

The Significance of Significant Fundamental Moral Disagreement¹

Richard Rowland

Australian Catholic University

Close to final draft, final version Published in *Nous* DOI:
10.1111/nous.12170

Please cite published version at:

<http://onlinelibrary.wiley.com/doi/10.1111/nous.12170/abstract>

Abstract

This paper has four parts. In the first part I argue that moral facts are subject to a certain epistemic accessibility requirement. Namely, moral facts must be accessible to some possible agent. In the second part I show that because this accessibility requirement on moral facts holds, there is a route from facts about the moral disagreements of agents in idealized conditions to conclusions about what moral facts there are. In the third part I build on this route to show that (*) if there is significant moral disagreement in idealized conditions, then our understanding of morality is fatally flawed and we should accept relativism over non-naturalism and quasi-realism. So, if, like many, you think that there would be significant moral disagreement in idealized conditions, you should hold that our understanding of morality is fatally flawed and reject non-naturalism and quasi-realism. In the fourth part of this paper I show that (*) undermines the plausibility of non-naturalism, quasi-realism, and the view that our understanding of morality is not fatally flawed even if we do not have sufficient reason to believe that there would be significant moral disagreement in idealized conditions.

0. Introduction

Fundamental moral disagreements are moral disagreements that would survive in ideal conditions. That is, conditions in which the parties to the disagreement are (at least) informed of all the non-moral facts and are fully rational.² Many have argued that if there are fundamental moral disagreements, then moral realism is false.³ However, it has not been made clear how fundamental moral disagreement could have any metaphysical consequences at all or pose a problem for moral realism. After all, moral realists hold that moral truths are independent of us. And it may well be that there would be disagreement in ideal conditions about other truths that are independent of us. But this would not cast doubt on there being such truths. If there were disagreements in ideal conditions about whether dualism holds or whether we have free will,

¹ I would like to thank Zoë Cosker, Billy Dunaway, Brad Hooker, Guy Fletcher, Wouter Kalf, James Kirkpatrick, Jussi Suikkanen, several anonymous referees, and audiences at the University of Reading, the University of Cardiff, and the University of Oxford for extremely helpful comments on previous versions of, and aspects of, this paper.

² See Stich (2009), Doris and Plakias (2008), and Tersman (2006, ch. 2).

³ See, for instance, Bennigson (1996), Brandt (1959: 281–4), Machery et al (2005), Doris and Plakias (2008), Rawls (1999: 290, 301), Smith (1994: 5–6), Stich (2009), and Wright (1995: 223).

this might show that we cannot know whether we have free will or whether dualism holds. But such disagreement would not show that there is no truth of the matter about whether we have free will or show that we should construe truths about free will in an anti-realist fashion.⁴

In this paper I argue for four key claims. Firstly, I argue that moral facts and truths are distinct from facts and truths about, for instance, metaphysics and history in that there is a certain accessibility requirement on moral facts and truths. Namely, in §1 I argue that

- (1) If ϕ -ing has a moral status, then the moral status of ϕ -ing is accessible to some possible agent.

In contrast, there is no similar accessibility requirement on historical or metaphysical facts: we have free will even if no possible agent could know this; and it rained (suppose) on September 2nd 1408 in Edinburgh even if no possible agent could know this. (1) is interesting in its own right. For instance, it is sometimes taken as definitive of realism about a domain that facts and truths about that domain are not seen as epistemically constrained in any way.⁵ If we should accept this understanding of what it is to be a realist about a domain, then by showing (1) I will have shown that we should reject realism about moral facts and truths.

In §2 I show that given this accessibility requirement on moral facts and truths, that is (1), fundamental moral disagreement has moral consequences. In §2 I show that we should hold views about the epistemology of peer disagreement in idealized conditions and about epistemic accessibility that, together with (1) entail

- (2) If there would be disagreement about the moral status of ϕ -ing in ideal conditions, then ϕ -ing does not have a moral status.

In §3 I argue that (2) has meta-ethical consequences. (2) entails that if there would be disagreement in ideal conditions about the moral status of most of the actions that we currently disagree about the moral status of (that is, the actions that we disagree about in the applied ethics literature and in public culture), then most of the actions that we currently disagree about the moral status of do not have a moral status. Vegetarianism, bombing Syria, torture, the death penalty, breaking promises, pushing the fat man off of the bridge, lying: (2) entails that if there is disagreement about the moral status of these actions in ideal conditions, then it's not just that these actions are neither right nor wrong, they're not permissible either. In §3 I argue that, as a result of this entailment of (2),

- (3) If there is significant fundamental moral disagreement, then our understanding of morality is fatally flawed and we should accept relativism over non-naturalist realism and quasi-realism.

Given (3), if you believe that there would be significant moral disagreement in ideal conditions, as many do, then you should reject non-naturalism and quasi-realism and hold that our understanding of morality is fatally flawed. And the truth of (3) shows that those, such as David Brink, David Enoch, Sarah McGrath, Derek Parfit, and Russ Shafer-Landau, who have

⁴ See Shafer-Landau (2003: 228) and McGrath (2010: 64, 76).

⁵ See Kelly and McGrath (2010: 341).

argued that there is no route from facts about moral disagreement to consequences for moral realism are mistaken.⁶

Finally, in §4 I show that

- (4) Even if we do not have sufficient reason to believe that there is or would be significant fundamental moral disagreement, (3) weakens the plausibility of non-naturalism and quasi-realism.

I show that significant fundamental moral disagreement is a salient possibility. And that, since (a) we should reject non-naturalism and quasi-realism if there is significant fundamental moral disagreement and (b) significant fundamental moral disagreement is a salient possibility, it follows that (c) we should reduce our confidence in non-naturalism and quasi-realism. And (c) holds even if we do not have sufficient reason to believe that there is or would be significant fundamental moral disagreement.

1. Morality Must be Possibly Accessible

In §1.2 I will argue that

- (1.2) If ϕ -ing has a moral status, then the moral status of ϕ -ing makes a difference to what at least some possible agent ought to do.

And in §1.1 I will argue that

- (1.1) If the moral status of ϕ -ing makes a difference to what some possible agent ought to do, then the moral status of ϕ -ing is accessible to some possible agent.

The combination of (1.1) and (1.2) entails

- (1) If ϕ -ing has a moral status, then the moral status of ϕ -ing is accessible to some possible agent.

To argue for (1.1) and (1.2) I first need to introduce the idea of a fact p 's making a difference to what an agent \mathcal{A} ought to do. A fact most obviously makes a difference to what \mathcal{A} ought to do by providing a reason for \mathcal{A} to perform/not perform an action. For instance, the fact that failing to turn the steering wheel of my car will result in the death of a pedestrian makes a difference to what I ought to do by giving me most reason to turn the steering wheel of my car. But a fact can make a difference to what \mathcal{A} ought to do in other ways too. A fact can make a difference to what \mathcal{A} ought to do by disabling or silencing other facts that would otherwise make a difference to what \mathcal{A} ought to do. For instance, the fact that a promise that you made was made under duress does not itself provide you with a reason to refrain from keeping the promise. Rather the fact that a promise that you made was made under duress stops the fact that you made a promise from providing you with a reason to keep the promise, as it otherwise would.⁷ And the fact that a promise that you made was made under duress makes a difference to what you ought to do by stopping the fact that you made a promise

⁶ See Brink (1989: 197-209), Enoch (2009) (2011, ch. 8), McGrath (2010), Parfit (2011b: 543-569), and Shafer-Landau (2003, ch. 9).

⁷ See Dancy (2004: 38-42).

from providing you with a reason to keep the promise, as it otherwise would. So, p makes a difference to what some agent ought to do if p makes it the case that some agent ought to perform or not perform an action by providing reasons for that agent to perform or not perform that action. Or by disabling/enabling facts that would have/wouldn't have otherwise been reasons to perform an action.

In order to argue for (1.1) and (1.2) I also need to make clear what I mean by 'access'. In §2 I will argue that being able to justifiably believe that p is not sufficient to have epistemic access to p . But for the purpose of §1.1-2, this does not matter. Assume that A has epistemic access to p only if A is in a position to know or justifiably believe that p .⁸

1.1 If Morality is Normatively Relevant, then Morality Must Be Possibly Accessible

It seems that

Normative Relevance Entails Accessibility. If p makes a difference to what some agent ought to do, then p is accessible to at least some possible agent.

The following familiar type of case supports *Normative Relevance Entails Accessibility*. Consider

Mineshafts. A hundred miners are trapped underground with floodwaters rising. We are rescuers on the surface who are trying to save them. We know that the miners are in one of two mineshafts but we do not know which shaft they are in. In fact we have no reason to believe that they are in one shaft rather than the other. There are three floodgates that we could close by remote control. And depending on which gate we close the results will be as follows:

		The miners are in	
		Shaft A	Shaft B
We close	Gate 1	We save 100 lives	We save no lives
	Gate 2	We save no lives	We save 100 lives
	Gate 3	We save 90 lives	We save 90 lives

It is clear that we ought to close Gate 3 in this case. There is a fact, F , regarding which shaft the miners are in and so there is a gate that we could close, the closing of which would save more lives than our closing Gate 3. But this fact, F , does not seem to make a difference to what we ought to do because we cannot know this fact and we have no reason to believe that the miners are in one shaft rather than the other.⁹

So, our intuitions about *Mineshafts* support *Normative Relevance Entails Accessibility*. And it follows from *Normative Relevance Entails Accessibility* that

⁸ It will also be helpful to clarify what I mean by 'some possible agent'. By 'some possible agent' I mean a possible agent other than the trivially logically possible agent who has access to p in virtue of the fact that p . It might seem ad-hoc to make this restriction. But it is not clear to me that it is ad-hoc. For instance, an account of a property P will count as a judgment-dependent account of P if it gives an account of P in terms of the dispositions of possible agents other than such trivially possible agents; see Miller (2003: 130-131). So this distinction between some possible agent in the non-trivial sense and some possible agent in the trivial sense is already significant and utilized, and thereby not an ad-hoc distinction.

⁹ See Parfit (2011a: 159).

- (1.1) If the moral status of ϕ -ing makes a difference to what some possible agent ought to do, then the moral status of ϕ -ing is accessible to some possible agent.

There is an obvious line of objection to my argument for (1.1) from *Mineshafts*. Objectivists about obligation hold that the fact about which shaft the miners are in fact in does make a difference to what we ought to do. According to objectivists about obligation, what we ought to do is determined by all the normatively relevant facts. Suppose that, in *Mineshafts*, the miners are in Shaft A. According to objectivists, even if we have no reason to believe that the miners are in Shaft A rather than Shaft B, so long as the miners are in Shaft A, we ought to close Gate 1 rather than Gate 3. That is, according to objectivists, the fact that the miners are in Shaft A makes a difference to what we ought to do by making it the case that we ought to close Gate 1 rather than Gate 3 even if this fact is not accessible to us. So, it seems that in order to show that (1.1) holds, I need to show that objectivists about obligation are mistaken.

However, in order to show that (1.1) holds I do not need to show that objectivists about obligation are mistaken. Although objectivism about obligation is normally understood as the view that what we ought to do is a function of all the normatively relevant facts, in the rest of this section I will argue that the arguments that objectivists use to argue in favour of objectivism only count in favour of the view that

Objectivism (Possibly Accessible). What we ought to do is a function of all the normatively relevant facts that are accessible to some possible agent.

And that the arguments that objectivists use to argue in favour of objectivism do not count in favour of the stronger view that

Objectivism (Inaccessible). What we ought to do is a function of all the facts *including the facts that are not accessible to any possible agent*.

In the rest of this section I will argue that even if objectivists can give us reason to hold *Objectivism (Possibly Accessible)* and so can give us reason to believe that we ought not close gate 3 in *Mineshafts* they cannot give us reason, or at least have given us no reason, to hold *Objectivism (Inaccessible)*. But *Objectivism (Possibly Accessible)* does not conflict with (1.1); only *Objectivism (Inaccessible)* conflicts with (1.1). And so objectivists cannot give us, or at least have not given us, any reason to reject (1.1).

It might seem that even if I am successful in my argument in the rest of this section for the view that objectivists have given us no reason to hold *Objectivism (Inaccessible)*, and so have given us no reason to reject (1.1), I will not have shown that we have any reason to accept (1.1) but only that there is no reason for us to reject (1.1). But this is not the case. Even objectivists agree that our intuitions about cases like *Mineshafts* give us good reason to reject objectivist views. Objectivists just claim that we have most reason to accept objectivism given the other arguments for objectivism that they give. However, our intuitions about cases like *Mineshafts* provide us with a reason to accept the view that (i) facts that are not currently accessible to us do not make a difference to what we ought to; and *Mineshafts* provides us with a reason to accept the view that (ii) facts that are not accessible to *any possible agent* do not make a difference to what we ought to; *Mineshafts* gives us *pro tanto* reason to accept both (i) and (ii) because (ii) is entailed by (i). Even if, because of the good arguments for *Objectivism (Possibly Accessible)*, we should in the end reject (i), this does not show that there is no reason to accept

(i), but only that this reason is outweighed by reasons to reject (i). If my argument in the rest of this section is successful, we have no reason to reject (ii). And since *Mineshafts* provides reasons for us to accept (i) and (ii) and we have no reason to reject (ii), we should accept (ii) even if we should reject (i) all things considered. (ii) entails (1.1). And, so, if my argument in the rest of this section is successful, we should accept (1.1).

There are two main arguments for objectivism about obligation.¹⁰ According to the advice argument,

Advice. When we advise someone about what they ought to do we take into account all of the normatively relevant facts and not only the facts that the person we are advising currently knows about or has evidence for.

But, according to the advice argument, alternative views to objectivism cannot make sense of *Advice*. Because these alternative views hold that what we ought to do is a function of our current beliefs or our current evidence rather than all the normatively relevant facts.¹¹

Advice is underspecified. Either by accepting *Advice* we are accepting *Advice (Believed Inaccessible)* or we are not:

Advice (Believed Inaccessible). When we try to advise people we try to advise them in light of all the facts including the facts we believe to be inaccessible to all possible agents.

Advice (Believed Inaccessible) could be false consistent with

Advice (Believed Accessible). When we try to advise people we try to advise them in light of all of the facts including facts that we believe to be accessible to some possible agent but which are not in fact accessible to any possible agent. But not including facts that we do not believe to be accessible to any possible agent.

If *Advice (Believed Accessible)* holds but *Advice (Believed Inaccessible)* does not, then there is still a very clear way in which *Advice* holds; indeed, when most objectivists claim that *Advice* holds I do not believe that they mean that *Advice (Believed Inaccessible)* holds. And if *Advice (Believed Inaccessible)* were false but *Advice (Believed Accessible)* held, then we would try to advise others in light of all the facts that we thought relevant including facts that we thought someone could access. The combination of *Advice (Believed Accessible)* and the negation of *Advice (Believed Inaccessible)* only entails that if we were to come to believe that some fact *F* is not accessible to any possible agent, we would stop advising others in light of *F*.

It seems to me that perhaps we should accept *Advice (Believed Accessible)* but we should certainly reject *Advice (Believed Inaccessible)*. Because we do not advise others in light of facts that we believe to not be accessible to anyone. Although we may advise others in light of facts that we take ourselves to have access to but which in fact are inaccessible. It would be very odd for us to try to advise others in light of facts that we believe to not be accessible to anyone. For if we thought it reasonable to advise others in light of facts that we believe to not be accessible to anyone, then we would think it reasonable to say ‘the miners are in shaft *A*, so you should shut gate 1, but it’s impossible for me or anyone else to know that the miners are in shaft *A*’. And it is not reasonable to say this. (This claim breaches many plausible candidate norms of

¹⁰ See Graham (2010: 91-93) and Lord (2015: §4.1).

¹¹ See Graham (2010: 91-92), Thomson (1986: 179), and Lord (2015: §4.1).

assertion such as the knowledge norm and the reasonable belief norm. Furthermore, proponents of the truth norm of assertion do not believe that it is reasonable to assert claims like this. And argue that the truth norm can be combined with other theoretical tools to explain why it is not reasonable to assert claims like this one).¹² So, we should reject *Advice (Believed Inaccessible)* even though we should accept *Advice*; the plausibility of *Advice* does not extend to, and lend plausibility to, *Advice (Believed Inaccessible)*.

However, in order for the advice argument for objectivism to count in favour of *Objectivism (Inaccessible)* (rather than just *Objectivism (Possibly Accessible)*) it would have to be the case that *Advice (Believed Inaccessible)* holds. This is because we can explain *Advice (Believed Accessible)*—and so *Advice* so long as *Advice* is not taken to involve *Advice (Believed Inaccessible)*—without holding *Objectivism (Inaccessible)*. For so long as we hold *Objectivism (Possibly Accessible)* we can explain why people advise others in light of all the facts other than the facts that they believe to be inaccessible to everyone. (According to *Objectivism (Possibly Accessible)* remember, what we ought to do is a function of all the normatively relevant facts that are accessible to some possible agent). So, the advice argument does not favour *Objectivism (Inaccessible)* because we should not accept *Advice (Believed Inaccessible)*.

The other main argument for objectivism—the new information argument—stems from how we understand what happens to our obligations when we obtain new information that changes our understanding of our obligations. Namely,

when we change our view about our moral obligations as a result of acquiring information, we don't take it that our moral obligations have changed. Rather, it seems to us that we, at last, come to see what our moral obligations were all along. If this seeming is not delusory, then objectivism must be correct.¹³

But if p is not accessible to anyone, then it is impossible to acquire p , or rather to acquire information that shows that p . And in this case this argument for objectivism only counts in favour of *Objectivism (Possibly Accessible)* rather than *Objectivism (Inaccessible)*. Furthermore, we cannot have reliable intuitions about impossible scenarios (if we can have intuitions about them at all). But in this case it is no response that if we were to acquire an inaccessible fact, or information that showed that such an inaccessible fact obtained, that changed our views about our obligations, we would think that we have come to see what our obligations were all along. It is impossible to acquire a fact that is not accessible to any possible agent or to acquire information that shows that such a fact obtains.

It might seem that the new information argument can be revised in a way in which it would still count in favour of *Objectivism (Inaccessible)*. According to such a revised argument:

Seemingly Accessible New Information. When we acquire information that seems to show that p , even if p is in fact not accessible to any possible agent, p (or the fact that is evidence that seems to show that p) can still make a difference to our obligations. And when we acquire information that seems to show that p , where p is or would be relevant to our obligations, we don't think that our obligations change; we think that we have just discovered what our obligations were all along.

And only *Objectivism (Inaccessible)* can accommodate *Seemingly Accessible New Information*.

¹² See Weiner (2007).

¹³ Graham, (2010: 91). Cf. Ross (1930: 32), Lord (2015: §4.1), and Williams (1981: 22-26).

But it is not the case that only *Objectivism (Inaccessible)* can accommodate *Seemingly Accessible New Information*. Suppose that no one can access p , but we can access q , and q within our epistemic context appears to show that p . It is consistent with the negation of *Objectivism (Inaccessible)* that in this case the fact that (q within our epistemic context appears to show that p) makes a difference to what we ought to do even though p , because it is inaccessible, does not make a difference to what we ought to do. Furthermore, suppose that the fact q that seems to show that p makes it the case that we ought to ϕ . It is consistent with the negation of *Objectivism (Inaccessible)* that the fact that q within a particular epistemic context appears to show p determined that we ought to ϕ all along.

According to,

- (1.1) If the moral status of ϕ -ing makes a difference to what some possible agent ought to do, then the moral status of ϕ -ing is accessible to some possible agent.

I have shown that objectivists about obligation give us no reason to accept *Objectivism (Inaccessible)*, but only give us reason to accept *Objectivism (Possibly Accessible)*. And I've shown that (1.1) is only inconsistent with *Objectivism (Inaccessible)*. So, it is no objection to (1.1) that objectivists about obligation have arguments that militate against (1.1) because this is not the case. And, as I argued above, our intuitions about *Mineshafts* count in favour of (1.1) even if we should accept *Objectivism (Possibly Accessible)*. So, I've shown that we have reason to accept (1.1) and no reason to reject (1.1).¹⁴ So, we should accept (1.1).

1.2. Morality is Normatively Relevant

In this section I'll argue that

- (1.2) If ϕ -ing has a moral status, then the moral status of ϕ -ing makes a difference to what at least some possible agent ought to do.

Many hold that it is essential to our understanding of morality that moral rightness and wrongness entail categorical reasons. That is, that, necessarily if ϕ -ing is wrong, then we all have categorical reason not to ϕ , and if ϕ -ing is right, then we all have categorical reason to ϕ .¹⁵ A categorical reason for A to ϕ is a reason for A to ϕ that she has regardless of her desires, aims, goals, or the roles that she finds herself in.¹⁶ And it seems very plausible that rightness and wrongness necessarily entail categorical reasons. If it's wrong to fail to save a drowning

¹⁴ Sorensen (1995: 252-254) argues against the view that if A ought to ϕ , then A must be able to know that she ought to ϕ . Sorensen claims that: (i) if A 's obligations must be accessible to A , then A could infer that ϕ -ing is supererogatory from the fact that she does not know whether ϕ -ing is obligatory or supererogatory and this is counterintuitive; and (ii) if obligations must be knowable, then one could 'mute the call of duty by diminishing one's cognitive capacity'. But (1.1) only entails that if A is obliged to ϕ , some possible agent can know that she is obliged to ϕ . So, A could only infer that she is not obliged to ϕ if she knew that no possible agent could know that she is obliged to ϕ . And (1.1) only entails that A could mute the call of duty by doing the impossible and diminishing all possible agents' cognitive capacities.

¹⁵ See Olson (2010: 62-65), Joyce (2001: 5, 37-45), Bedke (2011), Mackie (1977: 29), Parfit (2011a: 172-174), Scanlon (1998: 11), and Smith (1994: 7).

¹⁶ Olson (2010: 64)

child when you could do so easily, then there's reason for anyone who could save a drowning child to try to save them regardless of whether they want to or not. And if it's right to give 20% of your income to charity, then we all have reason to do so, not just those of us who care about giving to charity.

So, it seems that there is a necessary connection between rightness and wrongness and categorical reasons. But necessary connections call for explanation; if there were no explanation of why necessarily, if ϕ -ing is wrong, then we all have categorical reason not to ϕ , it would be extremely puzzling why wrongness and categorical reasons are necessarily connected. (For instance, the supervenience of the mental on the non-mental and the moral on the non-moral, which are both necessary connections, require explanation. And even strong connections require explanation, this is why clairvoyance seems so puzzling).¹⁷ There are, as far as I can tell, only two deep explanations of why rightness and wrongness necessarily entail categorical reasons, namely:

- (a) Facts about rightness and wrongness (at least partially) consist in facts about reasons for action;
- (b) The reasons for action that rightness and wrongness entail consist (at least partially) in facts about rightness and wrongness.¹⁸

And both (a) and (b) entail:

- (1.2.1) If ϕ -ing is morally right or wrong, then the moral status of ϕ -ing makes a difference to what some agent ought to do.

Furthermore, it seems incoherent to say,

'It would be wrong for me to do it, but nonetheless I've got no reason at all not to do it'.

Of course, one might say 'it would be "wrong", for me to do it but nonetheless I've got no reason not to do it', where "wrong" means, 'wrong according to society's, or others', standards'. But this does not show that it does not seem incoherent to say 'it would be wrong for me to do it, but nonetheless I've got no reason at all not to do it'. And if it is incoherent to say, 'it would be wrong for me to do it, but nonetheless I've got no reason at all not to do it' (as it seems to be), then (1.2.1) holds.

Even if you are skeptical about whether we should really hold (a) or (b), it would be hard to understand how rightness and wrongness could necessarily entail categorical reasons if either (c) or (d) did not hold:

- (c) Necessarily, the rightness or wrongness of ϕ -ing provides us with reasons to ϕ /not- ϕ ;
- (d) Necessarily, the rightness or wrongness of ϕ -ing makes it the case that we have reasons to ϕ /not- ϕ .¹⁹

¹⁷ See Enoch (2011: 158) and Suikkanen (2005: 526-527).

¹⁸ Bedke (2011) argues for (a).

¹⁹ Scanlon (1998: 11) and Parfit (2011a: 172-174) argue for (c). On claims like (d), see Schroeder (2007: ch. 2).

And (c) and (d) entail

- (1.2.1) If ϕ -ing is morally right or wrong, then the moral status of ϕ -ing makes a difference to what some agent ought to do.

It might seem that a lot of the plausibility of the connection between rightness and wrongness and normative reasons could be preserved if we instead held:

- (e) Facts about rightness and wrongness, when they are possibly accessible, provide us (or some possible agents) with normative reasons.

And (e) does not entail (1.2.1).

But, firstly, to hold (e) would be to reject that rightness and wrongness necessarily entail categorical reasons. It seems that we should accept that rightness and wrongness necessarily entail categorical reasons. So, we should not accept (e). Secondly, if (e) holds, there would be a very strong connection between rightness and wrongness and categorical reasons that would not be explained. And, there would be a necessary connection between possibly accessible rightness and wrongness and categorical reasons, which would go unexplained. And strong connections and necessary connections require explanation.²⁰ Thirdly, it is very natural to think that rightness and wrongness just consist in facts about normativity; for rightness and wrongness necessarily matter, but how could they necessarily matter, if they outstripped normative import. Even those who do not hold (a) or (b) hold that rightness and wrongness necessarily matter and that rightness and wrongness do not outstrip their normative importance. For instance, Derek Parfit, who believes that the concept of something's being wrong in perhaps the most important sense cannot be analyzed—and that the notion of a reason cannot be analysed—holds that the concept of something being wrong in perhaps the most important sense is equivalent to the concept of something being such that it '*mustn't-be-done*'.²¹ Finally, no one seems to in fact hold (e). Those who do not hold that if ϕ -ing is wrong, then there is categorical reason for us not to ϕ still hold one of (a-d). For instance, those who hold that what we have a reason to do is constrained by our desires or our motivational set in some way, think that the same is true of what it's right/wrong for us to do.²² So, we should not hold (e) and we should instead hold one of (a-d). And in this case we should hold (1.2.1).

If we should hold (1.2.1), we should also hold that

- (1.2.2) If ϕ -ing is morally permissible, the fact that ϕ -ing is morally permissible makes a difference to what some agent ought to do.

This is because (1.2.1) essentially claims that, other things equal, if ϕ -ing weren't wrong, there would not be (some) reasons to ϕ that there would be if ϕ -ing were wrong. But although the permissibility of an action may not itself provide reasons to perform or not perform an action, the permissibility of ϕ -ing counterfactually affects the reasons to ϕ . That is, if ϕ -ing weren't permissible, then the reasons for us, or some other agent, to ϕ would be different than if ϕ -ing

²⁰ *Supra* note 16.

²¹ Parfit (2011a: 173)

²² See, for instance, Harman (1975: 7).

were permissible. ((a-d) and the view that rightness and wrongness necessarily entail categorical reasons all entail this). And in this case (1.2.2) holds.

The combination of (1.2.1) and (1.2.2) entails

- (1.2) If ϕ -ing has a moral status, then the moral status of ϕ -ing makes a difference to what at least some possible agent ought to do.

And in §1.1. I argued that

- (1.1) If the moral status of ϕ -ing makes a difference to what some agent ought to do, then the moral status of ϕ -ing is accessible to some possible agent.

And (1.1) and (1.2) entail

- (1) If ϕ -ing has a moral status, then the moral status of ϕ -ing is accessible to some possible agent.²³

²³ McGrath (2010: 72), Suikkanen (2008), and Wedgwood (2010a) argue that there are unknowable truths about wrongness because borderline cases tend to create in principle unknowable truths. For instance, suppose that it is wrong for us to fail to give at least 1% of our income to charity but it is not wrong for us to refrain from giving all of our income to charity. In this case, assuming classical logic, there will be a highest percentage n such that we are morally required to give n percent of our income to charity but we are not required to give $n+1$ percent of our income to charity. However, according to Suikkanen, Wedgwood, and McGrath even idealized reasoners would not be in a position to discern exactly what percentages n and $n+1$ are. Such idealized moral reasoners would not be able to know what percentages n and $n+1$ are because even if they correctly believed that, for instance, n is 20% and $n+1$ is 21% their beliefs would not be sufficiently safe to constitute knowledge. That is, they could have too easily believed that n is 19% and $n+1$ is 20% on almost the same grounds as the grounds on which they believe that n is 20% and $n+1$ is 21% for their true beliefs to count as knowledge. Cf. Williamson (2000: ch. 4).

However, whatever idealized moral reasoners know (and more on this in §2), they will know all of the (relevant) non-moral facts. And in this case we should not be so easily convinced that idealized moral reasoners would not be able to know what percentages n and $n+1$ are. Since idealized moral reasoners know all the non-moral facts, they will know, for instance, exactly how much each individual needs to give in order to collectively maximize pleasure. Now suppose, as might be the case, that n is the amount that it is and not $n+1$ because n is the exact percentage of income such that if each individual gave this percentage of their income, pleasure would be maximized. Idealized reasoners, since they know all the empirical facts, could not easily have been mistaken about the empirical facts that make it the case that the amount needed to collectively maximize pleasure is n rather than $n+1$. So, if the grounds on which we are required to give n but not $n+1$ percent of our income to charity are, for instance, that doing so would maximize pleasure, then idealized reasoners would have safe beliefs about what n and $n+1$ are.

It might be objected that there must be a similar case involving vague concepts or concepts that come in degrees and idealized reasoners will not be able to know exactly where the borderline is for something being M or not- M if M is a vague concept. However, this is not obvious. For idealized reasoners know all of the non-moral facts. And in this case it will follow that they will know the borderlines of these vague concepts. (Note that it is not logically impossible to know the borderlines of vague concepts, so it is not illegitimate to stipulate that idealized reasoners know these borderlines).

Furthermore, elsewhere Wedgwood (2010b: 222) argues that because moral truths are plausibly necessary truths safety-based arguments against moral knowledge cannot succeed. Safety conditions on knowledge are of the following general form: if S knows p , then S only believes p when p is true. Wedgwood plausibly claims that in order for a safety-based argument against S having knowledge of the moral proposition p to succeed it must be the case that there is a possible world in which p is false.

2. Morality Requires Convergence

The fact that morality, or rather moral facts, must be accessible to some possible agent has implications for the significance of fundamental moral disagreements. Fundamental moral disagreements are moral disagreements that would survive in ideal conditions.²⁴ I'll understand ideal conditions as conditions in which agents are fully informed of all the empirical and non-moral facts, are fully rational, are unaffected by cognitive biases, don't hold any conflicting beliefs, and have engaged in the very best reasoning process about normative ethics. Most philosophers claim or assume that the very best reasoning process about normative ethics is the most thorough search for wide reflective equilibrium. But I will not assume this; make the best reasoning process about normative ethics as broad as it needs to be. For instance, it might be that the very best reasoning process about normative ethics includes moral experiences and moral education.²⁵

Many, including Michael Smith, Crispin Wright, and John Rawls, have claimed that the amount of fundamental moral disagreements that there are affects the prospects of moral realism.²⁶ *Prima facie*, it is not clear why this would be the case. It may well be that there would be disagreement in ideal conditions about what happened on September 2nd 1408 but this would not show that there is no truth of the matter about what happened on September 2nd 1408 or that we should not be realists about what happened on September 2nd 1408. Similarly, there might be disagreement in ideal conditions about whether dualism holds. But this would not show that there is no fact of the matter about whether dualism holds.²⁷ My conjecture is that (a) there is no accessibility requirement on historical or (non-moral) metaphysical facts. But in contrast, as I showed in §1, (b) there is an accessibility requirement on moral facts. And (c) this contrast explains why whether there are/would be fundamental disagreements makes a difference to what moral facts there are and the prospects for (certain forms of) moral realism. In this section and the next I will sketch the precise route from facts about fundamental disagreement to conclusions about the moral facts that there are and the prospects of certain forms of moral realism.

There are good reasons to hold

- (2.1) \mathcal{A} has epistemic access to p at $T1$ only if there is a justification R that \mathcal{A} can justifiably believe p on the basis of at $T1$ that is such that \mathcal{A} 's belief that p on the basis of R would not be defeated in more ideal or idealized conditions.

(2.1) is supported by our intuitive judgments about cases. An agent with a Gettiered true justified belief that p does not seem to have access to p . For instance, consider Anna who justifiably believes that it is 8am because her clock tells her that it is 8am. But although it is in fact 8am, unbeknownst to Anna her clock has stopped. Suppose that Anna has no other way

And if moral truths are necessary, there is no such world. So, safety-based arguments against \mathcal{S} having knowledge of the moral proposition p cannot succeed. If this other argument of Wedgwood's is sound, then it is not the case that moral borderline cases ensure that there are inaccessible moral truths because even idealized reasoners could not have safe beliefs about such moral borderline cases. This is because—if this other argument of Wedgwood's is sound—given that moral truths are necessary truths, if an idealized reasoner has a true moral belief, there is no possible world in which they have the same belief but the moral belief that they have is false. So, idealized reasoners' moral beliefs could not be unsafe.

²⁴ *Supra* note 1.

²⁵ See DePaul (1993: ch. 4).

²⁶ *Supra* note 2.

²⁷ *Supra* note 3.

of knowing what time it is other than by looking at her clock; she is locked in a room with no windows for instance. In this case, Anna does not know that it is 8am even though she has a justified true belief that it is 8am. But neither does it seem that Anna has access to the fact that it is 8am. This is because there is a defeater for the only justification that Anna has access to for the belief that it is 8am: the fact that her clock has stopped working is a defeater for her justification for believing that it is 8am because her clock says that it is 8am. So this justification would be defeated in more ideal and idealized conditions. (Similarly, agents in fake barn cases do not seem to have epistemic access to the fact that they are looking at a real barn rather than a fake barn). (2.1) explains why Anna does not have access to the fact that it is 8am (and why agents in fake barn cases do not seem to have access to the fact that they are looking at a real barn). And if we do not hold (2.1) it is unclear how we can explain Anna's lack of epistemic access in this case.²⁸

Now, assume that one judges another to be an epistemic peer regarding p if one judges them to have the same evidence, intelligence, freedom from bias and cognitive functioning regarding p as one does.²⁹ It is uncontroversial that

- (2.2) Peer disagreement regarding p , when it makes one's justified belief regarding p no longer justified, provides a defeater for one's belief that p .³⁰

And in the rest of this section I will argue that we should accept

- (2.3) If idealized reasoners A and B hold conflicting beliefs about the moral status of ϕ -ing, know this, and know that they are idealized reasoners, then neither A nor B are justified in holding conflicting beliefs about the moral status of ϕ -ing; A and B are only justified in suspending belief about the moral status of ϕ -ing.

(2.1-2.3) entail that

²⁸ I have encountered an objection to (2.1). According to this objection, (2.1) entails that we have access to close to no facts, so we should reject (2.1). However, (2.1) does not entail that we have access to close to no facts. For instance, (2.1) would only entail that we have no access to the external world if skeptical hypotheses are defeaters for our justified beliefs about the external world. This is clearly up for debate in a way that it is not up for debate whether, for instance, the fact that her clock has stopped is a defeater for Anna's belief that it is 8am; skeptical hypotheses may be defeaters for our knowledge that there is an external world but the fact that her clock has stopped is certainly a defeater for Anna's justification for believing that it is 8am. And, furthermore, if you believe that skeptical hypotheses are defeaters on our justification for beliefs about the external world, then you should believe that we do not have *epistemic access* to the external world.

Similarly, it might seem that (2.1) entails that we do not have access to many everyday facts. For instance, it might seem that (2.1) entails that I do not have access to the fact that my bank opens at 9am tomorrow when my justification for this belief does not meet a high standard of justification. Because my justification is that I saw on the bank's website that the bank opens at 9am and it is possible that there was a typo on the website. But if there are no defeaters for this justification, if for instance there was no typo on the website, then it is consistent with (2.1) (though not entailed by (2.1)) that the fact that I saw that it says on the bank's website that they open at 9am tomorrow would be a justification for believing that the bank opens at 9am tomorrow even in more ideal or idealized conditions—even if I do not know that there was no typo. So, (2.1) does not entail that I do not have access to the fact that my bank opens at 9am tomorrow in this case.

²⁹ See Christensen (2009: 756-757).

³⁰ See, for instance, Goldberg (2013: 170, 181) and Weatherson (2013: 66).

- (2.4) If there would be disagreement about the moral status of ϕ -ing in idealized conditions, then no one has access to the moral status of ϕ -ing.

This is because, given (2.1-2.3), if there is such disagreement, then neither we nor any other agents, including agents in idealized conditions, have access to the moral status of ϕ -ing. But in §1 I established that

- (1) If ϕ -ing has a moral status, then the moral status of ϕ -ing is accessible to some possible agent.

And (2.4) and (1) entail

- (2) If there would be disagreement about the moral status of ϕ -ing in ideal conditions, then ϕ -ing does not have a moral status.

I'll give three arguments for

- (2.3) If idealized reasoners A and B hold conflicting beliefs about the moral status of ϕ -ing, know this, and know that they are idealized reasoners, then neither A nor B are justified in holding conflicting beliefs about the moral status of ϕ -ing; A and B are only justified in suspending belief about the moral status of ϕ -ing.

Firstly, if A and B know that one another are idealized reasoners, then they know that they have the same evidence, that they have not overestimated the force of any arguments, and that they are responding rationally to all the evidence and arguments that they have considered. So, A knows that she has no reason to privilege her own view about the moral status of ϕ -ing over B 's view, and A knows that she is no more likely than B to be right about the moral status of ϕ -ing and *vice versa*. And if A knows that she has no reason to privilege her own view about the moral status of ϕ -ing over B 's view and that she is no more likely than B to be right about the moral status of ϕ -ing, then it seems that A should suspend judgment about the moral status of ϕ -ing.

Secondly, if A is justified in holding onto her belief, then A should give extra weight to her view about the moral status of ϕ -ing when compared to the weight that she gives B 's view. But if A should give extra weight to her view over that which she should give to B 's view, then she should infer that she has gone through a better reasoning process than B about the moral status of ϕ -ing, has better evidence, or has understood arguments regarding the moral status of ϕ -ing better than B . But A *knows* that she has *not* gone through a better reasoning process than B and that she *does not have* better evidence and *has not* understood arguments regarding the moral status of ϕ -ing better than B . For A knows that both she and B are idealized agents that have engaged in the very best reasoning process about normative ethics. So, A should not infer that she has gone through a better reasoning process than B about the moral status of ϕ -ing, has better evidence, or has understood arguments regarding the moral status of ϕ -ing better than B . So, A should not give extra weight to her view. And so, A is not justified in holding onto her view.³¹

³¹ Cf. Elga (2007: 486-488).

Thirdly, (2.3) is entailed by a conciliatory view about peer disagreement and seems to be entailed by all plausible competitors to conciliatory views about peer disagreement. According to conciliationism, whenever one encounters a disagreement about some matter p with someone one judges to be an epistemic peer when it comes to p one ought to suspend belief about p .³² So if we should accept conciliationism generally, as many have argued that we should, then we should accept (2.3). However, we can also accept (2.3) without accepting a conciliatory view about all cases of epistemic peer disagreement. This is because there are relevant features of the type of case in (2.3) that are not shared by other cases of peer disagreements. Namely, A and B 's case is a case of a peer disagreement in which the epistemic peers know that there is no epistemic asymmetry between them. Proponents of the most plausible alternative views to conciliationism hold that only some peer disagreements about p should lead us to suspend belief about p . And proponents of these views hold that in cases in which epistemic peers should believe that they are epistemic peers regarding p , that they disagree regarding p , and that there is no epistemic asymmetry between one another regarding p , these epistemic peers should suspend belief regarding p .³³ But if epistemic peers should suspend belief regarding p so long as they should believe that they disagree regarding p , believe that they are epistemic peers regarding p , and believe that there is no epistemic asymmetry between them regarding p , then *a fortiori*, epistemic peers should suspend belief about p if they *know* that they disagree regarding p , *know* that they are peers regarding p , and *know* that there is no epistemic asymmetry between them regarding p . So, the most plausible views that compete with conciliationism entail (2.3) too. So, the most plausible views about epistemic peer disagreement entail (2.3).

The most common objection to (2.3) that I've encountered is that the antecedent of (2.3) is impossible. That is, that it is impossible for idealized reasoners A and B to hold conflicting beliefs about the moral status of ϕ -ing, know this, and know that they are idealized reasoners. One might believe that the antecedent of (2.3) is impossible because it is impossible for A to disagree with B about the moral status of ϕ -ing *and* judge that B is an epistemic peer regarding the moral status of ϕ -ing. But I see no reason to hold this view. I disagree with, for instance, esteemed Kantians about particular matters in normative ethics but I certainly believe that they are at least my epistemic peers regarding normative ethics. Although I disagree with luck egalitarians about redistributive justice I agree with them about most moral matters, judge most well known luck egalitarians to be extremely insightful and sophisticated moral thinkers, and to be at least my epistemic peers regarding distributive justice (and other matters).³⁴

Some think that the antecedent of (2.3) is impossible because

Impossible Disagreement. Reasoners that are so idealized could never disagree, for if they (always) knew that there was no cognitive asymmetry between them, idealized reasoners would never disagree about the moral status of ϕ -ing; they would always suspend about the morality of ϕ -ing.

However, it is not essential for my argument in this paper that A and B *in fact* disagree about the moral status of ϕ -ing *and then* come to suspend judgment about the moral status of ϕ -ing in light of this disagreement; if you think that A and B in virtue of understanding their competing dispositions would only ever suspend about the moral status of ϕ -ing, this is enough for my

³² See, for instance, Feldman (2006), Elga (2007), and Christensen (2009).

³³ See Kelly's (2010: 135-152) esteemed total evidence view and Lackey's (2008: 302-306, 316-318) justificationist view.

³⁴ See also Kornblith (2010: 47-51), McGrath (2007: 103-106), and Enoch (2014: 230-231).

purposes. Alternatively, if you think that *Impossible Disagreement* holds, then just modify the idealization conditions and (2.3) slightly such that: A and B are idealized reasoners but do not know the beliefs of one another at $T1$; they come to hold conflicting beliefs about the moral status of ϕ -ing at $T1$; and at $T2$ A and B come to know one another's beliefs and the fact that one another are idealized reasoners; at $T2$ A and B should suspend belief about the moral status of ϕ -ing.

Alternatively, it might be objected that the antecedent of (2.3) is impossible because if A and B disagree about the moral status of ϕ -ing, then one of them must have not engaged in the very best reasoning process about normative ethics. However, we should not understand the very best reasoning process about X as the one such that if we engage in it, we are guaranteed to arrive at the truth regarding X every time; indeed, at most, the standard view is only that the best reasoning process with regards to X must reliably enable one to arrive at the truth of the matter with regards to X .³⁵ And according to Michael DePaul, for R to be the best reasoning process about X it need only be the case that R is a better reasoning process about X than any other reasoning process about X .³⁶

So, I've shown that we should accept (2.1-3). And that (2.1-2.3) entail (2.4). But, I've also shown that (2.4) and (1) entail

- (2) If there would be disagreement about the moral status of ϕ -ing in ideal conditions, then ϕ -ing does not have a moral status.

So, I've shown that (2) holds.

3. The Meta-Ethical Consequences of Significant Fundamental Disagreement

In the last section I showed that

- (2) If there would be disagreement about the moral status of ϕ -ing in ideal conditions, then ϕ -ing does not have a moral status.

In this section I'll argue that, given (2),

³⁵ See, for instance, Kelly and McGrath (2010: 326).

³⁶ See DePaul (1993: ch. 2). Alternatively, it might seem that we should reject (2.3) because if idealized reasoners A and B disagree in virtue of having different moral perceptions, then the reasoner who has the correct moral perception is justified in their belief that is based on this perception and the reasoner who has an incorrect perception is not justified. However, this is a view that proponents of the view that we have moral perceptions that play a significant justificatory role do not accept. For proponents of this view do not accept that their view entails that we should be more steadfast in the face of disagreement than alternative views; see for instance McNaughton (1988: 60) and Audi (2013: 66). Furthermore, we should not accept that if agents have differing perceptions, then the person who has the correct perception is justified in sticking to the judgment that they formed on the basis of their perception when they still believe that the person whom they hold a conflicting judgment to is their epistemic peer. Even non-conciliationists hold this view; see, for instance, Kelly's (2010: 151-152) horse race case.

- (3) If there is significant fundamental moral disagreement, then our understanding of morality is fatally flawed and we should accept relativism over non-naturalist realism and quasi-realism.

(2) entails that

- (3.1) If in ideal conditions agreement would not emerge about the moral status of most of the actions and practices about whose moral status we currently disagree, then most of the actions and practices about whose moral status we currently disagree do not have a moral status.

(3.1) is meta-ethically important because

- (3.2) If most of the actions and practices about whose moral status we currently disagree do not have a moral status, then our understanding of morality is fatally flawed.

Although it might be that because of indeterminacy or vagueness a few of the actions and practices that we currently engage in moral disagreements about have no determinate moral status, if most of the actions that we currently engage in disagreements about the moral status of had no moral status, then our understanding of morality would be in significant error.³⁷ We think that pretty much all the actions and practices that we currently engage in moral disagreements about have *a moral status*. If we didn't think that these actions and practices have *a moral status*, we wouldn't be arguing about their moral status. And if the antecedent of (3.2) holds—that is, if most of the actions and practices about whose moral status we currently disagree do not have a moral status—then most of our discussions about the moral status of actions and practices are entirely (epistemologically) pointless.³⁸ If the antecedent of (3.2) holds, we would not be making a substantial first-order error about which moral status actions and practices have but rather an error about actions and practices in fact having such a status. If the antecedent of (3.2) holds, it's not just the case that taxing the rich to benefit the poor is neither right nor wrong, it's not permissible either. Vegetarianism, bombing Syria, torture, the death penalty, breaking promises, pushing the fat man off of the bridge, lying: none of these things are right, wrong, or permissible. So, it seems that we should hold (3.2). (I address the only two objections that I can imagine to (3.2) in a footnote³⁹).

³⁷ See Enoch (2009: 25, n. 39).

³⁸ See Tersman (2006: 1).

³⁹ It might be objected that (3.2) only holds if there is a presupposition or entailment of our particular token ascriptions of rightness, wrongness, and permissibility (i.e. ' ϕ -ing is wrong', ' ϕ -ing is right') that would be shown to be false if for most of the actions and practices Ψ about whose moral status we currently disagree, Ψ do not have a moral status. And no such presupposition or entailment of our claims about rightness, wrongness, and permissibility would be shown to be false if for most of the actions and practices Ψ about whose moral status we currently disagree, Ψ do not have a moral status. But this is not the generally accepted view. According to Michael Smith, for instance, our understanding of morality is fatally flawed if enough of our platitudes about morality are false. And many of the platitudes that Smith provides are not presuppositions or entailments of claims about rightness, wrongness, or permissibility. For instance, according to Smith (1994: 39-40), the following are platitudes about our moral concepts, 'if someone judges her ϕ -ing to be right then, other things being equal, she will be disposed to ϕ '; 'right acts are often concerned to promote or sustain or contribute in some way to human flourishing'; and 'right acts are in some way expressive of equal concern and respect'.

(3.1) and (3.2) entail

- (3.3) If in ideal conditions agreement would not emerge about the moral status of most of the actions and practices about whose moral status we currently disagree, then our understanding of morality is fatally flawed because most of the actions and practices about whose moral status we currently disagree do not have a moral status.

And (3.3) is important for the prospects of non-naturalist realism and for the prospects of quasi-realism. This is because, as I will argue in the rest of this section,

- (3.4) If our understanding of morality is fatally flawed because most of the actions and practices about whose moral status we disagree do not have a moral status, we should accept relativism rather than non-naturalism or quasi-realism.

All non-naturalists hold

Anti-Revisionism. Our understanding of morality is not in significant error.⁴⁰

(In contrast, not all naturalist realists are committed to *Anti-Revisionism*).⁴¹ If *Anti-Revisionism* is false, the main reasons to accept non-naturalist realism are seriously undermined. Because the main reason to accept non-naturalist realism is that it keeps our ordinary understanding of morality intact.⁴² But if *Anti-Revisionism* is false, this reason to accept non-naturalism falls by the wayside. And if we cannot preserve our ordinary understanding of morality by accepting non-naturalist realism, we might be less happy to live with the metaphysical and epistemological quirks and problems of non-naturalism.

If *Anti-Revisionism* is false because most of the actions and practices that we currently disagree about do not have a moral status, then the reasons to accept non-naturalist realism over relativist views are significantly blunted. Critics of relativism normally argue that relativist views are implausible because of their counter-intuitive and revisionist

Another objection to (3.2) is that it is just one platitude about our moral concepts that for most of the actions and practices Ψ about whose moral status we currently disagree, Ψ do have a moral status. But the fact that one platitude about morality is false does not show that there is no such thing as morality as we understand it. However, if this were the case, then we should not worry about the error theory. Error theorists hold that it is essential to our understanding of morality that there are categorical reasons, since our understanding of rightness, wrongness, goodness, and badness entail categorical reasons; see, for instance, Olson (2010). But there are no categorical reasons, so there is nothing that is right, wrong, good, and bad in the way that we understand these notions. But it is only at most one platitude about each of these concepts that there are categorical reasons to respond to things that these concepts apply to. So, the error theory relies on the view that it is possible that if one platitude about our understanding of C is false, then there is nothing to which C applies as we understand it. And if this is true, the falsity of one platitude about our moral concepts can—depending on the platitude—show that our understanding of morality is fatally flawed. So, either this objection is mistaken or we can refute the error theory very quickly. Since we cannot refute the error theory so quickly (in this way), this objection fails.

⁴⁰ See Enoch (2011: 8-10), Huemer (2005: xxi-xxiv), Shafer-Landau (2003: 12, 23), and Björnsson (2011: 368, 368 fn. 1 and the references therein).

⁴¹ See, for instance, Jackson (1998: 137) and Railton (1989: 157).

⁴² *Supra* note 39.

implications.⁴³ But (i) if *Anti-Revisionism* is false, then we must accept a counter-intuitive and revisionist meta-ethical view. And (ii) relativists have the resources to justifiably reject

- (3.1) If in ideal conditions agreement would not emerge about the moral status of most of the actions and practices about whose moral status we currently disagree, then most of the actions and practices about whose moral status we currently disagree do not have a moral status.

This is because (3.1) is entailed by (2) and relativists can hold that one of the key premises in the argument that I made for (2) in §2 is false. Namely relativists have the resources to reject (2.3)—the claim that if idealized reasoners find themselves in disagreement about the moral status of ϕ -ing, they should suspend judgment about the moral status of ϕ -ing and are not justified in making any positive judgment about the moral status of ϕ -ing. Relativists can hold that (2.3) is false because if idealized reasoner A believes that ϕ -ing is morally wrong and idealized reasoner B believes that ϕ -ing is morally right, A and B should not suspend belief about the moral status of ϕ -ing if relativism is true. Rather, if relativism holds, A should hold that ϕ -ing is morally wrong-for-him or morally wrong-by-his-standards and B should hold that ϕ -ing is morally right-for-him or morally right-by-his-standards.

Given that relativists can justifiably deny (3.1), if there is or would be significant fundamental moral disagreement, relativist and non-naturalist realist views have rather different revisionist implications, namely

Revisionist Implications of

⁴³ See, for instance, Olson (2010: 73-77), Rachels (2012: 14-21), and Schroeder (2010: 65-70).

Relativism

(a) There are no actions that are just wrong in a framework/standard-independent way

Non-naturalist Realism

(b) Most of the actions that we disagree about the moral status of have no moral status

It is not clear that (a) is a more counter-intuitive or revisionary consequence than (b). It is perhaps more counter-intuitive and more revisionary to hold (b) than to hold (a). Because if we hold (b), vegetarianism, bombing Syria, torture, the death penalty, breaking promises, pushing the fat man off of the bridge, lying, none of these actions and practices are right, wrong, or permissible. But if we hold (a), this is not the case. Furthermore, it is a familiar fact that many people hold that the morality of actions and practices is relative to particular standards such as the standards of a particular culture.⁴⁴ So, it is a familiar fact that many hold (a). But it is not clear that anyone except philosophers who understand and hold the moral error theory believe (b). And if (a) is at most as revisionist a consequence as (b), the main reason to accept non-naturalism over relativism was its anti-revisionism, and non-naturalism has other problematic metaphysical and epistemological consequences that relativism does not have, then we should accept relativism over non-naturalist realism (if there is or would be significant fundamental moral disagreement).

It might be objected that relativism entails that there is nothing that is wrong *tout court* and so relativism entails (b) as well as (a). However, relativism as a meta-ethical thesis is a thesis about what it takes for something to be wrong. Relativists hold that when we disagree about the moral status of the many actions and practices that we disagree about the moral status of our disagreements should be understood as disagreements about framework- or standard-relative rightness and wrongness.⁴⁵ So, relativism does not entail (b).

A comparison with contextualism in epistemology will further clarify the prospects of non-naturalist realism in comparison to relativism if there is significant fundamental moral disagreement. (i) Skeptical hypotheses *cast doubt on much of, but not all of our knowledge*; they do not cast doubt on our knowledge that we know that there is thought when we are thinking or our knowledge that we do not know everything for instance. But (ii) if we endorse contextualism about knowledge, skeptical hypotheses *do not cast doubt on as much of our knowledge*; we can know that we have hands in most contexts if we endorse contextualism.⁴⁶ Because of (i) and (ii) traditional *non-contextualist views about knowledge suffer a severe loss in plausibility relative to contextualist views* if they do not have a response to skeptical arguments. Similarly, I have argued that (i*) if in ideal conditions agreement would not emerge about the moral status of most of the actions and practices about whose moral status we currently disagree, then *most of the actions about whose moral status we currently disagree do not have a moral status*. And that, (ii*) if we endorse relativism *it is not the case that most of the actions about whose moral status we currently disagree do not have a moral status* even if in ideal conditions agreement would not emerge about the moral status of these actions. And because of (i*) and (ii*) *non-naturalism suffers a severe loss in plausibility relative to relativist views* if there is such significant moral disagreement in ideal conditions. (Rejecting this last claim would seem to commit one to denying the parallel claim that traditional non-contextualist views about knowledge suffer a severe loss in plausibility relative to contextualist views if they do not have a response to skeptical arguments).

⁴⁴ See Sarkissian et al (2011).

⁴⁵ See, for instance, Finlay (2014).

⁴⁶ See DeRose (1995).

So, I've shown that if our understanding of morality is fatally flawed because most of the actions and practices that we disagree about the moral status of have no moral status, then the plausibility of non-naturalist realism in comparison to relativism is undermined and we should accept relativism rather than non-naturalism.

Like non-naturalists, quasi-realist expressivists hold that our understanding of morality is not fatally flawed.⁴⁷ And, like non-naturalism, if there is/would be significant fundamental moral disagreement, quasi-realism will look far less attractive vis-à-vis relativism. As I've argued in this section, relativists have a way of justifiably denying

- (3.1) If in ideal conditions agreement would not emerge about the moral status of most of the actions and practices about whose moral status we currently disagree, then most of the actions and practices about whose moral status we currently disagree do not have a moral status.

Relativists can justifiably deny (3.1) because (3.1) is entailed by (2) and, as I've argued, relativists can deny (2) by justifiably denying a key premise in the argument that I made for (2), namely (2.3). But quasi-realists have no way of denying (3.1), (2), or (2.3) that non-naturalists would not also have, and non-naturalists have no way of denying (3.1), (2), or (2.3). And in this case, if there is significant moral disagreement in ideal conditions, like non-naturalism, quasi-realism will entail that most of the actions and practices about whose moral status we currently disagree do not have a moral status. But if there is significant moral disagreement in ideal conditions, relativism will not entail that most of the actions and practices about whose moral status we currently disagree do not have a moral status. And, as I've discussed, in this case, if there is significant moral disagreement in ideal conditions, quasi-realism will produce consequences that are at least as, and perhaps more, counter-intuitive or revisionary as the consequences that relativism produces.

Although quasi-realism does not face the metaphysical and epistemological problems that non-naturalism faces, the fact that endorsing relativism would lead to less, or only as, revisionist consequences as embracing quasi-realism might be sufficient to justify us in endorsing relativism rather than quasi-realist expressivism; especially if we see, as many do, quasi-realist expressivism as a view mostly motivated by the revisionist and counter-intuitive consequences of relativistic subjectivist and emotivist views.⁴⁸ Furthermore, although quasi-realist expressivism does not face the metaphysical and epistemological problems that non-naturalism faces, it does face other problems such as the Frege-Geach problem. And, if there were significant fundamental moral disagreement, and so if quasi-realism would have at least as counterintuitive and revisionist consequences as relativism, then those who are currently sympathetic to quasi-realism might be less willing to put up with the problematic consequences of quasi-realist expressivism that the Frege-Geach problem, for instance, makes clear.

So, I've shown that

- (3.4) If our understanding of morality is fatally flawed because most of the actions and practices about whose moral status we disagree do not have a moral status, we should accept relativism rather than non-naturalism or quasi-realism.

⁴⁷ See, for instance, Blackburn (1999).

⁴⁸ See, for instance, Schroeder (2010: ch. 4).

And I've shown that

- (3.3) If in ideal conditions agreement would not emerge about the moral status of most of the actions and practices about whose moral status we currently disagree, then our understanding of morality is fatally flawed because most of the actions and practices about whose moral status we currently disagree do not have a moral status.

And (3.3) and (3.4) entail

- (3) If there is significant fundamental moral disagreement, then our understanding of morality is fatally flawed and we should accept relativism over non-naturalist realism and quasi-realism.

4. The Significance of Significant Fundamental Disagreement

In the last section I showed that

- (3) If there is significant fundamental moral disagreement, then our understanding of morality is fatally flawed and we should accept relativism over non-naturalist realism and quasi-realism.

It is sometimes pressed to me that this conditional claim about the meta-ethical significance of fundamental moral disagreement may not be very important because it is far from clear that we have sufficient reason to believe that there is or would be significant fundamental moral disagreement. But in this section I'll argue that

- (4) Even if we do not have sufficient reason to believe that there is or would be significant fundamental moral disagreement, (3) weakens the plausibility of non-naturalism and quasi-realism.

We should hold that

- (4.1) There is a salient possibility, rather than a merely logical possibility, that there is or would be significant fundamental moral disagreement.

As I'll discuss below, there is a salient possibility that p if there is some non-trivial positive reason to believe that p . There is one quite direct reason to hold (4.1) that I will sketch. Many experts in the field believe that there is or would be significant fundamental moral disagreement. John Rawls, Bernard Williams, Terence Horgan, Mark Timmons, Nick Zangwill, Bart Streumer, John Doris, and Stephen Stich have all argued that there is or would be significant fundamental moral disagreement.⁴⁹ But merely logical possibilities are possibilities that pretty much no one, and certainly no experts in the field, believe obtain: no one, or no experts, believe that we are brains in vats or that all the zebras that we've seen are cleverly disguised horses. And I submit that if many experts—or at least many of those in at least as good a position as any others—regarding whether p believe that p , then we have some non-

⁴⁹ See Horgan and Timmons (2008: 231-232), Streumer (2011: 330-334), Williams (1985: 136), Zangwill (2000: 284) and *supra* note 2.

trivial positive reason to believe that p .⁵⁰ (Conciliationist and plausible non-conciliationist views about peer disagreement would seem to entail this too; see §2 above). And in this case we should take that p as a salient rather than a merely logical possibility. And in this case, because many experts in the field hold that there would be significant fundamental disagreement in ideal conditions, it follows that (4.1) holds.

Of course, many experts in the field also think that there is not, or would not be, significant fundamental moral disagreement. But this does not undermine (4.1). It is consistent with (4.1) that it is also a salient possibility that there is not or would not be significant fundamental moral disagreement. (Just as in an election between two parties there is a salient possibility that each will win).⁵¹

There is a second, less direct, argument for (4.1) too. Many hold that

- (a) The most thorough search for wide reflective equilibrium between our initial judgments about cases, more general principles that explain these judgments, and our best thought in other domains (such as metaphysics, economics, biology, etc) is the very best reasoning process about normative ethics.⁵²

But many hold that

- (b) If the best reasoning process about normative ethics is such a search for coherence between initial judgments and general principles, there will be several different conflicting wide reflective equilibria that some idealized reasoners settle on.⁵³

Now even if we do not have *sufficient reason* to hold (a) or (b) it seems hard to dispute that there are *positive non-trivial reasons* to believe both (a) and (b): even critics of wide reflective equilibrium agree that we have reasons to believe that the most thorough search for wide reflective equilibrium is the very best reasoning process about normative ethics.⁵⁴ And good arguments have been made, that no responses have been published to, which argue that

⁵⁰ See also Bostrom (2014: 21).

⁵¹ It has been put to me that because views are split about whether there would be significant fundamental moral disagreement, the appropriate response might be to suspend belief about this matter rather than to conclude that it is somewhat probable that there would be such significant fundamental moral disagreement. But even if we should suspend belief in this way, this does not show that the possibility of significant fundamental moral disagreement is not salient. Just as it might be that we should suspend belief about whether the UK will leave the EU but this does not show that the possibility that the UK will leave the EU is not a salient possibility.

⁵² See, for instance, Smith (1994: 40-41) and Scanlon (2002: 149).

⁵³ See, for instance Kelly and McGrath (2010: 338-340).

⁵⁴ Though critical of reflective equilibrium, Kelly and McGrath (2010: 352-353) seem to hold that the fact that (*) reflective equilibrium, seemingly uniquely amongst the views with which it competes, does not involve a seemingly 'cosmic' faculty of rational intuition is a reason to accept reflective equilibrium rather than competing views about the epistemology of morality. But in this case (*) is also a reason to accept that the most thorough search for wide reflective equilibrium is the best reasoning process about normative ethics; for we have reason to believe that there is no other method—consisting in our reason to believe that there are no 'cosmic' faculties of rational intuition. Scanlon (2002: 149, 151) argues that there are no alternative methods to a search for reflective equilibrium. It seems that we have reason to believe that this is correct. And in this case, we have a further reason to believe (a).

idealized agents involved in a search for wide reflective equilibrium would come to conflicting wide reflective equilibria.⁵⁵

If (a) and (b) hold, then

(c) There is or would be significant fundamental moral disagreement.

And since we have some reason to believe (a) and some reason to believe (b) we have some reason to believe (c).⁵⁶ So,

(4.1.1) We have some positive non-trivial reason to believe that there is or would be significant fundamental moral disagreement.

Now

(4.1.2) If there is some positive non-trivial reason to believe p , then there is not just a merely logical possibility that p , rather there is a salient possibility that p .

To see that we should hold (4.1.2) consider some cases. There are, right now in April 2016, some reasons to believe that the UK will leave the EU—provided by the stringency of support for such an exit from the EU—and the possibility that the UK will leave the EU is salient possibility. (Even though there is not, right now at least, sufficient reason to believe that the UK will leave the EU). Furthermore, it is possible that my train home tonight will never turn up. But I have no positive reason to believe that my train home will never turn up unless I am given some reason to believe that it will not turn up. And, because I have no positive reason to believe that my train will never turn up, it is not a salient possibility that my train home tonight will never turn up. In contrast, suppose that I'm driving home from a weekend away. And suppose that I know that the takeaway restaurant on the way home from which I would prefer—but do not need—to get dinner sometimes closes early on Sundays because of a lack of custom. In this case I don't have sufficient reason to believe that the takeaway has or will have closed early. But I have some reason to believe this, for I have some evidence that is such that the probability of the restaurant having closed early is higher on this evidence than it would have been without that evidence (some restaurants don't close early regardless, some restaurants don't have such droughts in customers). The takeaway's having closed is a salient possibility in this case rather than a merely logical possibility. And it seems that this is because I have some positive reason to believe that it will have closed.

So, we should hold (4.1.1) and (4.1.2). And it follows from (4.1.1) and (4.1.2) that

(4.1) There is a salient possibility, rather than a merely logical possibility, that there is or would be significant fundamental moral disagreement.

⁵⁵ *Supra* note 52.

⁵⁶ It might not follow that we have positive reason to believe (c) because we have reason to believe (a) and (b) if the truth of (b) would undermine (a). However, as I discussed in §2 this is not obviously the case. Furthermore, Kelly and McGrath (2010: 352-353) hold that there is *some* reason to hold (a) even though (b) holds and Scanlon (2002: 151-153) holds that the reasons that we have to hold (a) are independent of the reasons that we have to hold (b) because it is a virtue of the method of reflective equilibrium that it does not rule out moral skepticism. See also *supra* note 53.

We should also accept

- (4.2) If we should reject T if p , and there is a salient possibility that p , then T is less plausible than it otherwise would be and we should reduce our confidence in T (even if we do not have sufficient reason to believe p and even if there is a salient possibility that not- p).

To see that we should hold (4.2) consider the following case. We should reject the thesis that our experiences are genuine experiences of medium-sized dry goods if we come to believe that we are brains in vats (BIVs). Normally it is merely a logical possibility that we are BIVs. But suppose that we come to justifiably believe that we are in one of ten worlds. That everyone in each of these ten worlds believes that they are in a real world of medium-sized dry goods. All of the experiences of everyone in these worlds are pretty much the same. But in two of these worlds everyone in these worlds are in fact merely BIVs being manipulated by evil demons to experience a world of seeming medium-sized dry goods. In this case there would not be merely a logical possibility of our being BIVs. There would be a very clear and salient possibility that we are BIVs. And we would be worried by this clear and salient possibility even though we would not have sufficient reason to believe that we are merely BIVs (and even though there would be a salient possibility that *we are not mere* BIVs). And in this case we should reduce our confidence that we are experiencing an external world as we know it. In this scenario the thesis that our experiences are genuine experiences of medium-sized dry goods and that we are not BIVs would become less plausible than it otherwise would have been. This is because we know that there is a salient possibility, and not a merely logical possibility that our experiences are not veridical and that we are in fact BIVs.⁵⁷ (4.2) fits with and explains these judgments about this case. And if we do not accept (4.2) we cannot hold onto these judgments about this case.

Further similar cases support (4.2) too. For instance, when we are in the passenger seat of our friend's car looking out on rural surroundings, it is only a merely logical possibility that the barns that we see on our ride are fake barns that have been elaborately constructed by film-makers. But if we are riding in a piece of country that we know is sometimes filled with fake barns, the possibility that the barns that we see are fake becomes a salient possibility. For the likelihood on our evidence that the barns that we see are mere fake barns is much higher given that we are in fake barn country. And, even though we do not have sufficient reason to believe that the barns that we see are fake ones, we should reduce our confidence in the thesis that the barns that we see are genuine barns rather than fake barns constructed by film-makers if we come to know that we are in fake barn country. (Even though the possibility that the barns are genuine is still also a salient possibility).

So, I've shown that

- (4.1) There is a salient possibility, rather than a merely logical possibility, that there is or would be significant fundamental moral disagreement; and
- (4.2) If we should reject T if p , and there is a salient possibility that p , then T is less plausible than it otherwise would be and we should reduce our confidence in T (even if we do not have sufficient reason to believe p and even if there is a salient possibility that not- p).

But in §3 I showed that

⁵⁷ See Vavova (forthcoming: §4).

- (3) If there is significant fundamental moral disagreement, then our understanding of morality is fatally flawed and we should accept relativism over non-naturalist realism and quasi-realism.

And it follows from (4.1-2) and (3) that

- (4) Even if we do not have sufficient reason to believe that there is or would be significant fundamental moral disagreement, (3) weakens the plausibility of non-naturalism and quasi-realism.

That is, we should reduce our confidence in non-naturalism and quasi-realism given (3) (and (4.1-4.2)). And non-naturalism and quasi-realism are less plausible given (3) than they would have been if (3) were not the case.

5. Conclusion

I've shown that moral facts must be possibility accessible. And if like some you believe that realism is the view that moral facts need not be possibly accessible, then you should reject realism. I've shown that if there is significant fundamental disagreement, then our understanding of morality is fatally flawed and we should reject non-naturalist realism and quasi-realism. I also established that what is distinct about the moral domain, or moral facts, that makes it the case that fundamental disagreement is a problem for realist (and other) meta-ethical views in a way that disagreement in idealized conditions is not a problem for realist views about, for instance, history and metaphysics is that moral facts are epistemically constrained in this particular way, namely, that they must be possibly accessible.

Many think that there is or would be significant fundamental moral disagreement: Rawls, Williams, Horgan, Timmons, Zangwill, Streumer, Doris, and Stich just to name a few.⁵⁸ If you agree that there is or would be significant fundamental moral disagreement, then, you should not accept non-naturalism or quasi-realism. I've also shown that significant fundamental moral disagreement should matter to you even if you are not convinced that there would be such disagreement. I've shown that the fact that significant fundamental moral disagreement is a salient possibility weakens the plausibility of non-naturalism, quasi-realism, and the view that our understanding of morality is not fatally flawed.

David Brink, Derek Parfit, Sarah McGrath, Russ Shafer-Landau, and David Enoch have argued that facts about moral disagreement do not or could not undermine moral realism.⁵⁹ I've shown that this is a mistake. I've shown that significant fundamental moral disagreement is meta-ethically significant.

References

- Audi, Robert (2013). *Moral Perception*. Princeton, NJ: Princeton University Press.
Bedke, Matt (2011). 'Passing the Deontic Buck'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics Volume 6*. Oxford: OUP.
Bennigson, Thomas (1996). 'Irresolvable Disagreement and the Case Against Moral Realism'. *Southern Journal of Philosophy* 34 (4): 411-437.

⁵⁸ *Supra* note 48.

⁵⁹ *Supra* note 5.

- Björnsson, Gunnar (2011). 'Do "Objectivist" Features of Moral Discourse and Thinking Support Moral Objectivism?' *Journal of Ethics* 16: pp. 367-393.
- Bostrom, Nick (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford: OUP.
- Brandt, Richard (1959). *Ethical Theory: The Problems of Normative and Critical Ethics*. Englewood Cliff, NJ: Prentice-Hall.
- Brink, David (1989). *Moral Realism and the Foundations of Ethics*. Cambridge: CUP.
- Christensen, David (2009). 'Disagreement as Evidence: The Epistemology of Controversy'. *Philosophy Compass* 4/5: 756-767.
- Dancy, Jonathan (2004). *Ethics Without Principles*. Oxford: OUP.
- DePaul, Michael (1993). *Balance and Refinement: Beyond Coherence Methods of Moral Inquiry*. New York: Routledge.
- DeRose, Keith (1995). 'Solving the Skeptical Problem'. *The Philosophical Review*, 104, 1: 1-52.
- Doris, John and Alexandra Plakias (2008). 'How to argue about disagreement: Evaluative diversity and moral realism'. In Walter Sinnott-Armstrong (ed.), *Moral Psychology Volume 2: The Cognitive Science of Morality: Intuition and Diversity*. Cambridge, MA: MIT Press
- Elga, Adam (2007). 'Reflection and Disagreement', *Noûs* 41, 3: 478-502
- Enoch, David (2009). 'How is Moral Disagreement a Problem for Realism'. *Journal of Ethics* 13: 15-50.
- _____(2011). *Taking Morality Seriously*. Oxford: OUP.
- _____(2014). 'A Defense of Moral Deference'. *Journal of Philosophy* 111, 5: 229-258.
- Feldman, Richard (2006). 'Epistemological Puzzles about Disagreement'. In Stephen Hetherington (ed.), *Epistemology Futures*. Oxford: OUP.
- Finlay, Stephen (2014). *A Confusion of Tongues*. Oxford: OUP.
- Goldberg, Sandy (2013). 'Disagreement, Defeat, and Assertion'. In David Christensen and Jennifer Lackey (eds.), *The Epistemology of Disagreement*. Oxford: OUP.
- Graham, Peter (2010). 'In Defence of Objectivism about Obligation'. *Ethics* 121, 1: 88-115.
- Harman, Gilbert (1975). 'Moral Relativism Defended'. *The Philosophical Review* 84: 3-22.
- Horgan, Terry, and Mark Timmons (2008). 'Analytical Moral Functionalism Meets Moral Twin Earth'. In Ian Ravenscroft (ed.), *Minds, Ethics, and Conditionals*. Oxford: OUP.
- Huemer, Michael (2005). *Ethical Intuitionism*. Basingstoke: Palgrave Macmillan.
- Jackson, Frank (1998). *From Metaphysics to Ethics*. Oxford: OUP.
- Joyce, Richard (2001). *The Myth of Morality*. Cambridge: CUP.
- Kelly, Thomas (2010). 'Peer Disagreement and Higher-Order Evidence'. In Richard Feldman and Ted Warfield (ed.), *Disagreement*. Oxford: OUP.
- Kelly, Thomas, and Sarah McGrath (2010). 'Is Reflective Equilibrium Enough?' *Philosophical Perspectives* 24 (1): 325-359.
- Kornblith, Hilary (2010). 'Belief in the Face of Controversy'. In Richard Feldman and Ted Warfield (ed.), *Disagreement*. Oxford: OUP.
- Lackey, Jennifer (2010). 'A Justificationist View of Disagreement's Epistemic Significance'. In Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.), *Social Epistemology*. Oxford: OUP.
- Lord, Errol (2015). 'Acting for the Right Reasons, Abilities, and Obligation'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics Volume 10*. Oxford: OUP.
- Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*. New York: Penguin.
- Machery, Edward, Daniel Kelly, and Stephen Stich (2005). 'Moral Realism and Cross Cultural Normative Diversity'. *Behavioral and Brain Sciences*, 28, 6: 830.
- McNaughton, David (1988). *Moral Vision*. Oxford: Blackwell.
- McGrath, Sarah (2007). 'Moral Disagreement and Moral Expertise'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics Vol. 4*. Oxford: OUP.
- _____(2010). 'Moral Realism Without Convergence.' *Philosophical Topics* 38, 2: 59-90.
- Miller, Alexander (2003). *An Introduction to Contemporary Metaethics*. Blackwell: Polity Press.
- Olson, Jonas (2010). 'In Defence of Moral Error Theory'. In Michael Brady (ed.), *New Waves in Metaethics*. Basingstoke: Palgrave Macmillan.
- Parfit, Derek (2011a). *On What Matters Volume 1*. Oxford: OUP.

- _____(2011b). *On What Matters Volume 2*. Oxford: OUP.
- Rachels, James (2012). *The Elements of Moral Philosophy Seventh Edition*. New York: McGraw Hill.
- Railton, Peter (1989). 'Naturalism and prescriptivity'. *Social Philosophy and Policy* 7: pp. 151-174.
- Rawls, John (1974/1999). 'The Independence of Moral Theory'. In Samuel Freeman (ed.), *John Rawls: Collected Papers*. Cambridge, MA: Harvard University Press.
- Ross, W.D. (1930). *The Right and the Good*. Oxford: Clarendon.
- Sarkissian, Hagop, John Park, David Tien, Jennifer Cole Wright, and Joshua Knobe (2011). 'Folk Moral Relativism'. *Mind & Language* 26, 4: 482-505.
- Scanlon, T.M. (1998). *What We Owe To Each Other*. Cambridge, MA: Harvard University Press.
- _____(2002). 'Rawls on Justification'. In Samuel Freeman (ed.), *The Cambridge Companion to Rawls*. Cambridge: CUP.
- Schroeder, Mark (2007). *Slaves of the Passions*. Oxford: OUP.
- _____(2010). *Noncognitivism in Ethics*. Abingdon: Routledge.
- Shafer-Landau, Russ (2003). *Moral Realism: A Defence*. Oxford: OUP.
- Smith, Michael (1994). *The Moral Problem*. Oxford: Blackwell.
- Stich, Stephen (2009). 'The Persistence of Moral Disagreement'. *Leverbulme Lectures on Moral Psychology Lecture 2*, University of Sheffield, UK, May 11th 2009. Video available online at: <https://www.youtube.com/watch?v=mU0IrthDjV0>
- Sorensen, Roy (1995). 'Unknowable Obligations'. *Utilitas* 7, 2: 247-271.
- Streumer, Bart (2011). 'Are Normative Properties Descriptive Properties?' *Philosophical Studies* 154: 325-348.
- Suikkanen, Jussi (2005). 'Reasons and Value – A Defence of the Buck-Passing Account'. *Ethical Theory and Moral Practice* 7, 5: 513-535.
- _____(2008). 'Williamson and (Moral) Realism'. <http://peasoup.typepad.com/pea-soup/2008/03/williamson-and.html>.
- Tersman, Folke (2006). *Moral Disagreement*. Cambridge: CUP.
- Thomson, Judith Jarvis (1986). 'Imposing Rights'. In her *Rights, Restitution, and Risk*. Cambridge, MA: Harvard University Press.
- Vavova, Katia (forthcoming). 'Irrelevant Influences'. Forthcoming in *Philosophy and Phenomenological Research*.
- Weatherson, Brian (2013). 'Disagreements, Philosophical, and Otherwise'. In David Christensen and Jennifer Lackey (eds.), *The Epistemology of Disagreement*. Oxford: OUP.
- Wedgwood, Ralph (2010a). 'There are Unknowable Moral Truths'. <http://peasoup.typepad.com/pea-soup/2010/08/there-are-unknowable-moral-truths.html>.
- _____(2010b). 'The Moral Evil Demons'. In Richard Feldman and Ted A. Warfield (eds.), *Disagreement*. Oxford: OUP.
- Weiner, Matt (2007). 'Norms of Assertion'. *Philosophy Compass* 2, 2: 187-195.
- Williams, Bernard (1981). 'Moral Luck' in his *Moral Luck*. Cambridge: CUP.
- _____(1985). *Ethics and the Limits of Philosophy*. Abingdon: Routledge.
- Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: OUP.
- Wright, Crispin (1995). 'Truth in Ethics'. *Ratio* 8, 3: 209-226.
- Zangwill, Nick (2000). 'Against Analytic Moral Functionalism'. *Ratio* 13: 275-286.