



MORAL INTUITIONS, RELIABILITY AND DISAGREEMENT

BY DAVID KILLOREN

JOURNAL OF ETHICS & SOCIAL PHILOSOPHY

VOL. 4, No. 1 | JANUARY 2010

URL: WWW.JESP.ORG

COPYRIGHT © DAVID KILLOREN 2010

Moral intuitions, reliability and disagreement

David Killoren

Overview

THERE IS AN ANCIENT, YET STILL LIVELY, DEBATE in moral epistemology about the epistemic significance of disagreement. One of the important questions in that debate is whether, and to what extent, the prevalence and persistence of disagreement between our moral intuitions cause problems for those who seek to rely on intuitions in order to make moral decisions, issue moral judgments and craft moral theories. Meanwhile, in general epistemology, there is a relatively young, and very lively, debate about the epistemic significance of disagreement. A central question in that debate concerns peer disagreement: When I am confronted with an epistemic peer with whom I disagree, how should my confidence in my beliefs change (if at all)? The disagreement debate in moral epistemology has not been brought into much contact with the disagreement debate in general epistemology (though McGrath [2007] is an important exception). A purpose of this paper is to increase the area of contact between these two debates. In Section 1, I try to clarify the question I want to ask in this paper – this is the question whether we have any reasons to believe what I shall call “anti-intuitivism.” In Section 2, I argue that anti-intuitivism cannot be supported solely by investigating the mechanisms that produce our intuitions. In Section 3, I discuss an anti-intuitivist argument from disagreement which relies on the so-called “Equal Weight View.” In Section 4, I pause to clarify the notion of epistemic parity and to explain how it ought to be understood in the epistemology of moral intuition. In Section 5, I return to the anti-intuitivist argument from disagreement and explain how an apparently vulnerable premise of that argument may be quite resilient. In Section 6, I introduce a novel objection against the Equal Weight View in order to show how I think we can successfully resist the anti-intuitivist argument from disagreement.

1. Introduction

Consider two scenarios:

- (1) Jeb is deciding whether to do A . Jeb’s only goal is to do the morally right thing. To decide whether or not to do A , Jeb first carefully considers the non-moral (“brute”) facts about A and about the circumstance in which A would be performed. Then Jeb consults his intuition,¹ i.e., he asks himself whether, in light of

¹ What is a moral intuition? Some (e.g. Sinnott-Armstrong [2008a, p. 47]) have defined intuitions as immediate beliefs – immediate in that they are non-inferential. I agree that intuitions are immediate in this way, but I do not think the definition of intuitions as beliefs captures standard (philosophical) usage of the term “intuition.” One can doubt one’s intuitions (e.g.

the brute facts, A seems right or seems wrong. Jeb's intuition is that A is morally right. So Jeb does A .

- (2) Matt is deciding whether to do A . Matt's only goal is to do the morally right thing. To decide whether or not to do A , Matt takes a coin out of his pocket. Matt decides that "heads" means A is right and "tails" means A is wrong. Matt flips the coin. Heads is the result, so Matt does A .

My question is this: Are there any good reasons to think Jeb's way of deciding to act is no better than Matt's way of deciding to act? Matt's goal is to act rightly, but he has chosen a very bad way to aim at that goal – in fact, he could do no worse unless he would aim at acting wrongly. Is there good reason to think Jeb is in the same boat?

Why this question? Well, an affirmative answer would be disastrous for those who rely on their intuitions for guidance in moral decision-making. This fact, together with the fact that so many of us – philosophers and ordinary people alike² – do rely on our intuitions, seems sufficient to make my question interesting. Of course, a negative answer to this question would do little to vindicate reliance on intuitions. To vindicate reliance on intuitions, we would need to show that there are good reasons to think Jeb's way of deciding *is* better than Matt's. But I will not attempt to determine whether such reasons exist; the question I have chosen to examine is already difficult enough (as we will soon see).

Another question worth asking is whether Jeb is, or can be, *justified in believing* that A is the right thing to do. This is a question moral epistemologists have traditionally asked. But this question sets the bar higher than I (or Jeb) would want it. If I have a coin that is slightly biased toward heads, so that it comes up heads 50.1% of the time (and I know this), I am probably not justified in believing that the next flip will come up heads, but I might well be justified in *acting as if* the next flip will come up heads. For example, I ought to bet on heads in such a case (if I am not risk-averse). In general, it seems that it is easier to acquire justification to act as if p than it is to acquire justification to believe p . And for those of us who, like Jeb, want to know how to aim to do the right thing, justification to act as if A is right is sufficient; we

Peter Singer appears to have confessed to having intuitions that are inconsistent with his utilitarian beliefs). I think it is better to understand intuitions as *seemings* (so, on this point, I think I follow Bealer [1998] and Huemer [2006]). Just as a twig in water can seem bent even if I do not believe it is bent, we can intuit propositions we do not believe. Despite this, of course, most people do believe most of their intuitions. I will use "intuitive judgment" to refer to moral beliefs formed on the basis of one's intuitions.

² In fact, even though ordinary people are aware that they rely heavily on intuitions, they still may radically underestimate the degree of that reliance. A growing body of psychological research suggests that, even when it seems to us that we reach a judgment through some form of principles-based reasoning, it is often the case that the judgment was formed well in advance of conscious reflection. (See, e.g., Haidt [2001].) Thus, intuition both seems to be a large part of everyday moral thinking *and* is probably an even larger part of everyday moral thinking than it seems to be.

do not need justification to believe A is right (although we would like to have that too).

Of course, even if we were to establish that moral intuition is better (from the perspective of those aiming to do the right thing) than coin-tossing, this alone would not be enough to establish that one is justified in acting as if one's intuitions were correct. That is because there might be another source of information about morality that is even better than intuition. If so, then when information from that source conflicts with the results of intuition, one might not be justified in acting as if one's intuitions were true, even if intuiting (taken by itself) were better than coin-tossing. Still, the comparison of intuition with coin-tossing provides a baseline. I call the view that intuition is better than coin-tossing *intuitivism* and the view that intuition is no better than coin-tossing *anti-intuitivism*.³

I will sustain two assumptions throughout the course of this paper. The first assumption is that there is an intuition-independent moral reality: one's moral intuition that p does not by itself make p true; so any given intuition *could* be mistaken. (This is not equivalent with the somewhat more controversial claim that intuitions have *no role whatsoever* in making moral judgments true.) I think this assumption is highly plausible, but that is not the reason I make it. I make it because if it were false, intuitivism would certainly be true. So there is an interesting question about anti-intuitivism's prospects only if this assumption is made.

Secondly, I will assume that there is no source of information about morality other than intuition. This means that if intuiting is no better than coin-tossing, then (e.g.) there is no form of principles-based moral reasoning better than coin-tossing, either. I call this assumption *No Alternatives*. I assume No Alternatives simply because the stakes are highest, and therefore this paper's problem is most interesting, if it is true: Given No Alternatives, we would be completely in the dark about our moral obligations unless our intuitions have some degree of reliability. Of course, it might be that the assumption of No Alternatives will influence the outcome of my analysis. For example, it might be that, assuming No Alternatives, we can show that there are no reasons to believe anti-intuitivism, but without the assumption of No Alternatives, we could produce very good reasons to believe anti-intuitivism. This might be a reason not to assume No Alternatives. And, in fact, I do think that reasons for anti-intuitivism might be obscured by the assumption of No Alternatives. That is because, if there is an alternative source of information about the moral truth, and if intuition routinely clashes with that al-

³ I borrow the name "intuitivism" from Robert Audi [2005], but my usage differs somewhat from his (I think what he calls "intuitivism" implies what I call "intuitivism" but not vice versa). My main reason to use this name, rather than the more common name "intuitionism," is that I want to allow intuitivism to be free of associations with Intuitionism, the philosophical tradition given fullest expression in the philosophy of W.D. Ross, a tradition which typically goes far beyond the very minimal assertion I want to use "intuitivism" to pick out.

ternative source, this might provide us with very good reason to think intuition is completely unreliable. But this is simply obvious. The really interesting question, I think, is whether we might have reason to think intuition is completely unreliable if we *don't* have anything other than intuition on which to rely for moral guidance. To examine that question, we should assume No Alternatives.

2. Why anti-intuitivists should look beyond the mechanisms that produce our intuitions

One source of evidence that Jeb and Matt are in the same boat might be found in the persistence and prevalence of disagreement between intuitions.⁴ We may distinguish two kinds of disagreement. First, there can be *interpersonal* disagreement, as when Ellen intuits that slavery is wrong but Ellen's uncle intuits that slavery is permissible. Second, there can be *intrapersonal* disagreement, as when, on one occasion, Ellen intuits that slavery is wrong, while on another occasion, Ellen intuits that slavery is permissible. It has seemed to some that, if either or both of these kinds of disagreement are widespread, then intuition is obviously unreliable. Beginning in the next section, I will discuss in detail the nature of the problems posed for intuitivism by these forms of disagreement. But according to an alternative line of argument, we can show intuitions to be unreliable by investigating the psychological mechanisms that give rise to them. That is the line of argument I will examine in the present section. I will first say something about a particular instance of this line of argument, which proceeds from research conducted by Joshua Greene. Then I will say something more general about the prospects for this line of argument.

Joshua Greene [2001] has used fMRI brain-imaging techniques to investigate the way in which we make intuitive judgments about trolley cases introduced into the philosophical literature by Philippa Foot [1967] and Judith Thomson [1976]. As is well known, most people intuit that inaction is morally required in the so-called Push Case, but inaction is not morally required in the so-called Standard Case.⁵ Many philosophers have believed that our intui-

⁴ If intuitions are seemings rather than beliefs (see note 1), one might ask how disagreement between intuitions is possible. I will say that one person's intuition disagrees with another person's intuition exactly if things cannot (possibly) be the way they seem to be to the first person and (at the same time) be the way they seem to be to the second person. So, e.g., if I intuit that *A* is (all things considered) morally wrong, and you intuit that *A* is (all things considered) morally permissible, then presumably we disagree in the intended sense (assuming *A* cannot be, all things considered, both wrong and permissible).

⁵ In the Standard Case, a runaway trolley is hurtling down a track. The observer must choose whether to hit a switch or to do nothing. If she does nothing, then the trolley will continue along its current path and will kill five people who are tied to the track. If the observer hits the switch, then the trolley will go down a side track, avoiding the five, but killing one person who is tied to the side track. In the Push Case, a runaway trolley is hurtling down a track. The observer must choose whether to push a very large man from a bridge or do nothing. If

tions about these cases are obviously correct; the puzzle is that there is no obvious principled explanation why they are (both) correct. And so philosophers have spent decades looking for morally relevant differences between the two cases. However, Greene's research suggests that people's emotional reactions to the Push Case are different, and more intense, than their emotional reactions to the Standard Case, and Greene suggests that this difference in emotional response may explain why we typically have different intuitive judgments about these cases. Commenting on Greene's research, Peter Singer [2005] writes:

If...Greene is right to suggest that our intuitive responses are due to differences in the emotional pull of situations that involve bringing about someone's death in a close-up, personal way, and bringing about the same person's death in a way that is at a distance, and less personal, why should we believe that there is anything that justifies these responses? If Greene's initial results are confirmed by subsequent research, we may ultimately conclude that he has not only explained, but explained away the philosophical puzzle [posed by the trolley cases]. (p. 347)

Singer thinks that Greene has pointed the way to a debunking explanation of our intuitions about these cases. It is easy to imagine that this debunking project might be extended to our intuitions about many other cases. Perhaps it will eventually come out that all of our moral intuitions are explained by "emotional pull." If that happens, then it may seem that our intuitions would have shown themselves to be a quite poor guide to moral truth. We may then need to conclude that moral beliefs cannot be justified by intuitions. And – although Singer does not go quite this far – we might even need to conclude that our intuitions have no better chance of getting moral matters right than simply tossing a coin.

Singer is not the only one making this kind of an argument. Greene's work has inspired similar arguments,⁶ and Greene's is not the only work in moral psychology that has been marshaled to support a similar conclusion.⁷ But it is not entirely clear how a Greene-style debunking explanation successfully debunks. I see two different arguments that may be extracted from the paragraph excerpted above. First, Singer might intend to argue as follows: Clearly, the fact that we respond emotionally to a given action (e.g. pushing a

she does nothing, then the trolley will continue along its current path and will kill five people. If the observer pushes the very large man from the bridge, then he will land in front of the trolley, stopping its motion, so the five will be saved. But the very large man will die in the process.

⁶ People like Singer, including Greene himself (in e.g. his [2008]), seem to think that Greene's empirical research in the psychology of intuitions will produce some compelling evidence for some form of consequentialism. Broadly speaking, their idea seems to be either that consequentialism is motivated by something other than intuitions, or that consequentialism is motivated by intuitions for which no debunking explanation is forthcoming, or perhaps a mix of these.

⁷ Another line of argument that is in some ways similar has been offered by Sharon Street [2006]. I discuss Street's argument below.

person from a bridge in order to save five people) does not make that action right or wrong; therefore, if our intuitions are generated by our emotions, then they provide no good evidence as to what is right or wrong. But we should not attribute this argument to Singer. That is because this argument is horrible: the premise obviously provides no support for the conclusion. It is as if one were to say that a thermometer can provide no evidence of the temperature because the thermometer has no capacity to make it hot or cold.

A second line of argument might go as follows: We know that in many contexts, our emotional responses provide no evidence about right or wrong. For example, many people respond with disgust, even horror, at the sight of an open-heart surgery in progress; but that is no evidence that open-heart surgery is wrong. Therefore, if our intuitions are generated by our emotions, then they provide no good evidence as to what is right or wrong. I am not sure whether this is the sort of argument Singer intends; in any case, this argument appears to be stronger than the previous one. But it is still not very compelling. Here is why. When I observe an open-heart surgery in progress, I am fairly disgusted, but I experience no intuition that the surgeon is doing anything wrong. (That is why the example of open-heart surgery might seem useful for a defense of skepticism about intuitions generated by negative emotions.) By contrast, when I consider the act of pushing a man from a bridge to save five others, I respond with a negative emotion *and* I experience the intuition that the act is wrong. So, if we assume that Greene is right, then one of these emotions has a certain effect (it produces an intuition), while the other emotion lacks that effect. That means there must be a psychological difference, of some kind, between our emotional response to open-heart surgery and our emotional response to pushing a man from a bridge to save five others. Perhaps that difference – whatever it is – is epistemically important. Perhaps emotions of the kind we experience when we observe open-heart surgery provide no evidence about the moral truth, but emotions of the kind we experience when we imagine pushing someone from a bridge provide significant evidence about the moral truth. Until that possibility is ruled out, I do not see how the argument presently under consideration can go through. And obviously, one cannot rule out this possibility simply by pointing out that in *some* cases, such as the open-heart surgery case, our emotional response provides no evidence about the moral truth.

There are probably better readings of the passage from Singer excerpted above, although I do not know what they might be. So, rather than dwell on that passage any longer, I want to try to draw a general lesson from Singer's argument. At present, moral psychology has not given us a clear picture of the mechanisms by which our intuitions are generated. Perhaps our intuitions are generated by a fairly simple mechanism involving emotions, as Greene's research suggests. Perhaps there is a different mechanism. Perhaps the real mechanism, once discovered, will seem even less reliable, as a way to know moral truth, than the emotion-driven process posited by Greene. However, I suspect that, whatever the mechanism turns out to be, it will be quite difficult

to put together a compelling case for anti-intuitivism *simply* by inspecting that mechanism. This is because that mechanism must necessarily generate intuitions that we find to be intuitively plausible (since otherwise, it would not be the mechanism that generates our intuitions) and so it will seem to us that this mechanism is quite well-functioning indeed.

The fact that my own intuition-generating mechanism must produce intuitively plausible intuitions is not evidence that this mechanism reliably homes in on the moral truth. That fact merely precludes counterintuitive output of my intuition-generating mechanism, thus precluding one kind of evidence that my intuition-generating mechanism *does not* home in on the moral truth. And of course counterintuitive output of my intuition-generating mechanism is not the only possible kind of evidence that could support an anti-intuitivist's contentions. For example, if it turns out that my intuitions are produced by a malevolent demon, this would be evidence that my intuitions are not trustworthy.⁸ But a demon-controlled intuition-generating mechanism is obviously unreliable only because (we imagine) the demon knows the moral truth and wants to deceive us about it. Such a mechanism thus has a built-in truth-avoidance feature. It is much more difficult to imagine an intuition-generating mechanism that seems untrustworthy, solely in virtue of its psychological description, *and* (a) is describable in purely naturalistic terms, (b) does not involve the influence of an ill-intending intelligence in the functioning of human psychology, and (c) produces only intuitively plausible intuitions. Yet it is very likely that our intuition-generating mechanism, once fully revealed by psychological research, will be shown to have characteristics (a)-(c). Given this, it seems unlikely that a compelling case for anti-intuitivism will emerge solely from psychological research into the mechanisms that produce our moral intuitions.

Now I want to show that we will be able to say something similar even if we broaden our focus, and examine not just our intuition-generating mechanisms but also the processes by which our intuition-generating mechanisms are generated. I will show this by (very briefly) examining the case of biological evolutionary explanations for our intuitions. It looks as though some (but not all⁹) parts of our intuition-generating mechanisms are the

⁸ This sentence seems true, and I am willing to grant it for the sake of argument. But I believe it stands a very good chance of being false. I do believe that malevolent demons are very unreliable guides to the moral truth. But why do I believe this? I think it is because (a) there is a long list of moral judgments that I presently believe to be correct (e.g. I believe eating ice cream is usually morally permissible; burning cats for fun is usually morally wrong; etc.), and (b) I think the odds are very low that the output of a demon-influenced intuition-generating mechanism would match this list. Yet (b) would surely be false if my moral beliefs are shaped by intuitions produced by a malevolent demon. Thus, the revelation that my intuitions are produced by a malevolent demon might simply provide me with reason to revise my opinion about the reliability of malevolent demons (rather than provide me with reason to revise my opinion about the reliability of my intuitions).

⁹ This is evident from the fact that moral intuitions about many issues (including infanticide, slavery, gender roles, etc.) vary radically from culture to culture. Some of these differences

products of evolution through natural selection. Sharon Street argues that this natural history dims the prospects for moral knowledge, at least if realism about value is true (and, it bears emphasizing, I have already committed myself to realism).¹⁰ According to Street, if we think that realism is true, and we think that moral knowledge is possible, then we must endorse the “tracking account,” according to which “making certain evaluative judgments rather than others promoted reproductive success [and therefore were favored by natural selection] *because these judgments were true.*” Yet Street argues that the tracking account is not plausible. She thinks we should believe instead that making certain evaluative judgments rather than others promoted reproductive success, regardless of whether or not they are true, “because they forged adaptive links between our ancestors’ circumstances and their responses to those circumstances, getting them to act, feel, and believe in ways that turned out to be reproductively advantageous.” (This Street calls the “adaptive link account.”) So, Street concludes, there must be something wrong with a picture in which realism is true, moral knowledge is possible, and the evolutionary explanation of our moral intuitions is correct. (Street also thinks there is something wrong with a picture in which moral realism is true and moral knowledge is not possible. So she thinks that, given the evolutionary explanation of our moral intuitions, we should not be moral realists.)

Street’s issue is different from mine. Street wants to show (in the course of an argument against realism) that there cannot be moral knowledge unless realism is false. But we can readily grant this view without settling the question about intuitivism, the view that intuiting is better than coin-tossing. This is simply because, even if moral knowledge is impossible, intuiting might still be better than coin-tossing (just as I may be unable to know in advance who will win a very close election contest, even though a guess based on the latest poll numbers is better than a guess based on a coin toss). So I do not need to determine whether Street’s argument succeeds; the crucial question, for the

may be attributable to differences in non-moral beliefs, but not all. For example, Brandt [1954] describes Hopi children who torture their pet birds (c.f. Doris and Plakias [2008]). This behavior seems morally wrong to many of us, but apparently did not seem wrong to the Hopi children or to their elders, even though Brandt was unable to explain this moral disagreement by appeal to any disagreement about the brute facts (e.g. the Hopi seemed fully aware of the birds’ capacity to feel pain). So there appear to be cases of cross-cultural *fundamental* disagreement between moral intuitions. This kind of variation cannot be explained by natural selection simply because it is very unlikely that it can be explained by genetic differences between members of different cultures. Below, in Section 6, I will say a bit more about the epistemic significance of the role of culture in the shaping of our intuitions.

¹⁰ Street defines realism as the view that “there are at least some evaluative facts or truths that hold independently of all our evaluative attitudes.” Whether or not Street’s definition of realism is accurate or charitable is open to debate – a debate I will not enter here. For the limited purpose of evaluating the Street-inspired argument at issue here, I will regard realism as the view that there is an intuition-independent moral reality: one’s moral intuition that *p* does not by itself make *p* true, so any given intuition could be mistaken. This is a view I have assumed above; I will not evaluate the view here.

purposes of this paper, is whether Street-style evolutionary considerations provide any reasons to believe intuiting is no better than coin-tossing.

Let us simply grant Street's view that the adaptive link account is true and the tracking account is false. The tracking account would rule out anti-intuitivism, but the adaptive link account does not rule out intuitivism. For example, the adaptive link account is consistent with the view that intuitions and moral facts have a common explanans. The presence of such a common explanans could conceivably ensure that our intuitions reflect the moral facts more often than not. So the adaptive link account does not entail anti-intuitivism. Of course, the adaptive link account could provide a reason against believing intuitivism without entailing that intuitivism is false. But it is unclear what that reason might be. Natural selection, in Street's view, is responsible for our intuitions; therefore, natural selection has in our case produced a mechanism whose output is uniformly intuitively plausible, and thus immune from criticism by counterintuitive counterexample. I grant that there is nothing in what we know about natural selection that should make us think it a reliable way to produce creatures whose moral intuitions tend to be true. But this is just a lack of a reason to think intuitivism is true; it is not a reason to think intuitivism is false. Street's inspection of the evolutionary origins of our intuitions thus seems to be like Greene's inspection of the emotional infrastructure undergirding our intuitions: both lines of investigation can deprive the intuitivist of ammunition, but will probably not provide the anti-intuitivist with ammunition.¹¹

But none of this is to say that ammunition for the anti-intuitivist cannot be found. The phenomenon of moral disagreement might be fertile territory for the anti-intuitivist. An anti-intuitivist argument from disagreement would not need to posit an alternative to intuitions as a source of information about morality; its evidence for anti-intuitivism would consist entirely in observations about the frequency of conflict between moral intuitions. Such an argument also would not rely on inspection of the mechanisms by which our intuitions are produced. I think these are advantages of this line of argument. In the next section I will explain how such an argument might look.

3. An anti-intuitivist argument from disagreement

Suppose I consider some particular case in which slavery is practiced – say, a case typical of slavery in the antebellum American South, which I will hereafter call the Typical Case – and I intuit that slavery is wrong in this case, whereas you intuit that slavery is permissible in this case. Then one of us is mis-

¹¹ Below, in Section 6, I will argue that the situation is even worse than this for the anti-intuitivist: I will argue that evolutionary explanations for our intuitions might even provide an effective response to an anti-intuitivist argument from disagreement, thus depriving the anti-intuitivist as well as the intuitivist of some ammunition.

taken.¹² Once I discover the disagreement between us, it seems I have only three options: I can believe my intuition is correct; I can believe your intuition is correct; or I can suspend judgment.¹³ I do not think it is mistaken to say that I have only these three options, but I do think that saying this obscures a spectrum of possibilities. For instance, suppose that, prior to discovery of our disagreement, I believed with very high confidence that slavery is wrong. After discovering our disagreement, I might continue to believe this with the same high level of confidence, or I might continue to believe this with some significantly lower level of confidence. These seem like importantly different options, even though they are both ways of believing my intuition is correct. More generally: Suppose levels of confidence in a proposition p can be measured on an inclusive scale from 0 to 1, where 0 represents certainty that p is false, 1 represents certainty that p is true, and .5 represents suspension of judgment about whether p , i.e. agnosticism about whether p .¹⁴ Then we should say that, in the moment after I discover my disagreement with you, I have as many doxastic options as there are points in the interval from 0 to 1.¹⁵

If we are justified in issuing moral judgments on the basis of our intuitions, as intuitivists believe, then probably different intuitions justify different levels of confidence. I intuit that torturing innocent children for fun is morally wrong. I also intuit that pushing a man from a bridge to save five others from an oncoming train is morally wrong. But the first intuition seems to me to be somehow *phenomenally stronger*,¹⁶ and thus to justify much higher confi-

¹² This is so assuming (as I will assume here) that my intuition is that slavery is *all things considered* wrong, and your intuition is that slavery is *all things considered* permissible, and also assuming that slavery cannot be both all things considered wrong *and* all things considered permissible.

¹³ There might also be a fourth option: I might conclude that there is no fact of the matter about whether slavery is permissible. My own view is that taking this option would be every bit as troubling as (C), the conclusion of the anti-intuitivist argument from disagreement laid out below. So I believe that the intuitivist should want to avoid this option just as much as she wants to avoid (C). But I will not try to support this view here.

¹⁴ Some people – perhaps including Juan Comesana (personal correspondence) – would say that .5 should not be taken to represent suspension of judgment. They would say that each point along the spectrum from 0 to 1 represents a different judgment, so suspension of judgment is not found anywhere on the spectrum from 0 to 1. I do not need to dispute this point.

¹⁵ There is some debate among epistemologists about whether all belief-states fit on a continuous scale of the sort I am discussing here. The remarks I have made thus far do not imply a stance in that debate. I have not ruled out the possibility that one can believe p without having a precise level of confidence in p ; and I have not ruled out the possibility that believing p is something more (or something other) than having sufficiently high confidence that p .

¹⁶ A clear characterization of what I mean by “phenomenal strength” may prove elusive, but here are a couple of examples to illustrate the notion. “ $2+2=4$ ” seems true to me. “Modal realism is false” also seems true to me. But the first seeming is somehow more vivid and more compelling; it has more of what I am calling “phenomenal strength.” Similarly: Diet Coke seems like regular Coke to me. (I am not always able to tell the difference between Diet Coke and regular Coke, although I am able to tell the difference more often than not.)

dence, than the second intuition. Suppose that the degree of confidence justified by an intuition is determined solely by the phenomenal strength of that intuition (so that stronger intuitions justify higher levels of confidence). Call this the Strength-Confidence View. Those who accept this view (as well as others) may talk about my justified level of confidence, and distinguish this from my actual level of confidence. For example, it might be that, in light of my intuition that slavery is wrong, I have actual confidence 1 that slavery is wrong, whereas the phenomenal strength of my intuition only justifies confidence .9999 that slavery is wrong.

Intuitivism, I have said, is the view that issuing moral judgments on the basis of one's intuitions is better than issuing moral judgments on the basis of the outcome of a coin toss. The above points suggest a way to explain what is meant by "better" in that formulation. In the absence of any reason to believe that a given action A is or is not wrong, I should be an agnostic about whether A is wrong. A coin toss provides no reason to depart from this agnosticism. By contrast, the intuitivist should say, intuiting *does* provide some reason to depart from agnosticism. That is: For the intuitivist, if I intuit that A is wrong, then I am justified in having some level of confidence greater than .5 that A is wrong. This characterization of intuitivism leaves open the question *how much* of a departure from agnosticism my intuitions justify; different intuitivists will give different answers to that question. For example, intuitivists who accept the Strength-Confidence View will say that the extent of departure from agnosticism justified by an intuition is determined by the phenomenal strength of the intuition. But all intuitivists should agree (at a minimum) that some of my intuitions are *epistemically significant*, that is, some of my intuitions sometimes justify a departure from agnosticism.

Thus we have the first premise of the anti-intuitivist argument from disagreement:

- (1) If intuitivism is true, then some of my intuitions are epistemically significant.

According to the next premise in the argument, I should attribute equal epistemic significance to each intuition of which I am aware, and whether that intuition is mine or someone else's, except in cases where there is some reason to discount or disregard a given intuition.¹⁷ That is:

Regular Coke also seems like regular Coke to me. But I think the first seeming is typically less compelling than the second seeming, which is phenomenally stronger.

¹⁷ This idea is plausible because it is consistent with *any* view about what does or does not count as a reason to discount someone's intuition. For example, it is consistent with the view that one should discount one intuition, with respect to another intuition, when the first intuition is phenomenally weaker than the second; it is also consistent with the denial of this view.

(2) Your intuitions have the same epistemic significance as mine unless there is some reason to attribute less significance to your intuitions than to mine.

The question raised by premise (2) is: What constitutes a reason to attribute less significance to your intuitions than to mine? A natural suggestion is to answer this question in terms of *epistemic parity*. The idea would be that, if my intuition conflicts with yours, e.g. if I intuit that slavery is wrong while you intuit that slavery is permissible, then I may not discount or disregard your intuition if you are my epistemic peer; I may discount or disregard your intuition only if you are my epistemic inferior. (Something similar would be true about my past selves: I would be permitted to discount or disregard my past self's intuition only if I were epistemically superior to my past self.) If this is right, then we should agree to the following view:

(3) **Equal Weight View:** There is no reason to attribute less epistemic significance to your intuition than to mine if I correctly consider you to be my epistemic peer or epistemic superior.¹⁸

To evaluate the Equal Weight View, we will need a definition for epistemic parity. What does it mean for one person to consider another person to be an epistemic peer? If I disagree with you about a given question, I consider you to be my epistemic peer if I consider you to be *as good as I am* at evaluating whatever sort of question is at issue – or (equivalently) if I consider your *epistemic credentials* to be as good as mine with regard to whatever sort of question is at issue. For example, I might consider you to be as good as I am at arithmetic, i.e. my arithmetical epistemic peer. Then (according to the Equal Weight View in its standard formulation) if we disagree about a given sum, I should not attribute more epistemic significance to your result than to mine. Rather, in the event of such a disagreement, I must (pending further evidence) be agnostic about whether my result or your result is correct.

The example gives rough sense of what epistemic parity is, but this will not be sufficient for evaluating the Equal Weight View as it appears in premise (3). We may have a fairly good intuitive grasp of what it means for one person to be as good as another person at *arithmetic*, but it is not at all ob-

¹⁸ For comparison, two alternative versions of the Equal Weight View deserve mention here:

Version 1: There is no reason to attribute less epistemic significance to your intuition than to mine if I consider you to be my epistemic peer or epistemic superior.

Version 2: There is no reason to attribute less epistemic significance to your intuition than to mine if you *really are* my epistemic peer or epistemic superior.

The view in (3) is entailed by Version 1 and by Version 2. Thus the view in (3) should be acceptable regardless of whether one accepts Version 1 or Version 2. In the next section I will discuss the difference between Version 1 and Version 2.

vious what it would mean for one person to be “as good as” another at *moral intuition*. Thus, before we can even understand the claim made by the Equal Weight View, let alone evaluate it, we will need to do some work to show how epistemic parity ought to be understood in the context of moral intuition. That is what I will do in the next section (Section 4).

Now for a brief interlude: I want to show how a skeptical argument offered by Sarah McGrath [2007] can be improved. This exercise will help motivate and explain the next premise in the anti-intuivist argument from disagreement. McGrath works with a conception of epistemic parity developed by Adam Elga [2007]. For McGrath (as for Elga), “[y]ou consider someone your epistemic peer with respect to a given question just in case: in advance of either of you reasoning about the issue, you would have predicted that the person in question was just as likely as you to arrive at the correct answer.” And for Elga (as for McGrath), the Equal Weight View is the view that “one should give the same weight to one’s own assessments as one gives to the assessments of those one counts as one’s epistemic peers.”¹⁹ However, unlike Elga, McGrath thinks the Equal Weight View has skeptical consequences when applied in the moral domain. To show this, McGrath considers the example of Ann, a conservative Republican who thinks abortion is morally abhorrent in most circumstances, and Beth, a liberal Democrat who thinks abortion is morally permissible in most circumstances. Suppose Ann and Beth meet. After this encounter, could either Ann or Beth hold on to their opinions about abortion without falling afoul of the Equal Weight View? McGrath thinks not; she writes:

[W]e would expect Ann and Beth to agree about the answers to any number of moral questions. We would expect them to agree, for example, that slavery is morally abhorrent, that it is wrong to cause others pain for the sake of one’s own amusement, that lying is *prima facie* wrong, and about countless other issues. ... In short, Ann and Beth’s disagreements about abortion..., although substantial, almost surely take place against a relatively wide background of shared moral beliefs. ... [T]he relatively wide background of agreement seems to tell in favor of Ann’s taking it that Beth is more or less equally likely to get the hard questions right. (p. 105-6)

If Ann should think that she and Beth are about equally likely to get the “hard questions” right, then (given Elga’s definition of epistemic parity) Ann should regard Beth as her epistemic peer. But then the Equal Weight View would require Ann to attribute the same significance to her own opinion as she does to Beth’s opinion. In that case, it is hard to see how Ann could be permitted to hold on to her pro-choice convictions after she learns of Beth’s pro-life convictions. Rather, it seems Ann should become an agnostic about the ethics of abortion.

¹⁹ This is different in several respects from the formulation of the Equal Weight View I give above; one important difference is that Elga’s definition refers to “assessments,” whereas my definition is tailored specially for the epistemology of *moral intuitions*, which are seemings (not assessments). See Section 4 for discussion relevant to this difference.

I think McGrath is right that the Equal Weight View has skeptical implications, but the example of Ann and Beth does not by itself show this. According to the Equal Weight View, Ann ought to attribute the same significance to *each* of the opinions of *all* the people she considers to be peers. So Ann might not be required to give up her pro-choice convictions after learning of Beth's opinions even if Ann considers Beth to be her epistemic peer. Ann would be required to give up her pro-choice convictions only if she finds that, among *all* of her peers, pro-life convictions are as common as pro-choice convictions.²⁰ And perhaps Beth is not the only person whom Ann considers to be a peer.

Similar points will hold when it comes to disagreement among moral intuitions about slavery. To show that I ought to give up my anti-slavery intuitions, it would not be enough to show that I have one epistemic peer with pro-slavery intuitions. Instead, what is needed is to show something like the following:

- (4) Among those who are either my epistemic peers or my epistemic superiors, a majority intuit that slavery is permissible in the Typical Case.

To motivate (4), one might begin with a seemingly plausible historical claim: Pro-slavery intuitions have been very common throughout most of human history, and therefore anti-slavery intuitions represent a minority of all intuitions concerning slavery. I will discuss this historical claim below, in Section 5; it is not obviously true. But *if* it were true, then to deny (4) we would have to claim that the majority of my epistemic peers and superiors are concentrated among a minority of all the individuals who have ever lived. This apparently provides a *prima facie* presumption in favor of (4) – an obligation not to deny (4) unless there is a compelling objection to (4).

²⁰ This may be counterintuitive. Suppose that, in one scenario, Beth is the only peer Ann knows, while in another scenario, Ann knows ten peers: Beth and nine others who have pro-choice opinions like Ann. If I am correct, then (according to the Equal Weight View) Ann should be moved to agnosticism by Beth's pro-life opinion in the first scenario, whereas in the second scenario Ann should retain her pro-choice opinion. But this makes it look as if Beth's opinion becomes (relatively) less significant as the size of Ann's acquaintances increases – a strange result. Even so, I think it is going to be very difficult to avoid this result without giving up the Equal Weight View. Suppose that Ann is moved to agnosticism in the second scenario. Then Ann (tacitly) considers Beth's pro-life opinion to be ten times as significant as any of her peers' pro-choice opinions. This is because, in the second scenario, there are ten peers (including Ann herself) who have pro-choice opinions – so if Ann becomes an agnostic about whether abortion is wrong, she treats these ten pro-choice opinions as if they are precisely counterbalanced by Beth's lone pro-life opinion. It would be contrary to the Equal Weight View to treat Beth's opinion as if it is ten times as significant as any of the other peers' opinions. By contrast, in the first scenario, there are only two peers: Ann and Beth. So if Ann is moved to agnosticism in that case, she treats all known peers' opinions as if they are equally significant, just as the Equal Weight View prescribes.

The foregoing is obviously not a complete defense of (4). I offer these points here only to give an idea why someone might believe (4). In Section 5, I will discuss in depth the possibilities for defending and attacking (4).

We have thus far assembled the following premises: (1) If intuitivism is true, then some of my intuitions are epistemically significant. (2) Your intuitions have the same epistemic significance as mine unless there is some reason to attribute less significance to your intuitions than to mine. (3) Equal Weight View: There is no reason to attribute less epistemic significance to your intuition than to mine if I correctly consider you to be my epistemic peer or epistemic superior. (4) Among those who are either my epistemic peers or my epistemic superiors, a majority intuit that slavery is permissible in the Typical Case. These premises strongly support the following conclusion:

(C) If intuitivism is true, then I should not believe that slavery is wrong, and should not be an agnostic about whether slavery is wrong in the Typical Case; rather, I should have some confidence greater than .5 that slavery is permissible in the Typical Case.²¹

I think it is fairly clear that (C) makes intuitivism unacceptable. I am highly confident that slavery in the Typical Case is wrong – and highly confident that this confidence is justified. If intuitivism implies I must doubt whether slavery is wrong, then I must abandon intuitivism. So the intuitivist should want to deny one of premises (1)-(4). (1) and (2) seem obviously true, and I think should not be disputed. So the intuitivist ought to concentrate her attack on either (3) or (4). In Section 5, I will consider what can be said against (4); in Section 6, I will attack (3). But before we get to all that, it will be useful to address a few technicalities having to do with the concept of epistemic parity.

4. Epistemic parity in the epistemology of moral intuition

In trying to use the notion of epistemic parity in the context of this paper, I face two problems. The first problem is that the *standard* notion of epistemic parity is used in conventional epistemology and is therefore typically understood in terms of belief (or other doxastic states). But moral intuitions are not beliefs (or other doxastic states); they are, I have claimed, seemings (see note

²¹ It is not actually the case that (C) follows from (1)-(4). What follows from (1)-(4) is that, if intuitivism is true, then: If (a) I intuit that slavery is wrong, (b) my intuition that slavery is wrong (taken by itself, absent any other evidence) justifies a departure from agnosticism, (c) I can know that premise (4) is true and (d) I have no other evidence (beyond my own intuitions and my knowledge of others' intuitions) about the morality of slavery, *then* I should have some confidence greater than .5 that slavery is permissible. But (a) is certainly true, (b) seems eminently plausible from any viable intuitivist perspective, (c) is probably true if (4) can be defended and (d) is guaranteed by the assumption of No Alternatives from Section 1. Thus, if (1)-(4) are not contested, then (C) will prove to be very difficult to avoid.

1). Therefore, since moral intuitions are my focus in the present context, I will not be able to rely on an unmodified form of the standard notion of epistemic parity; I will need to develop an alternative notion. Ordinarily, the way to develop such an alternative notion would be to begin with the standard notion as a model and to adjust it in order to make it suit its new context. But this will not work, because there is in fact no such thing as *the* notion of epistemic parity. Rather, there are several different definitions in use by different epistemologists. Worse, these different definitions correspond with slightly different ways of understanding the Equal Weight View. This diversity of available models is the second problem I face.

I cannot here resolve existing disagreements about how we should understand epistemic parity, so I will have to remain neutral between several major approaches. In this section, I will discuss three different definitions of epistemic parity and will explain how we can use each of these three definitions as models for developing alternative definitions of epistemic parity for use in the epistemology of moral intuition. Then, in the next section (Section 5), I will use these results to explore the possibilities for a case against premise (4) of the anti-intuivist argument from disagreement.

Adam Elga (whose view I discussed briefly in the previous section) takes what I shall call *the probabilistic approach* to epistemic parity. That approach is represented in the following definition:

Definition 1: S_1 considers S_2 to be an epistemic peer, with respect to a question Q , exactly if, prior to knowing S_2 's opinion about Q , S_1 considers S_2 to be as likely as S_1 to answer questions like Q correctly.²²

An important feature of Definition 1 is that it says only what it is to *consider* someone to be my peer; it does not give conditions for someone to *actually be* my peer. Thus Definition 1 leaves open how to determine whether I *correctly* consider you to be my epistemic peer. In fact, some philosophers who take the probabilistic approach maintain that there is not much at stake in the question whether you actually are my peer. These philosophers may say that there is simply no fact of the matter about whether you actually are my peer; or they may say that you actually are my peer just in case I consider you to be my peer.²³ In their view, if I find that my opinion conflicts with your opinion,

²² To fully understand this definition (and the alternative definitions considered below), it would be necessary to have criteria by which to determine whether a given question is “like” another. Spelling out such criteria is a lengthy project I will not be able to undertake here. Fortunately, it will often be clear whether one question is relevantly like another: for instance, any given arithmetic question is usually relevantly like any other arithmetic question of similar difficulty; and any given arithmetic question is usually not relevantly like the question of whether Napoleon died in France or the question of whether the man who just passed us on the street is Tom Cruise. But there will be hard cases, and a complete elaboration of the concept of epistemic parity would need to provide a way to decide such cases.

²³ I take Juan Comesana (personal correspondence) to be inclined to endorse the second of these possibilities.

I can determine how (if at all) I ought to modify my opinion by determining whether I consider you to be my peer; I do not need to determine whether you actually are my peer.

An alternative view of epistemic parity is represented in what I will call the *epistemic-virtues approach*, which is expressed by Thomas Kelly [2006]:

[T]he class of epistemic peers with respect to a given question are equals, not only with respect to their possession of the sort of general epistemic virtues enumerated by Gutting [which are: “intelligence, perspicacity, honest, thoroughness, and other relevant epistemic virtues”], but also with respect to their exposure to evidence and arguments which bear on the question at issue. (p. 3)

Despite the specificity of the virtues listed here, those who accept a Kelly-style conception of epistemic parity will accept that the requirements for epistemic parity differ from one domain of inquiry to another. For example, if you and I disagree about a difficult question in mathematics, the fact that I am less intelligent than you might prevent me from being your epistemic peer. But if you and I disagree about whether the man who just passed us on the street is Tom Cruise,²⁴ the fact that I am less intelligent than you is unlikely to make much difference; it is much more important to know whether my eyesight is as good as yours, whether I have seen as many Tom Cruise movies as you, whether I am as good at recognizing faces as you are, etc. To put the point another way: intelligence is an important epistemic virtue in mathematics, whereas good eyesight is not; good eyesight is an important epistemic virtue in the recognition of faces, whereas intelligence is not. Presumably something similar will be true about epistemic vices: bad eyesight, or stupidity, will be vices in some contexts, not others. Thus, in the present view, to determine whether you are my epistemic peer, I must determine the epistemic virtues and vices relevant to whatever is the question at issue; and I must evaluate my mix of vices and virtues and compare it with your mix of vices and virtues.

Here is a definition of epistemic parity to represent the epistemic-virtues approach:

Definition 2: S_1 and S_2 are epistemic peers, with respect to a question Q , exactly if S_1 's mix of epistemic virtues and vices is just as good (for the purpose of answering questions like Q) as S_2 's mix of epistemic virtues and vices.

Those who take the epistemic-virtues approach will need to explain how to compare different mixes of vices and virtues. This question is going to prove difficult.²⁵ But according to one plausible criterion, I am more virtuous (and

²⁴ This example comes from Sinnott-Armstrong [2002].

²⁵ To illustrate one source of difficulty, consider an example modified from Christensen [2007]. Suppose I believe that you are just as intelligent, thoughtful and free from bias as I

less vicious) than you (with respect to a given question Q) exactly if my belief-forming process, when used to form beliefs about questions like Q , is more *reliable* than yours.²⁶ If we accept this criterion, we may be able to dispense with the language of virtues and vices altogether and adopt what I will call the *reliability approach* to epistemic parity, in which we may use a definition of epistemic parity like the following:

Definition 3: S_1 and S_2 are epistemic peers, with respect to a question Q , exactly if S_1 's belief-forming process, when used to form beliefs about questions like Q , is just as reliable as S_2 's belief-forming process.

Thus far I have laid out three different conceptions of epistemic parity. I think each is initially plausible. These definitions are not obviously inconsistent with one another, although I suspect that Definition 3 is inconsistent with Definition 1.²⁷ Now I want to note that these three definitions correspond with different ways of spelling out the Equal Weight View. On the one hand, there is an *internalist* version of the Equal Weight View, which says that if I *consider you to be* my epistemic peer, then I ought to attribute the same epistemic significance to your opinion as to mine. On the other hand, there is an *externalist* version of the Equal Weight View, which says that if you *really are* my epistemic peer, then I ought to attribute the same epistemic significance to your opinion as to mine. A definition like Definition 1 (the probabilistic approach) may be useful in articulating the internalist version; a definition like Definition 2 (the epistemic-virtues approach) or Definition 3 (the reliability approach) may be useful in articulating the externalist version. It is not easy to decide whether the externalist version or the internalist version of the Equal Weight View is better; both have apparent disadvantages.²⁸ But

am. We both compute 376×215 . My result is 80840. Your result is 2. It seems clear that, in this case, I should not attribute any epistemic significance to your opinion – but it is not at all clear whether I have found reason to think you lack some relevant vice. Those who accept Definition 2 will need to show how your result has revealed some epistemic vice that I did not previously know you possessed, or else they will have to abandon the Equal Weight View.

²⁶ Roughly: One belief-forming process is as reliable as another, with respect to questions like Q , if the first process produces true beliefs about questions like Q as often as the second process does. Thus reliability is a certain sort of measure of accuracy.

²⁷ Definition 3 would be consistent with Definition 1 if it were the case that: S_1 considers S_2 to be as likely as S_1 to correctly answer questions like Q *exactly if* S_1 judges that S_2 's belief-forming process is just as reliable as S_2 's belief-forming process when applied to questions like Q . Unfortunately, this biconditional is probably not strictly true (for reasons I will not go into here), although it might be approximately true.

²⁸ The internalist version implies that, if I consider you to be my epistemic peer for very bad reasons, I still ought to attribute equal significance to your opinion as to mine. The externalist version implies that, if in fact you are my epistemic peer, but I have no way of knowing this, I still ought to attribute equal significance to your opinion as to mine. Obviously, neither of these counterintuitive implications alone provides decisive refutation of either the internalist or the externalist versions.

note that these two versions of the Equal Weight View have conflicting implications primarily in cases where one person *mistakenly* judges another person to be her epistemic peer. Therefore it is mainly when dealing with such cases that we will need to pay special attention to the difference between the internalist and externalist versions. (Of course, as we have seen, it is at least arguable that such cases are not even possible.)

We can use each of the definitions I have reviewed as a model for developing a conception of epistemic parity for use in the context of moral intuition. Following Definition 3, which defines epistemic parity in terms of the reliability of *belief*-forming processes, we may wish to understand epistemic parity, in the context of moral intuition, in terms of the reliability of *intuition*-forming processes. In that case, we may use the following definition of epistemic parity:

Definition 3.I: S_1 and S_2 are epistemic peers, with respect to a case C , exactly if S_1 's intuition-forming process, when used to form intuitions about cases like C , is just as reliable as S_2 's intuition-forming process.²⁹

Similarly, if different epistemic virtues and vices are relevant in different kinds of truth-seeking endeavors, then perhaps too there are virtues and vices relevant in the domain of moral intuition. For example, it seems not unreasonable to think that, if you do not have enough intelligence to understand the details of a given case, then you cannot be my peer (provided my intelligence *is* sufficient to understand the details of the case). (I will discuss some other potentially intuition-relevant virtues in the next section.) So, in the domain of moral intuition, the possession of a minimum level of intelligence looks like an epistemic virtue. Such considerations inspire the following definition:

Definition 2.I: S_1 and S_2 are epistemic peers, with respect to a case C , exactly if S_1 's mix of epistemic virtues and vices is just as good as S_2 's mix of epistemic virtues and vices.

It is also possible to take a probabilistic approach to understanding epistemic parity in the domain of moral intuition. In this approach we would define consideration of someone to be a peer as follows:

Definition 1.I: S_1 considers S_2 to be an epistemic peer, with respect to a case C , exactly if, prior to knowing S_2 's intuition about C , S_1 considers S_2 to be as likely as S_1 to have correct intuitions about C .

²⁹ Roughly: One intuition-forming process is as reliable as another, with respect to cases like C , if the first process produces true intuitions as often as the second process does. Thus, as with belief-forming processes, reliability here can be considered to be a measure of accuracy.

I do not know which of these definitions we ought to adopt, so in what follows I will strive to remain neutral between them. In the next section I will discuss some arguments for and against premise (4) of the anti-intuivist argument from disagreement.

5. How many of my epistemic peers have pro-slavery intuitions?

Here again is premise (4):

(4) Among those who are either my epistemic peers or my epistemic superiors, a majority intuit that slavery is permissible in the Typical Case.

Here is a quick-and-dirty argument in favor of (4). Most people, during most of human history, have had pro-slavery intuitions; in fact, the intuition that slavery is wrong in cases like the Typical Case is a relatively recent phenomenon. Thus we can assume that the sheer number of people who intuit (or when alive did intuit) that slavery is permissible in the Typical Case is significantly larger than the number of people who intuit that slavery is wrong in the Typical Case. Given this, the only way to resist (4) is to argue that the *percentage* of people with anti-slavery intuitions who are my peers (or superiors) is much higher than the percentage of people with pro-slavery intuitions who are my peers (or superiors). That is: In order to resist (4), one must argue that the epistemic credentials of people with anti-slavery intuitions tend to be significantly better, on average, than the epistemic credentials of people with pro-slavery intuitions. But it turns out not to be possible to make that argument. No matter which definition of epistemic parity we prefer, people with anti-slavery intuitions turn out to be no more likely (on average) to be my epistemic inferiors than people with pro-slavery intuitions. Thus, we cannot find a good argument against (4). But in the absence of a good argument against (4), we should accept (4). As we have just seen, in order to reject (4) we have to claim that people in a certain group (i.e. the pro-slavery group) tend to have poorer epistemic credentials than people in another group (i.e. the anti-slavery group). The burden of proof is on anyone who makes a claim of this form, because in the absence of reasons to think otherwise, we should assume that the members of any two groups have equally good epistemic credentials; otherwise, we would be guilty of a certain kind of (epistemic) prejudice.

There are three main points at which the preceding quick-and-dirty argument for (4) can be criticized. First, there is the argument's claim that, in the scope of human history, the number of people who have or have had pro-slavery intuitions is greater than the number of people who have or have had anti-slavery intuitions. It is hard to see how one could verify this claim; in any event, this claim is not obviously true, and there are reasons to think it is false (as I will show below). Another point for criticism is the argument's claim that the burden of proof is on those who wish to deny (4). Indeed,

there is some question whether it ever makes sense to say that one side or the other in a given dispute has the “burden of proof.”³⁰ But set these two points to one side, for the moment. A third point for criticism is the argument’s claim that, no matter which definition of epistemic parity we prefer, people with anti-slavery intuitions turn out to be no more likely (on average) to be my epistemic inferiors than people with pro-slavery intuitions. I want to examine the case against this claim in some depth, with reference to each of the three definitions of epistemic parity discussed in the previous section.

Suppose we accept Definition 2.I, the definition of epistemic parity in terms of epistemic virtues and vices. Can we show that people with pro-slavery intuitions tend to lack certain key virtues possessed by those with anti-slavery intuitions?

Degree of consistency between one’s intuitions might be thought to be an epistemic virtue.³¹ For example, it might be that there is greater consistency between my anti-slavery intuition and my other intuitions than there typically is between the pro-slavery intuitions of my pro-slavery opponents and their other intuitions. If this were true, it would go a long way toward impugning the epistemic credentials of those with pro-slavery intuitions. But what would it mean for one intuition to be inconsistent with another intuition? Certainly, if I intuit that *A* is wrong, and at the same time I intuit that *A* is permissible, then my intuitions are inconsistent. But this will not help much, simply because this sort of intuitional inconsistency is probably quite rare – even among those who intuit that slavery is permissible in the Typical Case.

Another possibility is that inconsistency arises when one is unable to identify relevant differences to explain divergent intuitions. For example, imagine a slave-owner who intuits that it would be morally wrong for him to enslave his own mother but also intuits that it is morally permissible for him to enslave some particular woman – a woman who happens to be one of his slaves – even though he is unable to identify a relevant difference between his mother and the slave. I suspect that most slave-owners would have found themselves in this kind of position. Does this mean most slave-owners exhibited the vice of inconsistency? Unfortunately, no. Such a slave-owner’s position is fully consistent. To see this, recall the two trolley cases discussed above: Standard Case and Push Case. Many people intuit that inaction is required in Push Case but not required in Standard Case, despite being unable to find a relevant difference between the two cases, even after expending great effort looking for such a difference. It seems clear that such people’s

³⁰ A valuable discussion of this question recently occurred on the blog *PEA Soup*, in the comments to a post by Michael Cholbi: <http://peasoup.typepad.com/peasoup/2009/07/shifting-burdens-of-proof-in-practical-ethics.html>.

³¹ The following discussion has been significantly influenced by insights provided by an anonymous reviewer.

intuitions are not inconsistent.³² (Perhaps these people are criticizable on other grounds, but it is not obvious what those grounds might be.)

A third possibility is that inconsistency arises when one has divergent intuitions about cases one believes to be similar in all relevant respects. For example, suppose the slave-owner intuits that enslaving his mother is wrong, and intuits that enslaving the slave is permissible, while at the same time believing that there is no relevant difference between enslaving his mother and enslaving the slave. I grant that this would be a genuine form of inconsistency and a serious epistemic vice, but I wonder how many people with pro-slavery intuitions are really guilty of this form of inconsistency. My suspicion is that not very many are.³³ Thus I think it will be difficult to show that people with pro-slavery intuitions tend in large numbers to exhibit vicious inconsistency.

Perhaps *familiarity with relevant brute facts* is an epistemic virtue. For instance, suppose you believe that people of African descent have the intelligence and emotional capacities of chickens or some other farm animal, and (in large part because you hold this mistaken belief) you intuit that slavery in

³² It might be claimed that this analogy fails because, even if some people cannot quite explain why they think Standard Case and Push Case differ morally, most people are able to offer (or at least gesture in the direction of) a few *potentially* relevant differences between Standard Case and Push Case. But first of all, this sets a very low bar; I suspect that most slave-owners would be able to offer *potentially* relevant differences that *might* explain their divergent intuitions about the morality of slavery. Secondly, it seems to me that consistency simply does not require us to clear even this very low bar. Suppose Leroy has been thinking about Standard Case and Push Case for some time. Leroy has the usual intuitions about these cases. When Leroy first learned about these cases, he had a few half-formed ideas about why these cases differ morally, but he has found that none of his initial ideas stand up to scrutiny (so he has abandoned them) and he has not come up with any new ideas. Leroy suspects that, if he were to continue thinking, he might find a relevant difference between Standard Case and Push Case – but now Leroy has to leave for work, and he will not be able to think about these cases again anytime soon. Suppose that Leroy continues, after this point, to believe that inaction is required in Push Case but not required in Standard Case. Does this make Leroy inconsistent? Clearly not, I think.

³³ In claiming this, it might be thought that I am suggesting that a typical slave-owner would be a Hare-style fanatic, who would be willing to be enslaved if his own enslavement turns out to be required by whatever principle he believes justifies his enslavement of others. But I have not suggested this (and do not believe it). I suspect that a typical slave-owner would intuitively judge his own enslavement to be wrong. And I doubt that a typical slave-owner would be able to provide any defensible principle that would permit him to enslave others but would forbid others from enslaving him. But (for reasons I have already discussed) I do not see that this would make him inconsistent. Inconsistency *would* arise if a slave-owner endorses a principle that unavoidably implies either that his enslavement of others is wrong, or that others' enslavement of him would be permissible. However, I suspect that almost any principle endorsed by a slave-owner that *appears* to have such implications could be shown not to have such implications. For example, Jefferson's phrase "All men are created equal," which is often taken to represent a moral principle of some kind, and which has been endorsed by many slave-owners, does not seem to me to have any interesting implications at all (unless supplemented by auxiliary principles). But I do not have the space to make the case for this view here.

the Typical Case is permissible. Then it would seem very strange to think you are my epistemic peer (assuming I do not have similarly mistaken beliefs).

I grant that peers ought to have similar levels of familiarity with the relevant brute facts. But I am not sure that people with pro-slavery intuitions usually do worse, on this score, than people with anti-slavery intuitions. Consider that many of the people who have had pro-slavery intuitions lived in slave societies, so either owned slaves themselves, worked with slaves, or had seen and heard about slaves. I conjecture that wildly false beliefs about slaves would not survive very well in such an environment.³⁴ For example, in a slave society like the antebellum South, it would be well-known that slaves can do things like marry and speak English; such facts would prove to anyone with common sense that slaves have high levels of intelligence and emotional capacities. Of course, there can be little doubt that people living in slave societies would have had many false beliefs. Even so, it seems probable that a person living in such a society would have been aware of most of the brute facts that we would offer as reasons why slavery is wrong. That is: I believe that I could explain why slavery is wrong, to someone from such a society, in terms of beliefs widely held in such a society.³⁵ Indeed, people in slave societies might have been privy to relevant facts unknown to us: consider that a slave-owner would have personally witnessed many of the cruelest effects of slavery, whereas I can only dimly and imperfectly imagine these facts. Thus, it is at least not obvious that people with anti-slavery intuitions would normally have better familiarity with the relevant brute facts than people with pro-slavery intuitions.

Let us now turn to reliability. Suppose it turns out that the intuitions of people in the anti-slavery camp are generally more reliable than the intuitions of people in the pro-slavery camp. This could provide the basis for a very strong case against (4). Indeed, this would provide for an *airtight* case against (4) if we were to accept the reliability approach to epistemic parity exemplified in Definition 3.I. Further, the case against (4) might also be strong if we were to accept the epistemic-virtues approach of Definition 2.1, because (as we saw in Section 3) the fact that one person's intuitions are more reliable

³⁴ When I speak of a “wildly false” non-moral belief about slaves, I primarily have in mind a belief *B* such that, if *B* were true, then slavery would (by our lights) be morally defensible. For example, if it were the case that human slaves have the intelligence and emotional capacities of farm animals, then slavery might be defensible (although even this is not entirely clear). My conjecture is that this sort of belief would not be held by the vast majority of those who live in a slave society. See note 35.

³⁵ This is to say that I could simply grant all of a typical slave-owner's non-moral beliefs, including his mistaken beliefs, and still have sufficient reason to believe slavery is wrong. For example, it appears that in the antebellum South it was widely believed that people of African descent are innately childlike intellectually and emotionally. Even if I were to grant this mistaken belief, I would still have good reasons to consider the enslavement of people of African descent to be wrong. (These are the same reasons why I would say it is wrong to enslave anyone who is childlike – e.g. mentally disabled people, actual children, etc.)

than another's can, on some views, show that the first person is more epistemically virtuous (and less vicious) than the second person.

How could we support the view that the intuitions of people in the anti-slavery camp are generally more reliable than the intuitions of people in the pro-slavery camp? Here is an argument that seems too easy. The people in the pro-slavery camp intuit that slavery is permissible. But slavery is obviously wrong. Getting such questions wrong is strong evidence of unreliability. By contrast, the intuitions of people in the anti-slavery camp are obviously correct (at least when it comes to the ethics of slavery). Thus we have evidence that the intuitions of those in the pro-slavery camp are less reliable than the intuitions of people in the anti-slavery camp.

This argument seems too easy because it looks like it must somehow beg the question. But question-begging is not always bad.³⁶ For instance, this argument begs the question against those who believe that slavery is permissible; but I would think we should have little interest in addressing such people, so it is not bad to beg the question against them. However, the argument might beg the question against the anti-intuitivist. After all, we are presently considering ways to attack premise (4) of the anti-intuitivist argument from disagreement. The conclusion of that argument is that intuitivists are committed to the obviously wrong view that slavery is permissible. So perhaps it begs the question, in a bad way, for the intuitivist to attack (4) by means of an argument that relies on the obviously correct view that slavery is wrong.

Whether this charge of question-begging (of the bad kind) can stick to the foregoing argument is unclear. But we might be able to circumvent that charge altogether. Suppose we were to survey the intuitions of people in the pro-slavery camp about a wide variety of different moral issues – issues that have little to do with slavery. And suppose we were to find that, in each case, their intuitions conflict with the intuitions of those of us in the anti-slavery camp. This result could produce evidence of unreliability that does not assume that slavery is wrong, and therefore would not even arguably beg the question against the anti-intuitivist. Moreover, this result could be relevant even if we were to eschew the reliability approach to epistemic parity and adopt the probabilistic approach as exemplified in Definition 1.I. The fact that members of a certain group tend to have intuitions that conflict with mine might justify me in judging that the intuitions of members of this group are less likely to be correct than mine.

But I think this strategy will not work. In general, if the intuitions of people in one group routinely conflict with the intuitions of people in another group, such conflicts might be evidence that the intuitions of members of *both* groups are unreliable, but it is difficult to see why such conflicts should be evidence that one group has less reliable intuitions than the other. Similarly, such conflicts might justify me in judging that neither group is likely to

³⁶ Shafer-Landau [2009] includes valuable insights about when question-begging is and is not a mistake; some of what I say here follows him.

have correct intuitions very often, but it is difficult to see why such conflicts should be evidence that the intuitions of members of one group are less likely to be correct than the intuitions of members of the other group.³⁷

I have just considered a few ways to argue that people with anti-slavery intuitions are less likely (on average) to be my epistemic peers (or superiors) than people with pro-slavery intuitions. I have certainly not considered *all* ways to make that argument, but I hope I have shown that there will be some difficulty in defending this position. Of course, even if that argument cannot succeed, there are other ways to undermine the quick-and-dirty argument for (4) that I gave at the beginning of this section. One could argue, for example, that pro-slavery intuitions are very unusual in history. Perhaps it is really quite rare for slavery to seem morally permissible to anyone; perhaps, in slave societies, slavery is practiced for its material and economic benefits *despite* the fact that it seems wrong to most people in such societies – or would seem wrong to them, if they would allow themselves to consider the moral question. Mundane demographic considerations are also relevant here. Due to technological advances in the 19th and 20th centuries, anti-slavery intuitions have become widespread during a period in which the human population has exploded. Thus, even if *everyone* born prior to (say) the year 1800 had had pro-slavery intuitions, it still might be that anti-slavery intuitions are more common among all the people who have ever lived. This would be enough to show that (4) is false even if people with anti-slavery intuitions are no more likely to be my epistemic peers (or superiors) than people with pro-slavery intuitions.

Thus, it is far from obvious that the quick-and-dirty argument for (4) succeeds. But it is also not obvious that it cannot succeed. Some intuitivists will be satisfied with this. They will say that, in the absence of a clear argument for (4), we should not accept it – and therefore should not accept the anti-intuitivist argument from disagreement. But other intuitivists will be nervous. They will say that, in the absence of a decisive refutation of (4), we should worry that (4) might really be true.

Fortunately for the nervous intuitivist, one of the other premises in the anti-intuitivist argument from disagreement can be refuted – or so I will argue. The vulnerable premise is premise (3) – the Equal Weight View. In the next section I will show how this premise can be effectively undermined.

6. Against the Equal Weight View

The Equal Weight View is premise (3) of the argument considered in Section 3. Here it is again, for ease of reference:

³⁷ Of course, things would be different if we had an extra-intuitional source of information about morality and could appeal to this source of information in the event of conflicts between intuitions. But the existence of such a source of information conflicts with No Alternatives. See the Introduction for my reasons to assume No Alternatives.

(3) **Equal Weight View:** There is no reason to attribute less epistemic significance to your intuition than to mine if I correctly consider you to be my epistemic peer or epistemic superior.

I will begin my case against this principle with an analogy. Suppose I am a watchmaker and I correctly consider myself to be very good at watchmaking. I have been recruited to participate in an experiment that will be conducted by Alex, a psychologist. Alex has given me a list containing tens of thousands of watchmakers, all of whom have also agreed to be in the experiment. I have read through the entire list (I am a very fast and thorough reader) and it turns out that I know all of the people on the list (I am a very well-connected watchmaker). I believe, correctly, that each watchmaker on the list is as good a watchmaker as I am (but no better). Alex tells me that our experiment will proceed as follows. On Tuesday, I and all the other watchmakers on the list will make separate watches in our separate workshops. Then, on Tuesday evening, Alex will visit each of us to collect the watches we have made. On Wednesday, Alex will return to my workshop and will present me with two boxes. In the first box will be my watch. In the second box will be a watch made by one of the other watchmakers. This is all that Alex tells me; he refuses to say more.

On Tuesday I dutifully make my watch. That evening Alex comes, as expected, and collects the watch I made. On Wednesday, Alex arrives to my workshop carrying two boxes. Alex says:

Before I allow you to open these boxes, I will tell you how I decided what to put in them. The first box contains your watch, as planned. To select a watch for the second box, I used the following procedure. First, I arranged into a single row all the tens of thousands of watches made by the other watchmakers participating in this experiment. Then I made a guess about what time it was, and I set my wristwatch to that time. (I have previously shown experimentally that my guesses about the time are rarely exactly correct, although it is also rare that my guesses are more than a few minutes off.) Then I went down the row, one by one, looking for a watch consistent with my guess. I won't tell you how long it took me to find such a watch, but I did eventually find one. That is the watch I put into the second box.

After saying this, Alex gives me the two boxes. In the first box I find my watch. I am sure it is mine because it bears my engraved signature. It says the time is 11:56. In the second box I find a watch that says the time is 11:50. This watch bears the inscription *R. Smith*. I recognize this as Robert Smith's signature; I have high confidence in Smith's abilities and regard him as my equal when it comes to watchmaking. (In this respect Smith is no different from any of the other watchmakers who participated in the experiment.)

Alex now asks me to make a judgment about the time. Alex forbids me to consult a third watch. His forbiddance is unnecessary – I do not need any more evidence; I say confidently, without hesitation, that the current time is 11:56. Alex takes note of my answer and thanks me for my participation in the experiment. I ask him to stay for coffee, but he says he has no time. He

has to visit each of the other watchmakers on the list; they are all going to be presented with the same problem.

Is my confidence that it is 11:56 justified? It seems so. Smith's watch says it is 11:50, and Smith is *ex hypothesi* an excellent watchmaker – every bit as good a watchmaker as I am. But Smith's watch was selected only because it matches Alex's guess, and I should have no great confidence in Alex's guesswork. It might be argued that Smith's watch would not match Alex's guess, or would be very unlikely to match Alex's guess, unless Alex's guess were correct. This is, in one sense, true: In normal circumstances, I should think it very unlikely that Smith's watch would match Alex's guess unless Alex's guess were correct. But it does not follow from this that I should trust Smith's watch in the present unusual circumstance. After all, Smith's watch was one in a batch of tens of thousands. We might expect there to be quite a few bad watches in a batch of tens of thousands, even if each watch in the batch had been made by an excellent watchmaker. Thus, even if we assume that Alex's guess had been wrong, it would be no great surprise that Alex is able to find a watch made by an excellent watchmaker to match it. And it would be no great surprise if Smith's watch just happened to be that watch.

What is going on in this example? I have been presented with two watches. One is made by an excellent watchmaker (me); the other is made by an equally excellent watchmaker (Smith). Aside from these two watches, I have no further evidence about the current time (e.g. I have not looked at a third watch to see whether it is consistent with my watch or with Smith's) and yet I am justified in placing much greater trust in my watch than in Smith's. I am so justified because of the way in which Smith's watch was selected by Alex. The selection of Smith's watch was, we might say, *rigged*: Smith's watch was selected from a large batch solely for its compliance with a guess in which I should have little trust. By contrast, the selection of my watch was not rigged; my watch was presented to me just because it is the watch I made, and the fact that a given watch is a watch I made is good reason for me to trust it, since I am *ex hypothesi* an excellent watchmaker.

It seems that similar points would hold in an example involving moral intuitions and epistemic peers. Imagine a crowd of ten thousand people, all of whom I correctly and justifiably believe to be my epistemic peers. If ten people are chosen randomly from this crowd, and all ten report pro-slavery intuitions, then perhaps this should lead me to seriously question my own anti-slavery intuitions. But suppose instead that ten people are selected from the crowd, not randomly, but by Alex, who has gone into the crowd searching specifically for ten people with pro-slavery intuitions. It would be strange to think that, when Alex presents me with these ten people, I am given any good reason to doubt my anti-slavery intuition. This is because Alex would surely have been able to find ten people with pro-slavery intuitions as long as no more than 99.9% of the people in the crowd have anti-slavery intuitions. There is not much evidence that slavery is permissible in a demonstration that no more than 99.9% of my peers intuit that slavery is wrong. Consider

also that, unless the crowd is very nearly uniform, Alex would have been able to find ten people with anti-slavery intuitions, if he wanted, just as well as he could find ten with pro-slavery intuitions. In that case, if I ought to attribute equal significance to the intuitions of anyone Alex selects from the crowd, then Alex would have control over how I ought to believe: Alex could freely choose whether I ought to relinquish my anti-slavery convictions or not simply by choosing whether or not to bring forth people with anti-slavery intuitions or people with pro-slavery intuitions. But this seems absurd.

In this example, a rigged selection process executed by Alex poisons any evidence that the pro-slavery intuitions of these ten people might otherwise have provided. Yet any evidence that might be provided by my own intuition is not poisoned in this way. It is not as if I am aware of my own intuition because Alex brought it to my attention; I am aware of my own intuition *automatically*, as it were, simply because it is *my own* intuition. This appears to be an epistemically relevant difference between my encounter with my own intuition and my encounter with the intuitions of the ten people selected by Alex. In light of this difference, it seems that I could be entitled to attribute greater epistemic significance to my own intuition than to the intuitions of any of these ten people.

If this is right, then the Equal Weight View is false: there can, after all, be reason to attribute less epistemic significance to your intuition than to mine even if you are my epistemic peer. If my encounter with your intuition is rigged, as in the Alex example just considered, then it seems I should not attribute the same significance to your intuition as to mine; rather, I should attribute considerably less significance to your intuition than to mine.

The Equal Weight View can be modified to deal with this problem, as follows:

Better Equal Weight View: There is no reason to attribute less epistemic significance to your intuition than to mine if (a) you are my epistemic peer and (b) my encounter with your intuition is not rigged.³⁸

Perhaps, as I have argued, the *mere* fact that an epistemic peer disagrees with me is not, by itself, epistemically significant. But it would seem significant if I am shown that a peer *could easily have different intuitions than mine*. If a peer could easily have different intuitions than mine, then perhaps I myself could easily

³⁸ What do I mean by “rigged” here? I do not think I need to provide necessary and sufficient conditions. The Alex examples show that if I encounter a peer *simply because* her opinion is consistent with a guess made by a bad guesser, or simply because she is a peer who disagrees with me, then my encounter with her is rigged. Perhaps, then, we should say that a rigged encounter with a given intuition is an encounter that occurs because of the intuition’s content. But I am not entirely satisfied with this definition, for reasons there is not space to discuss here. Alternatively, it might be suggested that a rigged encounter is any encounter that is not statistically random. But in a moment we will see reasons to think this is unsatisfactory.

have had different intuitions.³⁹ This would show that my intuition is *unsafe*.⁴⁰ If this is right, we may have a diagnosis of the second Alex case: the rigged presentation of a disagreeing peer is insignificant in that case because it provides no evidence as to whether a peer could easily have different intuitions than I have. But perhaps we should expect this problem would not arise in cases in which a peer is presented in a non-rigged fashion. Thus the Better Equal Weight View seems to stand a good chance of being correct.

I have no interesting objections to the Better Equal Weight View. And the Better Equal Weight View might still cause problems for intuitivism. Suppose that a historian shows that, in a random sample of all the people who have ever lived, a sizable majority would intuit that slavery is permissible in the Typical Case. It seems clear that, when I learn of this poll, my encounter with the intuitions of the people in the pollster's sample is not rigged. Then, unless we can refute premise (4) of the anti-intuitivist argument from disagreement, it seems the Better Equal Weight View will require me to abandon my anti-slavery convictions.

But the matter is complicated by the evolutionary history of our intuitions. I will defend two claims. First, I want to show that, given an evolutionary history of others' moral intuitions, my encounter with others' intuitions should *almost always* be considered rigged, even when they are randomly selected; this gives me a good reason to discount the significance of others' intuitions. Call this the **Other-Thesis**. Second, I want to show that an evolutionary history of *my own* moral intuitions would *not* provide reason to consider my encounter with *my own* intuitions to be rigged; thus, I am not given a reason to discount the significance of my own intuitions when I am given an evolutionary history of my own intuitions. Call this the **Self-Thesis**. If both of these theses were established, then it would follow that, given an evolutionary history of our moral intuitions, it is consistent with the Better Equal Weight View for an anti-slavery intuitivist to attribute high epistemic significance to her own anti-slavery intuitions while discounting, or even disregarding entirely, the pro-slavery intuitions of her pro-slavery peers.

Evolution is a multiple-generation selection process, and in a moment I will consider a process of this type. But I will begin my argument for the Other-Thesis and for the Self-Thesis by considering a single-generation selection process. Suppose (contrary to fact) that a single gene causes anti-slavery intuitions and a single (different) gene causes pro-slavery intuitions. Imagine two populations, Tribe A and Tribe B, living far apart from one another in separate areas. Due to a strange disease that takes hold in Tribe A's area, everyone with the pro-slavery gene suddenly dies; those with the anti-slavery gene are spared. At around the same time, a similar disease afflicts Tribe B's

³⁹ I am grateful to an anonymous reviewer for calling my attention to this point.

⁴⁰ *S*'s intuition that *p* is *safe* if and only if, "as a matter of fact, though perhaps not as a matter of strict necessity, *S* would not intuit that *p* without it being so that *p*." I take this concept from Ernest Sosa [1999], who defines safety in terms of *belief*, not intuition. The preceding quotation is Sosa's except for my substitution of "intuit" for "believe."

area, but it has a reverse effect: everyone in Tribe B with the anti-slavery gene suddenly dies; those with the pro-slavery gene are spared.

Suppose Lisa is a member of Tribe A, unaware of the existence of Tribe B. Lisa has anti-slavery intuitions, so she survives the epidemic. Suppose Lisa does not know even roughly how many members of her tribe were eliminated by the disease (because Lisa is reclusive and spends all her time in her hut). But Lisa does know that an epidemic killed all and only those with pro-slavery intuitions – and that it would have killed her, if she had not had anti-slavery intuitions. Should having gone through this selection process cause Lisa to revise downward any epistemic significance she previously attributed to her own intuitions? It seems not. The fact that those who disagreed with Lisa have been killed does not show Lisa's own intuitions to be less trustworthy; and the fact that she was spared death only because of the content of her intuitions is not a reason for her to distrust that content. These points, I want to emphasize, do *not* show that Lisa should attribute high (or any) epistemic significance to her own intuitions; rather, they show that the selection process to which Lisa has been subjected does not give Lisa a reason to attribute less significance to her intuitions than she would have done absent that selection process. This is so, I suggest, because in this case, after the disease has had its effect, Lisa's encounter with her own anti-slavery intuitions is *permitted* by a selection process – but it is not *rigged* by that selection process.

Now suppose that, shortly after the disease has had these effects on Lisa's tribe, Lisa takes a long walk and discovers the previously unknown Tribe B. Lisa learns about the effect of the disease that afflicted Tribe B: anyone in Tribe B who had anti-slavery intuitions is now dead, and everyone still living has pro-slavery intuitions. (No one will tell her how many members of Tribe B were killed by the disease, because Tribe B has a rule against counting dead people.) In that case it seems Lisa should attribute little or no epistemic significance to the intuitions of the still-living members of this tribe. This is so even if we assume that Lisa considers Tribe B to be a more enlightened society than her own and therefore considers the members of Tribe B to be her epistemic superiors. Lisa has, in that case, found that 100 percent of a group of her superiors have pro-slavery intuitions. But Lisa can know that she would have found this even if only 1% of the members of this enlightened society had had pro-slavery intuitions at the moment before the disease took effect. Lisa finds pro-slavery intuitions prevalent in this society *just because* a disease ensured she could find nothing else. So, although Lisa considers these tribe members to be her superiors, her encounter with their intuitions is rigged, not just permitted, by a selection process. Given this, it seems that Lisa should attribute little or no significance to their intuitions.

Thus, in the example of these single-generation selection processes, there appears to be a self-other epistemic asymmetry: a selection process permits, but does not rig, Lisa's encounter with her own intuition, but a precisely similar selection process rigs Lisa's encounter with others' intuitions – thus Lisa is provided a reason to regard others' intuitions as epistemically in-

significant without being provided a reason to regard her own intuitions as epistemically insignificant. Next I will argue that we find a similar asymmetry when we consider multiple-generation selection processes. But first I want to briefly mention an objection to my argument about Lisa. Note that, if my argument goes through, then a similar argument would work equally well for a person from Tribe B. Imagine that Rick, a person from Tribe B with pro-slavery intuitions, meets Lisa. Suppose that Rick and Lisa both know about the selection processes to which each has been subjected, and Rick and Lisa agree on all the relevant facts about these selection processes. I am committed to say that, even after agreeing about all these facts, Rick could be justified in attributing significance to his own intuition but not Lisa's, while Lisa could be justified in attributing significance to her own intuition but not Rick's. Thus, after Rick and Lisa meet, Rick could be justified in continuing to believe that slavery in the Typical Case is permissible, while Lisa could be justified in continuing to believe that slavery in the Typical Case is wrong. But this seems odd.

Of course, there is nothing particularly odd about two people being justified in believing differently. I may be justified in believing the world is round even if a person who has been subjected to misleading evidence, e.g. someone raised in a family of Flat-Earthers, is justified in believing the world is flat. But there is something odd about a case in which two people can be justified in believing differently *after having come to an agreement about all the relevant facts*. Such a case would apparently be a counterexample to the Uniqueness Principle, which says that "given one's total evidence, there is a unique rational doxastic attitude that one can take to any proposition" (White [2005]). Given the Uniqueness Principle, if two individuals have all the same evidence bearing on a given proposition, then they cannot rationally disagree about that proposition. It seems that, in the case just described, Rick and Lisa have all the same evidence – yet, I claim, they could still rationally disagree about whether slavery is wrong. This implication is counterintuitive, but it is defensible. Different parts of this shared body of evidence have been learned by Rick and Lisa in different ways. Rick and Lisa both know that Rick has a pro-slavery intuition, but Lisa learned this fact in a rigged fashion, whereas Rick did not. Likewise, Rick and Lisa both know that Lisa has an anti-slavery intuition, but Rick learned this fact in a rigged fashion, whereas Lisa did not. I have argued that this is an epistemically relevant difference. If that argument has been successful, then it explains why this case provides a counterexample to the Uniqueness Principle.

Let us now consider a multiple-generation selection process. Suppose that, after her visit to Tribe B, Lisa returns home to Tribe A. She has children, grows old and dies. The disease afflicting Tribe A is never cured and never dies out. Thus, generations later, Lisa's descendents, along with all other members of Tribe A, have anti-slavery intuitions. Rob is one of Lisa's descendents. Suppose Rob knows that he comes from a long line of people with anti-slavery intuitions, and knows that his lineage has been able to continue

because the disease that afflicts his tribe only kills those with pro-slavery intuitions. Should this lead Rob to consider his anti-slavery intuitions to be untrustworthy? It seems not. Perhaps it will be argued that the mere fact that Rob has inherited a genetic disposition toward anti-slavery intuitions should make Rob distrust those intuitions. But I have already responded to this view in Section 2, where I argue that it will be difficult to show that facts about one's intuition-generating mechanism provide reason to distrust one's intuitions. Perhaps it will be argued that Rob should distrust his intuitions, not just because those intuitions are due to an inherited disposition, but because this inherited disposition is the result of a selection process. But we already considered an argument like this – and it has not gotten any stronger in the generations that separate Rob and Lisa. The fact that, among Rob's tribe, pro-slavery lineages have been wiped out (or been prevented from existing in the first place) by a disease is not a reason for Rob to distrust the anti-slavery intuitions prevalent in his own lineage. Likewise, the fact that Rob's own lineage has been unaffected by the disease, and thus has been able to thrive, is not a reason for Rob to distrust the anti-slavery intuitions prevalent in his own lineage. Thus Rob's inheritance of anti-slavery intuitions is merely permitted, not rigged, by a selection process.

Suppose now that Rob takes a long walk and finds Tribe B, still in the same place it had been in Lisa's day. Rob finds that Tribe B continues to be afflicted by a disease that kills people with anti-slavery intuitions, so all the people of Tribe B continue to have pro-slavery intuitions. How much significance should Rob attribute to the intuitions of the people he encounters in Tribe B? Not much, I argue, for reasons identical with those that applied generations ago: Rob finds pro-slavery intuitions prevalent in this tribe just because a disease ensured he could find nothing else; thus his encounter with the tribe-members' intuitions is rigged, not just permitted, by a selection process.

The moral of this story is as follows. The multiple-generation selection processes I have just considered are obviously far simpler than any intuition selection process that might occur in the actual world. Our moral intuitions regarding slavery are not caused by a single gene; rather, they are probably the result of a complex interplay between genetic, cultural, idiosyncratic and other factors. Further, if (as seems probable) forces of biological and cultural evolution shape many of these factors, those forces are certainly more complicated than the disease described in this story. But our real-world predicament may be *epistemically* similar with Rob's predicament. If my own anti-slavery intuition is the result of a cultural or biological evolutionary selection process, and if others' pro-slavery intuitions are the result of different evolutionary selection processes, this may give me, like Rob, a reason to deny epistemic significance to others' pro-slavery intuitions without giving me a reason

to deny epistemic significance to my own anti-slavery intuitions.⁴¹ Then the Other-Thesis and the Self-Thesis would be true. And in that case, even if I consider *everyone* with pro-slavery intuitions to be my epistemic peers, and even if I believe that, in the scope of human history, pro-slavery intuitions vastly outnumber anti-slavery intuitions, I could still be an intuitivist without being committed to abandon my anti-slavery convictions. The anti-intuitivist argument from disagreement would then fail.

Conclusion

The question with which I began is whether we can give good reasons to think that trusting one's intuitions is no better than tossing a coin. I do not think we have found such reasons. But that may be because I have looked in the wrong places, or because I have looked in the right places but have overlooked the reasons there. Of course, even if I were to have conclusively shown that there are no good reasons to think that trusting one's intuitions is no better than tossing a coin, that may seem to be a small accomplishment. After all, it might be that there are no good reasons to think that trusting one's intuitions *is* better than tossing a coin. In that case we would have no reasons for or against; and in general, when one has no reasons for or against believing something, one had better not believe it. Still – I believe that there are reasons to think that trusting intuitions is better than tossing a coin (although I have not discussed any of those reasons here), and so I will be pleased if it turns out there are no reasons against that view. It is also worth pointing out that, even if one were to conclusively show that trusting one's intuitions is better than tossing a coin, it still might be that trusting one's in-

⁴¹ Here it is worth mentioning that some pro-slavery intuitions are probably the result of the same selection process that has issued my anti-slavery intuitions. For instance, suppose that cultures are subject to selection processes, and suppose that my anti-slavery intuitions are the result of my immersion in a particular culture. Some members of my culture have pro-slavery intuitions; e.g. there are Ku Klux Klan members, who are arguably part of my culture, yet who seem to have pro-slavery intuitions. If these Klan members are part of my culture, then my encounter with their intuitions is not analogous with Rob's encounter with the intuitions of Tribe B, so I have not provided any reason to think that my encounter with Klan members' intuitions is rigged. Thus it may be that I ought to attribute equal significance to Klan members' pro-slavery intuitions as to my own anti-slavery intuitions. I think this particular result would not be very worrying simply because Klan members are rare; I believe that the vast majority of the members of my own culture have anti-slavery intuitions. Therefore, even if it turns out that I ought to attribute as much significance to the intuitions of each member of my culture as I attribute to my own intuitions, I think I would still be justified, on intuitivist grounds, in maintaining an anti-slavery position. But this is just one example illustrating a more general concern. For all I have said here, it might turn out that the selection process that issues my anti-slavery intuitions (whatever that process turns out to be) usually issues pro-slavery intuitions. I do not have space here to determine whether this is so. The issue is partly empirical, but also partly conceptual (since it depends on how we ought to individuate selection processes). So I have to concede that, if it turns out that the selection process that issues my anti-slavery intuitions usually issues pro-slavery intuitions, then it is at least possible that a compelling anti-intuitivist argument is available.

tutions is only *very narrowly* better than tossing a coin. I worry about this. My view (which I have not defended here) is that No Alternatives is true; I think we have no alternative (other than agnosticism about morality) to trusting our intuitions. So I think we should hope that our intuitions are as trustworthy as they can possibly be. However, I suggest that, somewhat paradoxically, it is far better for our intuitions to be narrowly better than coin-tossing, than it is for our intuitions to be no better than coin tossing. If our intuitions are no better than coin-tossing, then we do not even know how to aim at doing the right thing. If our intuitions are narrowly better than coin-tossing, then our aim might not be very good – but at least we will know how to aim.⁴²

David Killoren
University of Wisconsin – Madison
Philosophy Department
killoren@wisc.edu

⁴² I am immensely grateful to Jeff Behrends, Thomas Carson, Juan Comesana, Dan Hausman, Casey Helgeson, Justin Horn, Holly Kantin, Joshua Knobe, Christian Koons, Jonathan Lang, Clayton Littlejohn, Peter Nichols, Jeff Prett, Carolina Sartorio, Russ Shafer-Landau, Rob Streiffer, Peter Vranas, two anonymous reviewers and others for invaluable feedback on earlier drafts and/or conversations which vastly improved my thinking about aspects of the material I have covered in this paper. Many of these people disagree with some or all of the argument I have offered here.

References

- Audi, Robert. 2005. *The Good in the Right*. Princeton University Press.
- Bealer, George. 1998. "Intuition and the Autonomy of Philosophy" in DePaul and Ramsey [1998].
- Bonjour, Lawrence. 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5.
- Brandt, Richard. 1954. *Hopi Ethics: A theoretical analysis*. University of Chicago Press.
- Conee, Earl and Feldman, Richard. 1998. "The Generality Problem for Reliabilism." *Philosophical Studies* 129.
- Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 119.
- DePaul, Michael and Ramsey, William. 1998. *Rethinking Intuition*. Lanham: Rowman & Littlefield.
- Doris, John and Plakias, Alexandra. 2008. "How to Argue about Disagreement" in Sinnott-Armstrong [2008b].
- Elga, Adam. 2007. "Reflection and Disagreement." *Nous* 41.
- Foot, Philippa. 1967. "The Problem of Abortion and the Doctrine of the Double Effect." *Oxford Review* 5.
- Greene, Joshua. 2008. "The Secret Joke of Kant's Soul" in Sinnott-Armstrong [2008c].
- Greene, Joshua et al. 2001. "An fMRI Investigation of Emotional Engagement in Moral Judgment." *Science* 14.
- Haidt, Jonathan. 2001. "The Emotional Dog and its Rational Tail." *Psychological Review* 108.
- Hare, RM. 1981. *Moral Thinking*. Oxford University Press.
- Huemer, Michael. 2006. *Ethical Intuitionism*. Palgrave Macmillan.
- Kelly, Thomas. 2006. "The Epistemic Significance of Disagreement." *Oxford Studies in Epistemology* 1.
- McGrath, Sarah. 2007. "Moral Disagreement and Moral Expertise." *Oxford Studies in Metaethics* 3.
- Shafer-Landau, Russ. 2009. "A Defence of Categorical Reasons." *Proceedings of the Aristotelian Society* 2009.
- Singer, Peter. 2005. "Ethics and Intuitions." *The Journal of Ethics* 9.
- Sinnott-Armstrong, Walter. 2002. "Moral Relativity and Intuitionism." *Philosophical Issues* 12.
- . 2008a. "Framing Moral Intuitions" in Sinnott-Armstrong [2008b].
- . 2008b. *Moral Psychology Volume 2*. Cambridge: MIT Press.
- . 2008c. *Moral Psychology Volume 3*. Cambridge: MIT Press.
- Sosa, Ernest. 1999. "How Must Knowledge Be Modally Related to What Is Known?" *Philosophical Topics* 26.
- Stich, Steven. 2006. "Is Morality an Elegant Machine or a Kludge?" *Journal of Cognition and Culture* 6.
- Street, Sharon. 2006. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127.
- Thomson, Judith. 1976. "Killing, Letting Die and the Trolley Problem." *Monist* 59.
- White, Roger. 2005. "Epistemic Permissivism." *Philosophical Perspectives* 19.