a minor role in the debunking debate.

Second, if there are target beliefs about whose content a Darwinist view plausibly implies disagreement, it has to be shown that the disagreement is defeating and thus of an epistemically relevant kind. Not all disagreement is defeating. For example, disagreement between experts and novices is not considered to be defeating for both parties, nor is merely hypothetical disagreement defeating in an obvious sense.¹⁹ The prospects of the disagreement view thus depend on whether counterfactual disagreement is undercutting in the first place and whether we should, as it were, take seriously (in an epistemic sense) the moral views of our counterfactual selves. I suggest that the notion of an epistemic peer, gleaned from the literature on peer disagreement, will be relevant here, as it allows us to assess whether the moral views of our counterfactual selves should give us pause about our own moral views. However, if we demand that our counterfactual selves must be our moral peers in an interesting sense, we face a problem for the disagreement view.

Suppose that the proponent of the disagreement view wants to show that all objectivist moral beliefs are undercut. We thus have to show that there is total disagreement about moral norms. However, in cases of total moral disagreement between ourselves and our counterfactual selves it becomes hard to see how that disagreement could still be defeating. The problem is that we have to identify our counterfactual selves as peers when it comes to moral beliefs to argue that we should take seriously their moral views. But in cases where there is total disagreement about moral norms, we cannot rely on the *content* of our ‘moral’ beliefs to identify our peers. There is no content-based agreement and so we should conclude that our counterfactual selves are not our moral peers, in which case there would not be epistemically relevant disagreement and the disagreement view falters. Alternatively, we could try to identify a *formal* criterion for moral norms so that we could identify moral peers even in the absence of content-based agreement. The problem is that it proves perennially difficult to distinguish moral from other kinds of norms on purely formal grounds, such as their

¹⁸ Cf. Barkhausen (2016) for a related discussion.

¹⁹ Cf. Tersman (2013) for a discussion of the significance of hypothetical disagreement.

mode of justification, or the emotional reactions they trigger, or the peculiar developmental pathway the allegedly follow (cf. Southwood 2011). If there is no purely formal criterion for distinguishing moral norms from non-moral norms, however, then some agreement about the content of morality seems required after all to identify our counterfactual selves as peers. In that case, however, there would be some agreement about morality and thus it would be false that evolution undercuts all objectivist moral beliefs.

Therefore, a dilemma arises for the debunker: if there is total disagreement, then we seem to lack a basis for counting our counterfactual selves as peers, in which case there is no undercutting disagreement. If there is partial agreement, then we have a basis for taking our counterfactual selves as peers, but then there is a subset of moral beliefs that is not undercut by Darwinist views of morality.

These questions about the appropriate set of target beliefs and the peerhood status of our counterfactual selves strike me as crucial for assessing the disagreement view. What the present article establishes with regard to the disagreement view is a partial conclusion. If debunking arguments raise a problem about explaining the reliability of moral beliefs, and if third-factor accounts are a promising response, then the epistemic significance of disagreement plays no role in assessing the strength of evolutionary debunking arguments.

8 Conclusion

Once the *prima facie* legitimacy of third-factor accounts is granted, as many scholars do, disagreement is not a problem for third-factor explanations, contrary to Tersman's suggestion.

In particular, plausible assumptions about the structure of justification and the nature of the debunking challenge show that there cannot be radical disagreement amongst the beliefs supported by the third-factor account. Moreover, even if it were possible, there is good empirical support for thinking that there is no actual radical moral disagreement.

In conclusion, worries about disagreement do not rationally constrain third-factor explanations in metaethics. Plausibly, this finding generalises to other domains, such as mathematics. More generally, Tersman's proposed constraints on third-factor explanations do not serve to support the debunking-disagreement thesis. Insofar as third-factor accounts are *prima facie* legitimate, this is good news for moral objectivists.

Funding This work is part of the research programme 'Evolutionary Ethics? The (Meta-) Ethical Implications of Evolutionary Explanations of Morality' with project number 360-20-340, which is financed by the Netherlands Organisation for Scientific Research (NWO).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Artiga, M. (2015). Rescuing tracking theories of morality. *Philosophical Studies*, 172(12), 3357–3374.
- Barkhausen, M. (2016). Reductionist moral realism and the contingency of moral evolution. *Ethics*, 126(3), 662–689.
- Behrends, J. (2013). Meta-normative realism, evolution, and our reasons to survive. *Pacific Philosophical Quarterly*, 94(4), 486–502.
- Bergmann, M., & Kain, P. (Eds.). (2014). *Challenges to moral and religious belief: Disagreement and evolution*. Oxford: Oxford University Press.
- Berker, S. (2014). Does evolutionary psychology show that normativity is mind-dependent? In J. D'Arms & D. Jacobson (Eds.), *Moral psychology and human agency philosophical essays on the science of ethics* (pp. 215–252). Oxford: Oxford University Press.
- Bogardus, T. (2016). Only all naturalists should worry about only one evolutionary debunking argument. *Ethics*, 126(3), 636–661.
- Bowles, S., & Gintis, H. (2011). *A cooperative species: Human reciprocity and its evolution*. Princeton: Princeton University Press.
- Brandt, R. B. (1954). *Hopi ethics: A theoretical analysis*. Chicago: University of Chicago Press.
- Brosnan, K. (2011). Do the evolutionary origins of our moral beliefs undermine moral knowledge? *Biology and Philosophy*, 26(1), 51–64.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research*, 81(1), 185–215.
- Clarke-Doane, J. (2016). Debunking and dispensability. In U. D. Leibowitz & N. Sinclair (Eds.), *Explanation in ethics and mathematics* (pp. 23–36). Oxford: Oxford University Press.
- Copp, D. (2008). Darwinian skepticism about moral realism. *Philosophical Issues*, 18, 186–206.
- Curry, O. S. (2016). Morality as cooperation. A problem-centred approach. In T. K. Shackelford & R. D. Hansen (Eds.), *The evolution of morality* (pp. 27–51). Dordrecht: Springer.
- Curry, O. S., Mullins, D. A., & Whitehouse, H. (in press). Is it good to cooperate? Testing the theory of morality-as-cooperation in 60 societies. *Current Anthropology*.
- Doris, J. M., & Plakias, A. (2008). How to argue about disagreement: Evaluative diversity and moral realism. In W. Sinnott-Armstrong (Ed.), *Moral psychology: The cognitive science of morality. Intuition and diversity (A Bradford Book)* (Vol. 2, pp. 303–331). Cambridge: MIT Press.
- Elga, A. (2007). Reflection and disagreement. *Noûs*, 41(3), 478–502.
- Enoch, D. (2010). The epistemological challenge to metanormative realism. How best to understand it, and how to cope with it. *Philosophical Studies*, 148(3), 413–438.
- Gibbard, A. (2003). *Thinking how to live*. Cambridge: Harvard University Press.
- Joyce, R. (2006). *The evolution of morality (Life and mind)*. Cambridge: MIT Press.
- Joyce, R. (2013). Irrealism and the genealogy of morals. *Ratio*, 26(4), 351–372.
- Joyce, R. (2016). Evolution, truth-tracking, and moral scepticism. In *Essays in moral skepticism*. Oxford: Oxford University Press (pp. 142–158).
- Kitcher, P. (2011). *The ethical project*. Cambridge: Harvard University Press.
- Klenk, M. (2017a). Can moral realists deflect defeat due to evolutionary explanations of morality? *Pacific Philosophical Quarterly*, 98(S1), 227–248.
- Klenk, M. (2017b). Old wine in new bottles. Evolutionary debunking arguments and the Benacerraf-field challenge. *Ethical Theory and Moral Practice*, 20(4), 781–795.
- Klenk, M. (2018). Survival of defeat. Evolution, moral objectivity, and undercutting. Ph.D. thesis, Utrecht, Utrecht University.
- Kölbel, M. (2004). Faultless disagreement. *Proceedings of the Aristotelian Society*, 104, 53–73.
- Mogensen, A. L. (2016). Contingency anxiety and the epistemology of disagreement. *Pacific Philosophical Quarterly*, 97(4), 590–611.
- Nisbett, R. E., & Cohen, D. (1996). *Culture of honor: The psychology of violence in the south (new directions in social psychology)*. Boulder: Westview Press.
- Parfit, D. (2011). *On what matters: Volume one. On what matters*. Oxford: Oxford University Press (Vol 1).
- Pollock, J. L., & Cruz, J. (1999). *Contemporary theories of knowledge. Studies in epistemology and cognitive theory* (2nd ed.). Lanham: Rowman and Littlefield Publishers.
- Rovane, C. A. (2013). *The metaphysics and ethics of relativism*. Cambridge: Harvard University Press.

- Schechter, J. (2018). Explanatory challenges in metaethics. In T. McPherson & D. Plunkett (Eds.), *The Routledge handbook of metaethics (Routledge handbooks in philosophy)* (1st ed., pp. 443–458). New York: Routledge.
- Shafer-Landau, R. (2012). Evolutionary debunking, moral realism and moral knowledge. *Journal of Ethics and Social Philosophy*, 7(1), 1–37.
- Skarsaune, K. O. (2011). Darwin and moral realism. Survival of the iffiest. *Philosophical Studies*, 152(2), 229–243.
- Southwood, N. (2011). The moral/conventional distinction. *Mind*, 120(479), 761–802.
- Street, S. (2006). A Darwinian dilemma for realist theories of value. *Philosophical Studies*, 127(1), 109–166.
- Street, S. (2008). Reply to Copp. Naturalism, normativity, and the varieties of realism worth worrying about. *Philosophical Issues*, 18(1), 207–228.
- Street, S. (2016). Objectivity and truth. You'd better rethink it. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (1st ed., Vol. 11, pp. 293–334). Oxford: Oxford University Press.
- Talbott, W. J. (2015). How could a 'blind' evolutionary process have made human moral beliefs sensitive to strongly universal, objective moral standards? *Biology and Philosophy*, 30(5), 691–708.
- Tersman, F. (2013). Moral disagreement. Actual vs. possible. In D. E. Machuca (Ed.), *Disagreement and skepticism. Routledge studies in contemporary philosophy* (Vol. 46, pp. 90–108). New York: Routledge.
- Tersman, F. (2017). Debunking and disagreement. *Noûs*, 51(4), 754–774.
- Vavova, K. (2014). Debunking evolutionary debunking. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 9, pp. 76–101). Oxford: Oxford University Press.
- Vavova, K. (2015). Evolutionary debunking of moral realism. *Philosophy Compass*, 10(2), 104–116.
- White, R. (2010). You just believe that because.... *Philosophical Perspectives*, 24(1), 573–615.
- Wielenberg, E. J. (2014). *Robust ethics: The metaphysics and epistemology of godless normative realism* (1st ed.). Oxford: Oxford University Press.
- Wielenberg, E. J. (2016). Ethics and evolutionary theory. *Analysis*, 76(4), 502–515.