



# A New Route from Moral Disagreement to Moral Skepticism

**ABSTRACT:** *Moral disagreement is sometimes thought to pose problems for moral realism because it shows that we cannot achieve knowledge of the moral facts the realists posit. In particular, it is ‘fundamental’ moral disagreement—that is, disagreement that is not due to distorting factors such as ignorance of relevant nonmoral facts, bad reasoning skills, or the like—that is supposed to generate skeptical implications. In this paper, we show that this version of the disagreement challenge is flawed as it stands. The reason is that the epistemic assumptions it requires are incompatible with the possibility of fundamental disagreement. However, we also present an alternative reconstruction of the challenge that avoids the problem. The challenge we present crucially invokes the principle that knowledge requires ‘adherence’. While that requirement is usually not discussed in this context, we argue that it provides a promising explanation of why disagreement sometimes leads to skepticism.*

**KEYWORDS:** metaethics, moral disagreement, moral skepticism, moral realism, adherence

## Introduction

One avenue for questioning theories that posit moral facts is to argue that we lack the ability to know those facts. If it could be shown that the assumptions the theories involve—for example, those about the nature of moral facts—leave no room for moral knowledge then that is commonly seen as a *reductio*. This line of reasoning helps explain why moral diversity is given such a central role in metaethics because the extensive disagreement that occurs over moral issues is often taken to reinforce the worries about our ability to know moral truths—assuming they exist at all.

The most common target of the strategy just indicated is moral realism. According to realism, there are moral facts that are objective in the sense that they exist independently of our attitudes, conventions, and the like (see Brink [1989], Cuneo [2007], Enoch [2011], and Shafer-Landau [2003] for moral realists who explicate

Thanks to Billy Dunaway, David Enoch, Zoe Johnson King, David Killoren, Michael Klenk, Brian Leiter, Don Loeb, Jonathan Matheson, Victor Moberger, and Richard Rowland for valuable comments on preliminary drafts. We are also grateful to two anonymous reviewers whose perceptive criticisms have led to significant improvements.

their position along these lines). That commitment is supposed to make it particularly hard to accommodate the possibility of moral knowledge in the face of moral disagreement. Thus, if the argument is successful, realists are stuck with the position that the objective moral facts they posit are unknowable. As Russ Shafer-Landau puts it, although that position seems logically coherent, it ‘contains about zero appeal’ (2012: 1). In what follows, we call this argumentative strategy ‘the disagreement challenge’ to moral realism (there are other ways in which one could argue that moral disagreement undermines moral realism, see Tersman [2006] for an overview).

Many instances of the disagreement challenge in the literature have been met with powerful objections. One of their vulnerabilities is that they often rely upon the premise that much moral disagreement is especially deep or, as we shall say, ‘fundamental’, in the sense that it cannot be attributed to some distorting factor such as ignorance of relevant evidence, bad reasoning skills, bias, or the like. As this premise raises complex empirical issues, critics of the argument may reasonably question whether the available evidence is sufficient to justify it (for relevant discussion, see Nisbett and Cohen [1996], Doris and Plakias [2008], Fraser and Hauser [2010], and Demetriou [2014]). Another prominent type of criticism focuses on the inference from the premise about disagreement to the conclusion that moral realism leads to moral skepticism. A common complaint is that this crucial inference is often left unexplained or unmotivated or that it presupposes epistemic principles that overgeneralize by also excluding knowledge in areas where skepticism or antirealism is thought to be independently implausible (for example, Shafer-Landau [2006: 56] urges that the challenge generalizes to all areas of philosophy).

Do these objections spell doom for the disagreement challenge? We believe that such a verdict is premature for there are alternative ways to develop the argument that have so far been overlooked. Our aim in this paper is to present a novel version of the challenge that differs from others in two ways. First, it relies on weaker assumptions about the nature of the existing moral disagreement than those that are typically invoked, and, second, it motivates the step from disagreement to skepticism by appeal to an epistemic principle that is independently supported and less susceptible to overgeneralization worries. Since our version employs a weaker premise about the nature of the existing disagreement, it goes some way toward addressing the first type of criticism mentioned above, concerning the plausibility of such assumptions. However, we shall not provide a sustained defense of that aspect of the challenge. Instead, our main aim is to focus on the second type of criticism, concerning the inference from facts about moral diversity to skeptical conclusions. What we shall primarily argue for is the hypothetical claim that *if* the weaker premise that our version of the challenge invokes is true, then realists are committed to the skeptical implications that are thought to undermine their position.

The epistemic principle we shall offer to support that claim is related to the attempts to characterize knowledge in modal terms, along the lines of the suggestions that safety or sensitivity are necessary for knowledge. However, our principle concentrates instead on the less discussed property called ‘adherence’. As

the ‘adherence requirement’ (as we shall call it) straightforwardly explains why disagreement raises skeptical worries, it deserves a more central place in the debate about disagreement and skepticism than it is commonly given.

## 1. Preliminaries and Plan

There are some aspects of the disagreement challenge that we shall not discuss. One of them is whether there are metaethical positions besides realism to which it is equally applicable. We shall ignore that possibility and pursue the discussion on realist assumptions about moral truths, thought, and talk.

Another aspect of the challenge that we shall set aside concerns the step from the claim that moral realism rules out moral knowledge to the conclusion that it ought to be abandoned. We merely wish to note that the plausibility of that step depends on the scope and the modal strength of the type of skepticism to which realists are (allegedly) committed. Antirealists usually try to associate realism with skeptical conclusions that are both highly general and modally strong because tying the realist position to more restricted forms may not be enough to undermine it. The conclusions in question do not entail just that we in fact lack knowledge about *some* moral questions but also that it is not even possible for us to obtain *any* such knowledge, at least given the limitations we human beings have. Our version of the challenge is construed along the same lines. One can envision even stronger forms of skepticism, such as the view that moral knowledge cannot be obtained by any conceivable thinkers, regardless of their epistemic resources. That type of skepticism is stronger because we could imagine beings whose resources far exceed ours. Still, we take the relevant type of skepticism to be the one that focuses on human beings, as it is hard to see how facts about human disagreement could establish the stronger view.

The element of the disagreement challenge that we shall concentrate on is the step that (under realist assumptions) takes us from the claim that there are fundamental moral disagreements to the conclusion that the answers to the disputed questions are unknowable. That step can tentatively be represented by the following argument:

- A. There are disagreements about the truths that moral realists posit that would persist even if the parties were not influenced by distorting factors.
- B. If there are such disagreements, then we lack the ability to know those truths.
- C. Therefore: We lack the ability to know the truths that moral realists posit.

In this paper we offer a revised version of this crucial part of the challenge. The plan is as follows. We begin by addressing a problem with **A** that, in our view, is fatal to the argument given its present formulation. The problem concerns the notion of a distorting factor. While advocates of the disagreement challenge need an interpretation of that notion that validates both **A** and **B**, we argue in section 2 that no such interpretation is available. The reason is that the epistemic

assumptions that are needed to defend **B** generate a notion of a distorting factor that in turn excludes the truth of **A**. However, we shall also argue that we can reach the conclusion of the argument (**C**) on the basis of a claim about the existing diversity that avoids the problem since it is weaker than **A**. The weaker premise is formulated in section 3. In sections 4–6, we introduce the epistemic principle that our argument invokes to get from that weaker premise to the relevant skeptical conclusion and explain how it allows us to take that step. In section 7, we further defend the principle by considering some of its implications. We also explain how our version of the disagreement challenge is less vulnerable to certain common objections to arguments against realism, including the worries about overgeneralization mentioned above. In section 8, finally, we make some concluding remarks.

## 2. Distorting Factors

The question of whether there are any fundamental moral disagreements depends on what counts as a distorting factor. What makes such factors significant in this context is that somebody who does not suffer from them is arguably in the most favorable position, relative to the aim of uncovering moral truths, that a human being could ever hope to attain. Realism implies that at least one of the parties of any moral disagreement is in error and has thus failed to achieve moral knowledge. Accordingly, if there are fundamental disagreements, then realism implies that one could fail to achieve moral knowledge even in a position that is maximally favorable from an epistemic point of view. It is that implication that, ultimately, is supposed to commit realists to skepticism.

In accordance with these remarks, we may say that a factor is distorting to the extent that its presence or influence means that we are not in the most favorable position in relation to the aim of acquiring moral knowledge. This criterion accounts for the distorting nature of bad reasoning skills and ignorance of relevant nonmoral facts as well as for factors such as bias, lack of imagination, and so on. The status of some factors, such as lack of empathy, may be more controversial, since their epistemological relevance might depend on substantive ethical issues. But our aim is not to offer a comprehensive list of distorting factors, and precision at this point is not required for assessing the argument we shall develop.

There is one potential candidate of a distorting factor that we shall address, however, as it is needed to bring out the problem with **A**. The concept of a fundamental moral disagreement is construed so as to ensure that the relation between the parties is one of symmetry or parity. Since none of the parties is influenced by any distorting factors, none of them is in conditions that are inferior, relative to the aim of determining the truth of the disputed claim, to the conditions of their opponent. In other words, the parties are supposed to be *peers* in the sense that is used in the peer disagreement debate in epistemology.

According to one view that has emerged in the peer disagreement debate, when learning that one disagrees with a peer, one should withdraw one's initial verdict and become agnostic about the target claim (or at least reduce one's confidence in

it). That view is known as ‘conciliationism’ and is congenial with the intuitions that motivate the disagreement challenge (see Christensen [2007], Elga [2007], and Feldman [2006], though note that Elga prefers the term ‘the equal weight view’ for his position). Since conciliationism does not permit continued confidence in the face of peer disagreement, what it suggests, in effect, is that such steadfastness should be seen as a distorting factor (or as a manifestation of one). The opponent’s dissent is perhaps best seen as a piece of evidence that the subject has not properly acknowledged.

If continued confidence is a distorting factor, however, A faces a problem that seems devastating. For a direct result of that assumption is that, contrary to A, there are no fundamental moral disagreements. What the assumption rules out is not that people who disagree morally can be peers but rather that the parties of such a disagreement can be in the conditions that are most favorable for acquiring moral knowledge.

The problem can be brought out as follows. We have assumed that a disagreement is fundamental only if the parties would continue to disagree even if they were free from distorting factors. The existence of fundamental disagreement thus requires that it is possible that people disagree in such conditions. And that possibility is ruled out by the idea that retaining one’s belief in the face of peer disagreement is a distorting factor. For suppose that *a* is not subject to any distorting factors. If continued confidence in the face of peer disagreement is a distorting factor, then the claim that *a* is free from such factors but still disagrees with *b* implies that *b* is not her peer. Since that result in turn implies that *b* (unlike *a*) is subject to some distorting factor, it follows that their disagreement is not fundamental. Hence, if continued confidence in the face of peer disagreement is a distorting factor, people who are not subject to distorting factors could not disagree, which means that A is false.

This would be a minor problem if an advocate of the challenge could simply reject conciliationism and instead adopt a view that permits steadfastness in cases of peer disagreement (see Kelly [2005] and van Inwagen [1996] for two proponents of such views). However—and this is the second important point—the problem with that move is that it leaves advocates of the disagreement challenge with no explanation of why fundamental moral disagreements commit realists to skepticism. In other words, even if steadfast views can be used to save A from the problem just presented, they instead leave us with no way of substantiating B (the second premise of the argument above).

To see how the adoption of a steadfast view undermines B, recall that the disagreement challenge is supposed to show that even if realists are right about the existence of moral truths, we cannot achieve knowledge of them. The challenge thus relies on the idea that the moral diversity rules out the satisfiability of some other necessary condition for knowledge besides truth and belief, such as, for example, justification. Now, part of the reason why fundamental disagreements are thought to generate skeptical implications is that, given the symmetry that obtains between the parties, it is plausible to assume that either both parties have a justified belief or that none do. Advocates of the challenge need to reach the result that neither party has a justified belief so they have to exclude the possibility

that both of them do. And the problem is that adopting a steadfast view deprives them of grounds for doing so. For what such views imply is precisely that the symmetry that obtains in a peer disagreement does not exclude that both parties have justified beliefs.

The point may be illustrated by William Tolhurst's version of the disagreement challenge. Tolhurst's argument aims to show that the truths that moral realists posit could not be known, even if they exist, because no moral beliefs are justified. (The conclusion of Tolhurst's argument is that moral facts 'if any there be, are epistemically inaccessible' [1987: 611]; for a similar argument, see Schiffer [2002: 288].) His argument relies on the following principle:

If a person's belief that *j* is justified, then it is not possible for there to be another person who is similar in all epistemically relevant respects and who is justified in believing that not-*j*. (Tolhurst 1987: 611)

Tolhurst argues that if there could be moral disagreements between subjects who have true nonmoral beliefs and adequate reasoning skills, and do not suffer from other deficiencies—i.e., if fundamental moral disagreement is possible—then his principle (in conjunction with some further assumptions that need not concern us here) entails that no moral belief could be justified. What is noteworthy here is that Tolhurst's principle is in direct conflict with steadfast views about peer disagreement, as those views imply that conflicting beliefs can both be justified even if there are no relevant differences between the subjects.

A similar line of reasoning has been suggested by Derek Parfit, who writes that:

If we had strong reasons to believe that, even in ideal conditions, we and others would have deeply conflicting normative beliefs, it would be hard to defend the view that we have the intuitive ability to recognize some normative truths. We would have to believe that, when we disagree with others, it is only we who can recognize such truths. But if many other people, even in ideal conditions, could not recognize such truths, we could not rationally believe that we have this ability. How could *we* be so special? And if none of us could recognize such normative truths, we could not rationally believe that there *are* any such truths. (Parfit 2011: 546)

Parfit assumes that if our moral views are opposed by someone who is not in inferior epistemic circumstances, then our claim to have moral knowledge is undermined. That is why we could not rationally believe that there are moral truths in the face of such opposition. And, again, Parfit's assumption about the epistemic significance of disagreement is hard to reconcile with steadfast views.

The upshot is that advocates of the disagreement challenge face a dilemma. If continued belief in the face of peer disagreement is a distorting factor, then there are no fundamental moral disagreements—hence A is false. If such steadfastness is not a distorting factor because the competing verdicts in a peer disagreement may both be justified, then that invalidates the inference from the occurrence of

fundamental disagreement to moral skepticism—hence **B** is false. Either way the challenge is flawed.

### 3. Divergence and Convergence

In response to the dilemma, our strategy is to reconstruct the challenge so that it relies on a weaker assumption about moral diversity than **A**. On our version, what matters is not whether moral disagreement would persist if people were not influenced by distorting factors but whether people would reach *agreement* in such conditions. One may say that the important factor is the absence of convergence rather than the existence of divergence.

To explain how this assumption differs from the claim that there are fundamental moral disagreements we need to use a certain notion of agreement. On that notion, we agree about *P* if and only if either we both believe *P*, or we both reject *P*. If we both suspend judgment about *P* or have not even considered whether *P*, then that does not qualify as agreement. Hence, on our understanding, two agnostics (who suspend judgment about whether God exists) are not in theological agreement. Instead, they neither agree nor disagree.

We may now state two alternative assumptions about the existing diversity:

**The divergence premise.** There are moral disagreements such that if the parties were not influenced by any distorting factors, they would nevertheless disagree about the disputed claims.

**The lack of convergence premise.** There are moral disagreements such that if the parties were not influenced by any distorting factors, they would nevertheless fail to agree about the disputed claims.

The lack of convergence premise is weaker than the divergence premise (i.e., **A**). If *a* and *b* would continue to disagree about moral issues even if they were free from distorting factors, it follows that they would not agree about those issues in such conditions. But the converse does not hold. For *a* and *b* could fail to agree even in the absence of disagreement, as they could fail to agree by suspending judgment. This is why replacing the divergence premise with the lack of convergence premise avoids the dilemma faced by the initial construal of the challenge. Fundamental disagreements, which the divergence premise posits, are excluded by the view that belief in the face of peer disagreement is a distorting factor. But that view does not exclude that the parties may fail to agree about the target claim, precisely because they may fail by suspending judgment (which is just what conciliationism mandates).<sup>1</sup>

Before we continue the reconstruction of our alternative version of the disagreement challenge, we want to get a possible misunderstanding out of the

<sup>1</sup> We are not the first to articulate the disagreement challenge in terms of the lack of convergence (see, e.g., Williams 1985 and Wright 1992). However, previous discussions have not noted how the distinction between divergence and lack of convergence is needed to address the problems identified in section 2.

way. One explanation of why people do not agree about some issue may be that they simply do not care enough about it to try to get to the bottom of things. However, failures to reach moral agreement that are due to that factor are not sufficient to secure the truth of the lack of convergence premise as we understand it. Recall that a factor is distorting relative to some claim if the factor is an obstacle to the aim of uncovering the truth about that claim. In light of that criterion, insufficient interest will either count as a distorting factor itself or be accompanied by such factors, such as a failure to see where one's evidence points in the case at hand. If someone has considered all the pertinent arguments but fails to form a view because she does not care about the matter, then she could surely do better in relation to the aim of discovering the truth.

#### 4. Modal Principles about Knowledge

The shift from the divergence premise to the lack of convergence premise represents the first way in which our version of the disagreement challenge deviates from the initial version. The second deviation is the epistemic principle that we shall use to derive skeptical conclusions from that premise. Since the lack of convergence premise is weaker than **A**, our argument requires a premise that is stronger than **B** in order for the argument's skeptical conclusion (**C**) to follow. In the rest of the paper, we elaborate the epistemic principle that our challenge invokes to motivate that stronger claim, and we defend it by noting some of its advantages over other principles to which antirealists have appealed.

The intuition that inspires our principle is the familiar one that knowledge excludes certain kinds of coincidence (cf. Unger [1968], Zagzebski [1994], Pritchard [2005], Pritchard [2015]). The idea is that if we have a true belief about some matter, then that belief constitutes knowledge only if our having arrived at the pertinent truth is not due to a pure fluke or sheer chance. That is the intuition Gettier examples are often taken to illustrate. For instance, suppose that you form a true belief by noting a clock that, unbeknownst to you, is broken, but just so happens to show the correct time at the moment of observation. On the anticoincidence intuition, although your belief is true, it is not knowledge, since the match that obtains between your belief about the time and the facts is owed to a fluke—i.e., to the fact that you happened to look at the clock at the only moment when it correctly represents the time.

On one way of articulating the anticoincidence intuition, a true belief about some matter is knowledge only if you could not easily have been wrong about that matter. This is not to say that all possible scenarios in which you are wrong are relevant, of course. Suppose, for example, that you believe that a distant bird is a male because you saw it through an advanced telescope. The fact that you would have been mistaken if you had used an inferior telescope, or none at all, does not exclude that you have knowledge about the bird's sex.

Anticoincidence theorists accommodate that point by focusing only on scenarios in which the subject uses the same 'belief-forming method' as the one that she actually uses (Nozick [1981, 179]; see also Sosa [1999], Sosa [2002], Pritchard [2005], Williamson [2009]). The question is whether you could easily have been



wrong about the sex of the bird *given* that you used the advanced telescope. This understanding of the anticoincidence intuition has the plausible implication that not all coincidences exclude knowledge. Suppose that you are equipped with the telescope because you randomly met a keen ornithologist who lent it to you. Since you met her by accident, there is a sense in which you could easily have been wrong about the bird, due to the fact that you could easily have failed to be equipped with the advanced telescope. However, on the present understanding of the anticoincidence intuition, that coincidence does not rule out that you have acquired knowledge of the bird's sex by using the telescope, for in the scenario where you did not meet the ornithologist, you presumably did not use the same method for ascertaining the bird's sex.

The anticoincidence intuition underlies attempts to characterize knowledge in modal terms. According to Nozick's so-called tracking account of knowledge, for example, a true belief that *P* amounts to knowledge only if it is 'sensitive', where this means that the subject would not have believed *P* if *P* had been false (and the subject used the same method that she in fact used; Nozick [1981]; see also Roush [2005]). Others have made the related suggestion that the belief constitutes knowledge only if it is 'safe', where it is safe if and only if the subject would believe *P* through the relevant method only if *P* were true (i.e., if and only if there are no nearby worlds where she believes that *P* on the basis of the method that she actually used, but where *P* is false; see, e.g., Sosa [1999], Williamson [2000], and Pritchard [2005]).

Both these principles explain why we lack knowledge in Gettier cases such as the one with the broken clock. If you acquired a true belief about the time because you looked at that clock, then your belief is neither sensitive nor safe, as there are sufficiently nearby worlds where you look at the clock at a slightly different moment and thus hold the same belief even though it is false. It is important to note, however, that the sensitivity and safety principles at best capture one aspect of the anticoincidence intuition. To see how the intuition helps justify inferences from disagreement to skepticism one also needs to consider another aspect.

The sensitivity and safety principles focus on possible worlds in which one is mistaken because one holds the same belief as one actually holds even though the facts are different so that the belief is false. By contrast, what is highlighted by the presence of somebody who disagrees with you is that even if the facts were the *same*, your *beliefs* could have been different, because you could just as well have ended up with the beliefs of your opponent. Worlds in which the facts are the same while your beliefs are different represent another way in which you could have been mistaken and another way in which your actually being correct could be coincidental in a sense that rules out that you have knowledge.

This observation is accommodated by a further, less discussed, element in Nozick's tracking theory. According to Nozick, a true belief that *P* does not only have to be sensitive to constitute knowledge. The belief also has to be 'adherent', where it is adherent, roughly, if the subject believes *P* also in all nearby worlds where *P* is true and where she uses the method that she in fact used (Nozick 1981: 176). Suppose, for example, that we have formed a belief about some complex mathematical theorem by tossing a coin. On the idea in question, this belief does

not constitute knowledge, even if it is true, since there are nearby worlds in which the mathematical facts are the same but where the coin landed differently (so that we reached a different, incorrect verdict). It is a version of this ‘adherence requirement’ that we shall use to derive skeptical conclusions from the lack of convergence premise.

## 5. The Adherence Requirement

In presenting our version of this view, it will help to compare it with a modified version of the safety requirement that is due to Justin Clarke-Doane ([forthcoming](#)) and Duncan Pritchard (2008). Clarke-Doane and Pritchard recognize the significance of adherence by defining ‘safety’ in a more demanding way. On their definition, the safety of a belief that  $P$  requires both that there are no nearby worlds in which the subject believes that  $P$  (using the same method) while  $P$  is false *and* that there are no nearby worlds in which the subject has reached a false belief (using the same method) while  $P$  is true.<sup>2</sup> Clarke-Doane has also noted that safety understood in this way is relevant to the question of the epistemic significance of disagreement (Clarke-Doane 2017: 852). However, the version of the adherence requirement that we shall offer and our argument to the effect that it generates a skeptical conclusion differ from what Clarke-Doane has suggested. Before spelling out the argument (in section 6), we shall state and clarify our requirement and explain the differences.

On an alternative formulation of the anticoincidence intuition that motivates modal principles, knowledge requires not only that you are correct about the relevant matter but also that your being correct is, in a sense, *inevitable* (and so not dependent on flukes). However, given the truth of your belief, there are two different ways in which you could have failed to be correct: either by having had a different, false belief instead or by having no belief at all. Pritchard’s and Clarke-Doane’s modified safety requirement focuses only on the first possibility. By contrast, our version of the adherence requirement acknowledges the relevance of both ways in which one could fail to be correct, as a belief in  $P$  satisfies it only if the subject could not easily have *failed to have a true belief* in  $P$  (given that  $P$  is true). This feature of our requirement will turn out to be crucial when deriving skeptical conclusions from the lack of convergence premise.

A second important feature of our requirement concerns the role played by the concept of methods of belief-formation in modal principles. In our view, the term ‘method’ is not optimal in this context. Suppose that a person with perfect pitch hears somebody play an F# on the clarinet. This puts her in a position to know that the note is an F#, it seems, even if there are nearby worlds where she makes a mistake because her pitch is less than perfect. However, if the subject heard the

<sup>2</sup> See Pritchard (2008) and Clarke-Doane ([forthcoming](#)). Clarke-Doane’s account in addition involves a notion of ‘relevantly like’ propositions. On his account, ‘our belief that  $P$  is not safe if we could have easily had a false belief as to whether  $Q$ , where  $Q$  is any proposition relevantly like  $P$  [holding the method fixed]’ ([forthcoming](#): 156; Clarke-Doane refers here to Williamson 2000). As he notes, however, what counts as a relevantly like proposition is ‘surely vague and context sensitive’ ([forthcoming](#): 156). We agree with this judgment and thus consider it an advantage that our account need not rely on such a notion.

same note in both the actual world and in those worlds, it is strained to say that she uses different methods in them. We shall therefore invoke the concept of a subject's 'epistemic position' instead of that of her belief-forming methods. On our construal, if someone believes that  $P$ , then whether her belief is adherent depends on what goes on in nearby  $P$ -worlds in which the subject's epistemic position relative to  $P$  is at least as good as her position in the actual world.

What is an epistemic position? We take the features of a subject and her situation that characterize her epistemic position to fall into three main categories. First, a person's epistemic position includes relevant abilities or skills of different kinds, such as her reasoning skills, her perceptual abilities, and her 'intuition skills' (if such skills exist).<sup>3</sup> A second set of relevant aspects are features of her situation that influence her capacity to use those skills (either positively or negatively), such as the lighting conditions or the fact that she is subject to various biases. The third category includes the evidence (in the form of other beliefs) on which the inquirer may have based the belief.

Notably, an item can be part of the subject's epistemic position in relation to one belief even if it is not a part of her epistemic position in relation to another. Whether it is a part of her epistemic position relative to a given belief depends on the nature and subject matter of the belief. If the target belief (whose status is to be assessed) is inferential, for example, then her reasoning skills are part of her epistemic position, and so are the other beliefs from which she inferred it. If the belief is obtained solely through observation, by contrast, then the relevant aspects of her situation consist of her perceptual recognition capacities and the conditions that are relevant for their exercise.

The third category of items that define a subject's epistemic position raises special questions. If a subject's belief that  $P$  ( $BP$ ) is inferred from another belief ( $BQ$ ), then we have assumed that  $BQ$  is part of her epistemic position. It may seem that this assumption makes adherence come too cheaply in some cases, at least if the subject is a competent reasoner. For suppose that the subject reasons competently and that  $P$  follows logically from  $Q$ . This makes it hard to see how there could be nearby  $P$ -worlds in which she believes that  $Q$  and without being in an inferior epistemic position in any respect yet fails to believe that  $P$ . That is so, moreover, regardless of the status of  $BQ$  and independently of whether  $BQ$  has been acquired through tossing a coin, for example. Yet, it goes against the spirit of the adherence requirement to claim that the status of  $BQ$  is thus irrelevant to the status of beliefs that are based on  $BQ$ .

The solution to this problem is to postulate that if a belief is derived from other beliefs, then it is adherent only if those other beliefs are also adherent. That is, the

<sup>3</sup> To have 'reasoning skills' is to have the ability to reason in accordance with valid inference rules. Whether something counts as a reasoning skill in this sense depends on general issues about the status of logic. If there is such a thing as the 'one true logic', then the rules of that logic presumably belong to the relevant ones (though logical disagreement might undermine our claims to know which those rules are). If there are instead different logics that can in some sense be equally correct, things get more complicated. However, we take it that even supporters of that view want to accommodate the intuition that there is *some* important normative difference between denying the antecedent and reasoning in accordance with modus ponens, for example. Thus, there seems to be room for a notion of reasoning skills on such views too.

adherence of *BP* requires that the subject believes *P* in all nearby *P*-worlds in which her epistemic position is not inferior to the position she actually has. But if *BP* is inferred from *BQ*, then *BP*'s adherence also requires that she believes *Q* in all nearby *Q*-worlds in which her epistemic position (in relation to *Q*) is not inferior. And if *BQ* is in turn inferred from some other belief (*BR*), then *BP*'s adherence requires in addition that the subject believes *R* in all nearby *R*-worlds in which her epistemic position (in relation to *R*) is not inferior. And so on.

Ultimately, there might be a belief (*BX*) from which the subject has inferred all the others but which is itself noninferential. If so, then *BP*'s adherence requires that the subject believes *X* in all nearby *X*-worlds in which her epistemic position is not inferior to the one she has in the actual world, but where her epistemic position is characterized solely by her skills and conditions. If *BX* satisfies that condition, then we shall say that it is 'basically adherent'. *BP*'s adherence accordingly requires that *BX* is basically adherent.

What if *BP* is based on a circular chain of inferences? As it is doubtful if such chains can confer a positive epistemic status, we shall make the additional assumption that circular chains cannot ground a belief's adherence. That is, we assume that if a belief is based on other beliefs then its adherence requires that some of those other beliefs are basically adherent.

The points we have now made can be summarized as follows. *BX* is basically adherent if and only if the subject believes *X* in all nearby *X*-worlds where her skills and conditions are the same as in the actual world. Let us further say that when the belief *BR* is inferred from *BX*, it is 'directly adherently inferred' if and only if the subject believes *R* in all nearby *R*-worlds in which she believes *X* and has equally good skills and conditions as she actually has. Finally, let us say that *BQ* is 'indirectly adherently inferred' from *BX* if and only if there is a sequence of beliefs, *BQ*, *B1*, *B2*, . . . , *BN*, *BX*, such that each belief in the sequence is directly adherently inferred from the next belief in the sequence. We can now state our definition of adherence as follows:

**Adherence.** *S*'s belief that *P* is adherent if and only if it is basically adherent or adherently inferred, either directly or indirectly, from beliefs that are basically adherent.

What we want to suggest, however, is not just that a subject knows that *P* only if her belief that *P* is adherent in the indicated sense. The belief must also be the result, in a causal sense, of the pertinent features of her epistemic position. This represents another difference between our principle and Clarke-Doane's and Pritchard's modified safety requirement.

Adding the causal condition is not just an ad hoc move. On the contrary, it is motivated by the same considerations that underlie the need to appeal to methods of belief formation (or, in our case, to epistemic positions) in the first place. Possible worlds in which the subject uses other methods are irrelevant to the evaluation of her belief in *P* because the important question is whether she has cognitive resources *in the actual world* that allow her reliably to determine whether *P* is true. Now, suppose that those resources played no causal role in the

formation of the belief but that her belief is nevertheless held in all nearby *P*-worlds due to other factors. We may imagine, for example, that the subject accepts some complicated mathematical theorem because of wishful thinking that is so deeply rooted that it ensures that she has the same belief in all nearby worlds. Then the modal robustness of the belief is immaterial to whether her cognitive resources have the relevant features. That is why our version of the adherence requirement includes an explanatory component.<sup>4</sup>

Accordingly, the principle we want to propose can be stated as follows:

**The Adherence Requirement.** A subject knows that *P* only if her belief that *P* is adherent and is the causal result of features of her epistemic position.

Note that given our definition of adherence, this requirement assigns a special status to noninferential beliefs, in that inferred beliefs constitute knowledge only if they are inferred from basically adherent (noninferential) beliefs. This feature might seem to make the position incompatible with certain internalist views, such as versions of coherentism that entail that no belief has such a special status. However, to the extent that there is such a conflict, we think that it simply manifests the difficulties that coherentist views have in accounting for the anticoincidence intuition that we have focused on (see section 4). Consider, for example, Hartry Field's example of a person who happens to have true beliefs about the goings-on in some remote village in Nepal even though she has never visited the village or been in contact with anyone who has (Field 1989: 26). The example is meant to illustrate that although a subject's beliefs about some matter are true, the correspondence between her beliefs and the facts that thus obtains might be purely coincidental in the sense that we have taken to be relevant as to whether she has knowledge. This impression remains, it seems, even if her beliefs hang together through internal coherence relations. Thus, we think that there is already independent reason to be suspicious of the relevant forms of coherentism in this context.

## 6. The Revised Argument

We turn now to the task of explaining how our version of the adherence requirement allows us to derive skeptical conclusions from the lack of convergence premise (given realist assumptions about moral truths). In the next section, we elaborate the advantages of our principle compared to other principles that antirealists have invoked, and that completes our defense of the version of the disagreement challenge that we are proposing.

<sup>4</sup> An anonymous referee has suggested that one may explain why wishful thinking does not generate knowledge by arguing instead that beliefs that are caused in that way are not even defeasibly justified. However, in our view, that alternative explanation is at best incomplete. For example, suppose, for the sake of argument, that a subject is defeasibly justified in believing that *P* if it seems to her that *P*, other things being equal (cf. Huemer 2007). If the subject's 'seeming' is in turn explained by wishful thinking, then this plausibly undermines her claim to have knowledge even given that her belief is defeasibly justified and that there are no nearby worlds in which she believes something else.

Suppose that *a* and *b* do not agree about *P*. This means that they cannot both have a true belief about *P*. If *a* and *b* are also in equally good epistemic positions in relation to *P*, then the fact that they fail to agree illustrates that neither of them has an epistemic position that is good enough to ensure that they could not easily have failed to be correct. For even if one of them in fact has a correct belief about *P*, there is also one who has failed to acquire such a belief. In other words, if *a* believes that *P* and *P* is true, then the fact that *b* does not accept *P* shows that *a*'s epistemic position nevertheless fails to ensure the adherence of her belief. In this sense, it is just a coincidence that she has a true belief about the matter, rather than being less fortunate as *b* was. If we instead assume that *P* is false and that *b* believes that *P* is false, the same reasoning applies to *b*. The fact that they cannot reach agreement shows that their epistemic capacities are simply not good enough to allow them to determine, in a robust way, whether *P* is true. Accordingly, on the adherence requirement, neither of them has knowledge about *P*.

This general connection between adherence and failure to agree shows why the lack of convergence premise, via the adherence requirement, saddles moral realism with skeptical implications. Recall the formulation of that premise:

**The lack of convergence premise.** There are moral disagreements such that if the parties were not influenced by any distorting factors, they would nevertheless fail to agree about the disputed claims.

This claim entails that even if some of us would arrive at true beliefs regarding the disputed moral issues, at least some of us would not in spite of the fact that none of us were subject to any distorting factors. In such circumstances, we would all have equally good skills and be in equally good conditions as concerns our ability to utilize those skills. What the absence of convergence thus shows is that even in such circumstances our epistemic positions would not be good enough to generate moral beliefs that are robustly connected to the facts they concern—that is, not connected in the way the adherence requirement requires. This is sufficient to conclude that we would not have knowledge about the disputed moral issues even if we were in the most favorable epistemic conditions we could ever hope to be in.

Note moreover that if we—that is, again, all of us human beings—would not have adherent moral beliefs even in such optimal epistemic positions, we cannot have adherent moral beliefs from suboptimal epistemic positions either. We adherently believe that *P* only if we could not easily have failed to be right about *P*, given that our epistemic position is at least as good as it in fact is; that is, equally good or better. Adherence failure in optimal conditions thus entails adherence failure in suboptimal conditions as well. That is why the adherence requirement and the lack of convergence premise together generate the stronger skeptical conclusion that moral facts are not only unknown but also unknowable for us.

## 7. The Adherence Requirement: Some Implications

By invoking the adherence requirement, our version of the disagreement challenge offers a straightforward explanation of the inference from disagreement to

skepticism. In this regard it differs importantly from other versions that either offer no account at all of the relevant step or accounts that seem to us to be problematic. The key feature of our version is that although the adherence requirement generates verdicts that are congenial with conciliationism and related ideas (such as Tolhurst's principle), the role it assigns to disagreement is indirect. On our view, moral disagreements have skeptical implications only through being a *symptom* of the fact that the competing verdicts are not adherent. Ultimately, it is the absence of adherence that explains why the target beliefs do not constitute knowledge, even if one of them happens to be true.

One of the advantages that this gives our argument compared to alternative construals is that it answers the worry that by appealing to disagreement in support of skepticism, we commit ourselves to thinking that we could acquire knowledge by eradicating the disagreement by whatever means. For example, Thomas Kelly notes that

whether there is any actual disagreement with respect to some question [. . .] might, in a particular case, be an extremely contingent and fragile matter. In particular, whether there is any actual disagreement might very well depend on factors that everyone will immediately recognize as irrelevant to the truth of the question at issue. (Kelly 2005: 181)

Kelly continues by asking us to imagine that 'there would be considerable disagreement with respect to some issue, but that all of the would-be dissenters have been put to death by an evil and intolerant tyrant' (Kelly 2005: 181), and he insists that if the absence of disagreement is due to such factors, this hardly provides any reassurance for a nonskeptic (see Tersman [2012] for a discussion).

The adherence requirement accommodates Kelly's plausible verdicts about the case in question. For, as we just stressed, even if the requirement implies that disagreement between peers is sufficient to conclude that neither party has knowledge, the disagreement itself is not the problem. The problem is the underlying fact that is revealed by the disagreement, namely, that the conflicting beliefs are not adherent. One cannot eradicate that fact by killing the opposition, just as one cannot eradicate the solubility of a substance by removing all the liquids in the vicinity.

This in turn allows us to address another common worry about arguments from moral disagreement, namely, the fact that besides those moral views that are disputed, there are also those that are shared virtually universally (cf. McGrath [2008]). This is not a decisive objection to our version. For the fact that some moral beliefs are undisputed does not ensure that they satisfy the adherence requirement, not even if they are caused by processes that ensure that they are unanimously held also in nearby possible worlds. The reason is that to satisfy the requirement, a belief must also be the result of some features of the subject's epistemic position that ensure its adherence. Otherwise its status could be challenged on exactly the same grounds as the belief in the mathematical theorem that is based on wishful thinking (see section 5). Hence, the objector must insist that the agreement obtains *because* some features of our epistemic position

ensure that we adherently believe in the relevant moral claims. Whether that explanatory hypothesis is plausible depends on further issues we have not addressed here but discuss elsewhere (Risberg and Tersman, *n.d.*). Our point in the present context is just that the existing patches of moral agreement do not in their own right pose a fatal problem for our argument.

Finally, let us turn to the worry that the disagreement challenge overgeneralizes by committing its advocates to implausible forms of skepticism or antirealism about other domains besides morality, such as the empirical sciences, mathematics, modality, and metaethics itself. Appealing to those areas as ‘partners in innocence’ is a common defensive strategy among realists, both in response to the disagreement challenge and to other objections. Consider, for example, the idea that knowledge requires a causal connection with the target facts, which has sometimes been used to argue that nonnaturalist moral realism excludes the existence of moral knowledge for purely a priori reasons. A popular reply is that those arguments fail because the relevant epistemological principle also rules out mathematical and modal knowledge (see, Clarke-Doane [2014: sect. 3] and Olson [2014: ch. 5.2]).

An especially striking version of the overgeneralization objection seeks to show that the epistemic principles that are invoked are so strong that they exclude knowledge across the board. That problem arises for the argument that takes moral disagreement to generate moral skepticism because the disagreement cannot be resolved by providing a justification that is independent of the beliefs in dispute (cf. Bennigson [1996]). The implicit assumption here is that knowledge requires the availability of a justification of the beliefs in the target area that does not ‘beg the question’ against a skeptic by invoking the contents of some of the beliefs whose status is challenged. As is often pointed out, though, that view does not only threaten to rule out knowledge in ethics but knowledge quite generally (see Vavova 2014: sect. 3).

The adherence requirement, by contrast, is more permissive, as the adherence of our beliefs in a domain does not require that they can be justified independently. More precisely, there may be domains where the skills we have for uncovering the truths in that domain ensure that our beliefs satisfy the adherence requirement even though the conviction that we have such resources can be justified only with reference to the contents of our beliefs in that domain. The absence of a non-question-begging justification may thus coexist with the adherence of the beliefs in question.

In addition, the adherence requirement does not entail that knowledge requires some form of causal contact. What it entails is just that a true mathematical belief constitutes knowledge only if the subject’s epistemic position relative to that issue ensures that she could not easily have failed to have that belief (unless her epistemic position was worse). Whether that is so is not settled by the assumption that mathematical facts are causally inert.

The question of whether the adherence requirement generates skeptical conclusions about ethics is accordingly independent of whether it also leads to such conclusions in other cases. That is a distinct and, so far, open question, one that we can settle only by investigating our beliefs in those other areas (for example, by examining the disagreements that occur there). Of course, this does



not provide strong reason to *deny* that the requirement leads to skepticism or antirealism about mathematics, modality, or the like. What we want to stress is just that unlike some other epistemic principles, the adherence requirement does not automatically *commit* one to such verdicts about those disciplines.

## 8. Concluding Remarks

Our aim in this paper has been to address a specific step of the disagreement challenge, one that goes from certain suppositions about moral divergence to the conclusion that realism leaves no room for moral knowledge. In particular, what we hope to have shown is that the adherence requirement helps to vindicate that step and thus generates a novel reconstruction of the challenge for moral realism.

On the adherence requirement, roughly, we have knowledge in an area only if our cognitive resources for discovering the facts in that area (be it evidence, inferential skills, our perceptual faculties, our abilities to grasp facts through intuition, or whatever) have securely steered us toward those facts. The important epistemological question is thus whether the resources we have in the moral domain are such that we could not easily have failed to be right unless we were less well equipped. And the point is that if people would fail to reach moral agreement even if they were in the most favorable position that they could attain, that fact reveals that we lack such resources in the case of ethics.

Like other modal requirements on knowledge, the adherence requirement raises a number of questions about how to delimit the set of possible worlds that are relevant to the epistemic status of a belief. For example, how is the concept of a subject's epistemic position to be characterized in more detail? And what features of a subject's situation, more specifically, count as parts of her epistemic position in relation to moral issues? Similarly, there are questions remaining about the other elements of the disagreement challenge, not least concerning the plausibility of the lack of convergence premise. While those questions call for more work, they do not undermine our proposal that the adherence requirement provides a promising explanation of why disagreement sometimes leads to skepticism.

OLLE RISBERG

UPPSALA UNIVERSITY

[OLLE.RISBERG@FILOSOFI.UU.SE](mailto:OLLE.RISBERG@FILOSOFI.UU.SE)

FOLKE TERSMAN 

UPPSALA UNIVERSITY

[FOLKE.TERSMAN@FILOSOFI.UU.SE](mailto:FOLKE.TERSMAN@FILOSOFI.UU.SE)

## References

- Bennigson, T. (1996) 'Irresolvable Disagreement and the Case Against Moral Realism'. *Southern Journal of Philosophy*, 34, 411–37.
- Brink, D. O. (1989) *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.

- Christensen, D. (2007) 'Epistemology of Disagreement: The Good News'. *The Philosophical Review*, 116, 187–217.
- Clarke-Doane, J. (2014) 'Moral Epistemology: The Mathematics Analogy'. *Noûs*, 48, 238–55.
- Clarke-Doane, J. (2017) 'Objectivity and Reliability'. *Canadian Journal of Philosophy*, 47, 841–55.
- Clarke-Doane, J. (forthcoming) *Morality and Mathematics*. Oxford: Oxford University Press.
- Cuneo, T. (2007) *The Normative Web*. Oxford: Oxford University Press.
- Demetriou, D. (2014) 'What Should Realists Say About Honor Cultures?' *Ethical Theory and Moral Practice*, 17, 893–911.
- Doris, J., and A. Plakias. (2008) 'How to Argue About Disagreement: Evaluative Diversity and Moral Realism'. In W. Sinnott-Armstrong (ed.), *Moral Psychology*, vol. 2, *The Cognitive Science of Morality: Intuition and Diversity* (Cambridge, MA: MIT Press), 303–31.
- Elga, A. (2007) 'Reflection and Disagreement'. *Noûs*, 41, 478–502.
- Enoch, D. (2011) *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Feldman, R. (2006) 'Epistemological Puzzles about Disagreement'. In S. Hetherington (ed.), *Epistemology Futures* (Oxford: Oxford University Press), 216–36.
- Field, H. (1989) *Realism, Mathematics, and Modality*. Oxford: Blackwell.
- Fraser, B., and M. Hauser. (2010) 'The Argument from Disagreement and the Role of Cross-Cultural Empirical Data'. *Mind and Language*, 25, 541–60.
- Huemer, M. (2007) 'Compassionate Phenomenal Conservatism'. *Philosophy and Phenomenological Research*, 74, 30–55.
- Kelly, T. (2005) 'The Epistemic Significance of Disagreement'. In J. Hawthorne and T. Gendler (eds.), *Oxford Studies in Epistemology*, vol. 1 (Oxford: Oxford University Press), 167–96.
- McGrath, S. (2008) 'Moral Disagreement and Moral Expertise'. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics 4* (Oxford: Oxford University Press), 87–108.
- Nisbett, R., and D. Cohen. (1996) *Culture of Honor: The Psychology of Violence in the South*. Boulder, CO: Westview Press.
- Nozick, R. (1981) *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Olson, J. (2014) *Moral Error Theory*. Oxford: Oxford University Press.
- Parfit, D. (2011) *On What Matters*. Oxford: Oxford University Press.
- Pritchard, D. (2005) *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, D. (2008) 'Safety-Based Epistemology: Whither Now?' *Journal of Philosophical Research*, 34, 33–45.
- Pritchard, D. (2015) 'Anti-luck Epistemology and the Gettier Problem'. *Philosophical Studies*, 172, 93–111.
- Risberg, O., and F. Tersman. (n.d.) 'From Evolutionary Theory to Moral Skepticism, via Disagreement'. Unpublished ms.
- Roush, S. (2005) *Tracking Truth: Knowledge, Evidence and Science*. Oxford: Oxford University Press.
- Schiffer, S. (2002) 'Moral Realism and Indeterminacy'. *Philosophical Papers*, 12, 286–304.
- Shafer-Landau, R. (2003) *Moral Realism: A Defence*. Oxford: Oxford University Press.
- Shafer-Landau, R. (2006) 'Ethics as Philosophy: A Defense of Ethical Nonnaturalism'. In T. Horgan and M. Timmons (eds.), *Metaethics After Moore* (Oxford: Oxford University Press), 54–62.
- Shafer-Landau, R. (2012) 'Evolutionary Debunking, Moral Realism and Moral Knowledge'. *Journal of Ethics & Social Philosophy*, 7, 1–37.
- Sosa, E. (1999) 'How Must Knowledge be Modally Related to What is Known?' *Philosophical Topics*, 26, 373–84.
- Sosa, E. (2002) 'Tracking, Competence and Knowledge'. In P. Moser (ed.), *Oxford Handbook of Epistemology* (Oxford: Oxford University Press), 264–87.
- Tersman, F. (2006) *Moral Disagreement*. New York: Cambridge University Press.
- Tersman, F. (2012) 'Moral Disagreement: Actual vs Possible'. In D. Machuca (ed.), *Disagreement and Skepticism* (London: Routledge), 90–108.
- Tolhurst, W. (1987) 'The Argument from Moral Disagreement'. *Ethics*, 97, 610–21.
- Unger, P. (1968) 'An Analysis of Factual Knowledge'. *Journal of Philosophy*, 65, 157–70.

- Van Inwagen, P. (1996) 'It is Wrong Everywhere, Always, and for Anyone, to Believe Anything upon Insufficient Evidence'. In J. Jordan and D. Howard-Snyder (eds.), *Faith, Freedom, and Rationality* (London: Rowman & Littlefield), 137–53.
- Vavova, K. (2014) 'Debunking Evolutionary Debunking'. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 9 (Oxford: Oxford University Press), 76–101.
- Williams, B. (1985) *Ethics and the Limits of Philosophy*. London: Fontana.
- Williamson, T. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2009). 'Replies to Critics'. In P. Greenough and D. Pritchard (eds.), *Williamson on Knowledge* (New York: Oxford University Press), 279–384.
- Wright, C. (1992) *Truth and Objectivity*. Cambridge, MA: Harvard University Press.
- Zagzebski, L. (1994) 'The Inescapability of Gettier Problems'. *The Philosophical Quarterly*, 44, 65–73.