4

Moral Disagreement and Moral Expertise

Sarah McGrath

1. INTRODUCTION

The phenomenon of persistent ethical disagreement is often cited in connection with the question of whether there is any "absolute" morality, or whether, instead, morality is in some sense merely "a matter of personal opinion". Citing disagreement, many people who hold strong views about controversial issues such as the permissibility of abortion, eating meat, or the death penalty deny that these views are anything more than "personal beliefs". But while there might be inconsistencies lurking in this position, it is not obviously at fault for according the facts about disagreement some epistemic weight.

This paper addresses the question of whether and to what extent moral disagreement undermines moral knowledge. The most familiar arguments from disagreement in the literature purport to establish conclusions about the metaphysics of morality: that there are no moral facts, or that there are no moral properties, or that the moral facts are relative rather than absolute. Of course, the conclusions of some such metaphysical arguments might be perfectly consistent with the existence of considerable moral knowledge. For example, even if there is some successful argument from disagreement to the conclusion that moral facts are relative rather than absolute, this might very well be consistent with our having just as much moral knowledge as we

ordinarily take ourselves to have. (Although of course, such an argument might alter our conception of what it is that we know.) On the other hand, a metaphysical argument from disagreement which successfully showed that there are no moral facts would presumably rule out the possibility of moral knowledge.

By contrast, epistemological arguments from disagreement purport to undermine moral knowledge by showing that, regardless of the metaphysics of the moral facts, we are not in a position to have anything like the amount of moral knowledge that we ordinarily take ourselves to have. For reasons that I explore below, there are various respects in which epistemological arguments from disagreement present a more formidable skeptical challenge than metaphysical ones. My main goal in this paper is to develop an epistemological argument that creates a difficulty for our controversial moral beliefs and to explore the extent to which it succeeds.

2. METAPHYSICAL ARGUMENTS

As a representative metaphysical argument, consider J. L. Mackie's well-known "argument from relativity" (1977: 36–8). According to Mackie, "radical differences between first order moral judgments" provide a compelling reason to doubt "the objectivity of values". While it is not entirely clear what Mackie means when he denies the objectivity of values, he does seem to mean, minimally, that all claims to the effect that something has a certain moral property are false. If Mackie is right about this, then we have very little moral knowledge—far less than we thought we had. Perhaps one could know that nothing is morally wrong, but one could not know of any particular action that it is morally wrong—for all claims to the effect that a particular action is right or wrong are false. (Just as, having learned that there aren't any witches, one could know that Marilyn Manson is not a witch. What one can't know is that anybody *is* a witch.)

Significantly, Mackie does not think that scientific disagreement supports an analogous conclusion about science. He argues that the skeptical inference is compelling in the moral but not in the scientific case because moral and scientific disagreements have different explanations. While scientific disagreement is best explained by the fact that scientists draw different conclusions from inadequate evidence, disagreement about moral codes is better explained by "people's adherence to and participation in different ways of life" (p. 36). In the moral case, "the causal connection seems to be mainly that way round: it is that people approve of monogamy because they participate in a monogamous way of life rather than that they participate in a monogamous way of life because they approve of

monogamy".[1] The hypothesis that moral codes are mere reflections of ways of life better explains the pattern of moral variation than does the hypothesis that different people have different "seriously inadequate and badly distorted" perceptions of objective values (p. 37). Thus, there are no objective values.

One immediate concern about Mackie's argument is that it seems to prove too much: it is not true that, in general, where differences in belief co-vary with differences in ways of life, we ought to draw similar conclusions. For example, within the United States, beliefs about evolutionary theory seem to satisfy the relevant criteria. According to a Harris Poll conducted in the summer of 2005, only one-fifth of Americans believe that human beings evolved from other species; only half think that other plants or animals did; 64 percent believe that "human beings were created directly by God". The poll shows that variation in these beliefs reflects differences in the ways of life of the individuals who hold them, in the sense of reflecting the religious, political, and cultural features of the communities to which they belong: individuals who embrace creationism are more likely to be from the South, to be Republicans, to be religious, and to lack college educations. By contrast, Democrats, those from the Northeast and West, and those with college educations are more likely to believe in evolutionary theory. But while it does seem that people's beliefs about evolutionary theory reflect their ways of life in this sense, this does not support any surprising metaphysical conclusions about the facts at issue. In particular, it does not support the conclusion that there are no truths about the origins of the human species, or that all claims about human origins are false. Perhaps Mackie is correct in holding that the pattern of difference in moral beliefs corresponds to a pattern of difference in the cultural norms prevailing in the communities in which individuals were raised. But even if that is true, it does not show that there are no moral facts.

Of course, more could be said on behalf of Mackie's argument. In particular, one might argue that moral controversy and the controversy about human origins are disanalogous in ways that ultimately prove crucial. I will not explore arguments to that effect here, since I do not claim that the present difficulty is decisive. My purpose in raising this *prima facie*

---

[1] This quote leaves out some details: on Mackie's view, people's moral beliefs typically reflect "idealizations" of the ways that they actually live rather than simply reflecting those ways of living. So, for example, "the monogamy in which people participate may be less complete, less rigid, than that of which it leads them to approve" (p. 36). And some revisions in moral beliefs are explained not by changing idealizations but by the fact that people aim for consistency: thus, someone might change her belief about whether same-sex marriage is wrong because it conflicts with her other beliefs about what features of a relationship are relevant to whether people ought to marry.

difficulty for Mackie's argument is to highlight a quite general challenge for those who would have us draw conclusions about the metaphysics of morality from the existence of moral disagreement: such arguments naturally invite the charge that they prove too much. In order to successfully respond to this charge, proponents of such arguments must explain why we should not draw the same surprising metaphysical conclusions wherever we find apparently similar phenomena. Why, for example, doesn't widespread religious disagreement show that there is no fact of the matter about whether any gods exist, or that such facts are relative? Notoriously, it is difficult to explain why moral disagreement cries out for the metaethicist's favored metaphysical conclusion while similar disagreement in other domains does not. (Just as we would not want to conclude that all beliefs about human origins are false, so also we would not want to conclude that such beliefs are all relative, or by nature are knowable only by some special faculty of intuition.) Of course, this is not to say that the relevant explanation cannot be provided; only that the task of providing it cannot be avoided, and is far from trivial.

A second potential vulnerability for metaphysical arguments from disagreement is that in general such arguments have the form of inference to the best explanation arguments, according to which the best explanation of the kind of disagreement that we find in the moral domain is the preferred conclusion of the proponent of the argument: that there are no objective values, or, alternatively, that moral facts are relative facts, or that what look like moral claims are really just expressions of emotion, and so on. Because metaphysical arguments are inference to the best explanation arguments, one who offers such an argument must show that her favored conclusion better explains the data than any alternative hypothesis does. One competing hypothesis is the perfectly mundane one that the questions with respect to which we disagree are difficult ones, and at least some of us are getting them wrong; the others include the wide range of surprising candidate metaphysical hypotheses familiar from the metaethics literature. Again, there is no guarantee that such a case cannot be made on behalf of some preferred explanation. The point is just that it is not enough to point to a hypothesis that would adequately explain the relevant features of moral disagreement if it were true: one must show that the hypothesis better explains those features than would any competing hypothesis if *it* were true.

Thus, any metaphysical argument for the skeptical conclusion that we have little or no moral knowledge immediately inherits two potential vulnerabilities. First, to the extent that parallel reasoning applied to other domains would lead to conclusions that we are unwilling to accept, it is potentially vulnerable to the charge that it proves too much or

overgeneralizes. Second, inasmuch as such an argument is an inference to the best explanation argument, it is vulnerable to the provision of formidable competing explanations of moral disagreement. In the next section, I consider a line of epistemological argument which possesses neither of these vulnerabilities.

3. AN EPISTEMOLOGICAL ARGUMENT

Consider the following passage from Henry Sidgwick's *The Methods of Ethics*:

[I]f I find any of my judgments, intuitive or inferential, in direct conflict with a judgment of some other mind, there must be error somewhere: and if I have no more reason to suspect error in the other mind than in my own, reflective comparison between the two judgments necessarily reduces me temporarily to a state of neutrality. (p. 342)

Moreover, according to Sidgwick, "the absence of such disagreement must remain an indispensable negative condition of the certainty of our beliefs" (p. 342).

Let us call a belief **CONTROVERSIAL** just in case it satisfies the condition to which Sidgwick draws our attention. Thus your belief that p is CONTROVERSIAL if and only if it is denied by another person of whom it is true that: you have no more reason to think that he or she is in error than you are. Of course, a belief might be controversial without being CONTROVERSIAL. This is the case, for example, when some view that you hold is disputed, but you have reason to think that those who dispute it are more likely to be in error than you are.

As we have noted, Sidgwick holds that no belief that is CONTROVERSIAL can be *certain*. But a parallel claim about *knowledge* also seems attractive. That is, it seems plausible that

If one's belief that p is CONTROVERSIAL, then one does not know that p.

Suppose that you and your friend Alice intend to take the train together but discover that you have different views about what time it is scheduled to depart: you think that the train departs at a quarter past the hour, while she thinks that it departs at half past. Perhaps you have some good reason to think that Alice is the one who has made a mistake. For example, perhaps you know that she arrived at her view by consulting a train schedule that is out of date, while you arrived at yours by consulting the current schedule. Or perhaps you know that Alice is prone to carelessness with respect to such matters, as she has a past history of having made similar mistakes.

But suppose instead that you have no such reason to think that it is Alice who has made the mistake: as far you know, it is just as likely that you are mistaken as that she is. In that case, it seems that your belief about what time the train leaves does not amount to knowledge.[2]

Of course, it's clear enough that your belief does not amount to knowledge if you are in fact the one in error, i.e., if your belief about what time the train leaves is false. But even if your belief is true, and Alice is the one who has misread the schedule, it seems that your belief does not amount to knowledge provided that you have no good reason to think that she is the one who has made the mistake. Even if your belief would amount to knowledge in the absence of Alice's holding a contrary belief, the fact that she believes as she does can preclude your knowing in the circumstances. For plausibly, this would be a case in which misleading evidence undermines knowledge.[3]

This suggests the following epistemological argument for a certain kind of moral skepticism:

P1 Our controversial moral beliefs are CONTROVERSIAL.

P2 CONTROVERSIAL beliefs do not amount to knowledge.

C   Therefore, our controversial moral beliefs do not amount to knowledge.

The first premise and the conclusion of the argument refer to "our controversial moral beliefs". By this, I mean our beliefs about the correct answers to the kinds of questions that tend to be hotly contested in the applied ethics literature as well as in the broader culture: questions about

---

[2] Cases broadly similar to this one have recently been discussed in the epistemology literature devoted to the question of how we should respond to "peer disagreement". This literature has not directly addressed the question of how disagreement affects knowledge, which is our primary concern here. Significantly, however, a number of contributors to this literature (notably Feldman 2006, Christensen 2007, and Elga 2007) either endorse or express considerable sympathy for the view that peer disagreement should lead the peers to suspend judgment about the disputed question. Presumably, if one ought to suspend judgment as to whether p, then one does not know that P. Kelly (forthcoming) explicitly argues against the view that one is rationally required to suspend judgment in the face of peer disagreement but holds that one should nonetheless become less confident of one's original opinion, and that, all else being equal, as the number of peers on both sides of the issue increases, the push towards agnosticism grows stronger.

[3] Does the fact that your true belief is denied by the relevant kind of person suffice to undermine its status as knowledge, or must one also be *aware* that it is denied by such a person? This is a special case of a substantive and disputed question in epistemology, the question of whether (or in what circumstances) the existence of misleading evidence undermines knowledge when it is not possessed by the would be knower. For discussion, see Harman 1973, Lycan 1977, and Ginet 1980. In what follows, I sidestep this issue by focusing on cases in which one is aware of the disagreement.

the circumstances (if any) in which it is morally permissible to administer the death penalty, or to have an abortion, or to eat meat, or about how much money we are morally obligated to donate to those in dire need, and so on. It is clear that our beliefs about the answers to such questions are controversial ones. It is of course much less clear that they are also CONTROVERSIAL, i.e., that P1 is true. A good part of what follows is devoted to scrutinizing this claim. I begin, however, with a few preliminary remarks about the argument.

First, one who endorses the argument might remain studiously agnostic about the metaphysics of morality, and in particular, about whether there are any moral facts. That is, one who endorses the argument need not take a stand on whether such facts exist, or even on what, if anything, the relevant kind of disagreement suggests about their existence. The contention of one who endorses the argument is rather that the kind of disagreement that we find with respect to controversial moral questions precludes our knowing the correct answers to these questions, *regardless* of whether such questions have correct answers.

Second, the conclusion of the argument is that our beliefs about controversial moral matters do not amount to knowledge. The conclusion is not that it is unreasonable to hold those beliefs in the face of disagreement, or that we are rationally required to suspend judgment with respect to controversial moral matters. However, if the argument is successful, then the skeptic would seem to have made significant headway towards establishing these apparently stronger claims. For it has been argued, with considerable plausibility, that if one is not in a position to know whether p, then the reasonable course is to suspend judgment about whether p until further evidence becomes available; that is, one should not believe when one is in no position to know[4]. Thus, if the above argument is sound, then this would at the very least seem to put considerable pressure on the idea that it is rational for us to maintain our controversial moral views.

Third, in the previous section, we noted that metaphysical arguments from disagreement generally take the form of inference to the best explanation arguments, and that this fact presents a potential line of resistance to such arguments. Notice that the epistemological argument presented here is *not* best reconstructed as an inference to the best explanation argument. The suggestion is not that the best explanation of the disagreement is that no one knows; rather, the suggestion is that the circumstances of the disagreement are inconsistent with one's knowing. Thus,

---

[4] This conclusion will be especially attractive to those who take knowledge to be the aim of belief; for defense of this claim, see esp. Williamson 2000.

the argument does not share at least one of the two potential vulnerabilities characteristic of metaphysical arguments.

It is less clear, however, that the argument avoids the second potential vulnerability of metaphysical arguments: that of susceptibility to the charge of overgeneralizing, or proving too much. This issue is the focus of the next section.

4. DOES THE ARGUMENT OVERGENERALIZE?

In Section 2, we noted that metaphysical arguments run the risk of overgeneralizing. On the face of it, Mackie's argument that there are no objective values would, if it succeeded in showing claims about value to be false, show the same for claims about human origins. Since we can confidently assume that it would be a mistake to draw that conclusion about human origins, it seems that we can conclude that Mackie's argument falls short of showing that there are no objective values. Does the epistemological argument from disagreement similarly overgeneralize?

One might suspect that the epistemological argument does prove too much. Thus, in responding to a similar line of argument to the one under consideration here, Russ Shafer-Landau poses the following dilemma:

Either intractable disagreement among consistent intelligent parties forces them to suspend judgment about their contested views, or it doesn't. If it does, then we must suspend judgment about *all* of our philosophical views, as well as our belief that there is an external world, that I am an embodied being, that the earth is older than a second, etc. All of these have been challenged by brilliant, consistent, informed skeptics over the millennia.

Alternatively, if we are warranted in any of our beliefs, despite the presence of such skepticism, then justified belief is possible, even in the face of persistent disagreement. And so we could retain our moral beliefs, especially those we have carefully thought through, despite an inability to convince all of our intelligent opponents. (2004: 108–9)

Here, the suggestion is that our controversial moral beliefs are in the same epistemic boat as our beliefs that *there is an external world* and that *the earth has existed for more than one second*. If this were the case, then we could safely conclude that the argument from disagreement does prove too much, since (I assume) we do know that there is an external world and that the earth has existed for more than one second.

However, the idea that beliefs of this kind and our controversial moral beliefs are equally jeopardized by disagreement seems dubious. After all, my belief that *the earth is older than one second* faces much less opposition than my belief that *the death penalty is morally impermissible*. Even if it is true that

brilliant skeptics have disputed the former[5], they are vastly outnumbered by reasonable people who disagree. By contrast, with respect to, say, the moral permissibility or impermissibilty of the death penalty, the division of opinion is not that of lone geniuses vastly outnumbered by the opposition.

As Shafer-Landau interprets the skeptical challenge, it is the *absence of unanimity* among the relevant class of people which suffices to generate the skeptical conclusion. Even the existence of a single formidable dissenter who cannot be won over would suffice to undermine whatever justification one's belief originally enjoyed. This interpretation allows him to plausibly suggest that such a requirement, if consistently applied, would yield a sweeping and global skepticism. However, this is not the most charitable interpretation of the skeptical challenge. On a more charitable construal of that challenge, it is the fact that there is a substantial division of opinion with respect to controversial moral questions that undermines the possibility of knowing the answers to those questions.

In short, the beliefs that *the earth is older than one second* and that *there is an external world* are not CONTROVERSIAL. Even if these beliefs have on occasion been denied by some, including some of formidable intelligence (etc.), it does not follow that one has no more reason to suspect error in such minds than in one's own. Plausibly, one does have such reasons, reasons provided by facts about the distribution of opinion among the relevant class of people. If you and Alice have conflicting beliefs about what time the train is scheduled to depart, then it might be that both of your beliefs are CONTROVERSIAL. However, if you and Alice subsequently discover that ten other people have independently arrived at your belief while none shares hers, your belief is no longer rendered CONTROVERSIAL by the fact that Alice denies it. For now you do have reason to think that she is the one who has made the mistake. On the other hand, her belief—supposing she maintains it—is CONTROVERSIAL: she lacks any parallel reason.

Of course, it is no objection to the skeptical challenge under consideration that it fails to single out our controversial *moral* beliefs. Parallel arguments might be constructed to show that one lacks knowledge with respect to a significant number of topics—for example, philosophy, public policy, and

---

[5] One might quibble with Shafer-Landau's choice of examples. Skeptics about the external world or the past are best understood, I think, not as disputing first order propositions about the world such as *there is an external world* or *the world is older than one second*, but rather epistemic propositions such as *we know that there is an external world* or *we know that the world is older than one second*. It is clear that some have held views inconsistent with the latter propositions; it is less clear that anyone has held views inconsistent with the former. Still, if it could be shown that our controversial moral beliefs are no worse off than the relevant epistemic beliefs, I would take this to constitute an adequate vindication of the 'companions in guilt' strategy for resisting disagreement-inspired moral skepticism.

religion.[6] But this does not show that the argument *over*-generalizes. For it is far from clear that the answers to much disputed questions in such domains are known; in any case, that some of us have such knowledge is not a *datum* to which one might appeal in attempting to discredit the argument.

5. IN SEARCH OF MORAL EXPERTISE

The previous section defended the argument against the charge of overgeneralization. This section addresses the question of why one should think that one's beliefs about disputed moral questions are CONTROVERSIAL in the first place. We have emphasized that, even if a belief is controversial, it might not be CONTROVERSIAL. That is, even if the truth of a given belief is contested, a person who holds that belief might have good reason to think that anyone who thinks otherwise is more likely to be wrong than she herself is. Indeed, some beliefs might be *extremely* controversial without being CONTROVERSIAL. Consider again our earlier example of evolutionary theory. The proposition that *human beings evolved from other species* is vigorously denied by many, but it would be a mistake to conclude that it is therefore not known by any of those who believe it. Indeed, the fact that it is denied by many does not even preclude its being known by some who are relatively unfamiliar with the scientific evidence in its favor. Crucially, the proposition in question is not controversial among those who are known to possess the relevant expertise. Certain scientific questions might be highly controversial among the population as a whole, but when a consensus or near consensus exists among those with the relevant expertise, one need remain in a state of agnosticism only for as long as it takes to discover the content of that consensus. Thus, despite the large number of people who deny that human beings have evolved from other species, awareness of the expert consensus on the opposite side of the issue provides good reason to think that those who deny it are in error.

It might be thought that there is a parallel defense of one's controversial moral beliefs. That is, it might seem plausible that, although many dispute these beliefs, they are not CONTROVERSIAL, because they are not controversial among "the moral experts". This raises two questions. First, are there genuine moral experts? And second, if there are, how can they be recognized—either by themselves or by others—as such? Let us set aside the first question, and concede for the sake of argument that individuals with genuine moral expertise exist. How might they be identified?

---

[6] Cf. van Inwagen 1996.

The task of identifying those with genuine expertise will be a much less straightforward matter in some domains than in others. For the most part, the epistemology literature devoted to the topic of disagreement has focused on the idealized case, in which facts about relative expertise and "epistemic peerhood" are treated as given; the question that has dominated that literature concerns how we should respond to disagreement with our epistemic peers or equals.[7] But in actual, real-life cases, others do not typically wear their relative levels of competence on their sleeves. Of course, on occasion they do: most of us have good reason to think that the person whose shirt reads "Expert Plumbers" is someone to whose judgment we should defer with respect to whatever plumbing questions might arise. But in other cases, facts about relative levels of expertise and competence are far from transparent.

In general, identifying those with genuine expertise in some domain will be most straightforward when we have some kind of *independent check*, one not itself subject to significant controversy, by which we can tell who is (and who is not) getting things right. In certain domains, it is relatively easy for us to acquire evidence which bears straightforwardly on questions about relative expertise. Consider, for example, weather forecasting. Two weather forecasters might offer what seem to be equally compelling cases for their conflicting predictions about what tomorrow's weather will be like. But once tomorrow's weather rolls in, we will have an answer to the question of which of today's two conflicting predictions was more accurate. Thus, in the weather forecasting case, inductive track record evidence about who is more reliable is relatively easy to acquire. Moreover, crucially, such evidence can be readily assessed and assimilated by the layperson: one need not be an expert weather forecaster in order to reliably identify those who possess genuine expertise with respect to weather forecasting.

But significantly, we possess no similar independent check for moral expertise. If moral expertise stands to morality as weather forecasting expertise stands to weather, then a moral expert would be someone who consistently arrives at the correct answers to non-trivial moral questions (or at least, someone whose reliability with respect to such questions significantly exceeds that possessed by the average person, when the average person does not form his moral opinions by deferring to a moral expert). Given such a straightforward understanding of moral expertise, there is nothing particularly problematic about the idea that some individuals possess such expertise. The difficulty lies in arriving at compelling grounds for attributing such expertise, either to oneself or to others. A natural suggestion is that the possession of certain academic credentials, or professional concern with

---

[7] See esp. the work referred to in n. 2 above.

ethics, is good evidence that one possesses reliable moral judgment. I am acquainted with the ethics literature in a way that my plumber is not; moreover, I have taught ethics classes to college students and attended conferences devoted to the subject. He can claim no similar experiences. But in the absence of an independent check on my relative ability to therefore get the answers right, such facts would seem to constitute a relatively meager basis on which to conclude that I am his superior with respect to the reliability of my moral judgment. Again, contrast a case in which we know that one of two weather forecasters is more reliable than the other on the basis of his superior past track record. It would be a mistake, I think, to suppose that in these circumstances I have anything like the kind of evidence for the superiority of my moral judgment that is available in the weather-forecasting case.

Simply put, there is no obvious way to locate oneself in the space of moral expertise relative to others. It is true that professional philosophers who work in applied ethics have thought about the arguments longer than the average person has. Here, as elsewhere in philosophy, this has not resulted in a convergence of opinion. Yet even if these professionals were to converge on the view that, say, killing is no worse than letting die, on the grounds that no adequate metaphysical basis for imputing moral significance to this distinction could be found, it is not clear that ordinary people of the opposite conviction need treat this as conclusive. For it is less clear in the moral case than in various other cases that reliable judgment with respect to the relevant domain is the typical upshot of formal training. Here again the lack of an independent check seems crucial. If a moral expert is someone who tends to get the hard questions right, then good moral training is presumably whatever confers the relevant capacity. That studying structural engineering at MIT is good training for solving the kinds of problem that confront structural engineers can be more or less readily checked by, for example, examining the stability of bridges built by MIT-trained engineers. But in the moral case, since it is unclear how to check who is getting things right, it is unclear how to check whether MIT is a good place for moral training. Thus, while one might think that good moral training would consist in taking a series of ethics courses devoted to the critical examination of arguments on both sides of divisive issues, an equally plausible answer might be that good moral training consists in being raised by virtuous people who devote relatively little time to scrutinizing arguments for and against their views. Similarly, one might think that the best training for appreciating the permissibility or impermissibility of causing animal suffering would involve, among other things, witnessing such suffering. But we could just as easily imagine that the judgment of those best acquainted with the

slaughterhouse tends to become artificially deadened to the thought that animals matter.

If the population is substantially divided about, say, the moral permiss-ibility of abortion in certain circumstances, then, assuming that there is some non-relative fact of the matter, a large number of us are wrong. Unfortunately, we possess no analogue to an eye exam, by which we might determine whose moral vision is askew and whose is in good working order. Thus, the truth about where one stands in the space of moral expertise might prove elusive, even for intelligent, thoughtful people.

The upshot of these considerations is that it is quite unclear how one might argue, in a way that is not transparently question-begging or circular, that one's controversial moral beliefs are uncontroversial among the moral experts. But if, for all one knows, there is no consensus among the moral experts in favor of one's controversial moral beliefs, then one cannot appeal to the existence of such a consensus in order to show that those beliefs are not CONTROVERSIAL.

6. MORALITY AND THE CASE OF UNIQUE GREEN

In some ways, moral disagreement seems to parallel the diversity of opinion as to which shade of green is *unique green*. Unique green is that shade of green that is neither bluish nor yellowish. When asked to select the shade which is unique green, different subjects with normal color vision will select different shades.[8] As in the case of our controversial moral views, opinion about which shade is unique green not only fails to be unanimous, but is substantially divided. Perhaps if there were relatively widespread agreement as to which shade is unique green, then the dissenting judgments of a few who possessed otherwise normal color vision could be dismissed. But the fact that the actual division of opinion is substantial suggests that human beings are not reliable detectors of the relevant property. That relevantly similar creatures—creatures with the same type of visual system—arrive at different verdicts when similarly situated seems to show that that kind of creature is simply not well equipped to detect the presence or absence of the property in question. That human beings are not, as a species, reliable detectors of unique green seems to tell against crediting any individual with

---

[8] See Hardin 1988 for detailed information about the phenomenon. Hardin himself draws a conclusion analogous to Mackie's, viz. that nothing is colored. Cohen 2003 draws a relativist conclusion similar to Harman's: that objects have colors only relative to perceivers and circumstances of viewing. Byrne and Hilbert (2003) conclude that we should suspend judgment about which particular things are unique green.

knowledge that a certain shade is unique green, particularly if the individual knows of this general lack of reliability and has no good reason to think that he is exceptional in this respect.

Note that although questions about which shade of green is unique green are hard questions for human beings, such questions do not present themselves to us as difficult ones. In fact, most subjects are quite confident of their initial judgments; each person's view strikes her as obviously correct. This seems parallel to the moral case: in the moral case too, many find that their own views about controversial moral questions strike them as obviously correct. Cases in which we are quite confident in our original judgments, only to discover there is a substantial division of opinion among those with relevantly similar cognitive capacities, highlight the fact that there is more than one way to discover the relative difficulty of a given intellectual task. While one might learn that a given problem is hard by attempting to think it through and finding oneself struggling or unable to come up with an answer, one might, alternatively, learn that a given problem is hard by discovering that beings with the same cognitive capacities have arrived at wildly different answers. One might learn that a particular philosophical problem—say, about what makes it the case that I am the same person over time—is difficult by attempting to answer it, and finding oneself at a loss. But alternatively, one could learn that it is hard by discovering that there is a great deal of controversy about it. In the case of unique green, learning about the disagreement is the crucial way of finding out that the question is hard: each individual finds a shade that seems straightforwardly neither bluish nor yellowish to her; it is only upon discovering the extensive variation in judgment among those with similar cognitive capacities that the intellectual task is revealed to be difficult. The fact that with respect to various controversial moral questions, many of those on *both* sides of the issue experience their own view as obviously true suggests that here, as in the case of unique green, the better route to appreciating the relative difficulty of the problem is the more indirect of the two.

Judgments to the effect that a given shade is unique green are controversial; are such judgments also CONTROVERSIAL? Suppose that, in fact, some among us *are* reliable detectors of unique green: the initial judgments of members of this subpopulation are quite accurate, and non-accidentally so. Relative to the general population then, the members of this group are "experts" in the straightforward sense employed above. Consider their position, once they come to learn that others—whose color vision otherwise resembles their own—make contrary judgments. Do the experts have more reason to think that it is the others who are in error than that they are? It seems that they do not. For although their judgments are, *ex hypothesi*, more accurate, they have no reason to think so. Thus, even true, reliably formed
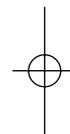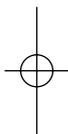
beliefs about which shade of green is unique green are CONTROVERSIAL. On the assumption that CONTROVERSIAL beliefs are not knowledge, neither are these beliefs.[9]

But while the case against attributing knowledge to even reliable detectors of unique green is quite strong, one might accept this conclusion while denying that the kind of disagreement that surrounds our controversial moral beliefs plays a similarly undermining role. The challenge, then, would be to point to some compelling difference between the moral and color case. Of course, there are some potentially relevant disanalogies. Here I will mention two, and argue that neither is sufficient for a successful defense of moral knowledge.

First, in the case of unique green, the subjects arrive at their judgments completely *independently* of one another. You select a particular shade as unique green, and those who select some other shade as unique green are neither influenced by one another, nor by some common influence. This stands in sharp contrast to the moral case, in which individuals do not arrive at their views about controversial moral issues in isolation. One might claim, then, that disagreement in the moral case creates significantly less skeptical pressure than it does in the unique green case, inasmuch as those on the other side of a particular moral controversy did not independently converge on their view.[10]

Now it is certainly true that, in some cases, learning that those who think otherwise did not arrive at their view independently can substantially reduce the skeptical pressure on one's own view. Consider an extreme case: the population is more or less evenly divided with respect to some question, with large numbers of people on both sides. Suppose you learn that all or almost all of those on the other side of the issue believe as they do because they unquestioningly defer to the judgment of a single charismatic individual whom they regard as a guru. If most of those on your side of the issue arrived at their view independently of one another, then it seems that you might reasonably conclude that your belief is not CONTROVERSIAL despite being controversial. After all, this would be a case in which there is substantial convergence on your view among individuals who made up their minds independently of one another. (Even if the guru himself arrived at the opposite opinion on his own, he is greatly outnumbered by those on the other side.)

---

[9] Notice that this conclusion can be accepted even by those drawn to reliabilist accounts of knowledge. On sophisticated versions of such views (e.g. Goldman 1986), true, reliably produced belief does not amount to knowledge when the subject is in possession of undermining evidence. (See Goldman's discussion of this point, 1986: 109–13.)

[10] On the importance of independence, see esp. Goldman 2001 and Kelly (forthcoming).

The difficulty with this response is just that actual moral controversies do not seem to exhibit the kind of asymmetry that might make the disagreement less threatening for either of the two sides. Granted, the many people who contest some controversial moral opinion of yours did not converge on their contrary view independently: the correct explanation of why that view is held by a substantial number of people will undoubtedly attribute a great deal to mutual influence, influence of common sources, and the like. But it is not the case that those who share *your* view have independently converged upon it. Thus, although it is true that people arrive at their judgments about unique green in relative isolation as compared to the moral case, it is far from clear that this disanalogy helps to defuse the skeptical challenge facing our moral beliefs.

A second potentially relevant disanalogy between moral controversy and the unique green case is the following. Judgments to the effect that a particular shade is unique green are *non-inferential* judgments. But for many of us, that is not how things are with our controversial moral beliefs. Many of us can provide reasons or arguments in favor of our controversial moral beliefs and against those of our opponents. Suppose, for example, that I disagree with Alice about whether abortion is morally permissible. She says that the fetus has a right to life, and that it follows that abortion is impermissible. But I have read Judith Jarvis Thomson's "In Defense of Abortion", and so I can supply an argument that even if the fetus does have a right to life it does not follow that abortion is impermissible. I take myself to have rebutted her argument and thus to have more reason to think that she is in error than that I am; I thus conclude that the fact that Alice disagrees with me about abortion does not render my belief CONTROVERSIAL.

More generally, one might think that non-inferential beliefs face skeptical pressure from disagreement that is not faced by beliefs that are based on arguments or discursive considerations.[11] Those who find this line of thought plausible will think that the key disanalogy between moral disagreement and the diversity of judgments about unique green lies here: many of us take ourselves to have compelling arguments for our controversial moral convictions while judgments about unique green are non-inferential, brute judgments.

Of course, the mere having of reasons cannot be sufficient to defuse the skeptical challenge. Presumably, if one can only offer bad reasons for one's

---

[11] For example, Walter Sinnott-Armstrong argues that moral intutionism leads to moral skepticism inasmuch as the moral intuitionist maintains that moral beliefs are non-inferentially justified. According to Sinnott-Armstrong, "disagreement creates a need for inferential justification" (2002: 312).

view, that is not sufficient to break the symmetry between oneself and one's opponent. And neither is the offering of genuine reasons on behalf of one's controversial beliefs, when those genuine reasons can be matched by similar reasons on the other side. After all, in the case of unique green, each of us can at least cite as a reason that this shade appears unique green to *me*. As each person's reason seems equally compelling, the symmetry remains, along with the skeptical pressure.

But suppose that one correctly recognizes that the argument on which the other person bases her belief is fallacious. One is then in a position to conclude that one's own belief is not rendered CONTROVERSIAL by the fact that *this* person holds a contrary view. If one could do this more generally, then one could establish that one's belief was not CONTROVERSIAL. But often those who engage in moral debate are dialectically skilled proponents of the rival views. In such cases, there will be non-fallacious arguments on both sides. The disagreement will then effectively reduce to one about the relative plausibility of the fundamental premises from which the arguments proceed. However, once the disagreement has been reduced to the question of whose premises are more compelling, the gap between the case of moral disagreement and the case of unique green seems to close. Of course, the premises of my argument seem more compelling to me than the premises of Alice's argument; but by the same token, the premises of Alice's argument seem more compelling to her than the premises of my argument.

Alice might show me pictures that motivate a premise of her argument for the conclusion that abortion is impermissible. But the pictures might not move me to agree that that the premise is true. Can I break the symmetry, then, by assuring myself that the reasons that I have are more compelling than hers? This seems no better than simply privileging my judgment about a given shade of green over Alice's contrary judgment.

Once again, it seems that I need some principled line of reasoning by which to privilege my judgment over that of those with whom I disagree. The final section of this paper examines a recent account of such reasoning due to Adam Elga.

7. ELGA'S PROPOSAL

In his recent paper "Reflection and Disagreement", Elga defends a view known in the epistemology literature as the "the equal weight view". According to the equal weight view, one is required to give equal weight to the judgment of an *epistemic peer* as to one's own judgment. You 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. For example, if I would have predicted that you and I would be equally likely to arrive at the correct solution to some mathematical problem in advance of our actually performing the calculation, then I consider you my epistemic peer with respect to that problem. According to the equal weight view, if you and I arrive at different answers, I am required to suspend judgment.

On the face of it, the equal weight view seems to have far-reaching skeptical consequences, requiring us to suspend judgment with respect to countless controversial questions. Elga, however, argues that this is not the case. His general strategy is to show that one's circle of epistemic peers includes only those with whom one is in substantial agreement on issues closely related to the one under dispute. Thus:

Consider Ann and Beth, two friends who stand at opposite ends of the political spectrum. Consider the claim that abortion is morally permissible. Does Ann consider Beth a peer with respect to this claim? That is: setting aside her own reasoning about the abortion claim (and Beth's contrary view about it), does Ann think Beth would be just as likely as her to get things right?

The answer is "no". For (let us suppose) Ann and Beth have discussed claims closely linked to the abortion claim. They have discussed, for example, whether human beings have souls, whether it is permissible to withhold treatment from certain terminally ill infants, and whether rights figure prominently in a correct ethical theory. By Ann's lights, Beth has reached wrong conclusions about most of these closely related questions. As a result, even setting aside her own reasoning about the abortion claim, Ann thinks it unlikely that Beth would be right in case the two of them disagree about abortion … The upshot is that Ann does not consider Beth an epistemic peer with respect to the abortion claim. (pp. 492–3).

It is clear enough how this general line of thought might be adapted so as to apply to the argument with which we are concerned. Again, a belief of yours is CONTROVERSIAL if and only if it is denied by another person of whom it is true that: you have no more reason to think that he or she is in error than that you are. But in what circumstances do you have no more reason to think that the other person is in error than that you are? Perhaps: exactly when, in advance of either of you reasoning about the case at hand, you would have predicted that the other person was just as likely as you to arrive at the correct answer. That is, we might take a CONTROVERSIAL belief to be one that is disputed by someone who is an epistemic peer in Elga's sense. In that case, one might appeal to the line of reasoning Elga provides and hold that one has significantly fewer CONTROVERSIAL beliefs than one might have thought, since those with whom one frequently disagrees over controversial moral questions are outside one's circle of peers. Following Elga's lead, one might say: even though Ann knows that Beth disagrees

with her about abortion, this has no tendency to make Ann's view about abortion CONTROVERSIAL.

Elga anticipates a natural objection that runs as follows: Ann cannot legitimately take her own views on the surrounding issues for granted and use them as a basis for concluding that Beth is more likely to get things wrong with respect to abortion. Rather, Ann should think of the entire cluster of related issues as a single compound issue, and take into account Beth's disagreement about this single compound issue. Once she does this, Ann will no longer be in a position to penalize Beth for having, by Ann's lights, false views about the surrounding issues. Hence, the case for skepticism is restored.

In response, Elga offers the following:

Consider the cluster of issues linked to abortion. Contrary to what the objection supposes, Ann does *not* consider Beth a peer about that cluster … That is because there is no fact of the matter about Ann's opinion of Beth, once so many of Ann's considerations have been set aside … To set aside Ann's reasoning about all of these issues is to set aside a large and central chunk of her ethical and political outlook. Once so much has been set aside, there is no determinate fact about what opinion of Beth remains. (pp. 495–6).

He motivates this claim with an example: plausibly, there is no determinate answer to the question of what your opinion of Jennifer Lopez is, setting aside your views that humans have bodies and that the Earth exists (p. 25).

However, while it seems right to say that there is no fact of the matter about your opinion of Lopez setting aside your beliefs about human embodiment and the existence of the Earth, the same maneuver seems less plausible when applied to the case of Ann and Beth. Recall that Elga characterizes Ann and Beth as "two friends at opposite ends of the political spectrum". We might then think of Ann as a conservative Republican who takes abortion to be morally abhorrent in most circumstances, and Beth as a liberal Democrat who thinks that it is morally permissible in most circumstances. No doubt, we would expect Ann and Beth to disagree about a wide range of moral issues. Significantly, however, this is perfectly consistent with a very substantial amount of moral agreement between the two. Indeed, we 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. Moreover, notice that many of the issues on which they are likely to agree are highly non-trivial, at least when judged by world-historical standards. (Consider, for example, their shared belief that slavery is morally abhorrent.) With respect to moral sensibility, we

would expect Ann and Beth to resemble one another far more than either resembles, for example, a committed Nazi, or an ancient Hittite lord for whom the thought that slavery is abhorrent was simply not on the moral map. In short, Ann and Beth's disagreements about abortion and related matters, although substantial, almost surely take place against a relatively wide background of shared moral beliefs. It seems wrong, then, to say that there is no fact of the matter about Ann's opinion of Beth's moral judgment, setting aside abortion and the many related issues about which they disagree. Indeed, once these disagreements are bracketed, the 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.

Once again, a comparison with the case of unique green seems apt. Subjects with normal color vision make contrary judgments as to which shade is unique green. Thus, those contrary judgments take place against a relatively substantial background of shared color judgments. Plausibly, this fact tends to increase the skeptical pressure: a shared background of agreement strengthens the case for counting the conflicting judgments as CONTROVERSIAL. The same would seem to be true in the moral case. Elga's proposal suggests that I can simply rank myself above those who disagree with me about controversial moral issues on the grounds that our disagreement is substantial. But this seems like a dubious procedure for locating myself relative to others in the space of moral expertise—all the more so when those who disagree do so against a wide background of *agreement*. After all, according to me, they usually get it right.

REFERENCES

Byrne, A., and Hilbert, D. (2003) "Color Realism and Color Science" *Behavioral and Brain Sciences* 26: 3–21.

Christensen, D. (2007) "Epistemology of Disagreement: the Good News" *Philosophical Review* 116: 187–217.

Cohen, J. (2003). "Perceptual Variation, Realism, and Relativization, or: How I Learned To Stop Worrying And Love Variations In Color Vision (Open peer commentary on Byrne and Hilbert, *Color Realism and Color Science*)" *Behavioral and Brain Sciences* 26: 25–6.

Elga, A. (2007) "Reflection and Disagreement" *Noûs*, LXI(3): 478–502.

Feldman, R. (2003) *Epistemology* (Upper Saddle River, NJ: Prentice-Hall).

—— (2006) "Epistemological Puzzles about Disagreement" in *Epistemology Futures*, ed. S. Hetherington (Oxford: Oxford University Press).

Ginet, C. (1980) "Knowing Less by Knowing More" *Midwest Studies in Philosophy.* 5: 151–62.

Goldman, A. (1986). *Epistemology and Cognition*. (Cambridge, MA: Harvard University Press).

———(2001) "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63.

Hardin, C. L. (1988) *Color for Philosophers* (Indianapolis: Hackett).

Harman, G. (1973) *Thought* (Princeton: Princeton University Press).

———and Thomson, J. J. (1996) *Moral Relativism and Moral Objectivity* (Cambridge, MA: Blackwell).

Harris Poll (2005) "Nearly Two-thirds of U.S. Adults Believe Human Beings Were Created by God." Harris Interactive, Inc. http://www.harrisinteractive.com/harris_poll/index.asp?PID=581

Kelly, T. (2005). "The Epistemic Significance of Disagreement" in *Oxford Studies in Epistemology*, ed. J. Hawthorne and T. Gendler Szabo (Oxford: Oxford University Press), 167–96

———(forthcoming) "Peer Disagreement and Higher Order Evidence" in *Disagreement*, ed. R. Feldman and T. Warfield (Oxford: Oxford University Press).

Lycan, W. (1977) "Evidence one does not possess" *Australasian Journal of Philosophy* 55: 114–26.

Mackie, J. L. (1977) *Ethics: Inventing Right and Wrong* (New York: Penguin).

Rosen, G. (2001) "Nominalism, Naturalism, Epistemic Relativism" *Philosophical Perspectives* 15: 69–91.

Shafer-Landau, R. (2004) *Whatever Happened to Good and Evil?* (Oxford: Oxford University Press).

Sher, George. (2001) "But I Could Be Wrong" *Social Philosophy and Policy* 18: 64–78.

Sidgwick, H. (1907/1981) *The Methods of Ethics* (Indianapolis: Hackett).

Sinnott-Armstrong, W. (2002) "Moral Relativity and Intuitionism" *Philosophical Issues* 12: 305–28.

van Inwagen, P. (1996). "It is Wrong, Always, Everywhere, and for Anyone, to Believe Anything, Upon Insufficient Evidence" in *Faith, Freedom and Rationality*, es. J. Jordan and D. Howard-Snyder (Lanham, MD: Rowman & Littlefield).

Williamson, T. (2000). *Knowledge and its Limits* (Oxford: Oxford University Press).