



<http://www.diva-portal.org>

Preprint

This is the submitted version of a chapter published in *Disagreement and Skepticism*.

Citation for the original published chapter:

Tersman, F. (2012)

Moral Disagreement: Actual vs Possible

In: Diego Machucha (ed.), *Disagreement and Skepticism* (pp. 90-108). London:

Routledge

Routledge Studies in Contemporary Philosophy

N.B. When citing this work, cite the original published chapter.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-188788>

This is a pre-peer-reviewed version of the article. It is published in final form in D. Machuca (ed.), *Disagreement and Skepticism*, London: Routledge, 2012 (pp. 90-108).

MORAL DISAGREEMENT: ACTUAL VS POSSIBLE

FOLKE TERSMAN

1. Introduction

Suppose that the disagreement over an issue is such that it justifies a skeptical response of some sort, such as the suspension of judgment regarding how to answer it. If so, had the mere possibility of the pertinent disagreement been enough to justify the same response? This question has recently been the object of some controversy. In his paper “The Epistemic Significance of Disagreement,” Thomas Kelly writes:

[I]t is extremely implausible that *actual* disagreement is always more epistemically significant than certain kinds of *merely possible* disagreement. After all, whether there is any actual disagreement with respect to some question as opposed to merely possible disagreement 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)

For example, Kelly continues, suppose 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” (2005, 181). Surely, if the absence of disagreement is due to such facts, it does not provide any reassurance for a non-skeptic.

In a paper that explicitly addresses Kelly’s position, Hilary Kornblith expresses a different view. According to him, the actual existence of dissenters, as contrasted with the mere possibility of opponents, *does* provide a special worry:

When I find that others disagree with me on a certain question, this gives me, *ceteris paribus*, reason to be less confident than I was that I am right. ... It is clear, moreover, that the mere possibility that someone might disagree with me does not have the same epistemic significance. Indeed, if it did, then since there might always be people who disagree with one on any question at all, treating merely possible disagreement as on a par with actual disagreement would result in total skepticism. (Kornblith 2010, 34)

Kelly shows some caution in stating his position. What he denies is that actual disagreement is *always* more threatening than “certain kinds of” merely possible disagreement. Kornblith, by contrast, seems to focus on the possibility of disagreement in general, leaving it open whether the possibilities he refers to are the same as those that Kelly has in mind. So one might wonder about whether they really disagree. In any case, Kelly’s caution is wise, in my view. There are several different arguments aimed at showing that the existence of disagreement justifies skeptical conclusions. Maybe they have different implications for the question of whether the actual existence of disagreement is more “intellectually threatening” than the mere possibility of disagreement?

In the present paper, I’m going to illustrate that that is indeed the case. More specifically, I aim to show that one can accept a skeptical conclusion about an area on the basis of the actual existence of disagreement and still deny that the mere possibility of such disagreement warrants a similar conclusion. Given the qualifications Kelly makes in stating his view, this conclusion does not strictly speaking conflict with this thesis. My point is that there are skeptical arguments that appeal to the existence of disagreement such that, even assuming that they are successful, they don’t justify drawing similar conclusions from the mere possibility of disagreement. This doesn’t exclude that there are also disagreement-based skeptical arguments for which the same does *not* hold. Still, if correct, my conclusion sheds light on the relative importance of actual and merely possible disagreement, and to do that is

my main purpose. Notice that I shall particularly focus on ethics and moral disagreement, although much of what I'm going to say applies, I submit, equally well to other discourses.

I also have another aim, namely: to explore Kornblith's suggestion that, if one takes the possibility of disagreement to justify a skeptical conclusion, then one is committed to a universal or global form of skepticism. I'm going to challenge that suggestion, on the ground that there is an argument for skepticism about ethics that does appeal to the mere possibility of certain kinds of disagreement but that doesn't, arguably, have the implication in question. As in the case of Kelly's view and my first conclusion, it is not clear if my second conclusion conflicts with Kornblith's view, since the type of skepticism he considers is different from the type that is pertinent in the context of the argument that I shall examine. However, again, the conclusion does have relevance to the more general issue about the relative importance of actual as contrasted with merely possible disagreement.

The plan of the paper is as follows. In Section 2, I shall define the types of disagreement relevant to the arguments that I shall consider. In Section 3, I shall distinguish between the different versions of skepticism that the arguments are supposed to support. In Section 4, I explain why actual disagreement, in some contexts, is more epistemically significant than merely possible disagreement. Finally, in Section 5, I shall turn to the question of whether one can be a skeptic about a given area on the basis of the mere possibility of certain types of disagreement in that area and still avoid global skepticism.

2. Peer Disagreement and Radical Disagreement

The exchange between Kelly and Kornblith is located within the context of the currently intense debate about "peer disagreement." The debate concerns what is the appropriate response to (the recognition of) the fact that one's views are contested by someone who qualify as an "epistemic peer"—i.e., a person who is just as well (or at least not less well)

equipped as oneself from a cognitive point of view (not less smart, or impartial, or informed about the relevant evidence, and so on).¹ Some in the debate hold, while others deny, that the proper response to peer disagreement is to suspend judgment about the truth of one's initial belief, or at least to be less confident about it.

Philosophers have of course pondered the phenomenon of disagreement long before the recent discussion about peers, and arguments that appeal to disagreement have been pursued and developed in other contexts, such as for example in the metaethical discussion about the status and function of moral thoughts and convictions. Just as in the peer disagreement debate, the focus in these other contexts has been on disagreements that meet certain conditions. It is usually not disagreement in general that is supposed to have any skeptical implications, but disagreement between parties who are not bad at reasoning, ignorant of relevant evidence, or unreasonable in other easily detectable ways.

Indeed, the parties to the disagreements that are considered to be most relevant in the metaethical context are often required to meet stronger conditions than those pertinent in the peer disagreement debate. In the peer disagreement debate, the question is whether I have a reason to revise my initial belief due to disagreement with a person who isn't cognitively inferior to me. And although it is common to focus on cases where none of the parties is obviously the subject to some cognitive shortcoming, such as sheer stupidity, it is not clear that the arguments for thinking that our disagreement with a peer gives us reason to reconsider our position presuppose that there is, from an external perspective, no general reason to doubt our trustworthiness whatsoever. In the metaethical discussion, by contrast, the relevant arguments usually *do* presuppose that there is no such reason

¹ The fact that someone who is *less* well endowed than oneself with regard to such virtues hardly gives one even a *prima facie* reason to revise one's views. The question is how to react to her dissent when there is no reason to dismiss it on such grounds.

Disagreements that belong to the pertinent type may be called ‘radical’. A disagreement is radical in so far as none of the parties is subject to some independently specifiable cognitive shortcoming, where a cognitive shortcoming is a condition that may plausibly be viewed as diminishing an inquirer’s chances of reaching a correct conclusion about the disputed issue. By requiring that the shortcoming be “independently specifiable,” it is meant that one can determine both that it qualifies as a shortcoming (relative to the disputed issue), and that an inquirer is subject to it, independently of determining which view about the disputed issue the inquirer actually holds or whether it is correct. Typical examples of such shortcomings, in the context of a moral disagreement, are ignorance of relevant non-moral considerations, bias, inferential error, lack of imagination, and the like.

3. Varieties of Skepticism

In the peer disagreement debate, we are often asked to imagine *ourselves* being in disagreement with other people. In the case of the metaethical debate, the perspective is more detached. We are asked to ponder the fact that other people disagree over issues that we ourselves might suspend judgment about, and consider the implications of this phenomenon regarding the nature of the issues and of the parties’ answers to them. The reason for this difference, as well as for the fact that the metaethical debate focuses on a narrower set of disagreements, has to do with the nature of the skeptical conclusions that are sought. Although there is some overlap, there are also differences.

The fact that the type of skepticism explored in the peer disagreement debate differs from the ones that provide the focus of more traditional discussions is stressed by another of the philosophers who helped to initiate the debate, namely, Richard Feldman. The position discussed in the peer disagreement debate is, he writes, “a kind of contingent real-world skepticism” that is contrasted with the positions philosophers address when considering

“whether it is possible for us to know about the existence and nature of things in the external world, or whether we can know about the future, or whether we can know about the past” (2006, 217). When disagreement is discussed in metaethics, by contrast, it generally *is* claims of the latter kind that are at issue—whether moral truths, if any there are, can be known—in addition to even more fundamental claims about the status of moral convictions, such as ontological and semantical claims of various kinds.

An example of a meta-ethical position that is defended with reference to disagreement is John Mackie’s ‘error theory’, according to which our moral judgments, though being in the business of stating facts, are all false (at least in so far as they actually ascribe a moral property).² Another example is the non-cognitivist view according to which our moral convictions cannot be true as they consist in conative attitudes of some sort, such as emotions or desires, rather than true/false beliefs. These views may be labeled ‘anti-realist’, as they conflict with the position known as ‘moral realism’, which, roughly, is the conjunction of the claim that moral convictions (our views about what is morally right and wrong, permitted and obligatory, and so on) do consist in beliefs and the thesis that there are independently existing facts in virtue of which such convictions can be true.³ However, the views in question can also be labeled ‘skeptical’, in a wide sense, as they entail epistemological claims of the type that are more traditionally associated with skepticism, such as the view that there is no moral knowledge. Notice that the skeptical conclusion sought in the peer disagreement debate, by contrast, is usually not supposed to have any anti-realist implications but is taken to be compatible with the view that the disputed issues have uniquely correct answers, and that there is evidence that, in principle, would allow the parties to determine their truth.

² See Mackie (1977, ch. 1).

³ For some contemporary realists, see Boyd (1988), Brink (1989), and Jackson (1998).

4. Convergence and Reliability

My aim in the present section is to illustrate that there are arguments to the effect that the actual existence of disagreement in an area warrants skeptical conclusions about it but that don't justify drawing similar conclusions from the mere possibility of such disagreement. Relative to these arguments, actual disagreement is indeed more "intellectually threatening," from a philosophical and epistemological point of view, than disagreement which is merely possible. Note that although I find the arguments initially plausible, I'm not going to provide an elaborate defense of them here. I discuss their plausibility elsewhere (Tersman 2006). I shall start by reconstructing an anti-realist argument that is familiar in the metaethical debate, and then consider an argument pertinent to the peer disagreement case.

4.1 Disagreement and Convergence

There is a consideration that is mentioned in support of the view that actual disagreement is not more epistemically threatening than merely possible disagreement that may seem compelling but isn't really. This is the point Kelly makes about killing off would-be dissenters. A similar point has been made by Markus Lammenranta:

Furthermore, is it well motivated to restrict the scope of serious disagreements to the actual ones? This would mean that we could gain knowledge by killing our opponents. Then there would no longer be actual disagreements. (Lammenranta 2011, 211)

If the thesis that actual disagreement poses an extra threat were to imply that the killing of opponents would by itself increase the level of justification by which we hold our beliefs, then that would be a reason to be worried about it. However, the thesis in question does not in fact presuppose that view, and there are arguments for it that do not entail such implications, as I shall now show.

Notice that some of the most prominent metaethical arguments that appeal to disagreement are best construed as a sort of indirect challenges. According to these arguments, the point is not that the existence of disagreement is somehow outright inconsistent with a realist view, but rather that it undermines arguments that may be offered in its support. An example is Mackie's "argument from relativity," a version of the argument from disagreement that has provided the source of inspiration of much of the contemporary discussion about disagreement in ethics.⁴

Those who appeal to disagreement in support of moral anti-realism are sometimes criticized on the ground that they thereby commit themselves to anti-realism also about the sciences, since there is plenty of disagreement there as well. In response to this claim, Mackie writes:

[T]he argument from relativity has some force simply because the actual variations in the moral codes are more readily explained by the hypothesis that they reflect ways of life than by the hypothesis that they express perceptions, most of them seriously inadequate and badly distorted, of objective values. (Mackie 1977, 37)

What Mackie does in this quote is, in effect, to offer a rudimentary theory about the genealogy of our moral convictions. According to this theory (which, incidentally, is logically independent of his anti-realism), we have the convictions that we have because they "reflect our ways of life," where this means, roughly, that they rationalize our way of life or perhaps the life we want to live. The point is not that this theory is very informative. The point is rather that it is neutral about the correctness of our moral convictions and therefore doesn't assume any moral facts. Thus, it provides support for the claim that is the central premise of

⁴ The argument is developed in Mackie (1977, ch. 1).

Gilbert Harman's well-known challenge to moral realism: the claim that the best explanation of our moral convictions does not entail that any of them is true.⁵

The significance of the (alleged) truth of this claim can be brought out as follows. If correct, it deprives realists of a crucial pro-argument—that which implies that we *do* have reason to believe in the facts posited by the realist because they *are* entailed by the best explanation of our moral thoughts. Of course, the (alleged) fact that *that* particular argument fails does not exclude that realism is true (that is what makes the argument indirect). Nor does it exclude that realism is justified, as there may be other arguments. But, unlike some other arguments, it would give realism crucial *independent* support, as the fact that we have moral convictions is ultimately an observable fact that can be established independently of what view we take on their status. And in the absence of other compelling pro-arguments for realism, the conclusion that this one fails does warrant a skeptical attitude towards realism.

In other words, given Mackie's way-of-life theory, the existence of our moral convictions provides no support for moral realism. Now, if Mackie is right in thinking that his way-of-life theory provides a better explanation of the "actual variation in moral codes" than theories that invoke realist assumptions (such as those that invoke the idea that our moral views are "perceptions of objective values"), then it obtains support from the diversity. Thus, the diversity gives us reason to accept the theory. This, in conjunction with the fact that the theory deprives the realist of a crucial pro-argument, is why the diversity poses a challenge to moral realism.

However, I am going to focus on an argument that constitutes a somewhat less complicated variation of the same theme. Its point of departure is the idea that one reason for being a realist or non-skeptic about an area is the existence of convergence and consensus

⁵ Harman's argument is developed in Harman (1977, ch. 1). For further discussion of this relation between Mackie's and Harman's arguments, see Tersman (2006, ch. 3).

among its practitioners. Such convergence is supposed to support a non-skeptical view at least in so far as it is the result of the fact that inquirers have gotten better at assessing the initially disputed claims, through, say, gathering more and better data, finding ways of articulating the target claims in a way that allows for more rigorous testing, and so on. More specifically, the idea is that such convergence is what a non-skeptical view predicts, and that it therefore helps to explain the tendency towards convergence in the area in question. This is, at least in part, why some philosophers have a realist and non-skeptical view about the natural sciences.

Given the cogency of the above reasoning, and given the extensive existence of radical moral disagreement, we may (again) argue that a non-skeptic about ethics is deprived of a crucial pro-argument. For, given the existence of such disagreement, the pertinent convergence is absent. Moreover, since the fact that people disagree entails that the parties *have* views about the disputed issues, this absence cannot simply be blamed on a lack of interest in finding out how to answer them. Also, since the disagreement is radical, we cannot blame it on a general lack of cognitive competence such that we have reason to expect that, if it went away, there would be convergence.

Notice that it is not assumed by the argument that I have just outlined that the existence of convergence in an area somehow *guarantees* that the views on which one converges are true or that the best stance about the status of the views articulated within that area is a realist one. Obviously, consensus doesn't always indicate truth. In some cases, it is due to social pressure, lack of courage to dissent, and similar psychological and social mechanisms.⁶ If convergence is due to such mechanisms, or to the slaughter of would-be dissenters, it obviously doesn't provide evidence for the truth or reliability of the convictions

⁶ The relevant mechanisms have been explored by social psychologists such as M. E. Sherif (1936) and S. E. Asch (1952).

on which people tend to converge (regardless of whether the survivors of the slaughter satisfy the conditions associated with cognitive competence).

This explains why Kelly's and Lammenranta's appeals to evil tyrants who create consensus by killing off dissenters have limited relevance to the issue of whether the mere possibility of disagreement is just as epistemically threatening as the actuality of it. For although the existence of a tendency towards consensus among competent inquirers is a necessary condition for the pro-argument that appeals to convergence to apply, it is not sufficient. The point about evil tyrants—the claim that, if one assigns special weight to actual disagreement, then one is committed to the view that killing off dissenters is a way of “gaining knowledge”—presupposes that it *is* sufficient. However, the argument for skepticism that is presently under consideration—the argument to the effect that the existence of radical moral disagreement deprives a non-skeptic of a crucial pro-argument—only requires that it is necessary.

Now, the crucial question is whether, assuming that the skeptical argument just sketched is sound, one could just as well have appealed to the mere possibility of radical disagreement. In other words, imagine that there *is* in fact agreement of the pertinent kind—people tend to resolve their moral disagreements as an effect of liberating themselves from the influence of cognitive shortcomings. Imagine also that an objector insists that, in spite of the fact that there is this agreement, it could have been absent: there is a possible world in which there is no tendency towards consensus, even among cognitively competent persons.

Whether the possibility of such a world is a cause for skeptical worries clearly depends on what else goes on in it. Thus, suppose that the views on which people tend to converge in the actual world are in fact true. If these views are contested only in worlds where they are false, the existence of those possible worlds obviously do not invite any skeptical

conclusions.⁷ Moreover, even if the moral facts are the same in the world in question, it is unclear how the diversity in it is supposed to undermine the pro-realist argument that appeals to the actual existence of convergence, as it is entirely consistent with the assumptions it involves. After all, the argument does *not* assume that the (alleged) truth of realism logically *implies* the existence of agreement, only that it is part of the best *explanation* of it.

The mere fact that disagreement is *logically* possible does evidently not undermine the thesis that the best explanation of the actual convergence is a realist one. However, if the possible world in which the disagreement occurs is relatively close-by, it might. But then one wonders how the existence of such a possibility could be established. Of course, if we are skeptics, we are likely to think that there is such a possibility, whereas if we are anti-skeptics we are likely to deny it. But the only way to show that there is this possibility seems to be to provide some argument to the effect that realism does *not* provide the best explanation of the existing agreement. And then the appeal to disagreement really drops out of the picture and has no independent significance. The *actual* absence of disagreement, by contrast, can be independently established, which accounts for its crucial role in the argument for a skeptical view that is under consideration here. Or, differently put, the point of the skeptical objection outlined above is that the existence of disagreement deprives a realist of a crucial pro-argument. The actual existence of disagreement means that both the soundness of the premises and the validity of the pro-argument can be questioned. If the disagreement is merely possible, there is only the second option. This is why the mere possibility of radical moral disagreement is less threatening than actual disagreement of the pertinent kind.

4.2 Peer Disagreement and Reliability

⁷ Of course, some hold that moral truths are necessary, in which case this comment may be ignored.

Let's turn to the peer disagreement debate. One argument to the effect that the actuality of a dissenting peer need be no more epistemically threatening than the mere possibility of such a person goes as follows. Why is the existence of peer disagreement supposed to justify a skeptical response, such as the suspension of judgment? Well, for someone to count as my peer the same evidence must be available to both of us. So, if I note that a belief of mine is believed to be false by a peer, I may plausibly conclude that the evidence to the same extent support opposite verdicts on it—i.e., that it underdetermines the belief—especially as I am to assume that the peer's reasoning capabilities are not inferior to mine. But to say that the evidence underdetermines the belief in this sense is just to say that it leaves any verdict on it unfounded, which should undermine my confidence in it. However, if that is the basis of the challenge, then disagreement (again) drops out of the picture. For then it is the *underdetermination* that motivates skepticism, regardless of whether it manifests itself in the actual existence of a dissenting peer. So the point is not so much, it seems, that merely possible and actual disagreement are on a par when it comes to epistemic significance, as that neither type ultimately has any significance at all.

However, the fact that the occurrence of peer disagreement ultimately carries no weight given *one* argument to the effect that the existence of such disagreement justifies skeptical responses does not exclude that it plays a crucial role in *other* such arguments.

Thus, notice that we treat some of our beliefs as justified on the ground that they are the results of the exercise of a capacity or skill of some sort or of some process such as perception. In other words, their justification is supposed to be due to the alleged reliability of the pertinent process, where a belief-forming process is reliable, roughly, if it generally leads to true beliefs. If one is confident in the truth of one of one's belief on the basis of the way it is formed, then one's ground for the belief essentially consists in the mere fact that it is *held*. Now suppose that another person has formed the opposite belief. This provides no worry if

one can explain away the dissent in a way that leaves the ground for retaining the view that one's *own* verdict was reliably formed intact, for example by attributing it to some lack of training, or some other shortcoming that we don't have. However, the point is that, if we assume that the dissenter is a peer, then we can't provide such an explaining-away. In such a case, maintaining one's confidence in the initial belief on the mere ground that one holds it manifests a kind of epistemic narcissism. It is like deciding to trust the readings of one thermometer over the contradicting readings of another on the mere ground that the first belongs to oneself.

In other words, we may say that, when we trust a belief on the ground that it was reliably formed, it is the fact that we *have* it that is our evidence for it. If another person has formed the opposite view, we are committed to viewing that fact as counter-evidence, unless we can explain it away in the sense suggested above. If the person is a peer, this is something we cannot do, which means that we rationally should conclude that the total evidence is ambiguous, and that we accordingly have a reason to lessen our confidence in the belief. And the strength of that reason increases if the number of dissenting peers multiply, at least if they have formed their views independently of each other.

Clearly, the role in this argument that is served by the actual existence of a dissenting peer could *not* be played by the mere possibility of such a person. The reason is not that there is no requirement to explain away that possibility in a way consistent with the assumption that our initial belief is reliably formed. The reason is rather that this is so *easy*. We may simply appeal to the fact that the assumption that our belief is reliably formed does not entail that it is true, since a process could be reliable without *always* resulting in true beliefs.⁸

Of course, there are possibilities such that, if they could be established, they would undermine our belief. For example, suppose that it is true that, if a highly reliable person were

⁸ Basically the same point is made by David Christensen. See Christensen (2007, 208).

to consider the topic on which we have a view, he or she would form the opposing belief. This would give us a reason to reconsider the view. But the question is how to establish that possibility. There seems to be no way to do so independently of taking a stand about the topic at issue or of offering a direct argument against our view or our reliability. The existence of an actual dissenting peer, by contrast, *can* be independently established, and a skeptic that appeals to that dissent can *in addition* appeal to all the considerations that are available to a skeptic who is left only with the possibility of disagreement. I conclude that a skeptic can coherently appeal to the actual existence of disagreement without committing herself to the view that skeptical conclusions obtain equal support from the mere possibility of such disagreement.

5. Skepticism and the Possibility of Disagreement

I shall now turn to the second aim of this paper, namely, to explore Kornblith's claim that, if one takes the mere possibility of disagreement in an area to justify a skeptical conclusion, then total skepticism looms. The type of skepticism Kornblith focuses on is different from the one that I shall concentrate on. Kornblith considers whether disagreement with a peer about a certain issue is a reason to have less confidence in one's initial view about it. An affirmative answer to that question is, as I pointed out in section 3, compatible with the realist assumption that the view in question is a true or false belief and that there is a uniquely correct answer to the disputed issue. However, it is the latter assumption that is the target of the (anti-realist) form of skepticism that I shall consider.

More specifically, the question I shall address is this: can one accept anti-realism about some areas on the basis of the mere possibility of certain types of disagreement in those areas and still, coherently and plausibly, deny anti-realism about others? That this is so, and in

particular that an *ethical* anti-realist who appeals to the possibility of disagreement need not be a global anti-realist, is the upshot of an argument developed by Crispin Wright.⁹

Wright formulated this argument in the context of the more general aim to find meaningful topics for the debates about realism and anti-realism in discourses such as mathematics, ethics, and aesthetics. He wanted more specifically to identify factors (“cruxes”) that can help us to distinguish those discourses about which we should be realists from those about which we should not. According to one of the suggestions he pursues, realism about a discourse is defensible only if it exhibits “Cognitive Command,” where this means that it is *a priori* that every disagreement that arises in that area (except those that can be attributed to vagueness) involves a cognitive shortcoming of some sort, such as inferential error or ignorance of relevant evidence.¹⁰ Thus, to use the terminology in this paper, the idea is that realism is defensible only if it is *a priori* that no disagreement is radical.

In support of this view, Wright appeals to an analogy between our belief-forming processes and certain devices whose function is to produce representations, such as cameras. It is incontestable, he thinks, that if two cameras function properly, and if conditions are suitable, then “they will produce divergent output if and only if presented with divergent

⁹ The argument is developed in Wright (1992, esp. 88–94 and 140–68), but see also Wright (1994, 1996). The argument is further elaborated in Wright (2001). Notice that Wright’s views on the matter have undergone some changes and modifications since his 1992 book. For others who appeal to the mere possibility of disagreement in support of anti-realist views on ethics, see, e.g., Goldman (1988, esp. 153–60) and Tolhurst (1987).

¹⁰ For the definition of Cognitive Command, see Wright (1992, 144). For the suggestion that it is a necessary (although perhaps not sufficient) condition for realism about a discourse to be defensible that it exhibits Cognitive Command, see Wright (1992, 148, 162, 175). In accordance with this suggestion, he stresses that, even if all *actual* disagreements happen to involve cognitive shortcomings, this does not settle the matter in favor of realism (e.g., Wright 1996, 15).

input” (1992, 91). Analogously, unless every instance of a “divergent output” in *ethics* can be attributed to “divergent input” (differences in evidence), or to malfunction (e.g., inferential error), or to unsuitable conditions (e.g., distraction), we have reason to doubt that moral convictions are the products of “a seriously representational mode of function” (1992, 94; see also 1996, 13–4). And since Wright thinks that the condition is not in fact met, he denies that moral views are the products of such a mode (1996, 15).

Wright acknowledges that a realist might respond by saying that every moral disagreement *does* involve a shortcoming, since at least one of its parties is *incorrect*.¹¹ However, Wright also thinks that, if this is the *only* shortcoming a disagreement is supposed to involve, then realists commit themselves to the view that the truth about the disputed issue may “in principle outreach the efforts of an ordinarily receptive, careful ... thinker” (1996, 9).¹² And although that idea may be available in other areas, Wright is skeptical in the case of ethics. Wright thinks that a realist about an area may appeal to this idea only if he can explain “what it is about its subject matter that potentially makes it so” (1992, 152). And he thinks that the prospects of providing such an explanation in the case of ethics are bleak. This is why he wants to pursue the debate on the assumption that, in ethics, “evidence transcendence is simply not in view” (1992, 82).¹³

Wright’s argument can accordingly be summarized as a dilemma: realists must assume either that it is *a priori* that every moral disagreement involves some cognitive shortcoming (*other* than the alleged fact that one of its parties is in error about the disputed issue) or that moral truths are somehow “epistemically inaccessible,” as I shall put it. Since Wright believes that neither option is believable, he rejects realism. That is, he rejects the view that moral

¹¹ For this response, see Williamson (1994).

¹² For this reasoning, see Wright (1992, 150–7; 1994, 333–6).

¹³ A similar view is taken by many others: see, e.g., Bennigson (1996, 413).

judgments are the products of “a seriously representational mode of function,” which is a view that is presupposed by realism. The argument can be reconstructed as follows:

- (1) There can be radical moral disagreements.
- (2) Realism implies that moral issues allow for uniquely and objectively true answers.
- (3) If a dispute over a moral issue can be radical, and if the issue has a uniquely true answer, then its truth is epistemically inaccessible.
- (4) If moral issues allow for uniquely and objectively true answers, these are not inaccessible.

This is an attempted *reductio ad absurdum*. Given (2) and (3), (1) commits realists to the existence of things that, according to (4), do not exist. However, all premises need clarification.

For example, consider the *scope* of (1). Realists usually concede that there *are* radical moral disagreements, although they stress that these are “few.” Some also concede that the proper response to such a case is to think that there is no correct answer to the issue in dispute (rather than that there is such an answer, but that it is unknowable). This is not to give in to anti-realism, however. For realism does not, they stress, entail that *all* moral issues admit of true answers, only that *most* of them do, where those that don’t can be attributed to, for example, vagueness and indeterminacy.¹⁴ This offers them a way to reconcile their realism with the conclusion that some moral issues do not allow for uniquely correct answers.

Therefore, in assessing the argument, I shall ignore cases that can be attributed to vagueness and indeterminacy (as Wright’s formulation of Cognitive Command suggests).¹⁵ For there are surely moral issues such that, if there were radical disagreements about them, it

¹⁴ See, e.g., Brink (1989, 202) and Sturgeon (1994, 95). For similar suggestions, see Shafer-Landau (1994, 1995).

¹⁵ Which cases are those? Obviously, that is a tricky question, and I shall not try to provide a general answer. Instead, I shall approach it on a case-to-case basis.

would be implausible to attribute it to such factors. Besides, a realist cannot consistently hold that there are no moral issues that allow for objective and uniquely correct answers. So, if people may disagree radically about those issues, realists are still, given (3), committed to (4).

The argument just outlined may obviously be questioned in a number of ways, and there are many points that deserve closer attention.¹⁶ One set of questions concern the locution “epistemically inaccessible.” Wright (1992) uses many different expressions to indicate the view he thinks that radical moral disagreement commits a realist to: that moral truths “can transcend all possibility of human knowledge” (151), that they may “transcend, even in principle, our abilities of recognition” (8–9), and that they may “elude the appreciation even of the most fortunately situated judge” (58). However, all these phrases are ambiguous in a crucial way. The difference could be brought out by construing it as a difference about the scope of the operator ‘possibly’. Since none of the parties to a radical disagreement are subject to any cognitive shortcoming, they are, we may say, in an optimal cognitive position (relative to the disputed judgment). By saying that a given truth p is inaccessible, we could mean two different things:

(A) Possibly, it is not the case that everyone who is in an optimal cognitive position relative to p knows that p .

(B) It is not the case that, possibly, there is a person who is in an optimal cognitive position relative to p and knows that p .¹⁷

(B) is stronger than (A), in that (B) implies (A), but not conversely. Let us say that, if a truth is inaccessible in the first sense, it is “weakly transcendent,” while, if it is inaccessible in the

¹⁶ I discuss the argument in detail in Tersman (2006, ch. 4).

¹⁷ What type of modality is relevant here? The type assumed by an advocate of the argument.

second sense, it is “strongly transcendent.” This distinction yields two interpretations of (3) and (4), respectively, and thus two versions of the argument reconstructed above.

If “epistemically inaccessible” means strong transcendence, then (3) is implausible. The reason why realists are supposed to be committed to the idea that some moral truths are inaccessible is that they cannot otherwise account for disagreements among people who are not subject to some cognitive shortcoming. But this seems already explained by assuming that they are *weakly* transcendent. There is no reason to go further and suppose that they are also strongly transcendent. I shall accordingly assume that “epistemically inaccessible” refers to weak transcendence. I shall ignore many of the objections that can be raised against the argument, however, since, in the present context, the crucial question is whether the argument is equally applicable to other areas.

Whether this is so depends on whether there are any areas where the possibility of radical disagreement can be ruled out. Note that Wright’s approach presupposes that there *are* such areas (i.e., areas for which it is *a priori* that every disagreement that arises in them can be attributed to some shortcoming other than being in error about the disputed issue). In relation to the Cognitive Command constraint, there are just two ways to defend realism about an area. One is to deny the version of (4) that applies to the discourse in question, i.e., to deny that it is implausible to assume that there are truths in that area that are inaccessible. The second is to deny the possibility of radical disagreement. Wright seems to think that the first strategy is available in the case of, for example, physics (1994, 335). However, there must also, presumably, be discourses for which the second strategy is available. For, if not, then the divide between discourses that should be construed realistically and those that should not could just as well be drawn on the basis of the question about whether the area in question allows for “evidence-transcendent” truths alone.

Still, Wright doesn't offer any suggestions about how to establish that the possibility of radical disagreement in an area is ruled out. In the rest of this section, I shall consider one such suggestion, a proposal that represents, as far as I can see, the only available option.

The problem with positing epistemically inaccessible truths in the case of ethics, according to Wright, is that there is no believable explanation of *why* they are thus inaccessible. In what follows, I shall assume that he is right. To explain why the truth about a given topic is inaccessible is (given that "inaccessible" refers to weak transcendence) to explain why an inquirer can be in error about it in spite of not being subject to any cognitive shortcoming (i.e., such that her cognitive position cannot be improved). What I assume, then, is that such errors are indeed inexplicable.

In defense of the view that appealing to the possibility of disagreement in support of skepticism leads to total skepticism, Kornblith points out that "there might always be people who disagree with one on any question at all" (2010, 34). Although this may hold for any particular topic, there are arguably limits to how extensive our disagreement with someone else could be. The reason has to do with the implications of that issue for matters of meaning and interpretation. Davidson has vividly expressed the pertinent point:

If I suppose that you believe a cloud is passing before the sun, I suppose you have the right sort of pattern of beliefs to support that one belief, and these beliefs I assume you to have must, in order to do their supporting work, be enough like my beliefs ... If I am right in attributing the belief to you, then you must have a pattern of beliefs much like mine. (1984b, 200; see also 1984c, 168.)

This observation is a source of inspiration for the set of constraints on interpretation that Davidson refers to as "the principle of charity." Concluding that someone disagrees with us on a set of reasonably well-defined topics requires attributing thoughts to her with specific contents. To be justified in attributing such thoughts requires ascribing determinate contents

or meanings to the sentences that are accepted and rejected by that person. Davidson thinks that knowledge of the meanings of a speaker's sentences may be captured by a finitely axiomatized "theory of interpretation" for the speaker. This is a theory that, for each of the sentences of the target idiolect, specifies its truth conditions. The principle of charity entails that such a theory is correct only if it has the following feature: it assigns truth conditions to the speaker's sentences such that, by and large, the sentences the speaker holds to be true *are* in fact true, by our, the interpreters' lights. That is, it entails that such a theory is correct only if it represents the speaker as having beliefs that to a significant extent overlap with ours.

Note that the point of the principle of charity is not to somehow eliminate disagreement. It does not rule out finding that a speaker is occasionally irrational or in error. Rather, as Davidson has frequently stressed, the point is to "make meaningful disagreement possible" (1984d, 196). Thus, he writes:

It isn't that any one false belief necessarily destroys our ability to identify further beliefs, but that the intelligibility of such identifications must depend on a background of largely unmentioned and unquestioned true beliefs. To put it another way: the more things a believer is right about, the sharper his errors are. Too much mistake simply blurs the focus. (1984c, 168)

One possible reason for holding on to a theory in spite of the fact that it represents the speaker as holding beliefs that one thinks are false is that alternative theories may be overly complicated (1984d, 196). Suppose that we find that a speaker accepts a sentence *s* when and only when it's raining (nearby). Then we try for a theory of interpretation that implies that *s* is true (in the speaker's idiolect) if and only if it's raining (nearby). Now, suppose that we at some later point find that the speaker, on one occasion, accepts the sentence when it is *not* raining. In this case, we have a choice: either to revise the theory so that it ascribes a meaning that makes *s* true also on that occasion, or to hold on to the initial version. If we choose to

revise it we must revise its axioms, and if the revisions result in a more complicated version, this may, according to Davidson, justify holding on to the initial one. Consequently, it may also justify attributing what, by our lights, is an error on the pertinent occasion.

However, Davidson has stressed that such considerations ultimately *do* justify attributing errors only if we can somehow explain the errors, for example by noting, in the case of the mistake about the weather, that she had misleading perceptual evidence. On this view, to attribute true beliefs is the default position, and in order to attribute an erratic belief we need a special reason in the form of a believable account of how the interpretee has acquired it in spite of it being false. Consider the following passage:

If you see a ketch sailing by and your companion says, ‘Look at that handsome yawl’, you may be faced with a problem of interpretation. One natural possibility is that your friend has mistaken a ketch for a yawl, and has formed a false belief. But if his vision is good and his line of sight favourable it is even more plausible that he does not use the word ‘yawl’ quite as you do, and has made no mistake at all about the position of the jigger on the passing yacht. (1984d, 196)¹⁸

Davidson has not been alone in advancing views along this line. David Lewis, Dagfinn Føllesdal, Steven Lukes, and many others have made similar suggestions. Wittgenstein was, perhaps, another advocate, as the following quote suggests: “Can we say: a *mistake* doesn’t only have a cause, it also has a ground? I.e., roughly: when someone makes a mistake, this can be fitted into what he knows aright” (1979, 11e).¹⁹

¹⁸ In another passage, he writes: “When we find a difference inexplicable, that is, not due to ignorance or confusion, the difference is not genuine: put from the point of view of an interpreter, finding a difference inexplicable is a sign of bad interpretation” (2000, 25).

¹⁹ See also Føllesdal (1975, 40), Lukes (1982, 262), and Lewis (1983, 112).

In other words, suppose that we believe that P and that the speaker rejects a sentence *s* of her idiolect that we initially suppose is true if and only if P. To explain an error, in the relevant sense, is to show that, given the subject's particular perspective, as defined by her evidence and her particular set of cognitive capacities, it is in a way reasonable (or at least not unreasonable) to hold on to the erratic belief. Given that the belief is false, however, and given that the idea that its truth is radically underdetermined by the evidence is ruled out, we must assume that there is some flaw in the speaker's evidence or in her cognitive capacities that has led her wrong. In other words, we must assume that she is subject to some cognitive shortcoming. So, if we suppose that there is no shortcoming, the initial supposition about the meaning of *s* is undermined. Sentence *s* must therefore be reinterpreted so that the apparent disagreement comes out as *merely* apparent and so that the situation is not seen to represent a genuine conflict of beliefs. We may accordingly conclude, on the basis of the Davidsonian program, that the possibility of disagreements that are both radical and genuine *is* excluded in areas where underdetermination is for some reason not supposed to be an option (i.e., that it is *a priori* that every (genuine) disagreement involves cognitive shortcomings).

Of course, there are competitors to Davidson's views about interpretation and about how content and meaning are determined. Still, I'm going to grant that the strategy just outlined is persuasive and that it does exclude the possibility of (genuine) radical disagreement in the areas to which it is applied. However, the assumption that there are areas on which it is properly applicable does not by itself justify the kind of mixed verdict I seek—the view that one can coherently reject realism about an area on the basis of the possibility of radical disagreement in that area without committing oneself to global skepticism—since one must also show that there are areas on which it is *not* applicable

There are arguments to that effect. That is, there are arguments to the effect that the Davidsonian views about belief attribution and interpretation are not applicable to ethics, even

though that they are applicable to other areas. A theory of interpretation, in Davidson's sense, captures the meanings of the sentences in the target idiolect by specifying their truth conditions. Whether ethical sentences *have* truth conditions is, however, controversial. It is denied by non-cognitivists. If they are right, it is farfetched to apply the Davidson views to determine the contents of thoughts expressed by such sentences, which leaves room for the view that we may correctly attribute to a speaker a moral view with which we disagree even if we cannot blame her acceptance of that view on some cognitive shortcoming.

Of course, a problem with appealing to a non-cognitivist view on ethics in the context of an argument for moral anti-realism is of course that non-cognitivism *is* a form of anti-realism. However, personally, I think that compelling, *independent* reasons can be offered for not applying the Davidsonian views to ethics, even assuming that they are applicable to other thoughts.²⁰ But the point I presently want to make is just that to provide such reasons is what establishing the mixed verdict requires. This means that it is in the debate about the plausibility and applicability of such principles about belief attribution and meaning where the real argumentative action is going to be. More generally, too much discussion about the philosophical significance of disagreement is in my view pursued under the assumption that the nature of instances of the phenomenon could be identified in a rather unproblematic way. More attention should be devoted to the deeper philosophical views on which assumptions about the types of disagreement that exist in an area are implicitly based.

6. Concluding Remarks

²⁰ I offer such considerations in Tersman (2006, ch. 6), where I call the view that we may correctly attribute to a speaker a moral view with which we disagree, as contrasted with (some) other states, even if we cannot blame her acceptance of that view on some cognitive shortcoming “the latitude idea.”

I have reconstructed and defended a skeptical argument that generates the conclusion that certain types of actual disagreements are more epistemically threatening than the mere possibility of such disagreements. I have also suggested that there are skeptical arguments worth considering that appeal to the mere possibility of disagreement in an area without inevitably leading to global skepticism. As I noted earlier, it is not clear if these results conflict with the views that Thomas Kelly and Hilary Kornblith express in the passages I quote in the introduction, partly because it is not entirely clear which possibilities they are considering. In fact, in the case of Kelly, it is clear that his thesis does *not* conflict with my conclusion, since he is merely denying that actual disagreements of a given type are *always* more significant than merely possible ones, and since that denial is compatible with the view that they sometimes are. What this illustrates is that the question of the relative importance of actual and merely possible disagreement is not simple matter or either/or. On some skeptical arguments that appeal to disagreement, actual disagreements are more important. On other arguments, they might not be. What I have achieved is therefore, at best, just a partial answer to the issue that provides the focus of this paper. A full answer requires a more systematic examination of all the skeptical arguments that may be adduced.²¹

References

Asch, S. E. (1952). *Social Psychology*. New York: Prentice Hall.

Bennigson, T. (1996). "Irresolvable Disagreement and the Case against Moral Realism," *Southern Journal of Philosophy* 34: 411–37.

²¹ I want to thank Zed Adams, Markus Lammenranta, and Diego Machucha, whose insightful and penetrating comments have greatly helped me to improve the paper.

- Boyd, R. (1988). "How to Be a Moral Realist," 181–228 in G. Sayre-McCord (ed.), *Essays on Moral Realism*. Ithaca: Cornell University Press.
- Brink, D. O. (1989). *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Christensen, D. (2007). "Epistemology of Disagreement: The Good News", *Philosophical Review* 116: 187-217.
- Davidson, D. (1984a). *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press.
- . (1984b). "The Method of Truth in Metaphysics," 199–214 in Davidson (1984a).
- . (1984c). "Thought and Talk," 155–70 in Davidson (1984a).
- . (1984d). "On the Very Idea of a Conceptual Scheme," 183–98 in Davidson (1984a).
- . (2000). "Objectivity and Practical Reason," 17–26 in E. Ullman-Margalit (ed.), *Reasoning Practically*. Oxford: Oxford University Press.
- Feldman, R. (2006). "Epistemological Puzzles about Disagreement," 216–36 in S. Hetherington (ed.), *Epistemology Futures*. Oxford: Oxford University Press.
- Føllesdal, D. (1975). "Meaning and Experience," 25–44 in S. Guttenplan (ed.), *Mind and Language*. Oxford: Clarendon Press.
- Goldman, A. H. (1988). *Moral Knowledge*. London: Routledge.
- Harman, G. (1977). *The Nature of Morality*. New York: Oxford University Press.
- Kelly, T. (2005). "The Epistemic Significance of Disagreement," *Oxford Studies in Epistemology* 1: 167–96.
- Kornblith, H. (2010). "Belief in the Face of Controversy," 29–52 in T. Warfield & R. Feldman (eds.), *Disagreement*. Oxford: Oxford University Press.
- Lammenranta, M. (2011). "Skepticism and Disagreement," 203–16 in D. Machuca (ed.), *Pyrrhonism in Ancient, Modern, and Contemporary Philosophy*. Dordrecht: Springer.

- Lewis, D. (1983). "Radical Interpretation," 108–21 in his *Philosophical Papers*, vol. I. Oxford: Oxford University Press.
- Lukes, S. (1992). "Relativism in its Place," 261–305 in M. Hollis & S. Lukes (eds.), *Rationality and Relativism*. Oxford: Blackwell.
- Mackie, J. L. (1977). *Ethics. Inventing Right and Wrong*. New York: Penguin.
- Shafer-Landau, R. (1994). "Ethical Disagreement, Ethical Objectivism and Moral Indeterminacy," *Philosophy and Phenomenological Research* 54: 331–44.
- . (1995). "Vagueness, Borderline Cases and Moral Realism," *American Philosophical Quarterly* 32: 83–96.
- Sherif, M. A. (1936). *The Psychology of Social Norms*. New York: Harper & Row.
- Sturgeon, N. (1994). "Moral Disagreement and Moral Relativism," *Social Philosophy and Policy* 11: 80–115.
- Tersman, F. (2006). *Moral Disagreement*. New York: Cambridge University Press.
- Tolhurst, W. (1987). "The Argument from Moral Disagreement," *Ethics* 97: 610–21.
- Williamson, T. (1994). "A Critical Study of *Truth and Objectivity*," *International Journal of Philosophical Studies* 2: 130–44.
- Wittgenstein, L. (1979). *On Certainty*. 4th edition. Oxford: Blackwell.
- Wright, C. (1992). *Truth and Objectivity*. Cambridge, Mass.: Harvard University Press.
- . (1994). "Realism: Pure and Simple?" *International Journal of Philosophical Studies* 2: 147–61.
- . (1996). "Truth in Ethics," 1–18 in B. Hooker (ed.), *Truth in Ethics*. Oxford: Blackwell.
- . (2001). "On Being in Quandary," *Mind* 110: 45–98.