

7-2015

The Problem of Epistemically Irrelevant Causal Factors

Derek L. McAllister

University of Arkansas, Fayetteville

Follow this and additional works at: <http://scholarworks.uark.edu/etd>

 Part of the [Epistemology Commons](#), and the [Religious Thought, Theology and Philosophy of Religion Commons](#)

Recommended Citation

McAllister, Derek L., "The Problem of Epistemically Irrelevant Causal Factors" (2015). *Theses and Dissertations*. 1299.
<http://scholarworks.uark.edu/etd/1299>

This Thesis is brought to you for free and open access by ScholarWorks@UARK. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of ScholarWorks@UARK. For more information, please contact scholar@uark.edu, ccmiddle@uark.edu.

The Problem of Epistemically Irrelevant Causal Factors

A thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Arts in Philosophy

by

Derek McAllister
University of Central Arkansas
Bachelor of Arts in Philosophy, 2010
Biola University
Master of Arts in Philosophy of Religion and Ethics, 2013

July 2015
University of Arkansas

This thesis is approved for recommendation to the Graduate Council.

Dr. Thomas D. Senior
Thesis Director

Dr. Jack Lyons
Committee Member

Dr. Eric Funkhouser
Committee Member

Abstract

The problem of epistemically irrelevant causal factors is an epistemological phenomenon that occurs when a person becomes aware of some non-epistemic, causal factor that threatens to adversely influence her present belief, yet this factor is irrelevant to her deliberation concerning that belief. While the problem itself is apparently relatively widespread, very few have given it a detailed analysis. This thesis is one attempt to improve that. The first part, and the bulk, of this thesis is an analysis and explanation of what exactly the problem is and how it differs from nearby, related epistemological phenomena. The second part is my attempt at providing a meaningful solution to the problem such that one can remain justified in one's beliefs despite becoming aware of an epistemically irrelevant causal factor.

Acknowledgements

Special thanks to Tom Senior, who provided tremendous, unflinching help during this entire process, even during his research leave at Notre Dame. Thanks also to Jack Lyons and Eric Funkhouser, who took out time during their summer to read and interact with this thesis, and to Tomás Bogardus for providing helpful answers to my questions. I'm also grateful to Doug Geivett, who gave me the green light to research this nascent topic in both his epistemology and philosophy of religion courses in Spring 2012.

Many thanks are also due to those along the way who have lent an ear and their time in department hallways and at conferences: Kathryn Zawisza Bailey, David Barrett, Zach Biondi, Wesley Chambers, Justin Clardy, Dawna Hendricks, Bobby Howard, David Justice, Eric Panicco, Jonathan Reibsamen, David Rice, Brandon Rickabaugh, Joshua Stuchlik, and many others I can't remember; as well as the participants of the Midsouth Philosophy Conference, the LSU Graduate Philosophy Conference, the UMSL Gateway Graduate Philosophy Conference, and the 2nd Annual Philosopher's Cocoon Philosophy Conference. Thanks to the Institute for Humane Studies, the University of Arkansas Philosophy Department, and University of Arkansas Graduate School for helping to fund many of those conference trips.

Finally, I am very grateful for continual encouragement from my wife, who ensured that this thesis would be completed on time.

Dedication

For Nakia, my patient and understanding wife, and Oliver, my beautiful little boy. And for my amazing mother. I hope I've made you all proud.

Table of Contents

I.	The Problem of Epistemically Irrelevant Causal Factors	1
A.	The Initial Puzzle	2
B.	A Description of the Problem	7
C.	What It is Not.....	24
D.	Some Additional Puzzles	27
	Rochester and Rutgers	27
	Random Choice Generators	28
	Accidents of Birth.....	29
	Evolutionary Debunking.....	34
E.	An Argument Without Appeal to Counterfactuals	37
II.	Toward a Solution to the Problem of Epistemically Irrelevant Causal Factors.....	42
A.	What is Needed for a Successful Counterexample	43
B.	Good Evidence Evaluator Principle.....	49
C.	Well-Foundedness Principle	55
III.	Conclusion	59
	Bibliography	61

I. The Problem of Epistemically Irrelevant Causal Factors

There is a special problem in epistemology that has received relatively little sustained attention.¹ It is the skeptical worry that arises when there is some contingent factor in the causal history of my belief formation that is irrelevant to the epistemic evaluation of that belief. The danger it seems is that, had things gone differently with this factor, I would have believed other than I now do, even though my evidence would have been just as good and I would have been just as competent, all of which seems to undermine my present belief. And of course the more important or fundamental this belief is, the more of a problem I appear to have.

In section I. of this paper, I give a lengthy exposition and analysis of the problem. I find that it is needed to tease out exactly what the problem is for at least four important reasons. First, there is surprisingly very little that has been written on it, especially when compared to longstanding issues in epistemology such as the analysis of knowledge or external world skepticism. Moreover, of the few authors who have touched upon it, many have dealt with it only tangentially or have pressed it into service to make a larger point, while only a handful give a rigorous analysis of the problem itself.²

Second, this problem is easily confused with and blurred with other, nearby epistemological problems such that it is difficult to pin down what it is. As we shall see, while

¹ To some degree or another, the problem can be found in a few recent papers: Ballantyne (2012), Bogardus (2013), Cohen (2000), Dworkin (1996), Elga (ms.), Loftus (2010), Rosen (2001), Schechter (ms.), Schoenfield (2014), Sher (2001), van Inwagen (1994), Vavova (2010) and (ms.), and White (2010). For work that addresses the problem specifically within the context of religious belief, see Bogardus (2013), Garber (2007) and (2009), Loftus (2010), Plantinga (2000), and van Inwagen (1994). Within the context of evolution and morality, see Bogardus (forthcoming), Joyce (2006) and Street (2006). Finally, from a Continental perspective, Leiter (2004) discusses the hermeneutics of suspicion in Marx, Nietzsche, and Freud.

² I can think of only a few who give the problem itself a sustained rigorous analysis: Ballantyne (2012), Bogardus (2013), Vavova (2010) and (ms.), and White (2010).

our present problem raises concerns with trusting one's own reasons for belief, it is not difficult to imagine many other things that can generate this uneasiness and yet fail to be our problem.

Third, as it will become clearer throughout this section and especially later in Section II., not all instances of an epistemically irrelevant causal factor are what I call "distorting," or epistemically problematic. As Roger White explains, "only a fraction of the myriad of such factors even appear to raise a challenge to the status of my belief."³ Since not all irrelevant factors are epistemically bad, it will be crucial to distinguish those that are either innocuous or philosophically uninteresting from those that are the subject of our investigation.

Fourth, there are various elements at work that generate the problem, concepts that themselves can be difficult to navigate, including counterfactuals, causal influence, disagreement, and so on. Therefore, any analysis of our problem will need to handle these issues with the appropriate care, without getting too involved in the details of competing views concerning them.

Not until Section II. do I attempt to take up an evaluation of the problem. Try as I may to reserve all my criticism and attempts at a solution until later, however, it is inevitable that some of my own opinions will be apparent in this section.

A. The Initial Puzzle

Take a standard example concerning a fundamental philosophical belief, found in Cohen (2000):

Scholar: A scholar tosses a coin to determine whether to attend Harvard or Oxford. He attends Oxford and ends up endorsing the analytic/synthetic distinction. But he finds out that if he had gone to Harvard, he would have rejected the analytic/synthetic distinction, even though he would have been exposed to the same relevant evidence.

³ White (2010), p. 2.

Presumably, Cohen chose these institutions because they housed representatives of a fundamental disagreement in the history of philosophy: during the 1950s and 1960s, W. V. O. Quine at Harvard rejected the analytic/synthetic distinction, while P. F. Strawson at Oxford affirmed it.⁴ But Cohen could have chosen from a number of examples in which there is widespread disagreement among those informed and competent. Even now there are philosophy faculties around the world that are home to representatives of fundamental philosophical disagreements just like this one. So it appears that **Scholar** is generalizable within the philosophical domain, and likely within other domains as well.

On the subject of disagreement, I must say a few words early on in our investigation. I will assume in all that follows that rational disagreement between epistemic peers is possible. In other words, it is possible for two individuals who have access to the same relevant evidence and are just as competent can reasonably disagree about whether p . While at first this might seem quite impossible, given that evidence is intimately connected with truth, there are several possible explanations for epistemic peers reasonably disagreeing. For one, the two individuals might be epistemic peers only broadly construed (not carbon copies of one another), which can allow for an asymmetry in how each one appreciates the evidence or fits the evidence in with prior convictions. Second, while the evidence might dictate what proposition is to be held, it might be the case that a spectrum of rational credences can be reasonable to hold with respect to that proposition. For example, the evidence might dictate p , but it may be rational to believe p anywhere from a 0.5 to a 0.8 credence level.⁵ If this is possible, then it might be plausible to think that not only can individuals reasonably disagree about which credences to hold, but also which beliefs to hold—although it will perhaps not include disagreements between p and $\text{not-}p$,

⁴ See Rey (2013) for an overview of this important philosophical debate.

⁵ See Senor (ms.) for more on this.

it may allow for reasonable disagreement between p and withholding from p . Third, we are not epistemic machines: “our evidence sets are large, their contents often hard to determine on reflection, and the relevant evidential relations opaque.”⁶ This may warrant a reconsideration in what makes a disagreement reasonable, given that we are only human and we are doing the best with what we can. Given our limitations, an individual might be highly rational in the way he is responding to the evidence, although that person might still be in error.

I say that to say this. When we consider the problem in light of a counterexample scenario such as **Scholar**, there are two instances of disagreement. The first instance is a general disagreement between communities. For Scholar, it is the disagreement between those at Oxford and those at Harvard as it concerns the analytic/synthetic distinction. The second instance is the merely *possible* disagreement between a subject and his counterpart. In the example, it is the disagreement between Scholar in the actual world who goes to Oxford and Scholar in the counterfactual scenario where he attends Harvard. It is vitally important to keep these two instances of disagreement distinct in order to get a full handle on the problem.

I will say more about disagreement itself in Part C of this Section, and it will appear throughout the rest of this paper, but for now it is enough to earmark it and set it aside. As will become clear, I will do my part to distinguish this problem of epistemically irrelevant causal factors from the problem of disagreement, although our problem in many cases (though not all) presupposes disagreement. For those cases, I say let us suppose that rational disagreement is possible and that it does not provide a defeater for my belief. There might be this further worry that some non-epistemic causal factor has an influence on my belief, which *does* present a

⁶ Senor (ms.), p. 29.

defeater. (As we shall see later in Part B of this Section, disagreement is not *necessary* to generate our problem.)

Back to our example, Scholar attends Oxford as a result of a coin toss, although he likely would say this is not his reason for coming to his eventual belief in which he endorses the analytic/synthetic distinction. Of course, the coin could have landed differently, in which case Scholar would have gone to Harvard and eventually would have developed the opposite view, a belief in which he rejects the analytic/synthetic distinction, even though he would have been exposed to the same relevant evidence.⁷ This troubling possibility may lead Scholar to doubt whether he has good reasons for his present belief.

What evidence Scholar has will come to play an important role in setting up the problem and any solutions to it. In order to draw out this feature in more detail, compare the above case to the following:

Scientist: A scientist believes that $E=mc^2$ on the basis of good evidence. But she becomes aware of a primitive society somewhere in the South Pacific and learns that they do not believe that $E=mc^2$. She surmises that if she had grown up in this primitive society, she would not have believed as she now does.

Cases like **Scientist** are not epistemically troubling, because in this case Scientist's own current situation is epistemically privileged over that of a person in the more primitive society. In other words, Scientist has better evidence in her actual situation than she would have in the counterfactual situation (or, has better evidence than those in a more primitive society), and this gives Scientist a good *prima facie* reason to think that her belief is epistemically justified and well-founded and that it was *not* formed as a result of her simply being in a community where

⁷ Some might be troubled here by the usage of "would": perhaps if Scholar had gone to Harvard, he *easily might have* ended up holding the opposite view, but it is too strong to say that he *would*. I sympathize with this intuition. Though I stay true to Cohen's case here in using the "would have" counterfactual, in my later analysis (Section I., Part B) I opt for a weaker yet more plausible "easily might have" counterfactual.

most members share this belief (or some other plausible irrelevant factor—more on this just below). Thus, as things stand right now in her actual situation, Scientist can disregard this as a trivial possibility and rest assured that her present belief is well supported by her evidence.⁸

Far more epistemically troubling are cases in which one's present evidence is *not* obviously better when compared to the counterfactual scenario. These can fall into one of three categories: (i) cases like **Scholar** in which the evidence is just as good, (ii) cases in which one would have had better evidence than what one now has, or (iii) cases in which it is not clear whether one would have better evidence than what one now has. It should be clear why (ii) is troubling. Suppose that Scholar realizes upon reflection that he would have had far better evidence at Harvard than he now is in at Oxford. This will likely lead Scholar to doubt whether he has seen all the evidence and arguments available for his beliefs. For similar reasons (iii) is also troubling. If it is unclear to Scholar whether he has better evidence now than he would have in the counterfactual situation, then he might reasonably question whether he has seen all the evidence and arguments for his belief. For example, Scholar might simply be unaware of what arguments and evidence those at Harvard traffic in, so it is unclear to him whether he would be in a better situation epistemically.

While I have been speaking loosely up to now, it is helpful to define some terms. When I speak of being in a "better position epistemically," I mean by that, for example, that Scholar has better evidence or is more competent. By a "similar degree of competency," I mean being at least as likely to be getting at the truth in the domain in question (perhaps as a function of epistemic virtues or intelligence) in both the actual and counterfactual situations. By a "relevant

⁸ If the rationale here is unclear or far from obvious, I hope to support and clarify this in the present section and in Section II. If it helps with clarity, I also offer a version of the problem of irrelevant factors without appeal to counterfactuals in Section I. Part E.

sameness of evidence,” I mean, at least in the case of **Scholar**, that Scholar would have been exposed to all the same arguments and counterarguments for the analytic/synthetic distinction that he now is exposed to. (Once we change the domain to something far less heady than analytic philosophy, to religious disagreements for example, the nature of evidence will differ to also include perhaps testimony or perception.) Finally, by one’s evidence being “just as good,” I mean to draw a distinction between the content and the quality of a body of evidence. In some counterfactual scenario, one might not have the exact same evidence (e.g., access to the same exact philosophical arguments), but for all intents and purposes, the evidence will be of similar quality (e.g., access to philosophical arguments that are of similar strength).⁹

One final important feature of **Scholar** worth discussing is what I consider to be the driving force behind the case—the irrelevant factor itself. What troubles Scholar is the etiology of his belief. Given the facts about Oxford and the facts about Harvard, coupled together with the possibility that he could have been *there* and formed *that belief* rather than *here* where he formed *this belief*, there is a significant worry whether Scholar has formed his present belief on something epistemically irrelevant. Is Scholar’s belief in the analytic/synthetic distinction influenced by his surrounding community? To what degree is his belief formed as a result of non-epistemic factors rather than the evidence itself?

B. A Description of the Problem

I refer to this phenomenon as the problem of epistemically irrelevant causal factors, but it has been called various other things: the problem of irrelevant factors,¹⁰ the problem of contingency,¹¹ and the problem of historical variability.¹² What has led philosophers to call it a

⁹ I go into further detail on this below in Section 1., Part B.

¹⁰ Elga (ms.), Vavova (ms.).

¹¹ Bogardus (2013).

problem is the implicit challenge that questions whether, or the explicit conclusion which aims to demonstrate that, some particular belief has been defeated.

There are a couple ways to go about presenting the problem. One way is to present cases like **Scholar**, then speculate as to what might be the appropriate reaction to such cases: should Scholar be worried, or feel lucky, or something else? This is the strategy taken by Cohen (2000), followed later by cases from Ballantyne (2012), Elga (ms.), Schechter (ms.), Vavova (2010) and (ms.), and White (2010). This strategy can be helpful for filling in the narrative details that a formal argument might overlook; however, it can also leave out much of the precision we seek in delineating the exact problem. The other way, of course, is to present the problem as a formal argument, a strategy which is taken up by Ballantyne (2012), Bogardus (2013), Loftus (2010), and Sher (2001).

Here, in my analysis, I think it is helpful to present it both ways. We have already seen an example of a standard narrative case. And if we like, we can please ourselves with inventing several similar cases if this one does not satisfy. (In Part D, “Some Additional Puzzles,” I do some of that.) In keeping with the idea that this is a problem, then, the conclusion to any formal argument will be something like “S’s belief has been defeated” or “S’s belief is irrational.” In what follows, I shall give it my best try at presenting the strongest argument in favor of such a conclusion, relying upon recent authors who have offered their take on it.

An early presentation of the argument (within the context of moral beliefs) can be found in George Sher’s (2001), who gives only two premises: one about moral disagreement,

(S1) I often disagree with others about what I morally ought to do.
and the other “about contingent origins,”

¹² Ballantyne (2012).

(S2) The moral outlook that supports my current judgment about what I ought to do has been shaped by my upbringing and experiences. For (just about) any alternative judgment, there is some different upbringing and set of experiences that would have caused me to acquire a moral outlook that would in turn have supported that alternative judgment.

Sher gives us no explicit conclusion, but he takes it that these two premises together present a “challenge to the authority of my moral judgments.”¹³

There are a couple things to notice about this argument. First, whether it is in the moral domain or not, the belief in question must fall within some domain where disagreement exists.¹⁴ Second, Sher’s premise (S2) claims that there is presumably a strong degree of dependency upon my unique upbringing and experiences that plays a role in forming my belief. He does not go into detail about what this dependency relation amounts to. While he does note that “some different upbringing and set of experiences *would have caused* me to acquire” (emphasis mine) an alternative outlook, he does not analyze specifically how this works, whether as a function of some straightforward physical causal relation, or some more complex psychological influence, or something altogether different.

John W. Loftus (2010) presents the problem in an argument quite similar to that above, although his concerns have more to do with religious beliefs.

(L1) Rational people in distinct geographical locations around the globe overwhelmingly adopt and defend a wide diversity of religious faiths... This is the *religious diversity thesis*.

(L2) [I]t seems very likely that adopting one’s religious faith is not merely a matter of independent rational judgment but is causally dependent on cultural conditions to an overwhelming degree. This is the *religious dependency thesis*.

¹³ Sher (2001), p. 4.

¹⁴ Some later authors do not state this explicitly, but it is important as we will later see in this same Section I. Part B.

I believe the conclusion could go a number of ways, each more compelling than what Loftus gives, but he opts for:

(L3) Hence the odds are highly likely that any given adopted religious faith is false.

Although the contents differ, clearly this parallels Sher's argument with both the claim that there is a diversity or disagreement within some domain and the claim that one's beliefs within that domain in some way depend upon one's upbringing, experiences, or other cultural conditions.

Thus we can generalize the argument to accommodate multiple domains to produce an initially attractive argument for the problem of epistemically irrelevant causal factors:

(1) Disagreement (or difference of opinion) exists within domain *d*.

(2) S's belief *b* concerning *d* is strongly dependent upon S's upbringing, experiences, or other cultural conditions.

(3) Therefore, S's *b* concerning *d* is defeated.¹⁵

The conclusion here is meant to undermine justification. It is not necessarily an anti-knowledge argument. As such, if it is successful, S would no longer be justified (or, perhaps, rational) in believing *b*.

But as it stands this is not much of an argument. A glaring oversight is that there is no explicit inference from (1) to (2). As such, premise (2) is dubious. We are at risk of confusing correlation with causation. In other words, we can plainly see that beliefs and opinions differ from society to society, person to person. But it is altogether another matter to claim that what causes these different opinions are underlying cultural factors, whatever those may be. They

¹⁵ I will use Pollock's (1986) terminology for defeaters. In this case, the belief in question will be met with an undercutting defeater.

may very well be caused by such factors, but the argument above does not demonstrate that; it merely assumes it.¹⁶

However, such an assumption might not be unreasonable. As such, Sher and Loftus are in good company. As Antony Flew explains:

One positive reason for being especially leery towards religious opinions is that these vary so very much from society to society; being, it seems, mainly determined, as Descartes has it, “by custom and example.” The phrase occurs, in Part II of his Discourse on the Method, almost immediately after the observation: “I took into account also the very different character which a person brought up from infancy in France or Germany exhibits, from that which...he would have possessed had he lived among the Chinese or with savages.” (Flew, 1976)

Here, Flew points out that Descartes makes a judgment of religious dependency immediately following an observation of religious diversity. Included in this observation of religious diversity is the additional observation that there seems to be a high correlation of religious beliefs to culture. Nevertheless, Flew still hedges his claim, noting only that “it seems” that these opinions are determined by cultural factors. This reveals what I think is apparent—there is *no immediate inference* from the fact that opinions differ from society to society in a highly correlated way to the claim that the cultural factors present within each society cause the beliefs of people within that society. Nevertheless, there is something to be said for it being a dangerous *possibility* that one’s beliefs were caused by some irrelevant, non-epistemic influences. So in future permutations of this argument, we shall try to capture that possibility.

So far we have only the bare bones of an argument that attempts to uniquely distinguish the problem of epistemically irrelevant causal factors from other nearby problems. We need something further to set it apart. One feature that is clearly lacking is a counterfactual claim, a subjunctive conditional statement that allows us to compare our situation with another. While

¹⁶ We can note our worry here without entertaining a Humean skepticism with regard to causation.

not a necessary feature of the problem,¹⁷ we can see an implicit counterfactual in **Scholar** and other narrative cases that describe the problem (Sher provides hints of a counterfactual claim in his (S2) but doesn't spell it out fully). Recall: *If Scholar had gone to Harvard, he would have rejected the analytic/synthetic distinction.* There is something peculiarly fascinating about such a claim that threatens to make us much more uneasy than did our previous, weaker argument. Without this counterfactual claim, we have merely a claim of diversity and a claim of causal dependence, claims that are not clearly connected.

Also lacking is any information as to what kind of evidence and degree of competency S has relative to some counterfactual situation. These two features are important, because if S has good reason to think that he actually is more competent than he would have been or has better evidence than he would have had in some counterfactual situation, it does not seem that S should be bothered by this counterfactual possibility. This is because these facts about S are a *prima facie* indication that S's actual belief is epistemically justified and well-founded, and that it was *not* formed as a result of a non-epistemic, or irrelevant, factor (Recall **Scientist**). On the other hand, as we shall see below, if S has reason to think that he would have had the same or better evidence than he now has, this may provide for a defeater for S's belief.

Comparing evidence sets is a tricky undertaking, especially when it involves counterfactuals. Let us begin by looking more closely at cases in which S's evidence would be just as good, then move to cases in which there is a disparity in evidence. In **Scholar**, Scholar would have been exposed to the "same relevant evidence" as he now is. I take it that in **Scholar**, relevant sameness of evidence consists of, at a minimum, the same arguments and counterarguments for the viability of the analytic/synthetic distinction. But sameness of

¹⁷ See Section 1, Part E.

evidence will differ depending on the content of *b*. If *b* is not amenable to formal arguments—perhaps *b* was formed as a result of perception, or on authority, or by a reliable process—then relevant sameness of evidence will consist in whatever grounds *b* for S. Relevant sameness of evidence does not require that S is an atom-for-atom duplicate in both the actual and the counterfactual situation, or that all of S’s beliefs (excluding *b*) are exactly the same. This is too strong. We need only to require something weaker yet strong enough to generate our problem.

I propose we require that S’s evidence would be *minimally just as good*. Let us see what this means by looking at an example. Suppose that I am a Christian, and someone presents a counterfactual scenario to me, “You could have been born in the Middle East and have believed in Islam just as fervently as you now believe Christianity.”¹⁸ It is plausible to think that such a situation would still present a challenge even if I did not read exactly the same literature and was aware of all the same arguments. That is, it plausibly would still present a problem if my evidence would have been *just as good*—not necessarily the same evidence set, but at least the same *quality* of evidence. Perhaps I am now aware of only two arguments for Christianity, but in the counterfactual situation I would be aware of two similarly strong arguments for Islam. Even though they are not the same arguments, and even though I do not have the same relevant evidence, this strikes me as being just as problematic.

Also epistemically troubling are cases in which S would have had better evidence than he now has. For example, if Scholar realizes upon reflection that he would have been in a far better situation epistemically at Harvard than he now is in at Oxford, this will likely lead Scholar to doubt whether he has seen all the available evidence and arguments for his beliefs. Requiring

¹⁸ I discuss such a case in more detail in Section I. Part D.

that S's evidence would be *minimally* just as good leaves open the possibility that S would have been in a superior epistemic situation.

As for similarity of competency, I will assume that S is minimally just as competent in the counterfactual situation as S is in the actual situation. This will differ from person to person. So, perhaps, the problem presents a stronger challenge to individuals who are less competent than to those who are more competent. To explain, suppose that S is not very skilled or competent with respect to domain *d*. In such a case, it is plausible to think that there will be a great many nearby worlds¹⁹ in which S is more competent than he now is, because in those worlds any minimal qualitative change to some irrelevant factor may bring about a boon to his competency and skill—given that S begins so low, he can only go up. In other words, someone who is less competent with respect to domain *d* may have a greater *prima facie* reason to worry that he easily would have been more competent had things been different (i.e., would likely have gotten the answer right, which means S's actual answer is likely wrong).

Back to constructing our argument. By introducing a counterfactual claim, we can also implement these two additional features (i.e., S's evidence and competency, which factor into the quality of S's epistemic position relative to the counterfactual situation), such that our new argument presents a significantly stronger challenge. Thus we shall strengthen our premise (2).

Let us take another look at how other philosophers have attempted to construct the argument along these lines. Consider the following two counterfactual claims. The first is from Tomás Bogardus (2013), who presents the problem within the context of religious beliefs.

(Bog) If you had been born and raised elsewhere, else when, and formed religious beliefs using the same method you actually used, then you easily might have had different religious beliefs.

¹⁹ I am assuming, with Lewis (1973), that primacy is given to qualitative similarity over holding fast the laws of nature.

This gets us much closer to **Scholar**. We get our counterfactual, along with a sameness of method—I would add “using *minimally* the same method.” I take it that, in the context of religious beliefs, “same method” most likely means either theistic arguments, for example, or that one formed these beliefs as a result of testimony from other members in the surrounding community. (However, I am open to there being other “methods” by which one forms a religious belief.) We do not get similarity of competency unless we assume it.

Interestingly, Bogardus offers an “easily might have” counterfactual rather than a “would have” counterfactual. His reasons for this stem from concerns with analyzing what the epistemic failure is in this problem. After entertaining a bare counterfactual version of the problem, he moves on to examine whether the problem might get its force from the counterfactual coupled with a failure of the Safety criterion or by way of Luck, respectively.²⁰ I prefer, as Bogardus does, the strongest version of the problem, the Argument from Symmetry,²¹ which pairs actual disagreement within some domain with a counterfactual in which we hold fast sameness of evidence (and a similarity of competency). Placing the problem in the light of safety and non-accidentality produces this “easily might have” language, and Bogardus retains this language in his later symmetry argument from which I draw (Bog).

Although Bogardus’ is an anti-knowledge argument, I opt to use this weaker, though more plausible, counterfactual for my argument against justification.²² For one, in foregoing a stronger “would have” counterfactual, one avoids the cost of defending a more implausible premise. Second, the resulting argument still presents a strong enough epistemic challenge.

²⁰ Bogardus (2013), p. 379-87.

²¹ Bogardus (2013), p. 388-91.

²² Bogardus has said (p.c.) that although the “easily might have” counterfactual offers a weaker and more plausible premise, it may come at the cost of making the conclusion less plausibly follow.

When presented with the possibility that I easily might have gotten some proposition wrong, even though this “might have” language invokes probability of error rather than clear-cut error, if I have no strong, independent reason to think that my situation is not one of the few that got it right, then there may be good reason to think I should revise my belief.²³

Nathan Ballantyne (2012) offers a similar counterfactual.

(Bal) You have reason to believe that p is such that if your background had differed in certain respects, then you would not have accepted p , even though you would have used the same evidence for p and the cognitive skills relevant to appropriately believing p that you actually used.

This version is generalized to include any domain. We get our counterfactual, along with both a relevant sameness of evidence and a similarity of competency (he calls it having the same relevant “cognitive skills”). Again, I would modify these by adding “minimally” to allow for cases in which S would have been in a superior epistemic position. Here, Ballantyne uses “would not have accepted p ” to accommodate either the denial of or withholding from p . And, unlike Bogardus, Ballantyne uses the stronger “would have” counterfactual.

Taking these into consideration, we can amend our premise (2) into something much stronger. Let us take the whole argument together:

(1) Disagreement (or difference of opinion) exists within domain d .

(2*) S’s belief b concerning d easily might have been different²⁴ if some factor in S’s history that is (i) causally relevant to S’s b , but (ii) epistemically irrelevant to S’s deliberation about b had been different, even though (iii) S’s evidence would have been *minimally* just as good and (iv) S would have been *minimally* just as competent.

(3) Therefore, S’s b concerning d is defeated.

²³ As we shall see in Section II., this is close to Vavova’s view.

²⁴ To say that S’s b “might have been different” is to include any of the following situations: S’s believing not- b , and S’s withholding b , and S’s never having considered b .

What has been lurking in the background up to this point is that this factor is epistemically irrelevant to S's deliberation about *b*. That shows up in (ii) of premise (2*). Of course the factor—whether it be precipitated by a coin toss, or is the event itself of attending some university, or simply is the fact itself of Scholar's having gone to Oxford rather than Harvard—is epistemically irrelevant to S's reasoning concerning *b*.²⁵ When S deliberates about *b*, S does not consider contingent factors such as these to lend any epistemic support to the truth of *b*.

It is not clear how the factor is supposed to be “causally relevant” to S's *b*, and whether this causal relevance is strong enough to be considered distorting. However, there is a real sense in which it is possible that the factor is partly responsible for S's coming to have the belief S does. It seems, then, that one way of criticizing of (2*) would be to explain away or mitigate the causal influence of this factor, and instead explain that S has the belief S does for other, more epistemically appropriate reasons.²⁶

We now have a much improved argument, but it is not complete. Something also must be done about premise (1). So far we have only made an innocuous observation, that difference of opinion exists, whether this entails actual instances of individuals disagreeing or just the possibility of disagreement given the disparate opinions individuals hold. More needs to be said.

Let us consider (1) in the first-person, all the while keeping in mind the distinction I made earlier concerning two instances of disagreement (this is the first I presented). The mere fact that someone disagrees with me is not reason enough to think that my belief is thereby in danger of being undermined, for that person could be a young child, or a novice on the subject, or someone who is simply feigning disagreement. However, if I have good reason to think that

²⁵ See Section 1., Part D.2 “Random choice generators” for clarification on coin tosses and other precipitating factors.

²⁶ I return to this point in Section II.

the person is an expert or is just as informed as I am on the subject, then I may have reason to reconsider my belief. So it matters whom I disagree with. But more than this, it is not really who the person is that worries me when we disagree. If a child disagrees with me, yet I know her to be very informed on the subject, perhaps this gives me reason to reconsider. But if my wife, whom I trust, disagrees with me, yet I know that she is clearly misinformed, perhaps I should not reconsider my belief simply in light of this disagreement (but maybe I should on other, non-epistemic grounds!). So it also matters whether I have reason to believe that my interlocutor is well informed. My evidence for thinking so might reveal some parity in our evidence, or in our intellectual virtues, or in our similar likelihood in getting the point in question right.²⁷ On the other hand, the more reason I have to think that my interlocutor is my epistemic peer, perhaps the stronger *prima facie* reason I have for revising (though it might not be an *ultima facie* reason). Perhaps not. Perhaps it offers no reason at all to revise.

I will remain neutral here as to which view in the epistemic peer disagreement literature is correct.²⁸ The point is, the problem of disagreement itself is much more philosophically interesting when cast in light of epistemic peers disagreeing, as opposed to disagreement between an epistemic superior and an epistemic inferior. So, we may say that (1) is best phrased as disagreement either between epistemic peers or, minimally, among those who are “close

²⁷ Philosophers have recently disagreed about what counts for epistemic peerhood. Some say it is captured by sameness of evidence, while others include this criterion along with sameness of epistemic virtues. Still others think it is a matter of two people who are just as idly to get the point in question right.

²⁸ This despite my earlier comment urging the reader to assume, for the sake of argument, that reasonable disagreement between epistemic peers is possible. If the reader is unwilling to make this assumption with me, that is quite fine, since it appears that disagreement within a community (such as in premise (1)) is not necessary for the argument to go through. See just below for my evaluation of this.

enough” to epistemic peerhood status so that it remains an interesting problem. That new premise, along with the rest of the argument, is as follows:

(1*) Disagreement (or difference of opinion) exists within domain d among those who are both informed and competent concerning d .

(2*) S’s belief b concerning d easily might have been different if some factor in S’s history that is (i) causally relevant to S’s b , but (ii) epistemically irrelevant to S’s deliberation about b had been different, even though (iii) S’s evidence would have been *minimally* just as good and (iv) S would have been *minimally* just as competent.

(3) Therefore, S’s b concerning d is defeated.

I assume in premise (1*) that the disagreement must be substantive or significant. It does not need to be widespread or universal, though cases having disagreements with that degree of disparity presumably might satisfy the premise more easily. To say that it need only be substantive or significant, and not widespread, leaves open the possibility that the disagreeing parties might be two individuals (e.g., me and my brother) rather than whole communities. For example, if things had gone differently, I easily might have believed as my brother does about the exact details concerning a shared memory we have from childhood. Perhaps my belief has been influenced by some completely unrelated later event which invoked PTSD in me. If that later event had not occurred, I would still have all the same evidence and be just as competent, but I easily might have shared the same belief as my brother.

Requiring that the disagreeing parties be epistemic peers is much too strong, but premise (1*) captures those cases plus some. The closer parties are to epistemic peerhood status, the stronger the problem is; the further away they are, the weaker it is. There is a concern with vagueness that threatens to be serious, but we will set that aside for now. In any case, I do not think vagueness should worry us in this context, because our problem trades on an epistemic

worry. Insofar as I have reason to believe that parties are close to epistemic peerhood status, this problem will, *ceterus paribus*, present a strong challenge; and if not, then not.

But is premise (1*) necessary? The short answer is no. However, it is important because it captures those cases, which happen to encompass a majority of cases, in which my belief might have been different. Let me explain. In many of the paradigm cases of this problem, we assume premise (1*)—we assume that there is a significant possibility that my present true belief could have been false, or that my present false belief could have been true. (Many recent authors do not explicitly include premise (1*) in their versions of the argument, but it is assumed nonetheless.) The motivation for this assumption lies in the fact that there actually is significant disagreement among those who are informed and competent with respect to certain beliefs I hold, which raises the probability that a change in my history would have led me to acquire a different belief. To see that this is so, suppose that premise (1*) is false. Then I would not expect there to be a significant possibility that my present true belief could end up being false had things been different. I might instead suppose that my present true belief would have been true nevertheless, just because there is no widespread disagreement on the matter.

However, even though many of the paradigm cases of this problem assume disagreement, it is not necessary to generate the problem. It is the non-epistemic, irrelevant causal factor that provides a defeater for my belief, not the fact of disagreement in some domain. Take an example in which the belief in question is virtually universally held, such as *that the earth is round*. I was raised in a certain community A in which I developed this belief, but later on I realized that my upbringing is contingent and I could have been raised in another community B. My belief would have been the same, so that does not worry me; however, the grounds for my belief could be in question. It is possible that I formed my belief as a result of a non-epistemic factor, such as the

cultural influence of community A, rather than on good epistemic grounds. This remains a problem whether the belief I hold is false or true.

Nevertheless, the interesting paradigm cases are those in which my belief would have been (or easily could have been) different. In these cases, both the grounds for the belief and the truth of the belief are in question. Given that this produces a stronger overall argument, we shall keep premise (1*).

Premise (1*) is important for another, related reason. We are after only those cases in which there is interesting disparity among opinions. Uninteresting disparity of opinion would be something like the following. Plantinga has the belief “I was born in Michigan.”²⁹ But his belief easily might have been different if he were born in another place.³⁰ This sort of scenario can take off as a result of individuals disagreeing about where “I” was born—where “I” is used as an indexical. But this is not an interesting disagreement. For that matter, it is not even real disagreement at all, since “I” refers to the speaker in whose mouth it is uttered. Another example is the disagreement between those who think a tomato is a fruit and those who think it is a vegetable. This is not a veritable disagreement, since different classification systems are being used: scientifically, a tomato is a fruit, but culinarily it is a vegetable. The same can be said for many linguistic disputes, which often appear to be veritable disagreements, when in fact different classifications are being applied. Having premise (1*) ensures that we are considering some domain *d* in which there is veritable and interesting disagreement among informed and competent individuals. Paradigm examples include, but are not limited to, the domains of philosophy, politics, religion, and ethics.

²⁹ Plantinga (1995).

³⁰ I have significantly simplified this counterfactual for readability. I will leave it to the reader to consider how all the components of premise (2*) fit in.

Our argument needs an additional criterion. Both premises (1*) and (2*) could be true and yet I could have no awareness of them being true. Perhaps I have not even reflected upon the presence of widespread disagreement concerning my belief, or perhaps I have never even considered this complicated counterfactual possibility. It seems right to say that if I am not aware of (1*) and (2*), even though they are true, I suffer no ill epistemic consequences as a result. This is because (1*) and (2*) together present an evidential defeater for my belief *b*. If I never possess this piece of evidence, then I never have a defeater for my belief. Thus, our argument needs a premise specifying the requirement that I have an awareness of (1*) and (2*). (We see a similar requirement in Ballantyne's (Bal), "You have reason to believe that...".) So we'll shift our original premise (3) down into a conclusion we will mark (C). We will replace it with a new premise (3*).

(1*) Disagreement (or difference of opinion) exists within domain *d* among those who are both informed and competent.

(2*) S's belief *b* concerning *d* easily might have been different if some factor in S's history that is (i) causally relevant to S's *b*, but (ii) epistemically irrelevant to S's deliberation about *b* had been different, even though (iii) S's evidence would have been *minimally* just as good and (iv) S would have been *minimally* just as competent.

(3*) S has reason to believe (1*) and (2*).

(C) Therefore, S's *b* concerning *d* is defeated.

One might be tempted to think that the best strategy is to include the awareness requirement along with the counterfactual. While this is one way of formulating the argument, there are good reasons for keeping the awareness requirement separate. Granted, it is true that the argument presents a challenge only in virtue of S being aware of (1*) and (2*), so one might think that we should instead simply modify those premises to include "S has reason to believe..." at the beginning of each one (or, much more simply, lump everything together in one big premise,

since (3*) seems to be doing all the heavy lifting). But this would allow for odd cases. For example, I could have reason to believe falsely, in the sense that I have a justified false belief, that widespread disagreement exists among those informed and competent concerning the shape of the earth, thinking that the population divides roughly in half between those who believe the earth is flat (“flat-earthers”) and those who believe it is round (“round-earthers”). And given that I have reason also to believe premise (2*), then my belief is defeated. But this example is not one worth being concerned about. We are interested instead in cases in which there is actual, real-world disagreement or difference of opinion among those informed and competent. Having premises (1*) and (2*) as states of affairs that obtain, rather than merely as something that S believes, ensures this.

We should add one final premise to our argument, so as to draw out explicitly the logical connection between the first three premises and the presence of defeater language in the conclusion. The resulting argument is this:

(1*) Disagreement (or difference of opinion) exists within domain d among those who are both informed and competent.

(2*) S’s belief b concerning d easily might have been different if some factor in S’s history that is (i) causally relevant to S’s b , but (ii) epistemically irrelevant to S’s deliberation about b had been different, even though (iii) S’s evidence would have been *minimally* just as good and (iv) S would have been *minimally* just as competent.

(3*) S has reason to believe (1*) and (2*).

(4) If (1*), (2*), and (3*), then S’s b concerning d is defeated.³¹

(C) Therefore, S’s b concerning d is defeated.

³¹ So as to enhance readability, I simply refer to the previous premises as providing a sufficient condition for defeat rather than pack this premise full with convoluted and unnatural wording. In doing so, I think this avoids a Lewis Carroll-type infinite regress (Carroll 1995).

We now have what I take to be the strongest argument for the problem of epistemically irrelevant causal factors.

Putting all these together, we can arrange a set of criteria which are individually necessary and jointly sufficient for being classified as an instance of the problem of epistemically irrelevant causal factors, in its current form:

- (a) veritable and significant disagreement within some domain among those who are both informed and competent,
- (b) some contingent factor, which is causally relevant but epistemically irrelevant, that easily might have been different,
- (c) S's evidence being minimally just as good in both the actual and counterfactual situations,
- (d) S being minimally just as competent in both the actual and counterfactual situations, and
- (e) awareness of (a) — (d).

I qualify my previous claim, because later (in Section 1, Part E.) I present a form of the argument that does not appeal to counterfactuals, so those criteria will obviously look very different from what I have listed here.

C. What It is Not

In order to clarify and delineate our problem from nearby related problems and skeptical worries, I must say a few words about what it is not. First, the problem must be distinguished from mere skepticism or undermining of one's beliefs in general. There are various ways to arrive at a skeptical conclusion or to have one's beliefs defeated, but it is this specific way we are interested in.

Second, it is not the mere reflection upon and reassessment of one's views. As Roger White points out, "One can be inspired to reassess one's beliefs in all sorts of ways. The interesting question is whether the causal background of these beliefs can have epistemological

relevance itself.”³² For example, you might become concerned about the status of some religious belief, but this concern may just come down to wondering whether your current justification is sufficient. In this case, it is (epistemically) good that you reassess your justification for that religious belief, but certain non-epistemic, causal influence(s) of the belief need not be playing any role here. We are concerned with cases of reassessment only insofar as they are those in which my primary motivation for reassessment lies in the realization that my belief is the result of a contingent, causal factor as I have described.

Third, the problem of epistemically irrelevant causal factors is a distinct problem from epistemic peer disagreement. Roger White has argued that this phenomenon “really just come[s] back to the issue of disagreement.”³³ What is driving the worry that Scholar might end up with a different belief “has nothing to do with the facts he has discovered about the etiology of his own beliefs. It is just the fact that he has evidence that there are apparently very smart, well informed philosophers who differ in their opinions.”³⁴ I have a quite different intuition. What drives Scholar, for example, to be concerned about his present belief is not the fact that there are individuals who are his epistemic peers and who disagree with him. Rather, it is the fact that Scholar realizes that there is a possibility that his present belief was formed on the basis of non-epistemic, or irrelevant, reasons.

As we have seen, there are multiple features that come together to produce the unique problem that we have been discussing. Disagreement among informed and competent individuals is just one of these features. Granted, if S becomes aware of only premise (1*), it may be enough to produce a defeater for S’s belief. Whether it does may in fact depend upon S.

³² White (2010), p. 3.

³³ White (2010), p. 30.

³⁴ White (2010), p. 30.

To demonstrate, suppose that when S hears of general facts like (1*), S is not sufficiently self-reflective to recognize what consequences it has for his beliefs. Or perhaps S is so confident that his belief is true that he does not worry about the possibility of being mistaken simply because there are other informed and competent individuals who disagree with him. Or perhaps, as some authors in the disagreement literature have suggested, the fact of disagreement does not provide any first- or second-order evidence that S is mistaken. It is reasonable to think, in light of these scenarios, that it might take an additional awareness of the counterfactual possibility in (2*)—that S might easily have believed otherwise—for S to grasp the severity of his situation or, depending on your view of epistemic peer disagreement, for S’s belief even to be defeated in the first place. I will not offer a judgment on any of these scenarios, but I will only point out that disagreement is not identical to our present problem, even though many cases of the latter presuppose disagreement and even though it is possible that in some cases disagreement provides a defeater for one’s belief.

There is one further reason for keeping our problem separate from disagreement. One might be tempted to present an analysis of the problem of epistemically irrelevant causal factors in terms of possible peer disagreement—in which my possible peer, hypothetical me at Harvard, would come to believe differently that I actually do (this is the second instance of disagreement I introduced early on)—but putting it in these terms may not be productive. While this may be an interesting and helpful way of framing the problem, we may not benefit from access to the same answers available to those trying to solve the disagreement issue, primarily because those answers are aimed at actual peer disagreement and not hypothetical peer disagreement. These are quite different phenomena. As John Stuart Mill writes, it is quite different to hear counterarguments “from persons who actually believe them, who defend them in earnest and do

their very utmost for them...in their most plausible and persuasive form” (*On Liberty*).

Speculation about what my hypothetical self would argue, given the same evidence yet believing oppositely, is just not as forceful as when I hear those arguments and reasons in the mouths of others. Thus it remains an open question whether the resources available in the disagreement literature are available to us here.

D. Some Additional Puzzles

Upon encountering a description of the problem of epistemically irrelevant causal factors, one might be suspicious as to how common or far-reaching the problem is in the real world. Perhaps it is only a conceptual problem in the abstract but rarely, if ever, rears its head in reality. Here, I present and discuss other kinds of cases of the phenomenon, because I think it is helpful to see just how pervasive the problem really is.

Rochester and Rutgers

This case is almost identical to **Scholar**, but it differs in a few interesting ways. Of course, the universities and the professors are different, but the proposition at hand is also different in that it pits evidentialism against reliabilism, so this case is of special, perhaps personal, interest to epistemologists. Moreover, this case is more contemporary, unlike the debate concerning the analytic/synthetic distinction which flourished during the 1950s and 60s, so epistemologists today seeing this newer, updated case are perhaps met with a more tangible sense of worry that may be absent in **Scholar**. Indeed, upon considering the case, one philosopher has said to me, “That really does bother me, and I’m not quite sure what to do about it.”

Grad Student: Student goes to Rutgers to study with Alvin Goldman and becomes a reliabilist. But he realizes that if he had gone to Rochester to study with Richard Feldman, he very likely would have become an evidentialist instead, rejecting reliabilism, even though he would have been exposed to the

same relevant evidence.

As before, we should be careful to point out that not only is there a sameness of relevant evidence, but there is also the implicit fact that there is a similarity in competency in both the actual and counterfactual situations. Moreover, in the actual situation Student has good, independent reason to think that what *that* department is doing is legitimate. In other words, he has no independent reason to judge the department, professors, or teaching quality at Rochester as inferior to Rutgers. When Student becomes aware of this counterfactual possibility, it seems his belief is met with a *prima facie* defeater.

Random Choice Generators

In **Scholar** a coin toss is used to make a decision regarding where to attend university. Thus, coin tosses and other random choice generators (e.g., darts, dice, Roshambo, etc.) may be used to (eventually) generate the belief in question. That is, it makes no difference whether Scholar tossed a coin in choosing between Harvard or Oxford, or whether he threw darts at a map instead. Eventually, as a result of the random choice generator, Scholar goes to some university and believes *p*, then comes to realize that if he had gone to another university he would not have believed *p*.

However, this might lead one to believe (wrongly) that the random choice generator is what is doing the work to generate the problem. It is not—or at least it is not by itself. We can easily see this in two ways. First, and most obvious, a random choice generator is not by itself sufficient to generate the problem. Recall that certain features are required to collectively produce the unique types of cases we are after:

- (a) veritable and significant disagreement within some domain among those who are both informed and competent,
- (b) some contingent factor, which is causally relevant but epistemically irrelevant, that easily might have been different,

- (c) S's evidence being minimally just as good in both the actual and counterfactual situations,
- (d) S being minimally just as competent in both the actual and counterfactual situations, and
- (e) awareness of (a) — (d).

One might think that random choice generators satisfy (b), but I think this is mistaken. They do not by themselves satisfy (b), of course, because what is needed is the additional fact of, for example, Scholar's having gone to Oxford rather than Harvard. This can more easily be seen in the next point.

Second, a random choice generator is not necessary. We can imagine cases like **Scholar** in which all of (a)—(e) are satisfied without appeal to a random choice generator. In fact, we just saw one in **Grad Student**. Student decides to attend Rutgers, not as a result of any random choice generator, but as a result of certain idiosyncratic reasons (like anyone making a big life decision). They can be good reasons: perhaps Student is previously inclined to accept reliabilism, or perhaps he has been exposed far more to reliabilism than to its rivals. Or they can be epistemically bad reasons: perhaps instead he simply admires Alvin Goldman's writing style, or he wants to attend a university with a good football team, or Rutgers is closer to family, or (God forbid) he is fonder of Rutgers' school colors. The point is, just about any reason, epistemically good or bad, can stand in the gap and fulfill the coin toss function as it appears in **Scholar**. So, while it might help with the expository presentation to include clear-cut examples in the form of rational decision-making tools such as darts or dice, it certainly is not necessary to include random choice generators.

Accidents of Birth

Perhaps the most commonly recognized trigger for the problem of epistemically irrelevant causal factors, whether the context involves philosophers writing on this specific

problem or whether it involves just a cursory inventory of the epistemic status of one's own beliefs, is when the irrelevant factor is (or is precipitated by) one's own accident of birth. By "accident of birth" I mean the contingent facts surrounding one's own time and place of birth. This too would qualify as triggering (or as being itself) an irrelevant factor like our argument describes. Take any belief that seems to be cultivated in your youth and about which there are disparate opinions among those informed and competent. Given that you easily could have been born elsewhere and else when, consider how it easily might have been different, even though your evidence would have been just as good and you would have been just as competent. That belief is subject to an irrelevant factor such that it is causally relevant but epistemically irrelevant to your deliberation concerning that belief.

While there is fertile ground here for application to a number of different domains, as we shall see, religious beliefs commonly receive the brunt of criticism. In fact, this specific type of challenge has received more attention in print than has the problem of epistemically irrelevant causal factors in general. Hence, the discussion that follows will center around religious beliefs, though by no means do I intend to convey that these are the only beliefs subject to our problem when it comes to accidents of birth.

It seems to me that many people consider the following type of reasoning a strong challenge to religious belief, and some have even used it to describe the loss of their own religious convictions:

Christian: If things had gone differently before my birth, I easily could have been born in the Middle East and have been raised Muslim, just as I was raised Christian here in the United States. As a consequence, I would have believed just as confidently and have defended just as fervently several religious beliefs which I now reject.

This should sound familiar. At some time or another, perhaps such thoughts cross all our minds, instilling doubt about our present beliefs. Even prominent philosophers have reflected on this possibility concerning their own beliefs, some even coming to a conclusion similar to what our argument recommends.

We have already seen Descartes' musings: "I took into account also the very different character which a person brought up from infancy in France or Germany exhibits, from that which...he would have possessed had he lived among the Chinese or with savages."³⁵ Descartes recognized that an individual's beliefs concerning religion and values would likely be very different had that individual been born and raised in a different cultural climate.

Peter van Inwagen writes, "[I]f I and some child born in Cairo or Mecca had been exchanged in our cradles, very likely I should be a devout Muslim."³⁶ Van Inwagen is unsure whether the same can be said for the other child, however, since he himself was not raised Christian. John Hick develops this point further and recommends a suspicious attitude as the appropriate response: "[R]eligious allegiance depends in the great majority of cases on the accident of birth: someone born into a devout Muslim family in Pakistan is very likely to be a Muslim, someone born into a devout Hindu family in India to be a Hindu, someone born into a devout Christian family in Spain or Mexico to be a Catholic Christian; and so on. The conclusion that I have drawn is that a 'hermeneutic of suspicion' is appropriate in relation to the beliefs that have been instilled into one by the surrounding religious culture."³⁷ On Hick's recommendation, one should adopt a skeptical attitude towards one's own religious beliefs to the

³⁵ Descartes (*Discourse* Part II).

³⁶ van Inwagen (1995, 238).

³⁷ Hick (1997, 281).

degree to which one's religious beliefs are affected by time and place of birth (assuming one can know this).

The observation that one's religious beliefs are often dependent on accidents of birth can be pressed into service for an argument against religious exclusivity. I quote van Inwagen again to offer a representative line of thinking: "Well isn't it fortunate for you that you just happen to be a member of this 'unique instrument of salvation'. I suppose you realize that if you had been raised among Muslims, you would make similar claims for Islam?"³⁸ Even the recent and popular atheist writers have gotten on board with this criticism against religious belief and exclusivism. John W. Loftus co-ops this argument for his "Outsider Test for Faith," calling on religious believers of all faiths to take the challenge.³⁹ Once the subject has been made aware of the widespread diversity among religions and the counterfactual possibility that she could have been raised to believe any one of them, Loftus is confident that she will do the rational thing and give up her present religious beliefs in favor of agnosticism or skepticism.

I consider this challenge to religious belief to be a special application of the problem of irrelevant factors involving accidents of birth. But we must take care here to distinguish genuine cases of the problem of irrelevant factors from those that are not. Roger White's evaluation is worth quoting at length:

It is common for people to describe the loss of religious conviction as precipitated by the thought "I suppose I just believe this because I was brought up this way. I could easily have believed something completely different had I been raised in a different home." But plausibly what is going on here in many cases is just that a thought like this occasions the reassessment of their beliefs on more general grounds. The doubter may be led to consider what grounds he has for his previously unquestioned convictions and find them wanting. But reflection on the *causes* of belief need not really be playing any epistemological role here. Perhaps the doubter is right to think that he never had any good reason to believe

³⁸ van Inwagen (1995, 237-8).

³⁹ Loftus (2013).

as he did. (White, 2010: 3. Emphasis mine.)

In other words, many cases in which an individual is led to doubt his belief upon reflecting on his accident of birth will not be the kinds of cases we are concerned with, because they are not instances of the problem of irrelevant factors. White has diagnosed exactly why these cases are found wanting: because they are not cases in which one's doubts concern the *causes* of one's beliefs. (That is, S is not aware of criterion (i) in premise (2*) of our argument.) It is not a question about whether one should reassess one's beliefs in general. "The interesting question is whether the causal background of these beliefs can have epistemological relevance itself."⁴⁰

Thus, we must set aside these kinds of cases. But we must also exclude certain other cases, albeit on entirely different grounds. When we consider a case like **Christian**, we do not commonly think of our evidence being just as good and of ourselves as being just as competent in the counterfactual situation. It is understandable why, especially when it comes to religious beliefs. For example, I might think that, while I believe Christianity is true because I have access to special revelation provided in the Bible, Muslims do not believe it is true because they lack this evidence. This could plausibly give me reason to think that my actual epistemic situation is epistemically superior, that I would be epistemically worse off in any counterfactual scenario in which I am raised Muslim. From this, perhaps we can conclude that a difference in evidence in the religious domain provides for as much explanation of disparity of opinions as it does in the scientific domain, as we saw in **Scientist**. But I will not offer a judgment on that here. I will only say that any case like **Christian** does not meet all the criteria set in our argument unless S also has some independent reason to think that S would have had at least the same relevant evidence and be at least just as competent. Thus we need to modify the case:

⁴⁰ White (2010), p. 3.

Christian*: If things had gone differently before my birth, I easily could have been born in the Middle East and have been raised Muslim, just as I was raised Christian here in the United States, *even though my evidence would have been just as good and I would have been just as competent.* As a consequence, I would have believed just as confidently and have defended just as fervently several religious beliefs which I now reject.

There are a couple interesting observations to make. The scenario mentioned by van Inwagen does not seem to be cause for concern in its current form. If we assume that, in some counterfactual scenario in which van Inwagen is switched at birth, his evidence is just as good and he is just as competent, then we can see why this sort of symmetry would be a problem. As a result, van Inwagen might reasonably wonder whether he is wrong about his present Christian belief, perhaps withholding from it until he can search for further evidence. However, if we assume that van Inwagen would have been far worse off epistemically than he now is, then this gives him no good reason to withhold judgment in his present belief. He might just as reasonably think that his being in a (relatively) superior actual epistemic position is good evidence for thinking that his belief was formed on the basis of good evidence and not as a result of a distorting irrelevant factor.⁴¹

Evolutionary Debunking

There is one final kind of case I should discuss, one in which the distorting irrelevant factor is or is triggered by one's evolutionary history.⁴² Tomás Bogardus calls this a "sub-type of a more general worry, which one might call 'the problem of irrelevant causal factors.'"⁴³

⁴¹ For a full analysis of this line of thinking, see Section II.

⁴² I do not mean to imply that the problem of irrelevant factors is limited to just these kinds of cases we have seen in this section

⁴³ Bogardus (forthcoming), fn. 1.

While evolutionary debunking arguments comprise a vast literature themselves,⁴⁴ it is worth looking at these arguments as they relate to our present problem. Consider the following case:

Evie: Evie holds a certain moral belief *b*. But she realizes upon reflection that, if evolution is true, her moral faculty was naturally selected to produce adaptive moral beliefs, and not naturally selected to produce true moral beliefs. She thinks to herself that if our species had evolved elsewhere at some other time, she easily could have come to have a different moral belief, using the same method she actually used. This is because there easily could have been different conditions for survival, which could have led to different beliefs being selected. Evie concludes that this provides a defeater for her present moral belief *b*.⁴⁵

Of course, this case precludes certain views of moral psychology. On a view we might call “Rational Insight Theory,” moral beliefs “*can be formed solely* on the basis of what’s been variously called rational insight, direct perception, direct apprehension, or simply *presentation*.”⁴⁶ And so evolutionary adaptive mechanisms might not play a role at all in the formation of at least some of Evie’s beliefs. On another view we might call “Divine Revelation,” moral beliefs “*can be formed solely* on the basis of divine testimony, a supernaturally-endowed conscience, the inward instigation of the Holy Spirit (as Aquinas might say), etc.”⁴⁷ Again, some of Evie’s beliefs easily might have been formed this way instead of through evolutionary natural selection. As Bogardus argues, these two views cause problems for **Evie**, since they involve no mental intermediaries which are naturally selected to produce adaptive moral beliefs and thus susceptible to being potentially misled from delivering true moral beliefs. However, the inferences in **Evie** are valid on the following view of moral psychology, what Bogardus calls “Representationalism.” “[A]ll our moral judgements come by way of a *mental intermediary*, indication, report, or representation, which is delivered by our

⁴⁴ For example: Bogardus (forthcoming), Clark-Doane (2012), Joyce (2006), Kahane (2011), Street (2006), and Vavova (2014) and (2015).

⁴⁵ See Bogardus (forthcoming), p. 21-6 for an evaluation of a case similar to this.

⁴⁶ Bogardus (forthcoming), p. 6. Italics in original.

⁴⁷ Bogardus (forthcoming), p. 7. Italics in original.

moral faculty and figures crucially into our formation of our moral beliefs.”⁴⁸ These mental intermediaries can be anything from sentiments to thoughts, all of which are subject to natural selection.

Insofar as **Evie** and cases similar to it rest upon this Representationalist view of moral psychology, in which all of one’s moral beliefs are susceptible to being formed on the basis of what ensures the organism’s survival, there is risk for a distorting irrelevant factor. As circumstances are changed in counterfactual situations, so too are the requirements for an organism’s survival, for these are not necessary in any logical or metaphysical sense of necessity. Thus in many counterfactual scenarios, it is plausible to think that Evie might come to have different moral beliefs. Upon realizing this, she may wonder whether she is in a better epistemic situation with respect to these counterfactual situations, which gives her a *prima facie* defeater for her belief.

Notice that the case mentions using the same “method,” rather than sameness of evidence and similarity in competency. Reliabilists should be comfortable with this talk, because what entails knowledge on reliabilism is having a reliable method that produces knowledge in the subject. The method in **Evie** is just the process of natural selection that produces adaptive moral beliefs. It is clear on this view that there would indeed be a symmetry in method between the actual and counterfactual situations. Furthermore, in Evie’s case, in the larger community there is actual disagreement or difference of opinion concerning certain moral beliefs she holds, and she recognizes this, along with recognizing that evolution presents a causal influence which is irrelevant to her deliberation about *b*. This kind of case meets all the criteria for being classified under our problem.

⁴⁸ Bogardus (forthcoming), p. 5. Italics in original.

If talk of methods is cumbersome, consider the case in terms of what *evidence* Evie has. Evie easily could have come to have a different moral belief, using at least the same relevant evidence and having at least the same level of competency as she now has. Suppose that Evie's belief *b* is *that all forms of war are morally wrong*. Her evidence for this includes the thought that war involves the taking of human life and is harmful to those associated with it, as well as certain associated sentiments and affections that support these thoughts (i.e., when people die, she experiences sadness). But perhaps Evie's beliefs were cultivated because she was already predisposed (e.g., through intuitions selected for by evolution) to believe that way. Plausibly, there are others who have access to evidence that is just as good and yet are predisposed to believe a different way. So it is possible that Evie could have been predisposed to believe a different way as well, all despite having evidence that is similarly good and being just as competent. And so we have our problem of epistemically irrelevant causal factors once more.

E. An Argument Without Appeal to Counterfactuals

It has been suggested to me that perhaps an argument for the problem of epistemically irrelevant causal factors can be made without appeal to counterfactuals.⁴⁹ There are at least three reasons why someone might want to have such an argument, not the least of which one avoids the potentially troublesome language and set-up as a result of talking in terms of counterfactuals. That is, one could become distracted or misled by the presence of counterfactuals, to believe, for instance, that the problem is an instance of possible peer disagreement between myself and my counterpart in the counterfactual situation. Assuming modal realism is false, there is just me, and my counterpart does not exist as a peer, so it appears it would be best to consider the

⁴⁹ This suggestion comes from Tom Senor (p.c.).

problem as explicitly concerning only one person, myself.⁵⁰ After all, as I argued before, our problem gains traction not from disagreement, but from the presence of an epistemically irrelevant causal factor. Secondly, one might be suspicious about the truth value of subjunctive conditionals; or if one does not doubt that certain counterfactuals are themselves true, then one might be suspicious about what grounds them. These reservations about speaking in terms of counterfactuals can be easily alleviated if we are able to construct an argument that leaves them out. Thirdly, it may be unclear in setting up cases for the problem of epistemically irrelevant causal factors how broad or narrow (or how far away or near) we should construe possible scenarios. With enough detail, we can see how a certain case is a token instance of our phenomenon, but it is not entirely clear how far away a possible world can be before it is considered an implausible scenario or perhaps better classified as some other phenomenon.

Thus, I will attempt to give an argument in this section which does not appeal to counterfactuals. To do so, let us recall the single premise in our argument that is loaded with counterfactual language:

(2*) S's belief *b* concerning *d* easily might have been different if some factor in S's history that is (i) causally relevant to S's *b*, but (ii) epistemically irrelevant to S's deliberation about *b* had been different, even though (iii) S's evidence would have been *minimally* just as good and (iv) S would have been *minimally* just as competent.

This premise will need to be modified, but hopefully not beyond recognition. It is a fair concern to note that modifying our present problem of epistemically irrelevant causal factors too much in order to arrive at the same conclusion might change the problem altogether. This was a worry from the start of our investigation. We wanted to delineate our problem from other, nearby epistemological problems that reach the same or similar conclusion of defeat for one's belief.

⁵⁰ Or yourself, or S, as the case may be. I do not think that the problem changes in any significant way by modifying it from first-person to second-person to third-person.

But if we cannot replicate our problem exactly, we can perhaps formulate a similar enough argument that bears a “family resemblance,” or is a different species within the same genus of epistemological problems.

Keeping this in mind, consider a case involving two universities. As before, let us stipulate that disagreement (or difference of opinion) exists within domain d , in this case in epistemology, among those who are both informed and competent.

Undergrad: Undergrad attends Rochester and takes classes from evidentialist Richard Feldman, instead of attending Rutgers and taking classes with reliabilist Alvin Goldman. She did not choose Rochester because of its stand on this or a related issue. However, Undergrad comes to learn of some troubling statistics. Upon graduation, 90% of those who studied at Rochester accept evidentialism and 90% of those who studied at Rutgers accept reliabilism.⁵¹ This is the case despite the fact that people at both schools read the same material, and she has no reason to think that the Rutgers students are any better or worse at philosophy than the Rochester students. As a result, Undergrad comes to doubt her current belief that evidentialism is true.

This case seems plausible, and it leads us to think that there must be something non-epistemic that accounts for the difference between the students at each university, since they are reading the same material and are similarly competent.

The epistemically irrelevant causal factor in this case is the fact that Undergrad attends Rochester, where there are a high percentage evidentialist graduates, instead of Rutgers, where there are a high percentage of reliabilist graduates. It plays no role in Undergrad’s deliberation concerning her belief. Also, Undergrad is aware of both the presence of disagreement in d and the relevant statistical correlations between communities and beliefs, which lead her to doubt the epistemic standing of her belief. The counterfactual situation has been removed. In its place is simply the actual presence of those who have the same relevant evidence and are just as competent as Undergrad, yet who disagree with her.

⁵¹ This is mere speculation. I know of no such data that exist.

One might think this seems suspiciously similar to the phenomenon of epistemic peer disagreement. In a way it does. But it is not quite disagreement. It is not the case that Undergrad takes these students to be her epistemic peers (perhaps they are, perhaps they are just close enough), and that that is what generates the resulting defeater for her belief. Instead, it is the fact that there is a high correlation between the university and the beliefs held by students at that university which makes trouble for her belief. She realizes that there is a high likelihood that she would have accepted evidentialism by attending Rochester, and she worries whether she is just another statistic, *influenced by causes that play no role in her deliberation concerning her belief*. In this way, then, this is a kind of problem of epistemically irrelevant causal factors, because it includes this causally relevant but epistemically irrelevant factor which easily might have been different.

Returning to our formal argument, we can make a moderate modification to the second premise:

(2**) S has a belief *b* concerning *d* and is a member of community C in which 90% of members believe *b*, whereas 90% of community C* believe *r* (where *r* entails not-*b*), even though members of both C and C* are similarly competent and have access to the same relevant evidence.⁵²

Again, this looks suspiciously like a description of disagreement among peers, but it is more than that. If S were not a member of either community C or C*, S could still puzzle over the disagreement between the two communities, but it would not worry S personally. It becomes a real problem for S because she is a member of community C, and now she must discern whether her belief is just a product of her environment.

⁵² It should be obvious, but the percentages used here are completely arbitrary. There is an issue of vagueness, in that it is unclear how many individuals in each community must believe *b* before it triggers the problem, but there is no space here to offer a detailed analysis, even if one were forthcoming.

This leaves open an interesting question that I am not able to answer fully at present: which way of framing the problem of irrelevant factors is more forceful?⁵³ Intuitively, perhaps, it is much more forceful when considered in the first-person, when my beliefs are at stake. But this may just be psychological. I am not sure that this makes the problem epistemically more dangerous. In fact, it appears to me that (2**) is just one way to rephrase (2*) in using different words and concepts. If it includes the unique feature of irrelevant, non-epistemic factors potentially providing a defeater for one's belief, then I am happy to include this alternative argument as a member of the same class of arguments we can refer to as the problem of epistemically irrelevant causal factors.

⁵³ Tom Senor has acknowledged (p.c.) that framing the problem in first-person terms, from the counterfactual perspective, is perhaps more forceful than when framed using statistics like this.

II. Toward a Solution to the Problem of Epistemically Irrelevant Causal Factors

In Section II. of this paper, I present some possible counterexamples to the argument for the problem of epistemically irrelevant causal factors. First, however, I wish to say more about “distorting” factors.

It is important to distinguish between instances when an epistemically irrelevant causal factor is distorting and when it is innocuous. As Katia Vavova tells it, “realizing that your belief reflects an irrelevant influence is not always cause for concern.”⁵⁴ This is because, as Roger White explains, “only a fraction of the myriad of such factors even appear to raise a challenge to the status of [your] belief.”⁵⁵ We shall call these types of factors “distorting” factors, in the sense that they distort or disrupt one's general reliability in getting the point in question right, or simply, that they provide a defeater for one's justification. As such, we are interested in determining, to the extent we are able, when in general a factor is distorting and when in general it is not.

As an unintended but welcome benefit of narrowing the scope of our argument to include only a specific type of epistemically irrelevant causal factor, much of this work has already been done. We have narrowed our problem to exclude those cases in which S is actually in an epistemically superior position. In **Scientist**, for example, we saw that Scientist simply has better evidence for her belief that $E=mc^2$. Scientist might admit that her belief that $E= mc^2$ would have been otherwise had she grown up in some primitive society, which of course is explained by the fact that she would have had different, albeit inferior, evidence leading her to believe otherwise. This shouldn't bother Scientist now, however, because in her actual situation she (i) has superior evidence which she (ii) recognizes as superior, both of which (iii) provides

⁵⁴ Vavova (ms.), p. 4.

⁵⁵ White (2010), p. 2.

her with a *prima facie* reason for thinking that her belief that $E=mc^2$ was formed on the basis of good evidence and *not* as a result of a distorting irrelevant factor. Since the worry, or driving force, behind the problem of irrelevant factors is the possibility that S's belief was formed as a result of some causal, non-epistemic influence, Scientist can rest assured if she determines, or has good reason to believe, that her belief was instead formed in the right way (e.g., on the basis of good evidence, by a reliable process, via a properly functioning faculty, etc.).⁵⁶

Now that have set aside these kinds of cases as involving non-distorting factors, we can turn back to our argument and discriminate further by noting instances when epistemically irrelevant causal factors are non-distorting even when they pass through the narrow requirements set by the argument. That is, we can attempt to identify counterexamples to the argument for the problem of epistemically irrelevant causal factors, cases in which all of the criteria set out in the argument are satisfied, yet intuitively S is still justified in believing *b*.⁵⁷ In Part A., I discuss in more detail the features of a successful counterexample. In Parts B. and C., I present possible counterexamples to the argument.

A. What is Needed for a Successful Counterexample

Recall what I take to be the strongest argument for the problem of epistemically irrelevant causal factors:

(1*) Disagreement (or difference of opinion) exists within domain *d* among those who are both informed and competent.

(2*) S's belief *b* concerning *d* easily might have been different if some factor in S's history that is (i) causally relevant to S's *b*, but (ii) epistemically irrelevant to S's deliberation about *b* had been different, even though (iii) S's evidence would have been *minimally* just as good and (iv) S would have been

⁵⁶ In saying this, I have perhaps tipped my hand for the kinds of counterexamples I have in mind.

⁵⁷ Or, "S's *b* may still be justified for S." There may be multiple ways to phrase this, but I do not think that what follows hinges on this distinction.

minimally just as competent.

(3*) S has reason to believe (1*) and (2*).

(4) If (1*), (2*), and (3*), then S's *b* concerning *d* is defeated.

(C) Therefore, S's *b* concerning *d* is defeated.

If the above argument is sound,⁵⁸ then all factors which are described by the argument should be distorting factors, those which defeat S's belief. However, if we are able to identify some factor which passes through this argument yet is intuitively non-distorting, then this would provide a counterexample to the argument. The relevant question, then, is this: *are there cases which have b's that satisfy all the conditions of this argument yet are still non-distorting?*

We must be careful in formulating our cases so that they satisfy all the criteria, as this will turn out to be a fairly difficult exercise. When thinking of examples of significant disagreements so as to satisfy premise (1*), we might bring to mind survivalists who stock up arms and food for doomsday, Holocaust deniers in the 1960s and 1970s, or the recent "anti-vaxxers" movement against vaccinating children, all of which have many proponents and many opponents. But these examples ostensibly fail to be cases in which all of those who disagree are informed or competent. In fact, they appear to be cases in which it is some small fringe group versus everyone else. Those in the minority appear to be missing all the evidence or not "all there." For example, the Holocaust denial movement was proved to be entirely discredited as more historical evidence was uncovered proving the existence of gas chambers and other atrocities. Similarly, we think of survivalists as loony and paranoid, to some degree out of touch with reality. We need rather some examples of significant disagreements which are perhaps fairly widespread and yet in which both sides are informed and competent.

⁵⁸ More accurately, since this is an argument schema rather than an argument, I should say, "if every filled-out instance of this argument is sound..."

Interestingly, many, though not all, of the kinds of cases we are looking for will involve disagreements in which the theory is *underdetermined* by the evidence. Underdetermination refers to a situation in which the evidence is insufficient to determine which theory is correct. Lawrence Sklar (1975) distinguishes between two types of underdetermination. Radical underdetermination is an underdetermination “in the evidential limit,” which is a highly contentious theoretical claim. It is to say, in our current (limited) epistemic situation, that no evidence will ever decide the issue. One might think this, for instance, as an agnostic about religious claims. On the other hand, there is a weaker, more plausible kind of underdetermination: transient underdetermination in which one is limited *given one’s current evidence*. This position allows for the possibility of improving one’s evidence to the point where it uniquely determines one theory over all the others. In these cases, there can be reasonable disagreement among those informed and competent.

One might reasonably object that cases of transient underdetermination permit us only to withhold from p rather than assent to or deny p . I sympathize with this, but I think it is reasonable to allow for disagreement even in these cases for two reasons. First, we must take into account how a theory (or proposition) coheres with a subject’s prior web of beliefs. Even though p may be presently underdetermined by S’s evidence, it may be perfectly reasonable for S to hold to p , much more reasonable than for her to withhold from p , because p coheres very well with or is a consequence of what she already knows. Second, I worry about the halt of progress if we prohibit subjects from entertaining not only p , but what follows from p . It can be difficult to see the consequences of a belief that I do not hold. If p is some great underdetermined theory of science (and it happens to be right), progress could be stunted if I cease to pursue p and the consequences it has for other theories.

Paradigmatic examples of disagreements in which each competing theory is transiently underdetermined by the evidence can be found within, but are not limited to, the domains of religion, morality, history, philosophy, and theoretical science. Consider, for example, the early- to mid-20th century debate in which several nonstandard cosmologies were proposed as alternatives to the Big Bang, all of which were underdetermined for quite some time. In 1923 Edwin Hubble measured the distance of Andromeda galaxy from earth (using a formula which takes into account the period of variability of a certain star and its intrinsic brightness), and determined that the further away objects appear to be from earth, the faster they appear to be receding from us. Based on this observational data alone, at least four cosmological models were viable. The Big Bang model posited a universe that once was at a very dense and very hot point and has since expanded and cooled. The tired light model suggested that light gets “tired” and loses energy the further it travels, and thus the universe’s expansion is merely an illusion; this change in the speed of light accounts for the discrepancy in apparent velocity between objects closer and farther from us. The Steady-State model accepted the universe’s expansion but denied that it was once denser and hotter; rather, the density has remained constant as the universe creates new matter as it expands. The Oscillating universe model proposed a universe that expanded, as the Big Bang model predicted, but also that had contracted prior to that. All four of these models were consistent with the empirical evidence, which left each theory transiently underdetermined by the evidence. It was not until 1964 when Robert Wilson and Arno Penzias detected cosmic background radiation, which was predicted by the Big Bang model alone, that the apparent symmetry in evidential support was broken.

Transient underdetermination of a theory by the available evidence often creates situations in which there are multiple parties to a disagreement, many of whom are informed and competent. So these cases will satisfy (1*).

Not all cases which satisfy (1*), however, will be cases of transient underdetermination. It may be that both sides to a disagreement are well informed and highly competent, yet perhaps one or both of them enters the disagreement with some prior, more fundamental belief. In the disagreement between functionalists and dualists in philosophy of mind, both sides may be completely informed of the arguments and empirical evidence, yet each may hold the view he does given some prior commitment to physicalism or anti-physicalism, respectively. Intuitively, neither is being irrational, despite both having access to the same relevant evidence yet coming to different conclusions. However, this intuition appears to be in direct conflict with what has been called the Uniqueness thesis, which Richard Feldman (2011) defines as follows:

[A] body of evidence justifies at most one proposition out of a competing set of propositions (e.g., one theory out of a bunch of exclusive alternatives) and...it justifies at most one attitude toward any particular proposition.⁵⁹

This would seem to cause trouble for the sort of scenario I just described. Uniqueness demands that the evidence justifies either functionalism or dualism, but not both. This seems highly intuitive as well. Given that functionalism and dualism are mutually exclusive, and assuming this is not a case of underdetermination, we should expect the evidence to fully support only one theory at most, or at least to support one theory much more strongly than the other.

These apparently conflicting intuitions are shown to be consistent if we recognize the following distinction put forward by Thomas Senor (ms.). Senor defines “Evidential Uniqueness” as:

⁵⁹ Feldman (2011), p. 148.

EU: For any proposition and evidence set E, there is a unique, specific, and objective fact of the matter regarding the degree to which E evidentially supports P.

It is this idea that Feldman captures in his thesis. EU supports our intuition that “evidence for P is necessarily evidence against not-P,” as Senor puts it.⁶⁰ We must, however, distinguish this thesis from what Senor calls “Rational Uniqueness”:

RU: For any proposition P and evidence set E, E makes rational a unique doxastic attitude (or credence level) regarding P.

This is a much stronger claim that suggests, for example, that only one doxastic attitude is rational in the disagreement between functionalists and dualists. It follows that at least one party to this disagreement is being irrational.

However, as Senor argues, RU is implausible. Although there is uniquely one credence level that correlates with the evidential support for a proposition, there may be a spectrum of credences, any of which is rational to hold with respect to a proposition P given one’s evidence E.⁶¹ This is in part due to fact that “our evidence sets are large, their contents often hard to determine on reflection, and the relevant evidential relations opaque.”⁶² If this is right, it makes sense of our example. It is true that the same relevant evidence set E justifies at most only one theory, either functionalism or dualism. But it is also true that both the functionalist and the dualist can be *rational* or *reasonable* even though they hold disparate beliefs upon the same relevant evidence E.

⁶⁰ Senor (ms.), p. 29.

⁶¹ While the spectrum may not be wide enough to cover P and not-P, it may be wide enough to cover P and withholding from P. See Senor (ms.), p. 31-32.

⁶² Senor (ms.), p. 29.

Furthermore, it seems that very complex disagreements will typically fall into this category, because they do not appear to be cases of straightforward contradiction as much as broad inconsistency. As Senor puts it:

As a consequence of what each takes the evidence to indicate, each will think that the other's perspective is wrong. But neither takes the evidence to directly show the falsity of the other's belief. It is not so much that the dispute has a "P versus not-P" structure but a "P versus R" structure where R entails not-P and P entails not-R.⁶³

Our disagreement between the functionalist and the dualist appear to be of this complex nature, as do many disagreements in philosophy, ethics, religion, and politics.

So, there are a number of cases that can satisfy (1*), but we must take care in choosing the type of case for a successful counterexample. The best candidates appear to be cases of transient underdetermination of a theory by the evidence, or cases in which there is no underdetermination but in which there is still reasonable disagreement among those who are informed and competent. Keeping this in mind, I attempt to give no less than two counterexamples to the above argument in the following two sections.

B. Good Evidence Evaluator Principle

Katia Vavova (ms.) has offered a principled way for determining when one should revise or withhold one's belief in the face of the problem of irrelevant factors. She calls this the:

Good Independent Reason Principle (GIRP): To the extent that your independent evaluation gives you good reason to think that you are unreliable with respect to matters like *p*, you must reduce your confidence in *p*.

This seems right. Suppose I hold a belief about which political theory is superior. Upon becoming aware of problem of irrelevant factors, if I have independent reason to suspect that I am unreliable with respect to politics, then I should reduce my confidence in my belief. It is

⁶³ Senor (ms.), p. 32.

important that my reason for thinking I am unreliable is *independent* of my reasons for holding my belief *b*. That is, for my reason to be independent, I must set aside what has been called into doubt (namely, *b* and my arguments for *b*). If I do not set these aside, I fall into the error of bootstrapping.

Despite its intuitive appeal, GIRP is a negative account in that it indicates to us when the problem of epistemically irrelevant causal factors is successful (i.e., provides a defeater for one's belief). We can, however, turn this idea of being unreliable on its head in an attempt to give a counterexample to the above argument, keeping in mind the need for one's support to be independent. Instead of discussing reliability, a loaded term and ambiguous when applied to individuals,⁶⁴ I prefer to speak in terms of being a good evaluator of evidence, generally-speaking. Consider the following case involving philosophical beliefs:

Reader: Reader is an upper-level philosophy major taking a seminar on the metaphysics of time in which participants read through several articles. The seminar instructor has structured the course such that students will first read through some articles on presentism, then through some articles on the static block, and lastly through articles covering the growing block. Although Reader has never covered the material for this course in depth, having taken a couple metaphysics courses before, she is familiar enough with the content to be able to map out the different views on time. In this respect, she is already ahead of many of her classmates. Once Reader gets through all the articles on presentism, she finds herself agreeing with many of the arguments. But she reads on. She next reads through all the articles on the static block and then the growing block. Towards the end of the semester, when class presentations are usually given, Reader presents one of her favorite articles—on presentism—the view that resonates with her the most. She delivers, in the few minutes allotted, a sincere appraisal of all three views, then reveals that she believes that presentism has the best case to offer. During the seminar break, a classmate points out that if Reader had read the articles in a different order, she easily might have come to a different belief instead. Perhaps she was attracted to presentism because it was the view she read about first. Perhaps this inculcated a sense of loyalty to that view at the expense of all the others. However, Reader is unfazed. She has independent

⁶⁴ In this context, “reliability” refers to a *subject's* being reliable, which seems to support a sort of virtue epistemology. In many other contexts, “reliability” refers instead to a suitable reliable *process* which leads to truth and which a subject uses, such as in reliabilism.

reason to think that, in general, she is a good evaluator of evidence. She is a model student who goes to office hours and receives high marks. Being an upper-level philosophy major, she has often seen similar dialectics play out in other debates prior to learning about this one. When writing papers, she discusses the content and her arguments with her professors to make sure she gets things clear. She usually does. Although she is aware that confirmation bias is always possible, she always strives to read each work evenly and thoughtfully, judging each argument on its own merit, as she believes she has done this time. All told, she does not think that the mere order in which she read the articles plays any deviant role in the formation of her belief.

This case satisfies all the criteria in our argument. Reader's belief, *that presentism has the best case to offer*, falls within a domain in which there is significant disagreement. And it is plausible that Reader's belief easily might have been different if some epistemically irrelevant causal factor had been different, in this case the fact that she read the articles in a different order.

Ostensibly, also, her evidence and competency would be the same no matter the order, since she is reading the same articles. Moreover, Reader is aware of both the presence of disagreement and the counterfactual possibility.

Intuitively, however, it appears that Reader is not being unreasonable in her judgment to disregard the fact that she has read one view before the others as having a distorting influence on her present belief. It would be one thing if she realized that she is generally a poor judge of arguments and a bad evaluator of evidence, frequently falling into error. If that were the case, then a reevaluation would be rational upon being made aware of her scenario, similar to what GIRP recommends. However, Reader has good independent reason to think that she is generally a good evaluator of evidence, able to distinguish concepts and critically analyze arguments on their own, all while being careful to watch for psychological biases. As such, **Reader** appears to be a counterexample to our argument.

In light of this counterexample, we may draw up a principle that captures what Reader satisfies in order to remain justified:

Good Evidence Evaluator Principle (GEEP): To the extent that S has good independent reason to think that S is generally a good evaluator of evidence in domain *d*, and that S is generally good at recognizing and avoiding cognitive error concerning *d*, S need not be concerned that *b* concerning *d* was formed as a result of an irrelevant factor.

By “S need not be concerned,” I do not mean to invoke imagery of S digging in her heels, staunchly avoiding all counterevidence. I mean simply that, in the face of this *prima facie* defeater, S has a defeater defeater, which puts to rest any epistemic anxiety she might have had due to the problem. Thus S is reasonable in continuing to believe as she does.

As GEEP indicates, in addition to having good reason to think herself a good evidence evaluator, S must also have good reason to think that she is generally good at recognizing and avoiding cognitive error concerning *d*. To see the relevance of this condition, consider the following example. Suppose that Thom is an excellent NFL football scout, having an excellent track record of finding “gems” and “sleepers” in the college ranks. Thus, Thom has good, independent reason to think that he is generally a good evaluator of evidence in the domain of college football talent scouting. However, Thom also has two sons who play college football. He assigns them high draft grades. Thom’s excellent track record notwithstanding, each of his sons are drafted high but do not meet expectations in the pros. A colleague might point out to Thom that he was likely led astray by wishful thinking or some other relevant cognitive bias. However, insofar as Thom has good, independent reason to think he is generally good at recognizing and avoiding cognitive error within the domain of college football talent scouting, it seems intuitive to me that Thom is not unreasonable to entertain other plausible explanations for his sons’ underwhelming pro performance (e.g., the player did not fit well with the system used by the coaches, the college preparation did not project well to the pros, the player was not given

sufficient opportunity, the player was unprepared for life in a big media market, etc.). Hence the need for the additional condition in GEEP.⁶⁵

One draw of GEEP is that it rewards those who are generally good at evaluating evidence and is of little to no help to those who are not. The poorer S is at evaluating evidence in some domain, the more S should be concerned that she might have erred due to the presence of an irrelevant factor. There is a problem of vagueness here, in that it is unclear how good an evaluator S must be before she can disregard an irrelevant factor. But the frustration lies in not having a precise account of vagueness; it is not a fault with evidence evaluation itself.

GEEP is a general principle, which means that S does not rely upon any evidence that she has evaluated *p* well *on this particular occasion* to support her belief that she is a good evaluator of evidence in general. (We will see in Part C a principle that captures the idea of handling the evidence well *on this particular occasion*.) This evidence also has to have independent support, as we see in Reader's case. Reader does not appeal to her arguments for presentism in her defense; rather, she points to her "track record" as a good evaluator of evidence in general when it comes to philosophical arguments. (As I note, she also happens to believe that she has, this time, evaluated the arguments well, but this conviction is not what is driving her reply to her classmate. It is instead the conviction that she knows how to evaluate arguments well.)

Returning to our case, despite the fact that Reader has read one views before the others, she can disregard this as having any distorting influence on her present belief, because she has good reason to think that she is generally a good evaluator of the evidence. Of course, this does not preclude the possibility, however a slight one, that reading the articles in a certain order has

⁶⁵ Thanks to Jack Lyons and Eric Funkhouser for drawing my attention to the need for an additional condition in GEEP. The example I give here is a modified version of one given by Tom Senor.

led her to form a false belief. But I believe that perhaps she is reasonable to disregard this irrelevant, or non-epistemic, factor as having an adverse effect based on what she knows about how she generally responds to evidence. Perhaps her classmate could show her where she has gone wrong, or offer some counterarguments with which she is unaware. But then this would be to bring in new evidence, which of course might easily defeat her belief. But the irrelevant, or non-epistemic, factor of the order in which she read the articles is just not strong enough of a defeater to take seriously.

In reply to this counterexample, one might object with the following dilemma.⁶⁶ On the first horn, **Reader** does not fully satisfy the criteria of our original argument, because if Reader is really that good at evaluating evidence, it is *not* the case that she easily might have believed otherwise. On the second horn, if Reader satisfies the criterion in that she *easily* might have believed otherwise, then we cannot say that she is generally a good evaluator of evidence. It may be surprising, but I think that this is a terrific objection. Of course, I do not want to take the second horn, because that would be to accept the success of the original argument. The criteria of the argument are satisfied, and there is no longer a counterexample. However, the first horn is appealing. Were I to take it, I would essentially admit that my counterexample here does not successfully go through, since all the criteria from the original argument are not satisfied. But that is not a terrible thing, since we are describing one way in which S can escape the argument—just be a good evaluator of the evidence in general, and you will not have easily believed otherwise. If this is right, then GEEP remains a viable principle.

⁶⁶ Thanks to Tom Senor for bringing this issue to my attention.

C. Well-Foundedness Principle

Let us return to our previous question: *are there cases which have b's that satisfy all the conditions of this argument yet are still non-distorting?* Consider another case which also involves philosophical beliefs.

Philosopher: Philosopher believes that free will is incompatible with determinism. Upon reflection, however, Philosopher wonders whether perhaps he has arrived at this belief in part due to non-epistemic factors. After all, he did not have the best relationships with his professors in graduate school, one of which held the opposite view in fact. Realizing that this could have caused him to believe the view he now does, Philosopher imagines whether he would have held the same view if things had gone differently, perhaps if he had gotten along better with this certain professor. He does not worry that he lacks some evidence that this professor has—being the good student that he is, he knows that he would have learned just as much and would have been just as competent as he now is. The worry, rather, is whether his relationship with his professor has had a distorting influence on his belief. All in all, Philosopher decides that he has no reason to be concerned about the etiology of his belief. He knows that he has very good evidence for his belief, since he has encountered a very many of the arguments for and against his position, and he is convinced that the available evidence supports his view better than it does the alternative.

This case satisfies all the criteria in our original argument. Philosopher's belief falls within a domain in which there is significant disagreement. And Philosopher's belief easily might have been different if some epistemically irrelevant causal factor had been different—in this case the sour relationship with a professor from graduate school who holds an opposing view—even though he would have had access to the same relevant evidence and would have been just as competent. Furthermore, Philosopher is aware of both the presence of disagreement and the counterfactual possibility.

Intuitively, however, it appears to me that Philosopher is not being unreasonable in his judgment to disregard the relationship with the professor from graduate school as having a distorting influence on his present belief. Even if Philosopher's belief was initially formed as a result of this influence, this does not mean that it is now unjustified. To say that a belief is

unjustified simply because of how it originated is to commit the genetic fallacy. Rather, he has since come to support his belief with good evidence and it is now epistemically justified. Hence we have another counterexample.

From this, let us draw up a principle that captures what is going on in **Philosopher**, such that it allows him to be justified in his belief.

Well-Foundedness Principle (WFP): To the extent that S's belief *b* is well-founded on good evidence, S need not be concerned that *b* was formed as a result of an irrelevant factor.⁶⁷

As we saw with GEEP, WFP will not provide much benefit in cases in which S has poor evidence for her belief. The poorer the evidence or the weaker the support, the more reason for S to be concerned. Second, WFP does not preclude that S's belief was *initially* formed on the basis of some non-epistemic, or irrelevant, factor. It may have been. But again, it matters what evidential support S now has and how S responds to it that determines whether S should consider this epistemically irrelevant causal factor to have any negative, or distorting, influence. As WFP indicates, to the extent that S's belief is in fact well-founded on good evidence, this entails that S bases her belief on good evidence and the belief not epistemically unjustified as a result of a distorting irrelevant factor.

To understand this claim more fully, we must look at what well-foundedness is. Richard Feldman and Earl Conee define well-foundedness as the following:

- WF S's doxastic attitude *D* at *t* toward proposition *p* is well-founded if and only if
- (i) having *D* toward *p* is justified for S at *t*; and
 - (ii) S has *D* toward *p* on the basis of some body of evidence *e*, such that
 - (a) S has *e* as evidence at *t*;
 - (b) having *D* toward *p* fits *e*; and
 - (c) there is no more inclusive body of evidence *e* had by S at *t* such that having *D* toward *p* does not fit *e*.⁶⁸

⁶⁷ This strategy is similar to Kelly's (2010) view in the disagreement literature, in which S has higher-order evidence that defeats S's first-order evidence of disagreement.

As Alvin Goldman notes, Conee and Feldman’s account of well-foundedness is “their way of expressing the notion of *doxastic*, as opposed to *propositional*, justification (as Conee indicates in personal communication).”⁶⁹ Propositional justification concerns whether the subject has epistemically good reasons for her belief, whereas doxastic justification concerns the way in which the subject arrives at her belief. For example, Nakia might have good evidence for believing that Oliver will enjoy the new *Llama Llama* book she bought for him; he has six *Llama Llama* books already, and Oliver often chooses those books over many others to read before bedtime. However, Nakia’s belief is based on wishful thinking instead of her evidence. Her belief is propositionally justified, but it is not doxastically justified, because her belief is not based on the good evidence she has.

Returning to our earlier case, Philosopher not only has plausible arguments for his belief that free will is incompatible with determinism, which justifies his belief, he also *appreciates* the evidence and its epistemic merit in supporting his belief such that his belief is based on these plausible arguments, which makes his belief also doxastically justified, or well-founded. Philosopher’s belief being well-founded on good evidence does not absolutely preclude that his belief is now influenced by a distorting irrelevant causal factor; however, it does give Philosopher good reason to think that it is not. This is just what WFP captures. To the extent that Philosopher’s belief is well-founded on good evidence, it is highly likely that Philosopher has gotten the point in question right—i.e., has knowledge (although, in keeping with fallibilism, there is always the epistemic possibility that he is wrong). Given all this, Philosopher need not worry about the possibility that his belief was influenced in an epistemically bad way by his sour

⁶⁸ Conee and Feldman (2004/1985), p. 93.

⁶⁹ Goldman (2011), Section 4.

relationship with his professor in graduate school. That is, all told, it is not strong enough to plausibly be considered a genuine defeater for his belief. If there are arguments he has not heard or evidence he has not seen, these would provide plausible *prima facie* defeaters because it is introducing new evidence. However, once he takes stock of his whole situation and realizes that his belief is well-founded, Philosopher need not worry that it was influenced deviantly by his relationship with his professor.

Additional support for this is found when we consider it explicitly in counterfactual language. Suppose S has at least the same relevant evidence in both the actual and counterfactual situations. So it would seem that if S forms his belief on the basis of good evidence in the actual situation, then he would do the same in the counterfactual situation. But this is not so. It is entirely plausible to think that the same set of evidence and arguments does not provide the same degree of epistemic support to both p and $\text{not-}p$. While p may be based on good evidence E, it does not follow that that same good evidence E for p will be good evidence for $\text{not-}p$. E may in fact be ill-suited for supporting $\text{not-}p$. Therefore, given that Philosopher has plausible arguments for his belief, and given that his belief is based on these plausible arguments, he can be reasonable in thinking that the same degree of support would not hold for the opposite belief in the counterfactual scenario, and he can continue believing as he has.

III. Conclusion

In this paper I have given an analysis and explanation as to what exactly is the problem of epistemically irrelevant causal factors. We have seen it in narrative cases as well as in a formal argument. Although there are plenty of related epistemological phenomena nearby, such as disagreement, the problem can be distinguished from those as *sui generis*. Furthermore, while it can be helpful to structure the problem as a counterfactual possibility, it is not necessary, since what primarily generates the problem is awareness of a non-epistemic irrelevant factor itself, and examples can be constructed that capture this criterion while avoiding counterfactual language.

I have also attempted to develop counterexamples to the problem of epistemically irrelevant causal factors, which allow a subject to meet all the criteria of the argument yet remain justified in his belief. First, being generally good at evaluating evidence with respect to the relevant domain will help—even if this principle ends up not being a true counterexample, it remains an avenue for escaping the problem. Second, realizing that one’s belief is well-founded will provide a way to avoid the defeater conclusion that the problem seeks.

What may not be apparent in the foregoing analysis is that, in my estimation of the problem of epistemically irrelevant causal factors, I do not think that it provides one with a *strong* reason to think one has erred. Of course it provides a *prima facie* reason, but I think that it is a weak one, easily defeasible. As such, I assume a sort of epistemic permissivism, akin to Miriam Schoenfield’s (2014) strategy, in which one largely is reasonable in continuing to believe as one does unless and until presented with strong evidence to the contrary. Much like those in the disagreement literature who do not consider an instance of epistemic peer disagreement to be very “weighty” when compared to, say, one’s higher-order evidence that one has evaluated the

evidence well on this particular occasion,⁷⁰ I do not consider an instance of an epistemically irrelevant causal factor to be very “weighty” when compared to instances of candidate principles such as GEEP and WFP.

⁷⁰ See Kelly (2010).

Bibliography

- Ballantyne, Nathan (2012). "The Problem of Historical Variability." In *Disagreement and Skepticism (Routledge Studies in Contemporary Philosophy)*. Edited by Diego Machuca. New York: Routledge, 239-59.
- Bogardus, Tomás (forthcoming). "Only All Naturalists Should Worry about Only One Evolutionary Debunking Argument." *Ethics*.
- (2013). "The Problem of Contingency for Religious Belief." *Faith and Philosophy* 30:4, 371-92.
- Carroll, Lewis (1995). "What the Tortoise Said to Achilles." *Mind* 104: 416, 691-3.
- Cohen, G. A. (2000). *If You're an Egalitarian, How Come You're So Rich?* Cambridge, MA: Harvard University Press.
- Conee, Earl (1992). "The Truth Connection." *Philosophical and Phenomenological Research* 52:3, 657-69.
- Conee, Earl and Richard Feldman (2004/1985). "Evidentialism." In *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press, 83-107.
- Dworkin, Ronald (1996). "Objectivity and Truth: You'd Better Believe It." *Philosophy and Public Affairs* 25:2, 87-139.
- Elga, Adam (ms.). "Lucky To Be Rational." *Bellingham Summer Philosophy Conference*, Bellingham, WA, August 6, 2008. URL = <http://www.princeton.edu/~adame/papers/bellingham-lucky.pdf>.
- Feldman, Richard (1988/2004). "Having Evidence." In *Evidentialism*. Edited by Earl Conee and Richard Feldman. New York: Oxford University Press, 219-41.
- (2000). "The Ethics of Belief." *Philosophy and Phenomenological Research* 60, 667-95.
- (2011). "Reasonable Religious Disagreements." In *Social Epistemology: Essential Readings*. Edited by Alvin Goldman and Dennis Whitcomb. New York: Oxford University Press, 137-57.
- Flew, Antony (1976). "The Presumption of Atheism." In *The Presumption of Atheism and Other Philosophical Essays on God, Freedom, and Immortality*. New York: Barnes & Noble.

- Garber, Daniel (2007). “*Religio Philosophi*: Some Thoughts on God, Reason, and Faith.” In *Philosophers Without Gods: Meditations on Atheism and the Secular Life*. Edited by Louise Antony. New York: Oxford University Press, 32-40.
- (2009). “What Happens After Pascal’s Wager? Living Faith and Rational Belief.” *Aquinas Lecture 73*. Marquette University Press.
- Goldman, Alvin (1999). “Internalism Exposed.” *Journal of Philosophy* 96, 271-93.
- (2011). “Reliabilism.” *The Stanford Encyclopedia of Philosophy*. Fall 2011 Edition. Edited by Edward N. Zalta. URL = <http://plato.stanford.edu/archives/spr2011/entries/reliabilism/>.
- Hick, John (1997). “The Epistemological Challenge of Religious Pluralism.” *Faith and Philosophy* 14, 281.
- Joyce, Richard (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press.
- Kelly, Thomas (2010). “Peer Disagreement and Higher Order Evidence.” In *Disagreement*. Edited by Richard Feldman and Ted Wardfield. Oxford: Oxford University Press, 111-74.
- Leiter, Brian (2004). “The Hermeneutics of Suspicion: Recovering Marx, Nietzsche, and Freud.” *The Future for Philosophy*. Edited by Brian Leiter. Oxford: Oxford University Press, 74-105.
- Loftus, John (2010). *The Outsider Test for Faith: How to Know Which Religion is True*. Amherst, New York: Prometheus Books.
- Plantinga, Alvin (1995). “Pluralism: A Defense of Religious Exclusivism.” In *The Rationality of Belief and the Plurality of Faith*. Edited by Thomas Senor. Ithaca, NY: Cornell University Press, 191-215.
- (2000). *Warranted Christian Belief*. New York: Oxford University Press.
- Pollock, John (1986). *Contemporary Theories of Knowledge*. Towota, NJ: Roman & Littlefield Publishers.
- Rey, Georges (2013). “The Analytic/Synthetic Distinction.” *The Stanford Encyclopedia of Philosophy*. Fall 2013 Edition. Edited by Edward N. Zalta. URL = <http://plato.stanford.edu/archives/fall2013/entries/analytic-synthetic/>.
- Rosen, Gideon (2001). “Nominalism, Naturalism, Epistemic Relativism.” *Philosophical Perspectives* 15, *Metaphysics* 69-91.

- Schechter, Joshua (ms.). "Luck, Rationality, and Explanation: A Reply to Elga's 'Lucky to Be Rational'." *Bellingham Summer Philosophy Conference*, Bellingham, WA, August 6, 2008. URL = <<http://philpapers.org/rec/SCHLRA-2>>.
- Schoenfield, Miriam (2014). "Permission to Believe: Why Permissivism is True and What it Tells Us About Irrelevant Influences on Belief." *Noûs* 48:2, 193-218.
- Senor, Thomas (ms.). "Perception, Evidence, and the Epistemology of Disagreement." Unpublished paper. (Cited with permission from the author.) URL = <<http://philpapers.org/rec/SENPEA>>.
- Sher, George (2001). "But I Could Be Wrong." *Social Philosophy and Policy* 18, 64-78.
- Sklar, Lawrence (1975). "Methodological Conservatism." *Philosophical Review* 60, 20-43.
- Street, Sharon (2006). "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127:1, 109-66.
- van Inwagen, Peter (1994). "*Quam Dilecta*." In *God and the Philosophers*. Edited by Thomas V. Morris. New York: Oxford University Press, 31-60.
- (1995). "Non Est Hick." In *The Rationality of Belief and the Plurality of Faith: Essays in Honor of William P. Alston*. Edited by Tom Senor. Ithaca, NY: Cornell University Press, 216-41.
- Vavova, Katia (2010). *Rational Humility and Other Epistemic Killjoys*. (Doctoral dissertation). URL = <<http://dspace.mit.edu/bitstream/handle/1721.1/62421/710986677.pdf?...1>>.
- (ms.). "Irrelevant Influences." Unpublished paper. (Cited with permission from the author). URL = <http://www.mtholyoke.edu/~evavova/files/vavova_ii.pdf>.
- White, Roger (2010). "You Just Believe That Because..." *Philosophical Perspectives* 24, 573-615.