JOURNAL OF
SOCIAL
PHILOSOPHY

Moral Caution and the Epistemology of Disagreement

Jonathan Matheson

In this article, I propose, defend, and apply a principle for applied ethics. According to this principle, we should exercise moral caution, at least when we can. More formally, the principle claims that if you should believe or suspend judgment that doing an action is a serious moral wrong, while knowing that not doing that action is not morally wrong, then you should not do that action. After motivating this principle, I argue that it has significant application in applied ethics. The application to applied ethics comes by way of the epistemic significance of disagreement. I argue that we are in an impoverished epistemic position with respect to a number of morally controversial actions resulting from the widespread and persistent disagreement on those issues among intelligent, informed, and open-minded individuals on disparate sides of the debate. In doing so, I build on the emerging literature in epistemology on the epistemic significance of disagreement. However, while disagreement should make us skeptical about the moral status of some actions, there needn't be disagreement about the permissibility of refraining from doing those very actions. When there is not, our principle instructs us to exercise moral caution and not take the controversial course of action. Though such a principle has perhaps a number of applications in applied ethics, my focus here is limited to the question of whether it is morally permissible to eat animals for pleasure.

1. The Moral Caution Principle

Consider the following principle:

> MORAL CAUTION (MC): Having considered the moral status of doing action A in context C, if (i) subject S (epistemically) should believe or suspend judgment that doing A in C is a serious moral wrong, while (ii) S knows that refraining from doing A in C is not morally wrong, then S (morally) should not do A in C.

MC is a principle that links epistemology to morality. While MC is a principle about what you morally should not *do*, your epistemic situation—what you *know* (and what you do not know) and what you (epistemically) should believe (and what you should not believe)—play a critical role in the principle. According to MC, your epistemic situation can affect what is morally permissible for you to do. In particular, MC claims that being in a particular epistemic situation places certain moral constraints on you.

JOURNAL of SOCIAL PHILOSOPHY, Vol. 47 No. 2, Summer 2016, 120–141.

MC is an objective moral principle—it is a claim about what one objectively should and should not do. The truth of MC does not depend upon what any individual or group believes about MC, nor does it depend upon any other attitude that anyone may have toward MC. That said, MC claims that what is true of you—your epistemic situation—places constraints on what you objectively (morally) ought to do. Not everyone is in the same epistemic situation, and thus not everyone (epistemically) should adopt the same doxastic attitudes toward the same propositions. What you (epistemically) should believe is based upon your epistemic situation; it depends upon features of you that may not be shared by others.[1] MC thus claims that certain epistemic features of you place certain objective moral constraints on you. If you meet the antecedent of MC, then it is morally wrong for you to A in C. In this way, MC is somewhat similar to the claim that you ought not act against your conscience. While there are differences between these principles, both principles have it that something that is true of the subject (and can differ from subject to subject) affects what the subject (objectively) morally ought to do.

MC's application is limited in a number of ways. For one thing, MC concerns a particular epistemic situation. First, MC is focused only on situations where, upon having contemplated the morality of doing a certain action,[2] the doxastic attitude that you (epistemically) ought to adopt toward the proposition that doing that action is a serious moral wrong is either belief or suspension of judgment.[3] That is, having considered the morality of the act, the subject is propositionally justified[4] in either adopting *belief* or *suspension of judgment* toward the proposition that the act in question is a serious moral wrong.

It may be thought that in situations where an individual has never contemplated the morality of doing a certain action, that he or she is justified in suspending judgment about the morality of that action.[5] For that reason, MC stipulates that the subject has contemplated the relevant moral proposition, and that upon contemplation one of two doxastic attitudes are epistemically rational for the subject to adopt toward that proposition: either belief or suspension of judgment. Even still, it might be thought that MC has it that actions are guilty until proven innocent.[6] If we imagine a case where the subject has no positive reasons to think the action is morally wrong and no positive reasons to think that it isn't, then it may look like the subject should suspend judgment as to whether the action is a serious moral wrong. If so, then the subject would meet condition (i) of MC. However, condition (i) is not so easily met. It is plausible that, in general, actions are morally innocent until proven guilty, and so such an epistemic situation (where one lacks positive reasons to believe an action is morally wrong and positive reasons to believe it is not morally wrong) is one where the agent (epistemically) should believe the action is not a serious moral wrong. On this way of thinking of things, to satisfy condition (i) the subject must have *some* positive reason to believe that the action is a serious moral wrong. If the subject only has positive reasons to believe the action is a serious moral wrong, then the subject (epistemically) should believe the action is a serious moral wrong. If the

subject also has positive reasons to believe that the actions is not a serious moral wrong, then if those reasons as equally strong the subject (epistemically) should suspend judgment as to whether the action is a serious moral wrong.[7]

Second, MC is focused on an epistemic situation where you also know that refraining from doing a particular action is morally permissible—cases where you know that some available alternative is morally permissible.[8] MC takes for granted that in at least some cases we can be in the know about the moral status of an action—it takes for granted that some moral knowledge is possible. While MC takes this claim for granted, I will not be defending it here.[9] In cases where you fail to know that it is morally permissible to refrain from doing the action in question, the antecedent of MC is not met and MC fails to make a prescription.

It may be that in some contexts *any* action that you could undertake may be a serious moral wrong, at least given your evidence or knowledge. In some situations significant moral risk may simply be inevitable because none of your options are known by you to be morally permissible alternatives. In such situations, MC does not recommend or forbid any course of action; MC simply fails to have any application in such contexts, and thus fails to make any moral prescriptions in those contexts. In such contexts, there is simply no opportunity to exercise moral caution, whereas MC only tells you to exercise moral caution *when you can*. Along these lines, MC does not claim that the *only* actions that you are morally permitted to do are those that you are justified in believing are morally permissible. It is consistent with MC that some actions are morally permissible, and even morally obligatory, for you to do even if you are not justified in believing that the action in question is morally permissible. MC tells you to exercise moral caution, *at least when you can*; it does not claim that a necessary condition of an action being morally permissible is that the subject is justified in believing it is a morally permissible action, nor does it claim that moral caution must always be exercised (sometimes it just cannot be).

MC is a plausible moral principle. You should not do things that are morally wrong to do—you have a moral duty to play it safe when you can. So, if you should believe or suspend judgment that a certain action is a serious moral wrong, while knowing that refraining from doing that action is to not morally wrong, you should refrain from doing that action. In doing so, you do your best to avoid serious moral wrongs. Such moral caution is noble. Morality demands that you take morality seriously, and taking morality seriously requires exercising moral caution, at least when you have the opportunity to do so.

MC is in the spirit of the precautionary principle. The precautionary principle is a widely endorsed decision-making rule or approach, particularly for matters involving the environment and health.[10] While it is plausible that there is no *one* thing that is the precautionary principle,[11] the principle roughly claims that in cases of suitable risk precautionary measures should be taken even in the absence of conclusive evidence linking the action under consideration to the potential harm. Bodansky (2004) outlines three broad ways of interpreting the principle: (i) as excluding ignorance as a justification for inaction; (ii) as permitting actions on

the basis of precaution; and (iii) as a duty to take precaution or preventative measures. Within each of these broad interpretations of the principle there remain ambiguities regarding both what counts as "suitable risk" and what counts as "precaution." That said, MC is in the broad spirit of the precautionary principle since it too values taking precaution and the "better safe than sorry" approach.

However, there are also some important differences between MC and the precautionary principle. First, MC is stronger than many formulations of the precautionary principle in that it does make prescriptions about what actions to take.[12] MC claims that you should not do certain actions. It gives clear verdicts regarding some actions where it is not always clear that the precautionary principle does so. Second, MC is focused on individual actions rather than more general approaches to policies. While policy creation can itself be seen as individual action, the application of MC extends beyond matters of policy creation, and this difference in scope between the precautionary principle and MC is worth noting. Relatedly, MC is concerned with individual subjects and their actions. While MC can also apply to states or collections of states and their actions, MC has application to individual agents and their actions as well. In this way, MC differs in its deployment from typical applications of the precautionary principle where state actors are the relevant parties. Third, while MC and the precautionary principle both involve acting in ignorance (where ignorance is something less than "scientific certainty"), the principles differ in the *nature* of the relevant ignorance. As the precautionary principle is typically understood, it is *empirical* ignorance (ignorance of the empirical effects of some course of action and the relevant causal connections) which are pertinent. MC is not qualified in this way. In fact, the ignorance relevant to MC is *normative* ignorance—ignorance regarding the morality of an act (even perhaps amidst knowledge of all the germane empirical facts).[13] Finally, there is an important sense in which MC, unlike the precautionary principle, is restricted in its application. While MC can make prescriptions to exercise caution, it only does so in certain situations—cases where there is an epistemic asymmetry and an uncontroversially moral course of action. MC does not make a prescription in every case where moral risk is present. This difference is significant since it allows for MC to avoid an objection leveled at the precautionary principle.[14] Harris and Holm (2002) argue that the precautionary principle results in a paradox since following the principle can actually result in *more* harm being done. As we have seen, some scenarios *require* taking a moral risk and do not allow for moral caution to be exercised. Whereas the precautionary principle can be seen as leaving agents paralyzed in such situations, MC avoids any such difficulties by not making prescriptions in cases where there isn't an uncontroversially moral course of action. In some cases moral risk is inevitable, but in such cases MC makes no prescriptions as to what is to be done. As such, MC avoids the problems Harris and Holm (1999) attribute to the precautionary principle.[15]

MC can be further explicated by examining how it applies to cases—by seeing how its antecedent can be met and by spelling out its consequent. In what

follows I will argue that MC has significant application to our lives. This application comes in two steps: first by seeing how the epistemic significance of disagreement has it that condition (i) of MC is often met, and second by seeing how in some cases condition (ii) of MC can simultaneously be met.

2. The Problem of Disagreement

While MC is plausible, its application may be thought to be much more restricted than I have let on. After all, how often are we actually in the dark regarding whether some action is a serious moral wrong? Even if it is plausible, does MC have any interesting consequences? In this section I argue that the antecedent of MC is met more often than may be thought. In particular, I argue that the epistemic significance of disagreement has it that we should be skeptical about the moral status of morally controversial actions.

The literature on the epistemic significance of disagreement has focused on disagreement between a kind of epistemic equals—*epistemic peers*. Epistemic peers are individuals who are in an equally good epistemic position on a matter, where one's epistemic position is determined by the quality and quantity of one's evidence and one's ability to accurately evaluate that evidence. Epistemic peers about p are equally well positioned with respect to p, and are thus equally likely to be correct about p.

Conciliatory Views of disagreement claim that having evidence that you are party to a peer disagreement gives you a defeater for your doxastic attitude calling for you to make doxastic conciliation.[16] Steadfast Views of disagreement claim that it can be rational for you to stick to your guns, even when you get evidence that an epistemic peer disagrees with you.[17] Of relevance here, however, is not merely a disagreement between two epistemic equals, but a more general state of controversy amongst the relevant experts, and what we should believe in such a situation.[18]

Let us first consider the epistemic significance of disagreement regarding a non-philosophical claim. Consider the following adaptation[19] of a case given by David Christensen (2007):

> RESTAURANT CHECK:
> A group of colleagues goes out for dinner. Upon getting the check, they decide to add 20% gratuity and divide the check evenly, regardless of who ordered what. Having agreed to this, each of the colleagues looks at the bill and independently calculates the shares. Suppose that this is a common practice among the members of the group. Each colleague is quite reliable at performing such calculations. While errors have been made in the past, each individual is about as likely to be correct as any other. When they each reveal their findings, they discover a disagreement—5 individuals believe the shares are $33 a piece, 4 believe that they are $27 a piece, and 2 believe that the shares are $31.

What should the colleagues believe about their shares of the bill upon this discovery? While they should double-check their calculations, and perhaps seek

out a calculator, what should they *believe* about the shares before they do so? What should they *believe* about the shares in the state of disagreement? It seems that they should suspend judgment about what the shares are. It would be irrational for any party to simply stick to their guns and continue believing that he or she correctly calculated the shares in the face of such discovered disagreement. It would similarly be irrational for any party to simply defer and adopt the belief of one of the other parties. Notice that this conclusion holds for even the members of the group who *correctly* calculated the shares since they are presented with good reasons to believe that they did not correctly calculate the shares. The shares are controversial among individuals who are good judges of these matters, and there is not any better reason to locate the calculation error with the other disagreeing parties. While each party was able to examine the check, hearing the disparate conclusions about the shares by the other parties (who are each just as likely to be correct) gives each party powerful higher-order evidence that he or she has misjudged the evidence or made a calculation error. Since we are fallible epistemic agents, we must take seriously such evidence that we have made an error. Such evidence undermines any justification we may otherwise have for our belief on the disputed matter.

Typically, in cases where there is great controversy about what our evidence supports, suspension of judgment is called for. If we should suspend judgment about what our evidence regarding p supports, we should also suspend judgment about p itself. This is evidenced by the irrationality exhibited in the following, "I have no idea what my evidence about p supports, but p!" In the Restaurant Check case, the parties should suspend judgment about what their evidence about the shares supports, and so they should also suspend judgment about the shares.[20]

For instance, let's suppose that external world skepticism is false and that our perceptual experiences do in fact provide good reasons to believe that things are the way that they appear. Now, suppose that an individual takes an epistemology class where skepticism is forcefully presented. Suppose that having taken the course she should suspend judgment about whether external world skepticism is true. If she should suspend judgment about whether external world skepticism is true, then she should also suspend judgment about whether she has hands. It would be illegitimate to think that while she should suspend judgment about whether external world skepticism is correct, she nevertheless has undefeated good perceptual evidence that justifies her belief that she has hands (even on the assumption that her perceptual evidence *is* good evidence). After all, her evidence (and justified suspension of judgment) about skepticism is about *that very evidence* (her perceptual evidence) and how good it is.[21]

Likewise, evidence of a sufficient controversy among the experts regarding p is evidence that calls us to suspend judgment about what to make of the evidence regarding p. So such evidence also calls for us to suspend judgment about p, the controversial claim, as well. When those best positioned (epistemically) are unable to come to anything like a consensus, this is good reason to believe

that either our evidence on the matter is insufficient or that our abilities to evaluate that evidence are insufficient. Either way, evidence of such a controversy provides a defeater for any non-skeptical attitude on the matter.[22]

Similar skeptical conclusions appear to apply in cases of philosophical disagreements in general, and *moral* disagreements in particular. Philosophy is difficult, perhaps ethics especially so. On controversial moral matters, the experts disagree. Intelligent, well-informed, and open-minded individuals have come to different conclusions on a number of moral matters. Here, too we lack good reason to believe that the parties in some opposing group are all mistaken. There are a number of philosophers in various moral camps that are aware of all the relevant arguments, are equally concerned with discovering the truth, and are roughly equal in intelligence. Roughly, these philosophers are equally likely to be correct about such moral matters. Nevertheless, they often disagree. Plausibly, once we are aware of the widespread disagreement among the experts on moral matters, we should be skeptical about the relevant moral propositions as well. Our awareness of such controversy is sufficient to satisfy (i) of MC. It is not as though we have any good reason to believe that we hold in our possession some decisive argument that would convince all (or even most) intelligent and open-minded people who consider it. Similarly, it is implausible to attribute some error in reasoning or illicit bias to all the various disagreeing parties. So, evidence of the state of controversy among the experts over contentious moral claims indicates that we should suspend judgment regarding the truth of those moral claims.

In fact, the state of disagreement with respect to moral claims is even more dire than in our Restaurant Check case. While we have recourse with respect to bill disputes that tends to quickly resolve the disagreement, hope for parallel resolutions in moral disagreements is ill-founded. We do not possess the equivalent of a "moral calculator" that we can simply use to settle our moral disputes. Moral disagreements persist, and they have for some time. If anything, *more* alternative positions are discovered, adopted, and defended as time passes, not fewer.[23] So our moral disagreements appear to be both stable and persistent.[24]

3. Application for Applied Ethics

Thinking about ethical disagreement is not novel.[25] However, typically what is at issue in such discussions is a *metaethical* issue—whether ethical disagreement raises a problem for certain metaethical views such as moral cognitivism or moral realism. Here, the application of ethical disagreement is different. Here, the argument is that ethical disagreement plays a role in establishing what we ought to do—that there are *normative* moral implications of ethical disagreement.

Let's examine how the epistemic significance of disagreement when coupled with our principle MC can be applied to ethical issues. While there are presumably a number of applications, I focus here on just one—the question of

whether it is morally permissible to eat animals for pleasure, where I use the phrase "eating animals for pleasure" to pick out eating animals when it is not necessary for survival (either for your own survival or for someone else's).[26]

The question of whether it is morally permissible to eat animals for pleasure is a controversial moral issue. Many intelligent, informed, and open-minded individuals believe that it is morally permissible for humans to eat animals for pleasure. While those in this camp may require that the consumed animals are treated well throughout their lives and that they are killed "humanely," they maintain that morality does not forbid eating animals even when it is not a matter of survival. Let's call this group "the omnivores."[27] On the other hand, many intelligent, informed, and open-minded individuals believe that it is not morally permissible for humans to eat animals when it is not necessary for survival. Let's call this group "the vegetarians."[28] Some within this latter group believe that it is also morally impermissible to eat animal products for pleasure. Let's call this group "the vegans."[29] Members of each of these groups are aware of the arguments made by the members in the other groups. No one group can plausibly be maintained to be less biased, more intelligent, or more informed than any of the other groups. Put differently, the members from each of these groups are roughly in an equally good epistemic position on the matter as members from any of the other groups. From our perspective, and given our evidence, the members in each group are roughly in an equally good epistemic position on the matter, yet the disagreement persists. Given this fact and the conciliationist position advanced above, we should suspend judgment about what our evidence supports on this matter. Since we should suspend judgment about which arguments in this debate are sound, we should also withhold judgment on whether it is permissible to eat animals for pleasure. As in the Restaurant Check case, skepticism is called for until we can gather a decisive consensus among the relevant experts. The groups in the animal ethics debate are analogous to the groups in the Restaurant Check case.[30] Unlike in the Restaurant Check case, it is doubtful that such a consensus on the permissibility of eating animals for pleasure is on the horizon given our lack of a reliable moral calculator. So, given the epistemic significance of disagreement, we should be skeptical about whether it is morally wrong to eat animals for pleasure.

In fact, we should be skeptical about whether it is a *serious* moral wrong to eat animals for pleasure. Those who believe that it is morally wrong to eat animals for pleasure do not maintain that it is a minor moral foible to eat animals for pleasure, but that it is a grave moral transgression. So, the disagreement about the permissibility of eating animals for pleasure can be seen as a disagreement about whether it is a serious moral wrong to eat animals for pleasure.

However, one alternative to eating animals for pleasure is to refrain from eating animals for pleasure. Here, there is no significant moral controversy. There is widespread consensus among the relevant experts that not eating animals for pleasure is a morally permissible course of action. While there is disagreement about whether that course of action is morally *required*, there is not

such disagreement about whether it is morally *permissible*. It is uncontroversially morally permissible to be a vegetarian or a vegan. While there are a few who have argued against the permissibility of such actions,[31] there is no significant controversy surrounding this issue like that surrounding the issue of the permissibility of eating animals for pleasure. So, regarding the moral permissibility of vegetarianism, there is insufficient disagreement to block our knowledge of its moral status—we know that this is a morally permissible action.[32]

Here, the analogous disagreement over the shares of a restaurant check would be a case where nine individuals all believe that the shares are $37 and one sole individual maintains that the shares are $39. While there is not universal consensus, in such a disagreement it is reasonable to locate the error with the sole disagreeing party.[33] That this single party made a mistake is a better explanation of the disagreement than that all other nine individuals made a mistake and came to the same conclusion. This shows that we needn't have universal consensus to have reasonable beliefs. All we need is a good reason to locate the error with the dissenting party. So, applied to our issue, while the relevant experts are not unanimous in their belief that vegetarianism is morally permissible, there is nevertheless an overwhelming agreement that it is. Given such agreement, we can know that vegetarianism is permissible.

Combining these considerations we can see that the antecedent of MC is met with regard to the permissibility of eating animals for pleasure: we should suspend judgment as to whether eating animals for pleasure is a serious moral wrong, and we know that not eating animals for pleasure is not a serious moral wrong. If so, then MC tells us that we should not eat animals for pleasure. Exercising moral caution in these epistemic circumstances has it that we should not eat animals for pleasure. This provides a quite strong argument against eating animals for pleasure; one that does not rely on it being a serious moral wrong to eat animals for pleasure without appeal to one's epistemic situation.

It will be helpful to contrast this issue with another neighboring issue in applied ethics. Consider the moral permissibility of conducting cancer research on animals. Such research is controversial among the experts. Some believe that this research is morally wrong in that it uses sentient creatures as a mere means for our benefit. Others do not see these actions as morally problematic due to the potential benefit to human lives. Given the state of the disagreement, we should be skeptical about whether such animal experimentation is morally permissible. However, with regard to conducting cancer research on animals, the relevant epistemic asymmetry does *not* obtain. Refraining from conducting these experiments is *itself* a topic of moral controversy. Many experts believe that in refraining from conducting these experiments on animals we would be doing something seriously morally wrong. After all, there is a possibility that we can make things much better for human lives. In contrast, we have already noted that others believe that such experimentation is morally wrong. So with regard to cancer research on animals, for all we know conducting the research is a serious moral wrong, but for all we know failing to conduct the research is *also* a

serious moral wrong. So, on this issue, MC fails to give a prescription since its antecedent is not met; on this issue moral risk is inevitable—we have no opportunity to exercise moral caution.

The question of the moral permissibility of conducting cancer research on animals is importantly different than the question of the moral permissibility of eating animals for pleasure. Only in the latter case does the relevant epistemic asymmetry obtain. What is important to the argument here is that refraining from eating animals for pleasure is uncontroversially permissible. Since, for all we know, eating animals for pleasure is a serious moral wrong, and we have an uncontroversially permissible alternative, we should not animals for pleasure.

4. Objections

4.1. Aren't There Problems with Appealing to Philosophical 'Experts'?

One worry for this view concerns the nature of philosophical "experts." The application of MC to our ethical issue came by way of noting the widespread and persistent disagreement amongst the relevant experts over the permissibility of eating animals for pleasure. One may question who the relevant experts are, as well as whether there even is such a thing as a "philosophical expert."[34] After all, whereas the experts in some fields can be plausibly held to have a good deal more true beliefs than false beliefs on the topic, this is plausibly not the case in philosophy. It would not be surprising if even the best philosophers still had an overall bad track record with respect to their philosophical beliefs (at least the interesting ones). It could be that those who are best positioned (epistemically) to determine the truth of various philosophical claims are still in the end not very well positioned to determine the truth of philosophical claims. The lack of philosophical progress may also be seen as an indication that even our best philosophical minds really aren't that good at doing philosophy (at least insofar as attaining true philosophical beliefs is the goal).[35] Further, sorting out the opinions of the moral experts (or the closest thing we have to them) is by no means straightforward.[36] Given the problems in identifying the relevant experts, and what they believe, one might worry that the skeptical verdict simply does not follow.

However, far from mitigating the skeptical impact of disagreement, our inability to rationally determine the state of disagreement and the lack of success of those best positioned (epistemically) in the field only *enhances* the skeptical worry. If there are no philosophical experts, if even the best of us are wrong more often than not concerning philosophical claims, then this only reinforces the motivation to be skeptical about a great deal in philosophy. So, even if we are aware of what the "expert" opinion on some matter is, if we can reasonably take the "experts" (those best positioned epistemically) to not be very good at things like this, then we have a powerful defeater for *any* belief on the matter (including our own). Further, it looks as though we often are not in a position to

determine what the philosophical "expert" opinion is. If we should suspend judgment about what the expert opinion on some matter is, then we should suspend judgment about that matter as well. So, neither our inability to sort through the expert opinion nor the lack of philosophical experts mitigates the skeptical worry. If anything, such considerations only highlight the epistemically precarious position we are in with respect to controversial philosophical claims.[37]

Now, here the objector might chime back in and raise the worry from a different angle, claiming that given what was just said we also cannot know of any of the alternative actions (like abstaining from eating animals for pleasure) that they are not morally wrong, and thus condition (ii) of MC is not met. After all, if we can't sort through the expert opinions very well, or if the relevant "experts" aren't very good, then don't we fail to have this knowledge as well? If so, then the antecedent of MC will not be met, and the principle will be left without real world application.

While there is a legitimate worry here, I do not believe that the skeptical threat extends to our knowledge that abstaining from eating animals for pleasure is morally permissible. While this too is a philosophical claim, it is one of those rare philosophical claims that enjoy a considerable consensus from the philosophical community. While we are certainly not infallible on such matters, it is plausible that many philosophical claims that enjoy such a widespread consensus among philosophers are among those few philosophical claims that we can rightly be said to have knowledge about. That is, we appear to have a good reason for treating claims like "refraining from eating animals for pleasure is morally permissible" as belonging to a special class of philosophical claims—one to which our skeptical worries do not apply due to the overwhelming consensus about them. Now, this circumstance (the philosophical consensus) is clearly a contingent matter. The philosophical landscape may change and become such that we no longer know that abstaining from eating animals for pleasure is morally permissible. Should that day come, our principle MC would no longer apply to the issue of whether it is morally permissible to eat animals for pleasure. However, for better or worse, it does not appear that we are in such a situation as of right now.

4.2. Isn't This View Self-Defeating?

Another worry for this view is that it is self-defeating. Such a charge has been leveled against Conciliatory Views of disagreement in general,[38] and the application of MC to applied ethics was motivated by just such a view. The self-defeating worry gets started by noting that views on the epistemic significance of disagreement are themselves controversial philosophical views. There is widespread and persistent disagreement among many intelligent, informed, and open-minded individuals on this topic as well. If we should be skeptical about claims that are significantly controversial amongst the relevant experts, then we should be skeptical about these claims regarding the epistemic significance of

disagreement as well. An objection, then, is that our inability to be justified in believing Conciliatory Views of disagreement prevents MC from having the real-world application claimed above.

There are several things to say in response to such a worry. First, the sense in which these views are self-defeating does not concern their truth.[39] At most, such views are self-defeating in that they do not allow for their proponents to be justified in believing their views by their own lights. So, at most, the upshot here is that we are not justified in believing these views about the significance of disagreement, not that such views are thereby false. If Conciliatory Views of disagreement are true, then we are not justified in believing them.

Second, it does not follow from this feature of Conciliatory Views that we are justified in believing other competitor views about the epistemic significance of disagreement. If Conciliatory Views of disagreement are true, then we are not justified in believing *any* sufficiently controversial claims. While Conciliatory Views of disagreement may fall into this camp, so do competitor views of the epistemic significance of disagreement. So, the truth of Conciliatory Views of disagreement, when coupled with the current state of controversy, would have it that we are not justified in believing *any* view about the epistemic significance of disagreement. Put differently, the truth of Conciliatory Views would undermine our justification for believing *all* of the candidate views on the epistemic significance of disagreement (at least given the current state of controversy about them).[40]

This is particularly relevant since, if we should suspend judgment about how to respond to such disagreement, then it seems that we should also suspend judgment about the controversial claims themselves. If we should suspend judgment about what the correct view of the epistemic significance of disagreement is, then we should suspend judgment about how to evaluate the higher-order evidence coming from the disagreeing parties. But, if we should suspend judgment about how to evaluate the evidence regarding disputed claims, then we should suspend judgment about those disputed claims as well. Recall our analogy with external world skepticism. After all, my higher-order evidence is evidence about the rest of my evidence and what it supports—it is evidence about what to make of the first-order evidence. In this way, my higher-order evidence is not merely one more piece of my evidence. So, if I should suspend judgment about what to make of my evidence, then my evidence fails to support a non-skeptical attitude. If I can't tell what to make of the epistemic significance of disagreement, but am aware that there is sufficient controversy about p, then I should suspend judgment about p. This is but one more application of our principle that if we should suspend judgment about what our evidence concerning p supports, then we should suspend judgment regarding p as well.[41] If Conciliatory Views of disagreement are false, then MC will fail to have the application I have claimed here, but the consequence that we are not justified in believing Conciliatory Views of disagreement fails to block the application of MC.

What about our principle MC? Can we be justified in believing it, or should we be skeptical about it as well? Given what has been said, we can only be justified in believing MC if it is not itself subject to significant controversy among the relevant experts. As we saw above, while perhaps most claims in philosophy are quite controversial, not all are. Whether MC is subject to significant controversy will be determined in part by how well the arguments in this paper are received. But, as mentioned above, MC appears to be quite plausible. If MC is rather uncontroversial, then the epistemic significance of disagreement would not threaten our application of MC by mandating that we be skeptical about MC itself.

Further, even if we ought to be skeptical about MC, this fact would not preclude our application of MC. If we should be skeptical about the truth of MC, then we should be skeptical about it being morally permissible to eat animals for pleasure, for the reasons considered above. But if we should be skeptical about whether it is permissible to eat animals for pleasure, then so long as we still know that refraining from eating animals for pleasure is permissible and MC is in fact true, then we should refrain from eating meat for pleasure. MC does not claim that we must be justified in believing MC for it to have application to us. So long as MC is true and its antecedent conditions are met, we should refrain from eating meat for pleasure. Thus, even if the epistemic significance of disagreement has it that we are not justified in believing MC, this would not prevent the application of MC argued for here.[42]

4.3. Moral Spinelessness

Perhaps the most pressing problem for MC comes from the charge that it leads to moral spinelessness. While Conciliatory Views of disagreement have been defended from the charge that they are spineless,[43] MC may be thought to face a deeper problem by connecting one's epistemic situation to what is morally permissible. After all, according to MC what it is permissible for you to do depends in part on others—in particular it depends upon what others believe is morally permissible. It might be thought that this leaves MC with some untoward consequences, particularly in cases where the relevant experts are terribly confused about morality. For instance, it is unfortunately not too difficult to imagine a case where given the epistemic significance of disagreement an individual ought to suspend judgment about whether miscegenation is a serious moral wrong. We can imagine a situation where the relevant experts are quite evenly divided about this issue, with numerous intelligent, good-willed, and equally informed individuals on either side of the debate. In such a context, MC appears to have the consequence that you morally ought not marry or have sexual relations with someone of a different race. We know that in general there is nothing morally wrong with either interracial marriages or interracial sexual relations, and thankfully no such controversy exists today, but such controversy is at least *possible* (and plausibly has been *actual* at various times in history).

This consequence may seem even worse if any kind of moral progress depends upon some individuals standing against the majority or "expert" moral opinions and living in such a way that many see as morally wrong at the time. If moral progress depends upon individuals living lives that flout the currently accepted norms, then MC also appears to stand in the way of moral progress by sanctioning those bold enough to act on the courage of their convictions.

So, what can be said in defense of MC? First, in cases where the subject is reasonable in believing that she is in such a situation, it appears that condition (ii) of MC is not met. If our subject has good reason to take it that standing up to the majority or "expert" moral opinions and acting on her convictions is the only way to achieve moral progress in her society, then it may be that she would not *know* that refraining from so acting is morally permissible. She may know that refraining from such actions is morally permissible in more "ordinary" circumstances, but if she has good reason to believe that her circumstances are not ordinary in this way (if it is reasonable for her to believe that true moral progress hinges on her actions), then it seems that she would not know that refraining from those actions is morally permissible for her. In some cases, it may be that moral progress requires "moral renegades"[44] and that the call toward moral progress is weightier than the call to exercise caution. But if so, the antecedent of MC fails to be met in such situations (namely condition [ii]) since you cannot know what is not the case. So, there is good reason to believe that MC does not stand in the way of moral progress in the alleged way.

Nevertheless, MC can still be seen to forbid actions that would otherwise be permissible, and this might be seen to be problematic enough.[45] MC claims that in some circumstances individuals can be victims of an impoverished epistemic environment, and as such, may have fewer moral paths open to them than they otherwise would have. Ethicists have long noted that one's context can make a difference for morality—that different contexts call for different actions. What MC points out is that one's *intellectual context* is also relevant for morality. Different intellectual contexts can call for different actions just as different social contexts can call for different actions. This highlights the importance of philosophy and good ethical thinking, since according to MC sometimes changes must occur along intellectual lines before certain options are made morally available. According to this line of thought, intellectual progress may sometimes need to precede moral progress—changes in thinking may need to precede changes in what behavior is moral.

Motivation for the claim that our epistemic situation can constrain our moral options comes from the following example. Suppose that we are having a bonfire and enjoying throwing various things onto the fire. A crate is then delivered to us marked "living organisms." We cannot see what is inside the crate, and we cannot determine what kind of living organisms are contained therein. It could be cacti. It could be bunnies. It could be a human child. It is clear that we should not simply throw the crate, or commit ourselves to throwing the contents of the crate into the fire without first having a justified belief that it would not be

morally wrong to do so. Suppose that when we open the crate we discover that it is full of diseased cacti. We can even suppose that given their disease, it is important that the cacti be destroyed by fire as soon as possible. Having become justified in believing that it is permissible to throw the cacti into the fire, it is now morally permissible for us to throw the crate and its contents into the fire. However, before opening the crate to uncover its contents (while we were justified in suspending judgment about the morality of throwing the contents into the fire), it would not have been morally permissible to throw the contents of the crate into the fire. To do so, to throw the crate onto the fire without being justified in believing that doing so was permissible, would be to undergo undue moral risk.

However, it is worth noting that nothing about the contents of the crate changed from its delivery to our opening of the crate. Further, in general there is nothing morally wrong with throwing diseased cacti into a fire. The relevant change occurred in our epistemic position regarding the morality of throwing what is in the crate into the fire, and this epistemic change made all the difference for what we are morally permitted to do with the crate and its contents. Our impoverished epistemic position regarding the morality of throwing the contents of the crate into the fire made a significant difference regarding what we were morally permitted to do with the contents of the crate (at least while we remain in that impoverished epistemic position). Taking morality seriously required waiting to find out the contents of the crate before we (morally) could throw it onto the fire.[46]

Being mistaken or confused about the moral significance of one's race is much worse than being confused about the contents of a crate. However, the lesson still applies. As unfortunate as it is, to be unaware of the moral status of throwing the contents of the crate into the fire, it is far more unfortunate to be in an epistemic situation where the moral significance of race is at issue. The world would be a much better place were there never to have been any such controversies. That said, the unfortunate existence of such controversies appears to also have the unfortunate consequences of limiting the range of morally permissible actions open to individuals in those unfortunate circumstances (at least so long as the subject knows that refraining from those actions is morally permissible). It is morally wrong to take undue moral risks, so in some situations it would be wrong to perform certain actions that would not otherwise be wrong to do.

5. Conclusion

In this article, I have put forward and motivated an epistemic principle for applied ethics, one that claims that we should exercise moral caution. I have argued that this principle also has some significant consequences in applied ethics. I have briefly examined how this principle applies to whether it is permissible to eat animals for pleasure. If the arguments here are correct, then there is a powerful argument against being an omnivore that does not rely on it being a

moral wrong to eat animals for pleasure regardless of one's epistemic situation. So, MC provides a powerful case against the permissibility of eating animals for pleasure.

Previous versions of this paper were presented at the Florida Philosophical Association, the Alabama Philosophical Society, and the Association for Practical and Professional Ethics. I am grateful to the audiences at each of these venues for their helpful feedback. I would also like to thank Bryan Bannon, Brandon Carey, Aaron Creller, Matt Frise, Jonathan Fuqua, Errin Gilson, Richard Greene, Sarah Mattice, Kevin McCain, Nathan Nobis, Ted Poston, Jason Rogers, Bill Rowley, Alissa Swota, Sarah Wright, and several anonymous referees for helpful comments on this paper at various stages of construction. Finally, special thanks to the Florida Blue Center for Ethics at the University of North Florida for funding this research.

Notes

[1] There is a significant literature on the ethics of belief, pertaining to when it is the case that one (epistemically) ought to believe something. For instance, see Clifford (1999), James (1956), Chisholm (1956), Feldman (2000), Kornblith (2001), and Wolterstorff (2010). MC takes no stand on these debates. MC can be coupled with one's favored account of when one (epistemically) ought to believe/disbelieve/suspend judgment in general. That said, the application of MC that I will argue for below does rely on the claim that you (epistemically) ought to suspend judgment about claims that are suitably controversial amongst the relevant experts. However, those claims about the epistemology of disagreement are themselves neutral with respect to a more general account of the ethics of belief since what is given there is an account of defeat rather than a positive account of what factors make it such that you should believe a proposition.

[2] I will leave it implicit that actions are individuated in part by the relevant context. So, when I mention a particular action, it is always with respect to a particular context as well.

[3] It is worth noting here that the antecedent of MC does not merely require that you don't know that the action in question is not a serious moral wrong. There are many reasons why you could lack such knowledge: you could be gettiered, you could lack knowledge-level justification, and so forth. MC concerns an epistemic situation different than this—a situation where you (epistemically) should either believe or suspend judgment that it is a serious moral wrong. In such cases you will also lack knowledge that p, but not all ways of lacking knowledge have the same implications.

[4] A subject can be propositionally justified in adopting a doxastic attitude toward a proposition whether or not she has in fact adopted that doxastic attitude toward that proposition, and whether or not she has adopted it on the basis of the reasons that justify it.

[5] Thank you to Jason Rogers for helping me see this potential problem with MC.

[6] Thanks to an anonymous reviewer for pressing this point.

[7] As stated, MC is concerned with all-or-nothing doxastic attitudes like belief, disbelief, and suspension of judgment. Some philosophers prefer to think of our doxastic options in a more fine-grained way, and instead speak of degrees of belief or credences. They think one's doxastic options can be as precise as single point values on a 0–1 scale (inclusive) or a range of probability functions. On this scale, 0 represents a maximal confidence that the proposition is false, and 1 represents a maximal confidence that the proposition is true. The numbers between 0 and 1 proportionally represent a less than maximal degree of confidence. For those so inclined, MC may require some supplementation to indicate how various degrees of belief (or credences) map

onto the more coarse-grained doxastic attitudes. That said, I will not do so here. For more on the connection between all-or-nothing doxastic attitudes and more fine-grained degrees of belief, see Christensen (2004), Buchak (2014), Easwaran (2015), and Schoenfield (2012).

[8] It might be thought that merely requiring knowledge that *some* alternative action to A is morally permissible is problematic. After all, there could be cases where the act that the subject *will actually choose* is much worse than A and so it is not the case that S should not A even if the antecedent of MC is met (thanks to an anonymous referee for raising this concern). However, it is hard to see why this would be the case. Suppose that I know that it is permissible to praise someone but I am in fact about to viciously slander that person. The worry here is that I ought to "gently" slander them instead of viciously slandering them (even if I should suspend judgment over whether "gentle" slander is morally wrong). While it is true that I should gently slander them *given that I will slander them*, it is by no means clear that I just plain ought to gently slander them. In fact I take it that I should not do so. Some morally wrong actions are worse than others, but this does not make it the case that we ought to do *any* of those wrong actions (they are, after all, morally wrong). If anything, what we ought to do is some conditional action such as—if I am to slander, slander gently—but from such a conditional it does not follow that I ought to slander gently. In addition, MC will prescribe against doing any alternative of A that itself meets the antecedent conditions of MC.

[9] It is worth noting that MC takes these claims for granted in the sense that it only has application to the extent that moral knowledge is possible. If we lack moral knowledge, MC would not be shown to be false, but only without application.

[10] For a history of the application of the precautionary principle, see Raffensperger and Tickner (2007).

[11] See Bodansky (1992, 2004) and Gardiner (2006) for laying out the conceptual space.

[12] In this way, MC is like the third interpretation of the precautionary principle given by Bodansky (2004).

[13] One might think that the relevant normative ignorance requisite in MC could provide an excuse for the subject and thus have it that her action is not morally wrong (contra the verdict of MC). While it is plausible that non-culpable ignorance excuses, such an excuse fails to have such a consequence. The morality of an act and the moral responsibility of the subject for committing the act are distinct issues (one ethical, one metaphysical). That the subject has an excuse for A-ing due to non-culpable normative ignorance entails that the subject is not morally responsible for A-ing, not that her A-ing was not wrong. For more on this distinction see Smith (1983), Ginet (2000), and Pereboom (2001).

[14] There is another related sense in which MC is more precise than the precautionary principle. While the precautionary principle does not explicitly state what level of risk should trigger precautions, MC is clear on this front. According to MC, when you should believe or suspend judgment that an act is a serious moral wrong while knowing that an alternative is permissible, it is sufficiently risky and thus it would be morally wrong to so act. So, MC also makes precise what level of risk is relevant.

[15] MC is also in the spirit of rigorism, or tutiorism, in claiming that we have a duty to play it safe (thanks to an anonymous referee for bringing this connection to my attention). See Pascal (1967) for a defense. According to rigorism, you are morally obligated to take the course of action that is the least likely to involve wrongdoing. Rigorism thus claims something stronger than MC. Rigorism gives verdicts about cases that MC is silent on. MC only makes moral prescriptions in cases where there is a certain level of moral risk—cases where you should believe or suspend judgment that doing A is a serious moral wrong while you know that an alternative to A is morally permissible. Rigorists would agree that we have such a duty, but they also go beyond MC in requiring even greater degrees of moral caution be exercised.

[16] Arguments for Conciliatory Views can be found in Elga (2007), Christensen (2007), Feldman (2006a,b), Frances (2010), and Matheson (2015a).

[17] Arguments for Steadfast Views can be found in Kelly (2005, 2010), Plantinga (2000), van Inwagen (1996), Foley (2001), Enoch (2010), Bergmann (2009), Lackey (2010), and Moffett (2007).

[18] For this reason, even if Steadfast Views of *peer* disagreement are correct, it is doubtful that they would present a challenge to the arguments offered here. While there may be important asymmetries to appeal to in two-party disagreements (self-trust, private evidence, etc.) it is doubtful that such appeals can be effective when the disagreement is broader (incorporating a large number of individuals), and between individuals who are in an even better epistemic position on the matter than you are.

[19] This adaptation alters Christensen's case from a two-party peer disagreement, to a broader disagreement of the sort relevant here.

[20] For further defenses of the claim that discovered disagreement should lead to skepticism see Christensen (2007), Elga (2007), Conee (2009), Feldman (2006a,b), and Matheson (2009, 2015a).

[21] Along these lines, Christensen gives us the following principle:

> *Independence*: In evaluating the epistemic credentials of another's expressed belief about P, to determine how (or whether) to modify my own belief about P, I should do so in a way that doesn't rely on the reasoning behind my initial belief that P. (2009, 758)

Similar principles are endorsed in Elga (2007), Frances (2010), Kornblith (2010), and Matheson (2015a). For a critical discussion see also Kelly (2010), Lackey (2010), Lord (2015), and Sosa (2013).

[22] For a detailed argument on this matter see Carey and Matheson (2013) and Matheson (2015a).

[23] For more on this point see Adams (2013).

[24] While some things that we take to be moral disagreements may only be merely apparent moral disagreements (where there is no moral proposition over which the parties disagree), I will be taking it that at least some of our apparent moral disagreements are genuine.

[25] For instance, see Adams (2013), Audi (2008), Wedgwood (2010), Sinnott-Armstrong (2006; 2012), Shafer-Landau (1994, 2003), Huemer (2005), Ballantyne and Thurow (2013), and McGrath (2008; 2010; 2011).

[26] Regarding this later disjunct, we can imagine a case where terrorists kidnap you and threaten to kill your family if you do not partake in some meat from animals.

[27] For a defense of this group, see Crisp (1988) and Rogerson (2002).

[28] For a defense of this group, see Singer (1975) and Regan and Singer (1976).

[29] For a defense of this group, see Hooley and Nobis (2015) and Engel (2002).

[30] One may question the aptness of the analogy between ethical disagreements and restaurant check disagreements. After all, it is plausible that conative aspects of practical judgments play a larger role in disagreements with practical import (thanks to an anonymous referee for pressing me on this point). However, while there are differences between the cases, it does not appear that any of the differences are salient. While conative aspects may more readily influence matters with practical import, this factor tells equally against all sides in the imagined ethical disagreements. It is not plausible to maintain, say, that vegetarians are being misled by conative influences in ways that omnivores are not. While the conative aspects of practical judgments may lead us to more readily expect disagreement about such matters, they do not help us resolve them.

[31] See Davis (2003) and Zamir (2004) for two such examples.

[32] Further, many of the relevant arguments against vegetarianism will not apply since they merely require that sufficiently many individuals need to consume animal meat. Presumably, enough individuals will consume animal meat even if any given individual ceases to consume animal meat. So, such arguments would fail to give a reason now for any one individual to give up eating meat.

[33] Notice that this is true even when the dissenting party is a peer of the other parties.

[34] For an extended discussion on these points see Goldman (2001), Fumerton (2010), and Coady (2006; 2012).

[35] For more on this point, see Kornblith (2013).

[36] This problem is further explored in Carey and Matheson (2013), Matheson (2015a), and Goldman (2001).

[37] Similarly, it could be that we are mistaken about who is best positioned to discover philosophical truths. It could be that professional philosophers are actually much worse at this than the folk (thank you to an anonymous referee for pushing this point). What matters here is who we are *justified in believing* is best positioned and what we are *justified in believing* that they believe on the matter. While we could be wrong about who is best positioned, unless it is rational for us to believe that a decisive majority of those best positioned on the matter agree, then the skeptical results will still obtain.

[38] For more on this objection to Conciliatory Views of disagreement see Plantinga (2000), Elga (2010), and Weatherson (2014).

[39] This point is emphasized in Matheson (2015b).

[40] For more on responses to the self-defeat worry, see Matheson (2015b), Christensen (2009), Frances (2010), Kornblith (2013), Littlejohn (2013).

[41] For more on this point see Feldman (2006b), Matheson (2009), and Matheson (2015a).

[42] One might wonder what the point of advancing MC is given that disagreement over it (which is perhaps likely) will have it that we are not justified in believing it. However, the conciliationist should not be precluded from advancing philosophical claims simply because according to the view (and given the contingent disagreement) such claims are not justified. Debates are advanced, debates are turned in novel directions, and philosophical progress is made by advancing ideas, defending claims, and thinking of new examples, even if given the current state of disagreement we are not currently justified in believing the relevant claims. It is in that spirit that this article is written.

[43] See Elga (2007).

[44] This parallels Frances's (2010) term.

[45] In addition, one might believe that the call toward moral progress is supererogatory rather than obligatory.

[46] Notice that this verdict holds even if by waiting to uncover the contents of the crate, further harm has been done. We can suppose that in having not burned the crate immediately, the disease has further spread. Nevertheless, without having a justified belief that throwing the crate on the fire is morally permissible for us to do, it would have been morally wrong of us to do it.

References

Adams, Zed. 2013. "The Fragility of Moral Disagreement." In *Disagreement and Skepticism*, ed. D. Machuca, 131–49. New York: Routledge.

Audi, Robert. 2008. "Intuition, Inference, and Rational Disagreement in Ethics." *Ethical Theory and Moral Practice* 11: 475–92.

Ballantyne, Nathan, and Thurow, Joshua. 2013. "Moral Intuitionism Defeated?" *America Philosophical Quarterly* 50: 411–21.

Bergmann, Michael. 2009. "Rational Disagreement after Full Disclosure." *Episteme: A Journal of Social Epistemology* 6: 336–53.

Bodansky, D. 1992. "Scientific Uncertainty and the Precautionary Principle: Commentary: The Precautionary Principle." *Environment* 34: 4–5.

——. 2004. "Deconstructing the Precautionary Principle." In *Bringing New Law to Ocean Waters*, ed. D. D. Caron and H. N. Scheiber, 381–91. Leiden, The Netherlands: Brill.

Buchak, Lara. 2014. "Belief, Credence, and Norms." *Philosophical Studies* 169: 285–311.

Carey, Brandon, and Matheson, Jonathan. 2013. "How Skeptical is the Equal Weight View?" In *Disagreement and Skepticism*, ed. Diego Machuca, 131–49. New York: Routledge.

Chisholm, Roderick M. 1956. "Epistemic Statements and the Ethics of Belief." *Philosophy and Phenomenological Research* 16: 447–60.

Christensen, David. 2004. *Putting Logic in Its Place*. New York: Oxford University Press.
——. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116: 187–218.
——. 2009. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4: 756–67.
Clifford, W. K. 1999. "The Ethics of Belief." In *The Ethics of Belief and Other Essays*, ed. Timothy J. Madigan, 70–96. Amherst, MA: Prometheus. (Original edition 1877.)
Conee, Earl. 2009. "Peerage." *Episteme* 6: 313–23.
Coady, David. 2006. "When Experts Disagree." *Episteme* 3: 68–79.
——. 2012. *What to Believe Now: Applying Epistemology to Contemporary Issues*. Malden, MA: Wiley.
Crisp, Roger. 1988. "Utilitarianism and Vegetarianism." *International Journal of Applied Philosophy* 4 (1): 41–49.
Davis, Steven L. 2003. "The Least Harm Principle May Require That Humans Consume a Diet Containing Large Herbivores, Not a Vegan Diet," *Journal of Agricultural and Environmental Ethics* 16: 387–94.
Easwaran, Kenny. 2015. "Dr. Truelove or: How I Learned to Stop Worrying and Love Bayesian Probabilities." Advance online publication. *Noûs*. doi:10.1111/nous.12099
Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41: 478–502.
——. 2010. "How to Disagree About How to Disagree." In *Disagreement*, ed. Richard Feldman and Ted Warfield, 175–86. New York: Oxford University Press.
Engel, Mylan. 2002. "Why YOU are Committed to the Immorality of Eating Meat." In *Social and Personal Ethics*, ed. William Shaw, 4th ed., 212–221. Belmont, CA: Wadsworth.
Enoch, David. 2010. "Not Just a Truthometer: Taking Oneself Seriously (but Not Too Seriously) in Cases of Peer Disagreement." *Mind* 119: 953–97.
Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60: 667–95.
——. 2006a. "Epistemological Puzzles about Disagreement." In *Epistemological Futures*, ed. Steve Hetherington, 216–36. New York: Oxford University Press.
——. 2006b. "Reasonable Religious Disagreements." In *Philosophers without Gods: Meditations on Atheism and the Secular Life*, ed. L. Antony, 194–214. New York: Oxford University Press.
Foley, Richard. 2001. *Intellectual Trust in Oneself and Others*. Cambridge: Cambridge University Press.
Frances, Bryan. 2010. "The Reflective Epistemic Renegade." *Philosophy and Phenomenological Research* 81: 419–63.
Fumerton, Richard. 2010. "You Can't Trust a Philosopher." In *Disagreement*, ed. R. Feldman and T. Warfield, 91–110. New York: Oxford University Press.
Gardiner, S. 2006. "A Core Precautionary Principle." *Journal of Political Philosophy* 14: 33–60.
Ginet, Carl. 2000. "The Epistemic Requirements for Moral Responsibility." *Philosophical Perspectives* 14: 267–77.
Goldman, Alvin. 2001. "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63: 85–110.
Harris, J., and Holm, S. 1999. "Precautionary Principle Stifles Discovery." *Nature* 400: 398.
——, and ——. 2002. "Extending Human Lifespan and the Precautionary Paradox." *Journal of Medicine and Philosophy* 27: 355–68.
Hooley, Daniel, and Nobis, Nathan. 2015. "A Moral Argument for Veganism." In *Philosophy Comes to Dinner: Arguments on the Ethics of Eating*, ed. Andrew Chignell, Matthew Halteman, and Terence Cuneo. New York: Routledge.
Huemer, Michael. 2005. *Ethical Intuitionism*. London: Palgrave.
James, William. 1956. *The Will to Believe and Other Essays in Popular Philosophy; Human Immortality*. New York: Dover.

Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology*, Vol. 1, ed. John Hawthorne and Tamar Gendler, 167–96. New York: Oxford University Press.

——. 2010. "Peer Disagreement and Higher Order Evidence." In *Disagreement*, ed. R. Feldman and T. Warfield, 111–74. New York: Oxford University Press.

Kornblith, Hilary. 2001. "Epistemic Obligation and the Possibility of Internalism." In *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, ed. Abrol Fairweather and Linda Zagzebski, 231–48. New York: Oxford University Press.

——. 2010. "Belief in the Face of Controversy." In *Disagreement*, ed. R. Feldman and F. Warėeld. New York: Oxford University Press.

——. 2013. "Is Philosophical Knowledge Possible?" In *Disagreement and Skepticism*. ed. Diego Machuca, 260–76. New York: Routledge.

Lackey, Jennifer. 2010. "What Should We Do When We Disagree?" In *Oxford Studies in Epistemology*, ed. T. Gendler and J. Hawthorne, Vol. 3, 274–93. Oxford: Oxford University Press.

Littlejohn, Clayton. 2013. "Disagreement and Defeat." In *Disagreement and Skepticism*, ed. D. Machuca, 169–92. New York: Routledge.

Lord, Errol. 2015. "From Independence to Conciliationism: An Obituary." *Australasian Journal of Philosophy* 92: 365–77.

Matheson, Jonathan. 2009. "Conciliatory Views of Disagreement and Higher-Order Evidence." *Episteme: A Journal of Social Philosophy* 6: 269–79.

——. 2015a. *The Epistemic Significance of Disagreement*. New York: Palgrave Macmillan.

——. 2015b. "Are Conciliatory Views of Disagreement Self-Defeating?" *Social Epistemology* 29 (2): 145–59.

McGrath, Sarah. 2008. "Moral Disagreement and Moral Expertise." In *Oxford Studies in Metaethics*, Vol. 3, ed. Russ Shafer-Landau, 87–107. Oxford: Oxford University Press.

——. 2010. "Moral Realism without Convergence." *Philosophical Topics* 38 (2): 59–90.

——. 2011. "Skepticism about Moral Expertise as a Puzzle for Moral Realism." *Journal of Philosophy* 108 (3): 11–137.

Moffett, Mark. 2007. "Reasonable Disagreement and Rational Group Inquiry." *Episteme: A Journal of Social Epistemology* 4: 352–67.

Pascal, Blaise. 1967. *The Provincial Letters*, ed. and trans. A. J. Krailsheimer. Harmondsworth, UK: Penguin.

Pereboom, Derk. 2001. *Living Without Free Will*. Cambridge: Cambridge University Press.

Plantinga, Alvin. 2000. "Pluralism: A Defense of Religious Exclusivism." In *The Philosophical Challenge of Religious Diversity*, ed. Philip L. Quinn and Kevin Meeker, 172–92. New York: Oxford University Press.

Raffensperger, C., and Tickner J. A. 2007. "To Foresee and to Forestall." In *Public Health Ethics: Theory, Policy, and Practice*, ed. R. Bayer, L. O. Gostin, B. Jennings, and B. Steinbock, 319–25. Oxford: Oxford University Press.

Regan, Tom, and Singer, Peter, eds. 1976. *Animal Rights and Human Obligations*. New York: Prentice-Hall.

Rogerson, Kenneth F. 2002. "On the Morality of Eating Animals." *Southwest Philosophy Review* 18 (1): 105–111.

Shafer-Landau, Russ. 1994. "Ethical Disagreement, Ethical Objectivism and Moral Indeterminacy." *Philosophy and Phenomenological Research* 54 (2): 331–344.

——. 2003. *Moral Realism: A Defense*. New York: Oxford University Press.

Schoenfield, Miriam. 2012. "Chilling Out on Epistemic Rationality: A Defense of Imprecise Credences (and Other Imprecise Doxastic Attitudes)." *Philosophical Studies* 158 (2): 197–219.

Singer, Peter. 1975. *Animal Liberation*. New York: Avon.

Sinnott-Armstrong, Walter. 2006. "Moral Intuitionism Meets Moral Psychology." In *Metaethics after Moore*, ed. Terry Horgan and Mark Timmons, 339–65. New York: Oxford University Press.

——. 2012. "An Empirical Challenge to Moral Intuitionism." In *The New Intuitionism*, ed. Jill Graper Hernandez, 11–28. London: Continuum Press.

Smith, Holly. 1983. "Culpable Ignorance." *Philosophical Review* 92: 543–72.

Sosa, Ernest. 2013. "The Epistemology of Disagreement." In *The Epistemology of Disagreement: New Essays*, ed. J. Lackey and D. Christensen. Oxford: Oxford University Press.

van Inwagen, Peter. 1996. "It Is Wrong, Always, Everywhere, and for Anyone, to Believe Anything, Upon Insufficient Evidence." In *Faith, Freedom, and Rationality*, ed. J. Jordan and D. Howard-Snyder, 137–54. Lanham, MD: Rowman and Littlefield.

Weatherson, Brian. 2014. "Disagreements, Philosophical and Otherwise." In *The Epistemology of Disagreement: New Essays*, ed. D. Christensen and J. Lackey, 54–76. New York: Oxford University Press.

Wedgwood, Ralph. 2010. "The Moral Evil Demons." In *Disagreement*, ed. R. Feldman and T. Warfield, 216–46. Oxford: Oxford University Press.

Wolterstorff, Nicholas. 2010. "Ought to Believe – Two Concepts." In *Practices of Belief*, ed. Terence Cuneo, 62–85. Cambridge: Cambridge University Press.

Zamir, Tzachi. 2004. "Veganism." *Journal of Social Philosophy* 35 (3): 367–79.