
METEORITE

THE UNDERGRADUATE JOURNAL OF PHILOSOPHY
AT THE UNIVERSITY OF MICHIGAN

University of Michigan
Ann Arbor

SPRING 2020

EDITOR-IN-CHIEF: Colton Karpman

EDITORIAL STAFF: Philip Caldwell, David Kamper, Andrew Kovacs, Samit Lamba, Alison Lo, Shihua Lu, Zuzanna Lutrzykowska, Rachel Mazzotta, Charlotte Meltzer, Veronica Sikora, Tristan Sirls, Trey Smith

FACULTY ADVISOR: David John Baker, Ph.D

Meteorite gratefully acknowledges the support of Judith Beck and the staff of the Tanner Library, and the faculty and staff of the University of Michigan Department of Philosophy, without whom this publication would not be possible.

All authors retain copyright on original work unless otherwise noted.

Please direct copyright inquiries and requests for back issues to:

Meteorite c/o
Department of Philosophy
2215 Angell Hall
435 S. State St.
Ann Arbor, MI 48109
USA

Contents

Dogmatism and Cognitive Penetration	3
<i>Bin Hui Kwon</i>	
<i>Harvard University</i>	
Affected Ignorance and Moral Justification	9
<i>Patrick Magee</i>	
<i>Harvard University</i>	
An Interview with Peter Railton	15
About <i>Meteorite</i>	21

Dogmatism and Cognitive Penetration

Bin Hui Kwon
Harvard University

Introduction

The dogmatist about perceptual justification — henceforth “the dogmatist” — is committed to two claims: first, that a subject’s perceptual experience with some content, p , *prima facie* justifies her beliefs that p ; and second, that this *prima facie* justification is immediate, requiring no additional justifiers. For instance, according to the dogmatist, if I see a bird in flight, then I have *prima facie* justification for the belief, “there is a bird in flight.” According to his critic, the dogmatist’s first claim faces serious challenges from certain cases of “cognitive penetration” — a phenomenon in which some non-perceptual state of the subject, such as a belief or mood, alters his experience to correspond to that non-perceptual state.¹ To illustrate: suppose that there is an over-eager birder trekking through the forest. His expectation that he will see a bird in flight causes him to see a bird in flight, leading him to believe, “there is a bird in flight,” even though the distal stimulus is in fact a bat.² According to the dogmatist, unless there is a defeater readily available to the birder which would undermine the justificatory connection, his belief is *prima facie* justified. The critic, however, would object: on her view, it is implausible for the birder’s belief to be justified at all. Indeed, she would object even if the birder’s belief were true; even if the distal stimulus were a bird, the birder’s belief would be unjustified since he would have seen a bird regardless of whether or not there is one. We may call all such cases of cognitive penetration such as these “epistemically perverting.” In what follows, I argue that the critic is wrong; epistemically perverting cognitive penetration poses no particular threat to the dogmatist’s first claim.

Section I

I begin by clarifying the meaning and significance of “*prima facie* justification.” Put simply, “*prima facie* justification” means justification absent defeaters. More specifically, since dogmatism is traditionally cast in internalist terms, it means justification absent “evidential” defeaters, or defeaters accessible to the subject.³ Hence, defeaters not readily accessible to the subject, or “propositional” defeaters, are excluded from undermining *prima facie* justification. Given that dogmatism is to be understood within an internalist framework, the distinction between *a priori justification* and other kinds of justification is irrelevant to this dispute between the dogmatist and the critic; the critic could easily stipulate cases in which there are no evidential defeaters available to the subject which would undermine the putative justificatory connection — for instance, by introducing severe repression or amnesia in the subject.⁴

¹Siegel (2011), p. 208. For empirical studies on cognitive penetration, see Eberhardt et al. (2004), Levin and Banaji (2006), Stefanucci and Guess (2009), Stokes and Payne (2010), and Banerjee et al. (2012). For criticism of these studies’ methods, see Firestone and Scholl (2013).

²I mean by “expectation” a kind of belief directed toward a future state-of-affairs.

³Pyror (2000), fn. 9; Siegel (2011), p. 215.

⁴E.g., see Siegel (2011), p. 216.

So, leaving out the “*prima facie*” in “*prima facie* justification” for now, we should consider the critic’s grounds for denying any justification whatsoever to the victims of epistemically perverting cognitive penetration. There are at least two objections available to the critic’s intuitions which might tempt the critic to do this. The first of these is that the victim of cognitive penetration is somehow implicated in a vicious circle of belief and experience. This objection can be unpacked in a couple different ways. The crude approach is to accuse the dogmatist of defying the non-circularity principle: the principle that a belief cannot be in its own prior justificatory chain. But the dogmatist is surely innocent of the charge. Cognitive penetration consists in a feedback loop of belief and experience composed of a causal relation and a justificatory relation, not justificatory relations going each way.⁵ A more sophisticated approach requires that there exist special kinds of causal relations between information, as well as some principle prohibiting justification of a belief when the belief arises from a certain kind of causal circle.⁶ Given this, the critic can argue that epistemically perverting cognitive penetration vitiates justification because it results in this causal circle. Setting aside whether or not all cases of epistemically perverting cognitive penetration do in fact result in a causal circle, we should focus our attention on whether or not there is such a causal non-circularity principle in the first place.

Siegel’s “gossip circle” analogy illustrates in rather simple and intuitive terms why one should endorse such a principle:

We can compare this [example of cognitive penetration] to a gossip circle. In a gossip circle, Jill tells Jack that *p*, Jack believes her but quickly forgets that she’s the source of his belief, then shortly afterward Jack tells Jill that *p*. It seems silly for Jill to take Jack’s report that *p* as providing much if any support for *p*, beyond whatever evidence she already had. On the face of it, this looks like a feedback loop in which no new justification is introduced. Similarly, when beliefs are formed on the basis of cognitively penetrated experience, it is as if your belief that *p* told you to have an experience that *p*, and then your experience that *p* told you to believe that *p*.⁷

It is clear that the “telling” relation here is causal, not justificatory. A belief does not justify your experience insofar as your belief “tells” you to have a certain experience; it is strange or absurd to say that beliefs justify experience. Similarly, an experience does not justify your belief insofar as your experience “tells” you to have a certain belief; after all, you are not justified in having a certain belief merely insofar as another person tells you to have that belief.

Since this gossip circle analogy seems to be doing a lot of the legwork for Siegel’s objection to justification⁸, it is worth close scrutiny. I, for one, think the conclusion of the analogy itself is wrong; Jill does in fact obtain new justification for her belief that *p*. To see this, it helps to fill out some of the details. Suppose, for instance, that Jill tells Jack that Bill, a mutual acquaintance, is a jerk — an opinion which Jack immediately echoes back. Jill believes that Jack knows Bill intimately, and that Jack is generally trustworthy and expects him to give her his honest and well-educated opinion. Unbeknownst to Jill, however, and indeed as in the original analogy, unbeknownst to Jack himself, Jack is just trying to cozy up to her and is more than willing to throw Bill under the bus to this

⁵Pryor (2000), p. 540.

⁶Siegel (2011), p. 212.

⁷Ibid., p. 202.

⁸E.g., see *ibid.*, p. 214.

end. Since, from Jill's perspective, Jack is a reliable person attesting to Bill's jerkiness, Jill seems to have all the more reason to believe Bill is, in fact, a jerk. Obviously, if Jill knows that Jack is a toady, she should suspend any additional justification. But her knowing this would be like the epistemic subject's knowing that his experience is cognitively penetrated, and we should assume that defeaters like this are inaccessible.

Now, even if the conclusion Siegel draws from her gossip circle analogy is false, one could try to advance the objection in more general terms. After all, there might not be an adequate analogy between a gossip circle composed of persons and a causal circle within the mind. In order for this more general objection to succeed, however, it must show that whenever cognitive penetration results in a causal circle, the resulting cognitive penetration is epistemically perverting.

But this is false: there are possible cases in which cognitive penetration results in a causal circle without being epistemically perverting. To show this, it is first worth pointing out that not all cases of cognitive penetration, whether etiologically circular or not, are epistemically perverting. For example, an ornithologist's knowledge about bird species may allow her to perceive certain blue birds as members of their respective species — California scrub jays, Steller's jays, blue jays — whereas the birding novice may manage only to perceive them as blue birds and infer from his perceptual experience, after checking his birding manual, that they are different species of American jay. As a result, the ornithologist's cognitive penetration could be said to be epistemically enhancing, rather than epistemically perverting.

With this distinction in mind, we could go on to generate examples of epistemically enhancing cognitive penetration with the allegedly perverse causal circle. Consider the following: an ornithologist trekking through Big Sur believes, correctly, that there will be a California scrub jay somewhere near his hiking trail. While sitting on a tree stump, he spots a medium-sized flash of blue and, as a result of this belief, accurately perceives a California scrub jay emerging from the brush. And in turn, on the basis of this experience, he believes or reinforces his belief that there is, indeed, a scrub jay near his hiking trail. If, however, he did not have this belief leading up to the sighting, and the environment, distal stimulus, and locus of attention were identical, he would see only a blue bird, and he would be left in wonder about the species. And if the flash of blue had in fact been some other medium-size blue bird, he would be puzzled about the species in this case as well; that is, he would not misidentify the unknown bird as a California scrub jay (because he could not identify the flash of blue at all). Although the relation between the ornithologist's belief and his experience is no less circular than that of the over-eager birder, his belief nevertheless enriches his experience without distorting it; without any loss in reliability, it allows him to perceive more specific objects than he otherwise would — in this case, scrub jays instead of mere blue birds. The circularity does not, then, make the ornithologist's belief unjustified and is therefore not the *per se* source of the epistemological problem.

We now move onto the second of the critic's two objections, namely that, since i) epistemically perverting cognitive penetration constitutes or is a kind of epistemic vice, and ii) every belief resulting from an epistemic vice is unjustified, beliefs which result from epistemically perverting cognitive penetration are unjustified.⁹ I accept premise i but reject premise ii. Although it is difficult to definitively refute premise ii in the space

⁹Applying Aristotle's nominal definition of a virtue in general (*Ethica Nicomachea* II.6 1106a15-24), we may define epistemic virtue in particular as a characteristic which makes its possessor perform her epistemic functions well. We may, accordingly, define epistemic vice as a characteristic which would make its possessor perform her epistemic functions poorly.

allotted, here is at least an argument for rejecting it.

Suppose that a man has a moral vice — say, dishonesty. On account of his dishonesty, he lies to the Gestapo, denying the presence of any Jews in his house. The man's dishonesty does not exclude the possibility that he would lie to the Gestapo for a good reason, e.g., disagreement with Nazi anti-Semitic policy. Rather, his being dishonest suggests he would lie, largely irrespective of the reason or of the circumstances, good or bad, appropriate or inappropriate. So, even supposing that the man's action is justified in the strong sense that he does what he does for a good reason (rather than in the weak sense that there is a good reason for him to do what he does), the action nevertheless arises from a moral vice. This shows that, more generally, one's action may be morally justified even though it arises from a moral vice. Similarly, I suggest, one's belief may be epistemically justified even though it arises from an epistemic vice. In other words: though there is something morally wrong with the dishonest man, it is not wrong for him to do what he does in the given situation; analogously, though there is something epistemically wrong with a victim of cognitive penetration, it is not wrong for him to believe what he believes in the given situation. This analogy does not, to be sure, suffice as a refutation of the critic's argument, since it depends, first, on the strength of analogy between the epistemic and the moral, and second, on the distinction between virtue and justification in the moral case— both of which are subject to debate. But this analogy does itself require refutation if the critic's argument is to succeed.

Section II

Thus far, I have presented negative arguments showing that the critic's objections fail to refute dogmatism's verdict on beliefs resulting from cognitive penetration. In what follows, I present a positive argument showing that, given some plausible assumptions about the critic's own epistemological commitments, she should in fact endorse the dogmatist's verdict. First, I assume that, like many dogmatists, the critic is an internalist, or, more specifically, an access internalist: she believes justification supervenes or depends principally on what is accessible to, or within the "ken" of, the subject. The rationale for this assumption is that dogmatism is, as I said above, traditionally formulated as an internalist position, access internalism is the most common form of internalism, and it is most illuminating to suppose the least possible disagreement between the dogmatist and the critic. Second, I assume that the critic is not a skeptic about perceptual justification (distinct from the skeptic about perceptual knowledge) since her argument is toothless without the possibility of procuring justified empirical evidence to undermine the conclusions from cognitive penetration.

Given these assumptions, the following thought experiment poses trouble for the critic: suppose that there are two people — Dreamer and Wakeful. Dreamer's experience is indistinguishable in both phenomenal quality and content from that of Wakeful, except Dreamer is dreaming and Wakeful is not. It is an alleged virtue of access internalism that Dreamer has a lot of justified perceptual beliefs — just as many justified perceptual beliefs as Wakeful does. It is a virtue of dogmatism in particular that both Dreamer and Wakeful would likely have many justified perceptual beliefs.

Now suppose that there is, in addition, a third person, Penetrator, whose experience is similarly indistinguishable from those of Dreamer and Wakeful. The only difference is that Penetrator's experience is as cognitively penetrated as possible. How does the justificatory status of Penetrator's perceptual beliefs differ from that of Dreamer's perceptual beliefs (or "dream- experience" beliefs)? Both Dreamer's and Penetrator's

minds are or are among the immediate proximate causes of their experience to the exclusion of the external world, and neither have access to defeaters which would clue them into some discrepancy between their experience and the external world. One difference, however, between Dreamer and Penetrator is that it is not necessarily Dreamer's *beliefs* which cause his dreams to be as they are. But this difference seems to be irrelevant to the justificatory status of Dreamer's beliefs; one could very well suppose that, just like in the case of Penetrator, Dreamer's experience are caused by his beliefs without, in any obvious way, changing whatever our verdict may be on the justificatory status regarding his perceptual beliefs. So, the justificatory status of Dreamer's perceptual beliefs ought not differ from that of Penetrator's.

As a result of this justificatory equivalence, the critic is forced into the following dilemma: if, on the one hand, Dreamer's perceptual beliefs are justified, then so are Penetrator's perceptual beliefs. Then *a fortiori*, so are the perceptual beliefs of others who have much less severe cognitive penetration. So, the critic concedes defeat. If, on the other hand, Penetrator's perceptual beliefs are not justified, then Dreamer's perceptual beliefs are not either. Then the critic must either deny that Wakeful and Dreamer have the same justified perceptual beliefs in principle, or otherwise reveal herself as a skeptic, in which case Wakeful and Dreamer are both equally unjustified. And if she were to deny that Wakeful and Dreamer have the same justified perceptual beliefs in principle, she would be breaking rank as an access internalist.

If the critic stands her ground against the dogmatist, it is clear which horn of the dilemma she should choose. She could not feasibly be a skeptic about perceptual justification for the reason given above. But she could take the alleged implausibility of the dogmatist's position as grounds for either an internalism distinct from access internalism or some sort of externalism.

Since it is unclear to me what this alternative internalism would be, and how it would avoid the problems with externalism, I am content to set it aside. As for externalism, the critic does not seem to gain much by it; not only would she have to deny that Wakeful and Dreamer have the same justified perceptions in principle, but she would also, in a certain respect, fail to obtain the hoped-for advantage against dogmatism. For externalism would not salvage the critic's objection to the dogmatist's first claim *per se* — again, the claim that a subject's perceptual experience with some content, *p*, *prima facie* justifies her beliefs that *p*. In the beginning of this essay, I set aside the "*prima facie*" because the supposition of internalism made it irrelevant; the critic could always construct theoretical cases without relevant defeaters. But the critic could no longer do so if she were to suppose that both she and the dogmatist were externalists. If the meaning of "justification" in the dogmatist's first claim is to be understood in externalist terms, then it would only be consistent for "*prima facie*" to be understood in externalist terms as well. In that case, propositional defeaters, not just evidential defeaters, could be invoked to undermine justification. Since epistemically perverting cognitive penetration would constitute a propositional defeater, there would be no real dispute between the "externalist critic" and the "externalist dogmatist" concerning the matter; both would agree that the resulting beliefs are unjustified. Thus, the externalist critic's issue is not with the dogmatist's first claim *per se*, i.e. irrespective of further epistemological commitments, but with his first claim insofar as dogmatism is an internalism. Or, in other words, the externalist critic is less a critic of dogmatism than of internalism.

To sum up and conclude: Of the critic's two objections to the dogmatist's position on cognitive penetration — the objection from circularity and the objection from epistemic vice — the former fails and the second needs to be further developed. Moreover, the

tenability of the critic's opposition to the dogmatist depends on their respective positions on internalism and externalism. If the dogmatist and the critic both presuppose access internalism, the dogmatist wins the dispute. If both presuppose externalism, there simply is no dispute. If the dogmatist presupposes access internalism and the critic is open to externalism or to alternative forms of externalism, the dispute bottoms out at a disagreement either between access internalism and externalism or between access internalism and alternative forms of internalism.

To be clear, even though this much is to concede that the critic's opposition to dogmatism is rationally permissible, it is not to concede that cognitive penetration affords the critic any special argument against dogmatism. The justificatory equivalence between Dreamer and Penetrator as well as the invalidity of the critic's cognitive penetration-specific objections — especially the objection from circularity — suggest that dogmatism is no more threatened by cognitively penetrated experience than it is by particularly lifelike dreams.

References

- Aristotle. (1894). *Ethica Nicomachea* Edited by Ingram Bywater. Oxford: Clarendon, 1894.
- Banerjee, Chatterjee, and Sinha.(2012). "Is It Light or Dark? Recalling Moral Behavior Changes Perception of Brightness." In Lucy O'Brien, Matthew Soteriou (eds.), *Psychological Science*. 23(4): 407-409.
- Eberhardt, Goff, Purdie, and Davies. (2004). "Seeing Black: Race, Crime, and Visual Processing." *Journal of Personality and Social Psychology* 87(6):876–893.
- Firestone and Scholl. (2013). "'Top-Down' Effects Where None Should Be Found." *Psychological Science* 25(1):38-46.
- Levin and Banaji (2006). "Distortions in the Perceived Lightness of Faces: The Role of Race Categories." *Journal of Experimental Psychology: General* 135(4): 501-512.
- Pryor. (2000). "The Skeptic and the Dogmatist." *Nous* 25(1) 34(4): 517-549.
- Siegel. (2011). "Cognitive Penetrability and Perceptual Justification." *Nous* 46(2): 201-222.
- Stefanucci and Geuss (2009). "Big People, Little World: The Body Influences Size Perception." *Perception* 38(12): 1782-1795.
- Stokes and Payne (2010). "Mental Control and Visual Illusions: Errors of Action and Construal in Race-Based Weapon Misidentification." *The Science of Social Vision* 7): 295-305.

Affected Ignorance and Moral Justification

Patrick Magee
Harvard University

In her paper “Culture, Responsibility, and Affected Ignorance,” Michele M. Moody-Adams challenges the view that moral responsibility for an action can be lessened by the action’s widespread cultural acceptance. While cultural acceptance of a practice seems to at least mitigate participating in it, Moody-Adams argues that this view misunderstands the relationship between a culture and its members: since cultural practices only persist insofar as a culture’s members decide to transmit them, the maintenance of a practice is ultimately an intentional, rather than unintentional, choice of the members. In general, we have little reason to believe that ignorance of the wrongness of a practice like ancient Greek slavery is anything but *affected ignorance*, motivated by a desire to believe that the practice is morally permissible. But while Moody-Adams’s position is attractive, it relies on questionable epistemological foundations. In overlooking the influence of experts on belief formation, Moody-Adams’s position places an unreasonable standard on moral knowledge, rejecting modes of justification that in other domains are considered legitimate. Not only do such modes of justification allow for culturally mitigated responsibility, but as I hope to show, they do not in general force us to absolve perpetrators of historic injustices. In the following, I will first outline Moody-Adams’s position regarding culture and responsibility. I will then argue that Moody-Adams’s position is epistemically implausible, placing unreasonable demands on moral justification. Finally, I will show that in cases of historic injustice, this epistemological approach allows us to retain our intuitions of moral responsibility, albeit redirecting its focus in some cases.

In her paper, Moody-Adams is concerned to argue against two versions of which she calls the “inability thesis” regarding culture and responsibility. According to the first, which Moody-Adams attributes to Alan Donagan and Michael Slote, an upbringing in a morally backward culture can severely curtail one’s ability to know that a practice is morally impermissible. Following this view, an ancient Greek slaveholder or member of Hitler *Jugend* may simply be “unable to know that certain actions are wrong”¹; having imbibed a culture in which slaveowning or violence against noncombatants was considered licit, such a person would have little basis on which to conclude that the practices were in any way problematic. The second version, which Moody-Adams attributes to Susan Wolf, makes an even stronger claim: given certain circumstances, it may be *inevitable* that individuals come to hold “values and beliefs... that we now condemn.”² Whether the strong version of Wolf or the weaker version of Donagan and Slote, the inability thesis makes a rather simple claim with regard to moral action: at least in some cases, at least to some extent, culturally induced ignorance can mitigate an individual’s responsibility for a given action.

To Moody-Adams, however, a culture’s mere condoning of a practice is insufficient to mitigate an adult agent’s guilt for engaging in it. The reason, argues Moody-Adams, is that such activity on behalf of the culture does not “render [the agent] unable to question the morality of their culture’s practices.”³ When we consider historical agents who failed to question the morality of some widespread injustice, we cannot directly

¹Moody-Adams (1994), p. 293.

²Ibid., p. 293.

³Ibid., p. 294.

conclude that they *could not* have done so; rather, we can only conclude that they *did not* do so, and only infer the former if the circumstances suggest it. And according to Moody-Adams, when we examine the social circumstances in which these historical agents are inevitably embedded, “it will be difficult, if not impossible, to make a rationally compelling case for the inability thesis.”⁴ For generally, we will find that these social circumstances conspire to *support* the immoral practice to the benefit of a particular group.⁵ In the case of ancient Greece, for example, an entire cultural system was directed towards the perpetuation of slavery, an institution that clearly benefitted an elite group at the expense of the enslaved. Rather than *inculpable* ignorance, the ignorance of the beneficiary group in such examples is almost certainly *affected*, comprising their strategic refusal to consider that the immoral practice is in any way wrong. In the case of Greek slavery, Moody-Adams concludes that “[t]he belief that slavery was justified was insufficiently examined by those . . . who wanted to perpetuate the culture of slavery.”⁶ Indeed, to Moody-Adams, there is “no convincing evidence”⁷ that the slaveholders’ actions be attributable to anything else.

A natural objection to this line of argument is that cultural trends are rarely formed in this manner. In particular, the majority of Greek slaveowners did not *devise* the cultural norms and practices that undergirded the institution of slaveholding; rather, they inherited them from prior generations. Hence, it seems unfair to use these structures as evidence that the slaveowners were in general ill-motivated. But as Moody-Adams argues, this view misinterprets the relationship between a culture and its members. For a culture “persist[s] only because individual persons capable of responsible action persist.”⁸ Even though individuals inherit a substantial portion of their beliefs and practices from the previous generation, they retain the ability to “modify[], reshap[e], and . . . radically revis[e]” cultural conventions through individual actions⁹; when they fail to eliminate some problematic practice, it is not due to some vague “cultural limitations,” but rather to their “failings as human beings who perpetuate cultures.”¹⁰ Owing to such considerations, Moody-Adams concludes that the claim that culture can impair moral knowledge “cannot withstand rational scrutiny.”¹¹

Now, as an analysis of the general relationship between culture and individual action, Moody-Adams’s account has much to recommend. The notion that many instances of “cultural blindness” are in fact affected ignorance is highly plausible, and the observation that individuals ultimately choose to perpetuate certain aspects of culture is surely correct. Indeed, as an argument against the “strong” version of the inability thesis — that cultural circumstances can render it “inevitable” that individuals hold the reprehensible beliefs that they do — Moody-Adams’s argument seems sound, demonstrating that culture’s influence on an individual’s beliefs and practices is far from monolithic. But to refute the *general*¹² inability thesis, Moody-Adams has a far larger task. She must show not only that culture is *usually* not an exculpatory factor with regard to responsibility,

⁴Ibid., p. 295.

⁵I omit Moody-Adams’s lengthy discussion of the “internal perspective” and its influence on these situations; for the purposes of this essay, I believe this fairly represents Moody-Adams’s argument.

⁶Moody-Adams, p. 296

⁷Ibid., p. 296

⁸Ibid., p. 293.

⁹Ibid., p. 306

¹⁰Ibid., p. 298.

¹¹Ibid., p. 298 and p. 304.

¹²By “general inability thesis,” I mean the claim that “the link between culture and agency [can] undermine the standard attributions of responsibility for action” (Ibid., p. 292), which I see as the

but that it *never* is — that the view that cultural limitations “at least sometimes”¹³ mitigate personal wrongdoing is ultimately incorrect. She must give us compelling reason to believe that apparent instances of culturally mitigated wrongdoing “cannot withstand rational scrutiny”¹³ and should not motivate us to doubt her highly general claim. And as I will argue, Moody-Adams’s paper does no such thing. While it may lead us to believe that many apparent instances of culturally influenced wrongdoing are in fact culpable, epistemological considerations provide strong reason to believe that at least sometimes, culture can severely limit an agent’s ability to act morally.

To see why this is the case, we must first examine the ways in which cultural norms can influence one’s moral beliefs. Throughout the essay, Moody-Adams seems to take for granted that culturally dominant false moral beliefs¹⁵ are objectively unjustified. In the case of ancient Greek slavery, for example, she concludes that slaveholders’ “belief that slavery was justified was insufficiently examined”¹⁶, implying that if the slaveholders were to *sufficiently* examine the belief, they would realize its falsehood. Similarly, she argues that alleged instances of cultural blindness are often merely “an unwillingness to consider that some practice might be wrong”¹⁷, implying that a willingness to consider the matter would inevitably lead to a correct belief. Yet these two possibilities — true belief or unjustified false belief — are by no means exhaustive. In particular, they omit the possibility of having a justified *false* belief. That is, it seems entirely possible that an agent consider the possibility that a given practice might be wrong, exercise sufficient care in weighing contrary opinions, and come to the conclusion that the practice is not wrong, owing entirely to cultural factors. And in that case, it is plausible that the agent’s responsibility for participating in the practice *would* be undermined, since she exercised every proper moral precaution before engaging in it.¹⁸

Given that we as modern people find it incredible that someone could carefully examine the practice of slavery and *not* find it problematic, I will first examine a case I take to be epistemically analogous. Imagine a society in which, owing to some historic accident, everyone believes that the Pythagorean theorem is false. Rather than think that $a^2 + b^2 = c^2$, the society has the collective belief that $a^2 + b^2 = c^{2.000001}$, a belief that, being close enough to the actual truth, has resulted in no discernible effect on technological innovation or scientific progress. The belief is instilled in schoolchildren from a young age, and the society’s best mathematicians even have “proofs” that the result is correct. Now, imagine that a particularly smart student is learning this result in school. She is perturbed by the arbitrariness of the value 2.000001, which seems to violate her deepest mathematical intuitions, and decides to undertake a thorough examination of the result. As a result of her studies, she discovers a geometric proof that $a^2 + b^2$ is in fact equal to

essence of both the “strong” and “weak” versions of the thesis.

¹³Ibid., p. 292

¹³Ibid., p. 298

¹⁵By “culturally dominant false moral beliefs” I refer to the sorts of false moral beliefs — like male chauvinism or a belief in the justness of slavery — that correspond to the predominant cultural attitudes of the society. The phrase is somewhat awkward, but others (culturally influenced falsehoods, inherited false beliefs) either seem less clear or imply a causal connection between the culture and the belief.

¹⁶Ibid., p. 296.

¹⁷Ibid., p. 294.

¹⁸Of course, one could argue that even in this case, the agent’s responsibility is not undermined. But even if one finds this objection plausible (and I don’t), the need for such extensive discussion surely gives lie to Moody-Adams’s assertion that claims of culture undermining responsibility “cannot withstand rational scrutiny” (Ibid., p. 298 and p. 304).

c^2 . Having shown her work to mathematicians, she has been consistently told her result is wrong, but she is unable to shake the feeling that her proof is nevertheless correct. Is she justified in concluding that $a^2 + b^2 = c^2$? In one sense, the result is entirely justified: she arrived at it by proper methods of reasoning, and she has exercised every proper precaution in ensuring the rightness of her result. But on the other hand, consider what she is committed to by assenting to the belief that $a^2 + b^2 = c^2$. She must accept that the society's greatest mathematicians consistently fell into an error that only she, a mere student, was able to uncover. She must accept that an *entire society* was consistently deceived with respect to one of the most fundamental mathematical results. And she must accept these results solely on the basis of her intuition and mathematical abilities. What is more probable: that these seemingly incredible results obtain, or that the student simply made a mistake she is unable to recognize? While I make no claim that it is *unjustified* for the student to believe that $a^2 + b^2 = c^2$, it seems more than epistemically reasonable for the student to conclude, despite her mathematical misgivings, that $a^2 + b^2 = c^{2.000001}$. Indeed, one could argue that it would be a feat of epistemic *arrogance* for the student to trust her own intuitions and arguments over those of experts, especially when the apparent consequences of her belief are so unbelievable. Certainly, it seems unreasonable to expect that *every* responsible epistemic agent would come to the correct belief in this scenario.

But if such an agent is justified in concluding that $a^2 + b^2 = c^{2.000001}$, then it seems equally plausible, at least in some scenarios, that an agent be justified in believing an entirely false *moral* belief. For unless we apply substantially different justificatory standards with regard to moral beliefs, there seems to be no reason why an agent could not be similarly deceived with regard to some ethical proposition. Parallel to the previous example, consider a scenario in which an entire society has come to hold that slavery is morally licit, a conclusion steadfastly supported by the society's "moral experts" (say, those most respected for understanding and upholding the moral law). If an agent in such a society were to entertain the idea that slavery was in fact immoral, and to come to this idea for all the right reasons, it is far from clear that, under all circumstances, the agent would be forced to conclude that slavery was in fact *immoral*. For once again, such a belief would force the agent to accept some seemingly incredible conclusions — indeed, arguably *more* radical conclusions than in the mathematical case. The agent would have to conclude not merely that society has systematically gone wrong with respect to some abstract proposition, but that it had done so with respect to a moral atrocity, universally condoning one of the worst practices in human history. The agent would be forced to admit that those she most respected as moral agents — moral experts who unfailingly understood and followed the moral law, even where our agent failed to — had nevertheless blundered with regard to one of the most obvious moral truths. And as in the mathematical case, while it is not necessarily *unjustified* for the agent to decide that slavery is wrong, it is surely reasonable for the agent to simply conclude that she has made an error. Epistemically, the two cases seem virtually identical; we have little grounds for concluding that the agent is justified in one and unjustified in the other.

But if such a scenario is possible, then Moody-Adams's case is severely undermined. For *prima facie*, the scenario seems to involve an agent inculpably coming to hold false moral beliefs owing to the influence of culture: despite exercising all due caution, consideration, and moral judgment, the agent is led by her cultural setting to hold a false moral belief. In such a case, Moody-Adams's allegation of "affected ignorance" would simply not apply, since the agent exercised all due care in examining her belief. Similarly, the agent's belief in slavery cannot be chalked up to her "failings as [a] human

bein[g] who perpetuate[s] cultures”¹⁹, since her choice to perpetuate a belief in slavery was motivated by inculpable ignorance of its wrongness. Were the agent to act this false belief, it would seem that her responsibility would be at least partially mitigated, given that her ignorance of slavery’s wrongness is not due to some personal failing. At the very least, such a case gives lie to Moody-Adams’s claim that allegations of culturally influenced moral ignorance “cannot withstand rational scrutiny.”²⁰

Now, one might object to this line of argument by pointing out the scenario’s implausibility. It seems incredible that an entire society could come to accept slavery without any voices of dissent, and it seems similarly implausible that our agent could deliberate over the morality of slavery without a trace of selfish motivation. Surely the agent in the scenario is not *really* undertaking a perfectly rational analysis of slavery; in practice, personal considerations would almost certainly influence her decision process. Such an objection, however, overlooks the high bar that Moody-Adams has set in rejecting the inability thesis. For Moody-Adams’s argument to succeed, it cannot merely be the case that culturally mitigated responsibility is rare or implausible — as I have stated, it must be the case that it could *never* occur. And even if it is *unlikely* that scenarios like the one described ever occur, it certainly seems possible; there is nothing in the scenario that seems to *rule out* its ever occurring, even if one finds the scenario fantastical. And given that such possibility is all that is needed to refute Moody-Adams’s argument, an appeal to implausibility has little force as an objection.

Similarly, one could object that the situation I’ve described illicitly relies on the idea of “moral experts.” Such an objection could take two forms: one could challenge the very *idea* of moral experts, arguing that it is wrong to defer to others’ judgment regarding moral matters, or one could argue that the moral experts deferred to in the scenario simply aren’t moral experts — given that these “experts” endorse an atrocity like slavery, it’s unreasonable for our agent to defer to their moral judgment. But these two objections have the same response. Granting that it is *always* wrong for an agent to defer to another’s moral judgment — a proposition I would dispute — it is implausible that the sin of deferral rises anywhere near the sin of slavery. Indeed, one could argue that the fault of the agent deferring to a moral expert is not moral, but *prudential*. The agent is in every way attempting to follow the moral law, yet she is factually confused as to what this duty entails.²¹ She believes that it is her moral duty to defer to some individual, yet for some reason or another, this duty is specious — either because the “moral expert” is not an expert at all, or because it is *always* wrong to defer in this manner. And even if one deems such moral confusion blameworthy, it is simply implausible to assert that the level of blame rises anywhere near the level of slavery. In such a case, the guilt of the agent would be substantially mitigated, in line with the inability thesis.

Finally, one could object that the view I’ve defended — that at least sometimes, cultural circumstances can mitigate one’s guilt for a particular action — violates our deepest intuitions about historical injustice. When we examine cultures where slavery was common, we feel strongly that *someone* must be responsible for the atrocity; it seems unreasonable, even immoral, to argue that the apparent guilty parties were in

¹⁹Ibid., p. 298

²⁰Ibid., p. 298 and p. 304.

²¹One could object that this illicitly relies on a sort of Kantian conception of moral duty, according to which care for what is moral *de dicto*, rather than what is moral *de re* morally (c.f. Johnson King, “Accidentally Doing the Right Thing” (2020)), ultimately determines whether an agent is acting morally. But again, it is implausible that this error in judgment — being *confused* about the moral thing to do in a situation — rises anywhere near to the level of slavery.

fact innocent. Yet this is far from what I have done here. First, I have nowhere claimed that instances of culturally mitigated responsibility are in any way *common*. Indeed, my view is consistent with the notion that the vast majority of slaveholders *are* responsible for their actions — all I have argued is that it is *possible* for an agent in their position to experience mitigated responsibility, assuming they were to thoroughly examine the practice without any trace of selfish motivation. And even in cases where cultural factors lead an agent to commit some wrongdoing, it is far from clear that *no one* is to blame. Indeed, in many such cases, this simply defers the blame to those *responsible* for the cultural factors: in the case of slavery, those who knew the practice was wrong, yet pretended not to for selfish reasons.²² By making possible an environment in which someone could be inculpably deceived as to the wrongness of slavery, these people carry a high degree of guilt, justifying our intuitions of responsibility even in cases where the initial party is inculpably ignorant. In general, then, the position that cultural situations at least *sometimes* engender inculpable ignorance of moral duties does not force us to dispense with the notion of historic guilt.

While Moody-Adams's position is initially attractive, epistemic considerations suggest that at least in some situations, cultural circumstances may substantially mitigate an individual's culpability for wrongdoing. Owing to the testimony of apparent "moral experts," an agent in these circumstances may be entirely justified in holding a false moral belief, remaining inculpably ignorant despite her best efforts to act morally. But while such considerations tend to support the inability thesis against Moody-Adams's critique, the *extent* to which one judges an agent's guilt to be mitigated may depend on what account of moral worth one endorses. If one endorses what Zoë Johnson King calls the Kantian view of moral worth, one may judge the agent's guilt to be almost entirely mitigated, since the agent's actions stemmed from a sincere belief that what she was doing was moral.²³ In contrast, one less sympathetic to a Kantian view may judge the agent's guilt to be only slightly mitigated, since they will not be as impressed by the agent's desire to do what is right in the *de dicto* sense.²⁴ Such debates, however, bear only tangential relation to my argument. So long as it is possible that cultural conditions even *partially* impair an agent's ability to do the right thing, the inability thesis holds, and Moody-Adams's argument fails.

References

Moody-Adams. (1994) "Culture, Responsibility, and Affected Ignorance." *Ethics* 104: 291-309.

²²Of course, it is not *necessary* that such people exist, but it is highly plausible that they would, at least in the vast majority of cases.

²³According to this view, an action has moral worth if and only if the agent is motivated by the very fact that it is right. I do not claim that a Kantian perspective would force someone to conclude that the agent was blameless, only that a Kantian account would be more congenial to this conclusion than the "new view." (c.f. Johnson King, "Accidentally Doing the Right Thing" (2020))

²⁴For a discussion of what is right *de dicto* versus what is right *de re*, again see Johnson King, "Accidentally Doing the Right Thing" (2020).

An Interview with Peter Railton

Gregory S. Kavka Distinguished University Professor; John Stephenson
Perrin Professor; Arthur F. Thurnau Professor
The University of Michigan

The staff of Meteorite conduct an interview with a contemporary philosopher for each edition of the journal. For this edition, we were grateful to be able to interview Peter Railton of the University of Michigan.

Meteorite (asked by Tristan Sirls):

*In a recent book, *Metaphysics and Cognitive Science*¹, you mention that there has been a shift in meta-ethical debate over the previous decades. You point out that moral realism has transitioned from what you would describe as a "fringe position," to a common and favored position. Regarding your own position as a moral realist during this time, how have your own views changed or remained the same?*

Peter Railton:

The first thing to say is that, people differ quite a bit in what they understand by moral realism. So, at the time I first started writing on it, it was thought, for example, that in order to be a moral realist, one would have to be a cognitivist of some sort — that is to say, think that moral judgments have ordinary truth value. You'd also have to think that at least some of them were true. And over time, views about moral judgments that include non-cognitivist positions — positions that moral judgments are doing something other than stating a proposition that could be true or false — gained a lot of currency, with great thanks to my former colleague Allan Gibbard, who developed a very powerful view of that kind. And that initially seemed to be a challenge to moral realism, because it seemed like it was a rejection of the kind of ideas about cognitivism and truth. Instead, someone like Allan now views himself as a quasi-realist. That is to say, he thinks that his account of moral language should make it be the case that you can perfectly well talk about moral judgments as true or false, that they can be objects of knowledge, and he certainly thinks that some of them are true. And so in that sense, he's a moral realist, even though the term means something rather different from what it meant back when I was first writing.

Similarly, when I was writing, and actually in my own view today, the approach I take to moral discourse is naturalistic. And I'd say a large number of contemporary moral realists are non-naturalistic in their approach. So moral realism can mean a lot of different things. The way I would say my view has changed is that, first of all, a lot has gotten clearer about what might or might not be important to moral realism, thanks to these debates. At that time, I was having to kind of invent the elements of psychology and learning, and social learning, that were part of my account. By invent I mean I was able to cite some evidence for those things but not very much. And since then, the development of psychology and neuroscience has been such that many of the key concepts I need have gotten systematic treatment. And in a way, some of the mechanisms I hoped

¹Goldman and McLaughlin (2019).

would exist, do seem to exist. And so, for example, the idea that our evaluative judgments have a cognitive component, that they are subject to learning, that learning takes place by a kind of feedback process, that knowledge can be socially distributed — those things have gotten much greater currency. And so, I guess I'd say I've been quite encouraged by the developments that have taken place since I first wrote.

Meteorite (asked by Tristan Sirls):

So, building off the more scientific role in your research, one thing that I noticed is that a lot of debate in meta-ethics right now seems to be shaped by research in cognitive science. What are your thoughts on this role of science kind of acting as a player in allowing these meta-ethical views to develop more substantially?

Peter Railton:

On the whole, I think it's a good development, because I think our philosophical views should be based upon a realistic picture of psychology. And also because I think philosophers can have the impression that they have considered all conceptual possibilities, and they therefore reach certain kinds of problems or dilemmas. Whereas, sometimes psychologists develop notions of processes or concepts that weren't imagined by philosophers. So, in that sense, I think not only can we try to keep ourselves realistic in our psychology, but we can also get a better picture of some of the range of possibilities through psychology. An example I'd give is that in meta-ethics there's been a big debate over the years between people who have a view about moral judgments that's called motivational internalism — namely that anyone who makes a moral judgment, if they do so sincerely, will be to some extent motivated to comply with it. That's a internalist position of the motivational kind. And it's been thought that, by philosophers, that psychological states divide roughly into those which have what's called a mind-to-world direction of fit, like beliefs, where they can be correct or incorrect. And those which have a world-to-mind direction of fit, like desires, which roughly put us in a position to try to change the world to fit what we want, and the thought is you can't have a state with both directions of fit. The two are mutually exclusive. And if you want to be a cognitivist about moral judgment, you then think they can be true or false — that's mind-to-world direction of fit. If you want to be an internalist, a motivational internalist, you think that they also motivate action, and that's world-to-mind direction of fit, and the thought is that that leaves you with this paradox. But if it turns out that there are states that have both directions of fit, and I think affective states are like that, then you don't have the same paradox. And I think it was easy for philosophers to think, well here's a dichotomy. It's a clear dichotomy. Any psychological state is going to fall on one side or the other of it, without recognizing that there can be hybrid and complex states which have both features.

Meteorite (asked by Tristan Sirls):

You claim that a naturalistic tracking moral realist, or a naturalistic moral realist who tracks moral facts, faces the large issue of accounting for what you say are non-hypothetical thought and action guiding reasons for our moral considerations. So, could you explain a bit on what exactly a tracking moral realist is, what are we tracking, and then maybe reasons for why someone with this view would face this problem, and how science might help solve this?

Peter Railton:

One way to think about tracking is to think of it, not just in the case of moral judgments, but in general. You might say, one way to think about how we would justify our beliefs, say, would be if we thought that we had capacities: perceptual capacities, inferential capacities, and so on, that tracked various features of the realm of facts. And, that would be a way of conferring justification that might not require that we have a kind of internal access to the reasons for which we believe. So a lot of people think, for example, that perceptual beliefs are justified. And they're prima facie justified — default justified. And, you might say, they're default justified even in the minds of those people who do not have an argument for why they'd be justified, and they're justified in virtue of the fact that they track certain features of the environment. So, that's a kind of tracking picture. And you might have the same idea of ethics. So, someone might say, how would you ever justify a moral judgment? What would count as justification? Do you mean you can show that it's self-evident. And you might say well, I can't show that self-evident, because it doesn't seem to be contradictory to deny it. Can you show that all rational people will necessarily accept it, no matter what else they think or feel? I'd say no, that's too strong, moral judgments are synthetic, they're not analytic. And so, you can't just demonstrate them. And so, some say well okay, if you can't demonstrate them rationally, and they're not self-evident, what could justify them? Well, if we can show that moral judgments track features of the world, in virtue of which the moral judgments hold, then that would be a form of justification, or a way that knowledge could be gained. And so the thought is, what would that look like? What's it like to track value? And, when I first started writing on this, I had hoped to think in something like the following way.

So, here's the way the Bayesian thinks about belief: he thinks we may not be able to give a justification for the beliefs we start with, because any justification we give would probably be question begging — it would assume something we already believe. And we certainly don't want to try to justify in terms of something we don't believe. So, I think we have a starting point in our beliefs which we can't give a *de novo* justification for. On the other hand, we have a notion of what it is to be rational in belief, which is a matter of how those beliefs respond to new evidence — how do we revise them in light of new evidence? And what rationality is really about is how you revise your beliefs in response to evidence. And, there are even some kind of nice results about various systems for updating your beliefs — that they, for example, tend to lead to convergence, even from different starting points. So that's a picture of what it is to have rationally justified beliefs, where you don't give an explanation of what it is to justify a belief from scratch, from nothing.

I thought maybe you could do the same thing for value. People start with certain values, and they have certain experiences, and given their values, and given those experiences, there might be rational or irrational ways for their values to evolve and respond to the experience. And that evolution would be like the case of belief, it would be a kind of response to the evidence of experience, and people could be more or less rational in that response. And so, in the Bayesian case, people could be tracking features of the natural frequencies in the environment, and in the ethical case or moral case, people could be tracking features of what's good or bad, or what's a virtue or vice. But you don't explain that by justifying from scratch, you explain that in terms of how you respond to new evidence. And you are revising in such a way that, if it were possible to track features of the external world, this would be a way of doing it. And so, that requires a theory of how we represent value mentally and how we change. And that's where I found, asking about science, that there's an interesting feature of the psychologist's picture of the mammalian brain, as opposed to the philosopher's, that perhaps the central activity that it has is forming evaluative representations. So, you are learning about reward, where reward is

present — how much reward there is. You're learning about risk and uncertainty. And those factors are regulating your behavior. But these are evaluations, and we now have a much better picture of how it is that the brain represents these evaluations, and how they get updated. And it does look a lot like a sort of rational updating process, so that doesn't show the view is right, but it shows that the view isn't ridiculous. So that's an example of how science could be relevant. There are lots of other examples: people use certain types of scientific evidence to try to argue for or against certain normative views and ethics. So, you may have followed the debate about trolley problems. And one camp has argued that the psychology of moral judgment is essentially a psychology of intuition and feeling. And they use that as an argument for defending certain meta-ethical positions. There are others who have looked at trolley problems and said: when people answer in some of these situations, they're making use of more of their deliberative or cognitive capacities, and in these other situations they are making more use of their affective or emotional capacities. And we might use that distinction to try to see or say which are more credible or plausible. So that's another use. And I find that interesting — I don't think any of those arguments are very conclusive, but they certainly have given us a new way to try to tackle some of these questions, and they don't replace the normative theory, but they are a useful kind of evidence. It's the sort of thing one would want to know in assessing the normative theories.

Meteorite (asked by Tristan Sirls):

We've talked a lot about the contents of different philosophical theories. So, moving away from that a little bit, you mentioned in another interview several years ago that you think philosophy, and those who write it and teach it, can make a difference by sharing the knowledge that comes from it. And that over time, you've tried to figure out different ways of accomplishing this. So, do you still hold this belief that sharing philosophical knowledge will create good change and, more importantly, has anything changed in your method of trying to share this type of knowledge?

Peter Railton:

Well one thing that's changed is that, as I've worked more with psychologists, I've been writing for a different audience, an audience that includes psychologists. Lately I've begun to work on artificial intelligence and ethics, and that is another audience. So one way for philosophy to make change is if it can be a useful partner of other disciplines in developing ideas, and so I worked together with some psychologists, and people in neuroscience, to write a book which is a sort of popular book on mental architecture, which has had some impact. And I've also tried in various ways to, by working on ethical problems, develop arguments that might enable people to find some justification for certain ethical positions, and I'd say that has not had a big influence on the world, except perhaps through teaching. But my thought is that some of the problems that we face arise from conceptual confusions, or a lack of confidence and understanding ourselves, or our beliefs. And if I can make some contribution to that, that might be a positive thing. So for example, I've participated in some applied ethics. I early on did some work on cost-benefit analysis that was taken up a bit in the policy literature. I've been trying to work lately on questions about well-being and trying to write in a vein that could be useful for people in policy areas, as well as philosophy. But I can't claim to have had great success.

Meteorite (asked by Tristan Sirls):

So building off that, you mentioned how you've kind of applied philosophy to these other

subjects— psychology, artificial intelligence. So there is this idea that undergraduate humanities have been in decline in the previous decade, where the amount of undergraduate students pursuing majors in subjects like philosophy have been dropping by pretty substantial amounts—one stat I found was about 9500 students in 2014 to about 7500 in 2017.²

Peter Railton:

This is in the country?

Meteorite (asked by Tristan Sirls):

I think it was country wide. I guess a question I have for you is, if it seems like the humanities are declining and philosophy might be losing students, and other students are going to more STEM related fields — computer science, especially, is a huge one. Would a way forward for philosophy be by integrating itself with these fields in a more definite, and applicable way?

Peter Railton:

Well, I'd like to think that that's part of what we're doing here, you know, we have this minor within the philosophy major that concerns cognitive science. We have strong participation in, and close affiliation with the program in cognitive science. My experience has been that, working with people in cognitive science and in psychology, is that philosophy has something to contribute to those conversations. And so they appreciate what philosophy can do. And so I guess I think that philosophy is actually very well positioned to establish its usefulness and relevance in the sphere of cognitive science, and so in my own philosophy classes I get a fair number of people who are going to be in cognitive science, but who recognize that, for that reason, they need to have a better understanding of the mind and its operation. There are value questions they'll face so they want to take ethics. And so I would agree that that's an area where I think philosophy can very profitably work together with other disciplines, make a distinctive contribution, be improved by that interaction, and my experience has been that the people have dealt with in cognitive science, and in psychology appreciate that. I'm not sure what our numbers look like, but this department has done reasonably well in terms of enrollments, partly because of that. The other thing that I find quite widespread right now among young people is concern about social justice and political questions, and that's another area where philosophy's been very active, with its role in philosophy, politics, and economics. And so I think philosophy is going to have to make connections, and I think it can do so, and I don't think it's going to be detrimental to philosophy, I think it can help the philosophy.

Meteorite (asked by Tristan Sirls):

I think, as one last question. Do you have any advice for undergraduates or graduate students wanting to pursue a career either in philosophy, or incorporate philosophy into their own career somewhere else?

Peter Railton:

So, thinking first about undergraduates, one piece of advice I can give is that philosophy is good training for a lot of pre-professional work: medicine, law, business. The kinds of skills that one could acquire in philosophy, but also the capacity to deal with systematic

²Weinberg (2018).

thought concepts, the ability to think about normative questions, the sensitivity to the quality of reasoning— those are all things that are benefits in a lot of professional training. So, the first thing I would urge is that, if you're thinking about law, or medicine, or business, or maybe some of the sciences, where philosophy's work in methodology is now much more relevant to what's going on in the sciences than it ever was. I'd say there's no reason why philosophy couldn't be part of that — that doesn't mean you have to be a philosophy major, but being a philosophy minor would probably be a good choice. And it would enrich your life in other ways too. I like the saying that, when you're in college, you're furnishing the mind that you'll live in for the rest of your life, and philosophy helps you do that well or better. As far as more specialized study in philosophy, like graduate study in philosophy, there are, right now, a large number of MA programs in philosophy. Taking an MA in philosophy as more advanced training, but then perhaps moving on to another area, that seems like a potentially valuable move for the same reasons. I think what people have been saying recently is that because of the declining representation of the humanities in universities, if you want to go on to a career in philosophy, you'd have to recognize that you're taking quite a risk. On the other hand, if you really love it, and you've got aptitude for it, it's a risk that's well worth taking, and it doesn't disqualify you for many other things, and may qualify you for other things. It'd be great if you could tell people, don't worry about the future, just follow what you love, and I know that's impractical for a lot of people, but to weave philosophy into your training and your work is quite practical. And so, you probably shouldn't think about grad school in philosophy unless you love it so much that you'd like to have spent those six years studying philosophy, even if you don't become a philosopher. And I think a lot of people will be in that situation — that it will have been six years spent doing what they love — under difficult circumstances, and equipping them for lots of options beyond philosophy.

***Meteorite* (asked by Tristan Sirls):**

Well, thank you so much for being here with me today.

References

- Goldman and McLaughlin. (2019) *Metaphysics and Cognitive Science*. New York, NY: Oxford University Press.
- Weinberg. (November 30, 2018) "Has the Sharp Decline in Philosophy Majors Hit Bottom? (Guest Post by Eric Schwitzgebel)." Daily Nous. <http://dailynous.com/2018/11/30/sharp-decline-philosophy-majors-hit-bottom-guest-post-eric-schwitzgebel/>.

About *Meteorite*

Meteorite is a student-run publication at the University of Michigan dedicated to recognizing valuable contributions in philosophy by undergraduates all over the world. The journal was founded by students at the University of Michigan in 1998, however, publication of the journal has been spotty and punctuated by large periods of inactivity.

In 2018, the journal was revived (for the fourth time!) and a new editorial staff was formed, and, since then, we've published three consecutive yearly issues. As well, this year, we had planned to host our first annual conference, with those selected for publication in this issue to present their papers at the University of Michigan. Unfortunately, this was not possible due to COVID-19, but I have the utmost confidence in our current team that we have many years of steady publication and conferences on the horizon.

This year, we have read and discussed hundreds of essays from undergraduates, selecting the most enlightening and clear of the bunch for publication. We are pleased to present you this edition, and we appreciate everyone who played a contributory role.

If you would like to submit a paper, please contact us using the information below below.

We prefer that submissions be formatted according to the Chicago Manual of Style guidelines. Papers should be submitted in electronic form only.

Visit our website for submission information:
<https://meteorite.philosophy.lsa.umich.edu/>

Please email us with any inquiries:
meteorite@umich.edu

Or write to:
Meteorite c/o
Department of Philosophy
2215 Angell Hall
435 S. State St.
Ann Arbor, MI 48109
USA