

Rigidity, Symmetry and Defeasibility: On Weisberg's Puzzle for Perceptual Justification

Juan Comesaña

1. Introduction

Joathan Weisberg has presented a puzzle for Bayesian epistemologists (Weisberg (2009), Weisberg (2015)). The puzzle highlights the inability to model a specific kind of defeasibility of evidence in a Bayesian framework. But the puzzle is of much wider interest: it arises not just for Bayesians, and not just for formal modelers, but for any theory which adopts three very plausible principles: the rigidity, symmetry, and defeasibility of evidence. In this paper I first present Weisberg's puzzle as originally developed, in the context of Bayesian epistemology. Then I show how the problem arises for normal epistemology as well.¹ I offer a solution for the normal puzzle and then carry over that solution to the formal side. I end up by suggesting that taking the puzzle seriously means facing up to the possibility that we cannot give an informative, general characterization of epistemic justification, and that we should instead rest content with a kind of particularism.

2. Bayesianism and Defeaters

Bayesianism in epistemology takes degrees of belief, or credences, as the doxastic attitude to be modelled. Credences are formally represented by a function, Cr , which takes propositions as arguments and delivers real numbers as values. Traditional Bayesianism imposes two normative constraints on credences: at any given time, rational credences satisfy the probability calculus (call this constraint "Probabilism"), and, upon receiving evidence E , rational credences evolve by conditionalization on E (call this constraint "Conditionalization"). Credences obey the probability calculus if and only if all propositions are assigned credences between 1 and 0, logical truths are assigned credence 1, and the credence assigned to a disjunction whose disjuncts are logically incompatible is the sum of the credences assigned to each disjunct. The conditional probability of a proposition P given another proposition E is defined as follows: $Cr(P|E) = \frac{Cr(E \wedge P)}{Cr(E)}$ (provided that $Cr(E) > 0$). In this traditional Bayesian framework, which propositions a subject has as evidence at a given time is taken as an exogenous input, on which the framework is silent. Once that input is set, Conditionalization is the norm according to which, for any proposition P , $Cr_{new}(P) = Cr_{old}(P|E)$, where Cr_{old} is the subject's credence function

¹ The use of "normal" as a contrast for "formal" epistemology is from James Van Cleve.

before acquiring evidence E and Cr_{new} is the subject's credence function after acquiring evidence E .

If Probabilism and Conditionalization are the only two constraints on credences, then once a proposition is part of a subject's evidence two consequences follow: that proposition receives credence 1, and it does so forever. These two consequences follow easily from the definition of conditional probability just given. Many Bayesians, most prominently Jeffrey, saw these consequences as bugs of the traditional Bayesian formalism. Jeffrey proposed to replace Conditionalization with a generalization of it (Jeffrey (1965)). A famous theorem of the probability calculus is the Law of Total Probability, according to which the probability of a proposition is a weighted average of its conditional probability on a partition. Thus, if the E_i are a partition (i.e., if one E_i is logically guaranteed to be true, and if they are pairwise logically incompatible), then: $Cr(P) = \sum_i Cr(P|E_i)Cr(E_i)$. Jeffrey's proposal involves two modifications to standard Bayesianism. First, instead of conceiving the impact of experience as making a proposition certain, he suggests that we conceive of it as changing the credences over a partition. Thus, whereas on the standard Bayesian account a glimpse at a cloth under sub-optimal lighting conditions must have as a consequence that some proposition is assigned credence 1 (with the attendant problem of figuring out which proposition that is), on Jeffrey's account the consequence can rather be a change over the partition constituted by the propositions that, say, the cloth is green, red, or some other color. If the cloth looks rather reddish-green, then maybe my credences in the propositions that it is red and that it is green should go up, whereas my credence in the proposition that it is some other color should go down, without any of these propositions receiving credence 1 or 0. Second, instead of propagating the impact of the newly acquired evidence by Conditionalization, it should be done by Jeffrey-Conditionalization. Jeffrey-Conditionalization on a partition E_i for any proposition P is defined as follows: $Cr_{new}(P) = \sum_i Cr_{old}(P|E_i)Cr_{new}(E_i)$. It is easy to see that Jeffrey-Conditionalization is a diachronic form of the Law of Total Probability, where the conditional probabilities are the ones the subject had before acquiring the evidence, but the weights are given by the subject's credences in the relevant partition after acquiring the evidence. It can also be seen that Jeffrey-Conditionalization coincides with traditional Conditionalization in the special case where the evidence does bring one member of the partition to 1 (and the remaining to 0). Call Jeffrey's version of Bayesianism "Liberal Bayesianism," in contrast with "Traditional Bayesianism".

Liberal Bayesianism lacks the two consequences we noted for Traditional Bayesianism. First, experiences need not make any proposition certain. Second, the impact of an experience on evidence can be changed by another experience. If I take a second, longer look at the cloth, and it looks now almost definitely green, then I can change my credences over the partition in question again. But both versions of Bayesianism are vulnerable to another problem: the problem of the defeasibility of evidence.

John Pollock gave the canonical account of defeaters in contemporary epistemology (see, for instance, *Knowledge and Justification* 1974)). Pollock's framework is not Bayesian: it deals with beliefs instead of credences. His primitive is that of *possible doxastic justification* (as we might call it): it's possible for a subject to be justified in believing a proposition on the basis of being in a certain mental state. Some of those mental states will be doxastic:

beliefs of the subject. Importantly, however, some mental states will be non-doxastic: experiences of the subject. Indeed, in Pollock's framework experiences are foundational: they (defeasible) justify certain beliefs, and then those beliefs can themselves justify further beliefs. Thus, the structure of justification for Pollock is recursive: the justification of an empirical belief can eventually be traced back to some experience. Experiences themselves are neither justified nor unjustified: they just are. A mental state is a reason for believing a proposition just in case it is possible to be justified in believing that proposition on the basis of being in that mental state. Defeasible justification (whether by experiences or by other justified beliefs) can of course be defeated. Defeaters, like reasons, are also mental states. Defeaters, according to Pollock, are always defeaters of mental states as reasons for believing. Thus, if M is a reason to believe P (i.e., if it is possible to be justified in believing P on the basis of M), then $M +$ is a defeater for M as a reason to believe P if and only if the combined mental state $M \wedge M +$ is not a reason to believe P (i.e., it is not possible to become justified in believing P on the basis of the combined mental state $M \wedge M +$).

Famously, Pollock distinguished between rebutting and undercutting defeaters. Rebutting defeaters (for M as a reason to believe P), in addition to being defeaters, are also reasons to believe $\neg P$. Thus, that Mary says that P is false is a rebutting defeater for Peter's testimony that P as a reason to believe P .² Undercutting defeaters (for M as a reason to believe P) need not be (at least not in the first instance—more on this momentarily) reasons to believe $\neg P$. Rather, they are reasons to *disbelieve* that M would not be true unless P were true—in other words, they are reasons to believe that it might be the case that M is true and P is false. Thus, when my experience as of a red wall in front of me is a reason to believe that there is a red wall in front of me, believing that the lights themselves are red is an undercutting defeater for that reason. It is not (at least not directly) itself a reason to believe that the wall is not red, but rather a reason to think that it might well happen that the wall looks red but is not red. Whereas rebutting defeaters “attack the conclusion,” undercutting defeaters “attack the connection between the premises and the conclusion” (the terminology is Pollock's). When undercutting defeaters do their work, however, they do have the effect of decreasing the subject's confidence in P . If a subject starts out disbelieving P , then believes P on the basis of M , and later acquires an undercutting defeater $M +$, the net effect will be that the subject goes back to disbelieving P . Thus, even though undercutting defeaters are perhaps not *direct* evidence against P , they are indirect evidence against it.

Can defeaters be accommodated in a Bayesian framework? The answer is complicated. We have already seen that Pollock distinguishes between rebutting and undercutting defeaters. But there are two more distinctions, giving us eight kinds of defeaters in total (we will deal

² Here I am referring to the content of the belief (namely that Mary says that P is false) rather than to the belief itself (my belief that Mary says that P is false) as the reason. Pollock himself does this repeatedly for the case of doxastic reasons and defeaters. Moreover, the content itself is a better candidate for being the reason, and Pollock's framework obscures the distinction between reasons and their possession. For more on contents as reasons and mental states as reasons-providers, see Comesaña (2020).

with six of them in this section, leaving the remaining two for the next). Like reasons, defeaters too come in doxastic (beliefs) and non-doxastic (experiences) varieties. An experience as of a cat in front of you can be a rebutting defeater for your friend's testimony that her pet is a dog. And an experience as of red lighting can be an undercutting defeater for your experience as of a red wall in front of you (as a reason to trust this experience). Finally, there are what I shall call endogenous and exogenous defeaters. Endogenous defeaters are the ones we are familiar with from Pollock's work. Exogenous defeaters are at the heart of Weisberg's puzzle. In this section, I deal only with endogenous defeaters.

Let us start by seeing how doxastic defeaters can be accommodated in a Bayesian framework. Both the basic notion of a reason as well as the notion of a defeater will be encoded in a Bayesian framework in conditional probabilities. Doxastic reasons are straightforward: we can capture the idea that belief in P is a reason to believe Q as follows: $Cr(Q|P) > Cr(P)$. For the generic notion of a defeater, we have it that D is a defeater for P as a reason to believe Q just in case $Cr(Q|P \wedge D) \leq Cr(Q|P)$. If D is a rebutting defeater, then in addition to satisfying the generic notion just defined it will also be the case that $Cr(P|D) < Cr(P)$. Undercutting doxastic defeaters are harder to translate to a Bayesian framework, but not impossible—the key lies in working with a measure of evidential support, and the interested reader is referred to Kotzen (2019).

Now, how do we represent non-doxastic reasons and defeaters into a Bayesian framework? The only non-doxastic state Pollock has talked about to any measurable extent are experiences, so the question is: how do we represent experiences (or their epistemic impact) into a Bayesian framework? The answer here is complicated. Perhaps the most straightforward way of incorporating experiences into a Bayesian framework is via a proposition to the effect that the subject has the experience. Thus, let e be the proposition that the subject has an experience with the content that there is a red wall in front of him. We can then represent the claim that the experience is a reason to believe that the wall is red (W) as follows: $Cr(W|e) > Cr(W)$, and we can also define defeaters (both rebutting and undercutting) in a way analogous to the doxastic kind. But this representation requires the subject to have credences about experiences. Many epistemologists think of experiences as providing justification *directly*, without the need for a mediating belief (Pollock himself thinks this, as does Pryor (2000)). But the Bayesian representation of non-doxastic reasons and defeaters via a proposition to the effect that the subject has an experience need not be thought of as *requiring* belief in that proposition for the experience to justify belief. Rather, the experience by itself does the justifying, but the subject is nevertheless justified in believing that he is undergoing that experience whenever it justifies him. Alternatively, we can leave the experience altogether outside of the realm of the Bayesian formalism, and start directly with the proposition belief in which the experience directly justifies.³

³ For more on the issue of how to represent non-doxastic reasons in a Bayesian framework, see Pryor (2013).

Suffice it to say that it is not obvious how to represent non-doxastic defeaters in a Bayesian framework. For the purposes of this paper, however, I am going to assume that that particular problem is solved. I will do this for two reasons. First, whether endogenous non-doxastic defeaters can be represented in a Bayesian framework or not, the exogenous defeaters which give rise to Weisberg's puzzle are a much more pressing problem for Bayesianism. And second, I will argue that Weisberg's puzzle generalizes, and that the generalization is a problem not only for Bayesians, but for everybody. Whether or not endogenous non-doxastic defeaters can be represented in a Bayesian framework, then, we have bigger fish to fry.

3. Exogenous Defeaters and Weisberg's puzzle

Endogenous defeaters are the ones we are familiar with from Pollock's pioneer work on the topic. They are defeaters for other mental states as reasons to believe some proposition. In both Pollock's as well as the two Bayesian frameworks we are considering, something will play the role of basic evidence at a time. In Pollock's framework, our basic evidence at a time consists of the experiences we have at that time. In Traditional Bayesianism, our basic evidence at a time consists of the (empirical) propositions that are assigned credence 1 at that time. In Jeffrey's Liberal Bayesianism, our basic evidence at a time consists of credences in the partition affected by our experience at that time. Exogenous defeaters, then, are defeaters for our basic evidence itself, rather than defeaters for that evidence as reasons to believe other propositions. How plausible it is that there are exogenous defeaters depends on what one takes our basic evidence to be.

The following kind of story must be accounted for in any plausible epistemology. You open your eyes and you have an experience as of a red wall in front of you. On that basis, you believe that the wall in front of you is red. Some time later, you learn that there are red lights shining on the wall. On that basis, you cease to believe that the wall in front of you is red. Let us start by assuming that your evidence is constituted by a partition over the possible colors of the wall. In that case, Bayesianism (of either kind) cannot account for the story.

It is easy to see that Traditional Bayesianism won't be able to account for the story under that assumption. According to Traditional Bayesianism, when you open your eyes and have the experience, you acquire as evidence the proposition that the wall in front of you is red. Thus, you assign it credence 1. By virtue of Conditionalization, you keep on assigning credence 1 to this proposition no matter what else you learn. Thus, even once you learn that there are red lights shining on the wall, you remain fully confident that the wall in front of you is red, and thus Traditional Bayesianism gives us the wrong result.

An obvious reaction to this is to say that what is responsible for the wrong result is not Traditional Bayesianism, but rather the assumption about what our basic evidence is in that case. If our basic evidence is not that the wall is red but rather, say, that the wall looks red to us, then the problem doesn't arise. I am going to assume that the way to capture the idea that our basic evidence is that the wall looks red in Traditional Bayesianism is to say that the proposition that the wall looks red gets assigned credence 1. As I said before, if that

is not the right way to proceed here, then it is not at all clear what is the right way to proceed, and so Traditional Bayesianism has even harder problems in this area. With that assumption in place, it is true that Traditional Bayesianism can account for the story. For now, once you open your eyes you assign credence 1 to the proposition that the wall looks red. You conditionalize on that and come to assign some high (but less than 1) credence in the proposition that the wall is red. Later, you acquire as evidence the proposition that there are red lights shining on the wall, and so you conditionalize on the conjunction *the wall looks red and there are red lights shining on the wall*, with the result that your credence in the proposition that the wall is red goes back down. Thus, the story is accounted for.

That particular story is indeed accounted for, but it should be clear that the problem remains. For just as it was easy to find a proposition that works as a defeater for the claim that the wall is red, it is also easy to find a proposition that works as a defeater for the claim that the wall looks red to you. Following Weisberg, let that proposition be that you just received word from the cognitive psychology lab that you are very bad at telling how things look to you. When you open your eyes, you have as evidence that the wall looks red, but after reading the report you should no longer treat that proposition as evidence.⁴ For the same reason that Traditional Bayesianism couldn't account for the original story when we assumed that the evidence was the proposition that the wall is red, it cannot account for this modified story when we assume that the evidence is that the wall looks red.

One could, of course, keep internalizing the evidence. Maybe the basic evidence is not that the wall looks red to you, but rather that it seems to you that the wall looks red to you. From the point of view of someone like myself, for whom even the first internalizing move was a mistake, these further moves look just like a desperate attempt to save the theory. And they are desperate because it is pretty obvious how to construct defeaters even for these extra-internalized conceptions of evidence: let the report from the lab tell you that you are a bad judge of how it seems that things look to you. For the internalizing reply to succeed, it needs to find propositions which satisfy two constraints: first, it is plausible that they are our basic evidence, and second, they need to be indefeasible. I won't argue here that there are no propositions meeting these two constraints, but I will follow the consensus that there aren't.

Maybe Liberal Bayesianism can help? After all, it looks as if the problem for Traditional Bayesianism lies in the fact that, according to it, evidence is forever: once a proposition is part of your evidence, it stays part of your evidence. Liberal Bayesianism was designed precisely in order to get rid of this aspect of Traditional Bayesianism. But the problem still arises for Liberal Bayesianism. This is because Jeffrey-Conditionalization, just like Conditionalization, obeys a **Rigidity** principle. Informally, **Rigidity** is the claim that what is

⁴ Maybe if you have no opinions as to where your belief that the wall is red came from, learning that there are red lights shining on it should not affect your view of its color (for a complaint about Pollock's notion of undercutting defeater along these lines, see Sturgeon (2014), and cf. Casullo (2018)). But then the problem will arise for those cases where you *do* have opinions as to the origins of your beliefs.

evidence for what is independent of whether anyone has that evidence. In a Bayesian framework, **Rigidity** can be stated in terms of conditional credences. For Traditional Bayesianism, this is the claim that if $Cr(P|E) = n$ before the subject acquires E as evidence, then $Cr(P|E) = n$ after the subject acquires E as evidence. It is easy to see that Traditional Bayesianism satisfies **Rigidity**: once the subject acquires E as evidence, he conditionalizes on it, and so $Cr(P) = n$ and $Cr(E) = 1$. But the conditional credence of any proposition on a certain proposition is the same as its unconditional credence, and so $Cr(P|E) = n$ after receiving E as evidence. Jeffrey-Conditionalization too obeys **Rigidity**, this time by brute force. Indeed, it is part of the statement of Jeffrey-Conditionalization that the experience affects the subject's credence distribution over a partition E_i *and nothing more*—in particular, the experience must leave the conditional probabilities $Cr(P|E_i)$ the same as they were before, for every P and each E_i .

Here's the central insight of Weisberg's puzzle: **Rigidity** itself is incompatible with the defeasibility of evidence in a Bayesian framework. To see why, we need only note one additional feature of that framework: independence is symmetric. Independence is defined in terms of conditional probabilities, where P is independent of Q if and only if $Cr(P|Q) = Cr(P)$. The symmetry of independence guarantees that evidential connections, encoded in conditional probabilities, are also symmetric. For if E is evidence for P , this is encoded as the following inequality: $Cr(P|E) > Cr(P)$. And if that is true, then given symmetry $Cr(E|P) > Cr(E)$, which means that P is evidence for E .

With **Rigidity** and the symmetry of independence at hand, we can prove that evidence is indefeasible even in the Liberal Bayesian framework. Remember that, before opening your eyes, you have a certain credence distribution over the possible colors of the wall—maybe you are pretty sure that it is either red or white, but leave some credence for the possibility that it is some other color. Before you open your eyes, these credences of yours over this partition are independent of the color of the lights. Let us suppose that you are pretty sure that the lights are white, but leave some little credence for the possibility that they are red, and even less credence for the possibility that they are some other color. Your credences in the color of the lights is independent of your credences in the color of the wall. In particular, you think that whether the wall is white or not has nothing to do with whether the light is red or not (perhaps you know that the color of the wall was decided by the flip of a fair coin, and the color of the lights by another such flip). So:

1. $Cr_b(LR|E_i) = Cr_b(LR)$, for each E_i ;

where Cr_b is your credence distribution before you open your eyes, LR is the proposition that the lights are red, and the E_i are the propositions which form the partition that will be affected by the experience you will have when you open your eyes (in our example, the E_i are the propositions that the wall is white, that the wall is red, and that the wall is some other color). Given symmetry, 1 implies:

2. $Cr_b(E_i|LR) = Cr_b(E_i)$, for each E_i .

1 and 2 together say that, before opening your eyes, you take the color of the wall and the color of the light to have nothing to do with each other. Now, given **Rigidity**, 1 implies:

3. $Cr_a(LR|E_i) = Cr_b(LR)$, for each E_i ;

where Cr_a is your credence distribution after you open your eyes. Remember that the Law of Total Probability establishes that the probability of any proposition is a weighted value of its conditional probabilities over a partition. But 3 tells us that each of the $Cr(LR|E_i)$ has the same value, namely $Cr_b(LR)$. Therefore:

4. $Cr_a(LR) = Cr_a(LR|E_i)$, for each E_i .

Applying symmetry once more gives us:

5. $Cr_a(E_i|LR) = Cr_a(E_i)$, for each E_i ;

which means that learning that lights are red cannot affect your view about the color of the wall even after you opened your eyes.

That is Weisberg's puzzle: **Rigidity** entails the indefeasibility of evidence itself. But evidence obviously is defeasible. Given that rigidity is built right into Bayesian (and other formal⁵) frameworks, they are structurally unable to represent the defeasibility of evidence.

4. Generalizing Weisberg's puzzle

We saw that Pollock's framework differed from the Bayesian one in a manner of ways. Two of the more important ones are that Pollock deals with full beliefs rather than credences, and that he gives a prominent place to experiences. However, Pollock's account shares with Bayesianism a structural feature: basic empirical evidence is taken as an exogenous input to the system. In the Bayesian framework that empirical evidence takes the form of credences over a proposition or a partition, whereas for Pollock it is constituted by the experiences themselves. But what Pollock's framework shares with the Bayesian framework is that in neither is the basic empirical evidence itself defeasible. We saw how that happens on the Bayesian framework. For Pollock, the indefeasibility of basic evidence is enshrined in the definition of defeater: a defeater is always a defeater for some mental state as evidence for a belief, so it is not possible for basic evidence itself to be defeated. This is also mirrored in Pollock's "semantics" for defeaters, where "initial nodes" (i.e., the basic evidence) are by stipulation always undefeated.

Both Bayesianism and Pollock's epistemology, then, have structural features that make basic evidence indefeasible. But a generalization of Weisberg's puzzle shows that there are conceptual obstacles to respecting the defeasibility of evidence, even if one's epistemology does not prohibit it by fiat (as with Pollock).

⁵ Weisberg (2015) shows that the same problem arises for Dempster-Shafer theory and ranking theory as well.

In this section I will generalize Weisberg's puzzle in two ways. First, I will show how it applies beyond the confines of formal epistemology. Second, I will show that the puzzle arises for exogenous rebutting defeaters as well as for exogenous undercutting defeaters.

My first generalization of Weisberg's puzzle takes the form of an inconsistent triad. The first member of that triad is the very idea that basic evidence is defeasible:

Defeasibility: Basic evidence is defeasible—it is possible for one's basic evidence to lose its status as such in virtue of the acquisition of further evidence.

The other two members of the triad we already encountered in the formal version of Weisberg's puzzle. The first one is **Rigidity**:

Rigidity: Acquiring E as evidence cannot change what E is evidence for;

and the second one **Symmetry**:

Symmetry: E is evidence for/against P only if P is evidence for/against E .

One view of basic empirical evidence has it that it is constituted by the content of experiences when those experiences justify belief in their contents (see Comesaña (2020)). So I will show that the triad is inconsistent with this conception of evidence first. But a natural reaction to the incompatibility when evidence is assumed to be propositional is that it is the very propositionality of evidence that creates the problem. I will then illustrate the incompatibility with a non-propositional conception of basic evidence, in the hopes of forestalling this natural reaction.

Suppose, then, that I have not yet opened my eyes. In a moment I will open my eyes and have an experience as of a red wall in front of me. This experience is not itself my basic evidence, but rather *gives me* the proposition that the wall is red (WR) as my basic evidence. The experience is the evidence-provider, and its content is the evidence provided. Now, before I open my eyes and acquire that evidence, I think that the color of the wall is independent of the color of the lights: whether the lights are red (LR) or white has nothing to do with whether the wall itself is red or white. In particular, before opening my eyes, (1) WR is not evidence against LR . After opening my eyes, however, (2) LR is evidence against WR . If, following **Defeasibility**, we assume that basic evidence is at all defeasible, then this is a case in point. By **Symmetry**, it follows from (2) that, after opening my eyes, WR is evidence against LR . By **Rigidity**, then, it follows that even before opening my eyes WR is evidence against LR , contradicting (1). Therefore, **Defeasibility**, **Rigidity** and **Symmetry** are an inconsistent triad.

As I said, a natural reaction to this presentation of the problem is to complain that the theory of basic evidence assumed is wrong (just as it was a natural reaction to the same assumption in a Bayesian framework). My basic evidence, the complaint goes, is the experience itself, not its content. If we assume this alternative theory of basic empirical evidence, then the problem doesn't arise. For even if WR is not evidence against LR before I open my eyes, what LR defeats is the power of my experience itself to provide evidence for WR . The extent to which my experience justifies me in believing WR is constrained by the extent to which I am justified in believing that LR —the more justified in believing LR I am,

the less my experience justifies me in believing *WR*. Thus, before I open my eyes I think that the justification that my experience can provide for *WR* is *not* independent of my justification for believing *LR*, and so the first step of the proof of the inconsistency is blocked.

But, as we also saw before when examining Weisberg's puzzle for the Bayesian framework, this reaction simply pushes the problem one step back. What matters is simply that we assume **Defeasibility**, for *whatever* we take our basic evidence to be. So, let us suppose now that we take our basic evidence to be the experience itself (*e*), not its content. If we are to still hold on to **Defeasibility** (and, of course, we need to at least assume that we are still holding on to it to prove its inconsistency with **Rigidity** and **Symmetry**), then we need to find some information that will make my experience lose its status as basic evidence. As we did with the Bayesian case, let us assume that this information takes the form of a report from the cognitive psychology lab informing me that I am an unreliable judge of what experiences I have (*UJ*). Before I open my eyes, I take the possibility of having *e* to be independent of *UJ*. In particular, (1) before opening my eyes, *e* is not evidence against *UJ*. After opening my eyes, however, (2) *UJ* is evidence against *e*—when I read the lab report, I can no longer rationally treat *e* as part of my evidence. By **Symmetry**, it follows from (2) that after opening my eyes *e* is evidence against *UJ*. By **Rigidity**, then, it follows that even before opening my eyes *e* is evidence against *UJ*, contradicting (1). Therefore, **Defeasibility**, **Rigidity** and **Symmetry** are an inconsistent triad, even if we assume that our basic empirical evidence is constituted by our experiences and not by their contents. Moreover, it is easy to see now that the triad is inconsistent no matter our conception of basic evidence. Of course, if we assume that our basic evidence is indefeasible, then the problem doesn't arise—but this obviously does not go against the inconsistency of the triad, it just removes one of its members.

My second generalization of Weisberg's puzzle consists in pointing out that exogenous rebutting defeaters represent as much of a problem as exogenous undercutting defeaters for some theories. Suppose that instead of learning that the lights are red my reliable friend tells me that the wall is not red. In some variations of that case, the proposition that the wall is red should no longer be part of my evidence. Traditional Bayesianism has as much trouble accommodating exogenous rebutting defeaters as it does accommodating exogenous undercutting defeaters, because once a proposition is part of a subject's evidence Traditional Bayesianism has it that it retains that status forever. Pollock's epistemology cannot handle rebutting exogenous defeaters either, simply because defeaters are defined in Pollock's framework as reasons for doubting that the evidence supports some belief, never as reasons for doubting the evidence itself. Liberal Bayesianism can accommodate exogenous rebutting defeaters, as long as the initial evidence does not consist in assigning credence 1 to any proposition. For instance, if my initial experience makes me pretty, but not totally, confident in the proposition that the wall is red, then my friend's testimony can lower that confidence. Later on I will show that a different version of Bayesianism can handle rebutting (and, indeed, undercutting) exogenous defeaters while insisting that evidence gets credence 1. So, while exogenous rebutting defeaters are not as problematic as exogenous undercutting defeaters, they do raise some interesting issues that we will get back to below.

Weisberg's puzzle, then, is not relegated to attempts to formalize epistemic relations: its generalization applies very widely indeed.

5. Solving the Generalized Puzzle

As with any inconsistent triad, there are three ways out of the puzzle, by denying each of the members of the triad. In this section I examine those three possibilities, arguing that the right solution is to deny **Rigidity**. A full solution, however, must explain why **Rigidity** is false in this case, and must in addition solve a residual puzzle regarding **Symmetry**. In this section I also make some gestures regarding these two issues.

First, one could deny **Defeasibility**. The internalization moves examined before only work if accompanied by a denial of **Defeasibility**. Historically, internalization and indefeasibility did indeed march in tandem, for instance in (textbook accounts of) Cartesian epistemology. In contemporary epistemology, few are willing to uphold the indefeasibility of our empirical evidence explicitly. And yet we have seen that, implicitly, both Standard Bayesianism as well as Pollock's account of reasons and defeaters go for this kind of indefeasibility. I side with the majority of contemporary epistemologists in avoiding this solution to the problem.

Symmetry might look suspicious when we think about it in Pollock's framework. For while it might make sense to say that my experience as of a red wall in front of me is a reason for me to believe that there is a red wall in front of me, it doesn't seem to make sense (let alone be true) to say that my belief that there is a red wall in front of me is any kind of evidence for my experience.⁶ While I agree that **Symmetry** looks better in an alternative framework I have alluded to, I think that it does make sense in the traditional Pollockian framework as well (and that is a good thing, for it would be an indictment against that framework if **Symmetry** couldn't even be meaningfully formulated in it). I will therefore make more explicit my alternative framework as well as explain why I think that **Symmetry** has true formulations in it as well as in the Pollockian framework.

Although it is customary in some circles to treat the evidential relation as holding between doxastic attitudes, I am here thinking of it as holding between propositions. Thus, whereas epistemologists such as Pollock would say that a belief that p (or an experience as of p) can be evidence for a belief that q , I am assuming that the evidential relations hold in the first place between propositions—thus, that in that case it is p itself that is evidence for q itself. The underlying picture is one where epistemic justification can be thought of as being constituted by three levels. The first one is the level of objective evidential connections between propositions, which exist independently of any subject. At this level, for instance, we find the fact that the proposition that all observed emeralds are green is evidence for the proposition that all emeralds are green—but we do not find the fact that the proposition that all observed emeralds are grue is evidence for the proposition that all

⁶ Thanks to Kurt Sylvan for bringing this worry to my attention.

emeralds are grue. The second level corresponds to the traditional notion of propositional, or *ex-ante*, justification. This second level incorporates subjects into the picture, for propositional justification requires evidence possession, and subjects are the ones who can possess evidence. Propositions can be possessed as evidence by a subject either by being the content of a justified belief or an undefeated experience of the subject. Thus, I can possess as evidence the proposition that all observed emeralds are green by being justified in believing this proposition.⁷ Having that proposition as evidence, I now have propositional justification for believing that all emeralds are green. The third level corresponds to the traditional notion of doxastic, or *ex-post*, justification. If I do believe that all emeralds are green, and if this belief of mine is properly based on the evidence I have (that all observed emeralds are green), then I am doxastically justified in believing that all emeralds are green. In terms of this picture, then, **Symmetry** is a thesis about the first level, the level of objective evidential relations between propositions. It says that if p is evidence for q , then q is evidence for p .

Let us go back to the question of whether **Symmetry** makes sense when formulated in the more traditional, Pollockian framework. The worry was that, whereas it might make sense to say that an experience as of a red wall in front of me is evidence for (or a reason for) believing that there is a red wall in front of me, it does not make sense to think of a doxastic attitude as evidence for a reason. But remember that for Pollock beliefs themselves can be reasons for further beliefs. Given that beliefs can stand in both ends of the *reason-for* relation, **Symmetry** should make as much sense as the very idea that beliefs can be reasons for other beliefs. In both cases, an obvious way to interpret the idea that a belief is a reason for a belief is to say that the propositions that are the contents of the corresponding beliefs are evidentially related (and I have already remarked that Pollock himself talks in these terms routinely). But what about the particular case where an experience is the reason (which will be the basic case according to Pollock)? Does it make sense to say that a belief is a reason for an experience? Again, I think that it makes as much sense as saying that an experience is a reason for a belief. In both cases there is some awkwardness, but that awkwardness derives from taking experiences themselves (rather than their propositional content) to be the relata of the reason-for relation. The awkwardness derives from thinking that a *thing* (such as an experience) can be a reason. Despite the common use of the notion of evidence as applying to things (such as bloody knives), in its philosophical use it is almost universally understood to hold between propositions, or at any rate between truth-bearers, and many suggest that even that common use is better understood as elliptical (it is some proposition regarding bloody knives which is the real evidence). If so, one way to interpret the idea that experiences can be reasons for belief is as in the alternative picture I just presented, where it is the content of the experiences that are properly speaking the

⁷ More precisely, evidence possession requires undefeated *basic* justification for believing the proposition in question—i.e., justification not based on justification for believing any other proposition. This makes it unlikely that the proposition that all emeralds are green is part of my evidence, because it is most likely inferentially justified, but it simplifies the example in other respects, so I assume that it is basically justified.

reasons (which are had by a subject in virtue of having an experience with that content). I conclude, therefore, that there is no reason to be suspicious of **Symmetry** on the basis that it cannot even be properly formulated in a Pollockian framework. To the extent that there is some awkwardness in the formulation of the principle in that framework, that is a problem for the framework, not the principle.

Notice, moreover, that **Symmetry** as defined is a qualitative, not a quantitative, principle. It says that if E is evidence for/against P to *any* extent, then P must be evidence for/against E to *some* extent. Suppose, for instance, that some scientific hypothesis entails an observation. In that case, the hypothesis provides the strongest possible evidence for that observation. According to **Symmetry**, the observation itself must also provide at least some evidence for the hypothesis, although not necessarily very strong evidence. And this is indeed widely accepted, from the Hypothetico-Deductive model of confirmation to Bayesian accounts. To see the appeal of **Symmetry** we can think about what we would have to hold to reject it: we would have to say that, in some cases, even though P is evidentially completely independent of E , E is evidence for or against P . How could that be? If learning P shouldn't make you change your mind to any extent whatsoever regarding E , how could learning E make you change your mind regarding P ? To regard P as completely evidentially independent from E is to think that whether P obtains makes no difference as to whether E obtains. Thinking about it in geometrical terms, that means that the amount of logical space occupied by E is the same in the P region as in the $\neg P$ region, and the same goes for $\neg E$. But if that is so, then obviously the amount of logical space occupied by P is the same in E region as in the $\neg E$ region, and the same goes for $\neg P$, which means that E is completely evidentially independent from P .

The remaining option is to reject **Rigidity**, and indeed this is the option I will argue for. But it is important to concede at the outset that **Rigidity** is not just a feature of formal systems like Bayesianism, but a rather intuitive constrain on evidence. To violate rigidity is to say that, at least sometimes, whether E is evidence for P depends on whether E is had as evidence. This sounds very bad. However, rejecting **Rigidity** is the option that is mandated by Weisberg's puzzle. For remember the setup of that puzzle: before I open my eyes, the color of the wall and the color of the lights are evidentially independent of each other, but after I open my eyes the color of the lights is evidentially relevant to the color of the walls. Given **Symmetry**, this is just the denial of **Rigidity**.

The correct solution to the puzzle, then, is to deny **Rigidity**. But we need an explanation as to why that very plausible principle is false in this case. The explanation has to do with the relationship between evidence-providers and the evidence they provide. I will illustrate the explanation assuming that experiences provide their contents as evidence, and then see how the explanation would go for experiential conceptions of basic evidence.

An experience, then, defeasibly provides its content as evidence. But that provision can be defeated, both in a rebutting and in an undercutting way. Thus, my experience as of a red wall defeasibly provides me with the proposition that the wall is red as evidence. But my friend's testimony that the wall is not red is a rebutting defeater for that proposition, and the information that the lights are red is an undercutting defeater for it. So, rebutting and undercutting defeat applies to evidence-providers as much as it does to evidence itself.

That is to say, we need to generalize Pollock's definition of a defeater. Pollock's definition of a defeater has it that a mental state $M +$ is a defeater for E as evidence for P if and only if E is evidence for P and $E \wedge M +$ is not evidence for P . Call that defeat *by* evidence. We need to countenance defeat *of* evidence as well. We can say that a mental state $M +$ is a defeater for e as providing evidence P if and only if e provides P as evidence but $e \wedge M +$ does not provide E as evidence. To cover both cases we can generalize as follows: $M +$ is a defeater for ϕ as a justification to believe P if and only if ϕ justifies the subject in believing P but $\phi \wedge M +$ does not justify the subject in believing P . Two different sorts of things can justify a subject in believing a proposition: evidence, and evidence-providers. Weisberg's puzzle arises because evidence and evidence-providers justify in different ways. Evidence justifies by providing reasons for belief, whereas evidence-providers justify by providing evidence itself. What evidence-providers justify is belief in propositions which are part of one's evidence, whereas what evidence justifies is belief in further propositions which are not part of one's evidence.

Rigidity fails, then, because of the difference between objective evidential relations and the possession of evidence. Objective evidential relations exist in the Platonic heavens and do not care about whether anyone has any evidence (indeed, they do not care about whether anyone at all exists). In this pure evidential realm, the color of the wall and the color of the lights are completely independent of each other. But when someone possesses a proposition as evidence, new evidential relations are created. If I have as evidence the proposition that the wall is red, and if what provided me with that evidence is my visual experience as of a red wall in front of me, then the color of the wall and the color of the lights are now evidentially linked, so that learning that the lights are red should make me lower my confidence in the proposition that the wall is red (and so that proposition should no longer be part of my evidence). **Rigidity** fails because possession of evidence introduces new evidential relations, in virtue of the fact that the justification that evidence-providers give to the evidence thereby provided can itself be defeated.

Now, how would the explanation go if we assume that experiences are themselves evidence, rather than evidence-providers? Admittedly, things are messier in this case, but ultimately I take that mess to be a symptom of the fact that the view that experiences are evidence is misguided. I will not argue for that point here, however. Rather, I would point out the messy way in which proponents of that view can try to solve the puzzle along lines similar to the one I already gave.

If experiences are evidence, then what gives us that evidence? This is the messy part, for on a view on which experiences are evidence it is not clear that there is a distinction between evidence and its possession. One possibility is to say that the evidence is the proposition that I am having an experience with such-and-such a content, and that I have that evidence when I am justified in believing (or perhaps when I am in a position to know) that I am having that kind of experience. Notice that, on this interpretation, my evidence is not my belief that I am having a certain experience, but rather the proposition itself that I am having it. Indeed, I need not have any belief at all. Rather, I must have *ex-ante* (or propositional) justification for believing the proposition (or perhaps be in a position to know it), whether I actually believe it or not.

On a different interpretation of the view, it is the experience itself which is the evidence, and not a proposition to the effect that I am having the experience, and to have the experience is to have the evidence. On this interpretation of the view, justification to believe that I have the experience is not what provides me with the experience as evidence. Still, it is plausible to think that if the experience itself is to be part of my evidence, then I have to be justified in believing that I have the experience. Suppose that I have the experience, but I am justified in believing that I do not, or I am perhaps justified in being agnostic as to whether I have the experience or not. In those cases, it seems that, even if I do in fact have the experience, it cannot justify me in believing anything else—that is to say, it cannot function as evidence. Suppose, for instance, that I have an experience as of a red wall in front of me, but I am justified in believing that I have a different kind of experience—say, I have been convinced by scientists with access to mind scanners that I actually have an experience as of an orange wall in front of me, and that I am confusing red experiences with orange experiences. In that case, my experience can hardly justify me in believing that there is a red wall in front of me—or, indeed, anything else. The same goes if instead of being justified in believing that I do not have the experience I am merely justified in suspending judgment about whether I have it. Again, note well that I am not saying that for the experience to justify I have to actually believe that I have the experience—rather, I must be justified in so believing.

So, whether justified belief that I have the experience is what provides me the experience as evidence or not, being justified in believing that I have the experience is a necessary condition of my having it as evidence. So, anything that takes away that justification (such as the report from the lab saying that I am a bad judge of which experiences I have) will defeat the evidential power of the experience. The corresponding generalization of Pollock's definition of defeater will in this case be the following: $M +$ is a defeater for my justification for believing that I have e as a necessary condition for my having e as evidence if and only if I cannot be justified in believing that I have e on the basis of $M +$. And in this case what explains the violation of **Rigidity** will be, not the fact that evidence-providers introduce new evidential relations, but the fact that necessary conditions on evidence-possession introduce those new evidential relations.

I hasten to repeat that I prefer the first view about empirical evidence and its possession: the view according to which experiences defeasibly provide their contents as evidence. But I offer the solution in the previous paragraphs to those philosophers who prefer the alternative view on which experiences themselves are evidence.

6. Back to Bayesianism

Having presented that solution to the generalized version of Weisberg's puzzle, I explain now how to model that kind of solution in a Bayesian setting.

We saw that **Rigidity** and **Symmetry** are built into Bayesianism, so it might look as if it is impossible to give a Bayesian model where **Rigidity** fails. But it isn't, as long as as we recognize the limitations of the model and we let the lack of rigidity show up as a side-effect of the fact that we will not be modeling *all* the interesting epistemic relations.

Let me start with a case that can be used to object to Traditional Bayesianism. Suppose that I show you two balls, a black one and a red one, and then put them inside an urn. It would be natural to say that, at this point, it is part of your evidence that there is a black ball and a red ball inside the urn. I then let you sample with replacement from the urn—i.e., you can take out one ball, note its color, and return it to the urn. You do this 10,000 times, and every time you take out a black ball. At this point (and, indeed, way before this point), it is natural for you to start suspecting that I am playing a trick on you—or, at any rate, that for some reason there is no red ball in the urn. Traditional Bayesianism cannot model this, for if it is ever part of your evidence that there is a red ball in the urn, then it is forever part of your evidence. Liberal Bayesianism, with its rule of Jeffrey-Conditionalization, can model this, as long as we have it that you do not initially have credence 1 in the proposition that there is a red ball inside the urn. But there is another way of modeling this interaction which does not consist in moving from Traditional to Liberal Bayesianism. Rather, we can appeal to the notion of an ur-prior.

If a subject abides by both Traditional Bayesian normative constraints—i.e., if at any time his credences are probabilistically coherent and, whenever he acquires a proposition as evidence, he conditionalizes on it—there will be (infinitely many) ur-priors for him. An ur-prior for a Traditional Bayesian subject is a probability distribution over a set of propositions with the following property: at any given time t , the subject's credences are the result of conditionalizing his ur-prior on every proposition which is part of his evidence at t . Now, the reason that there will be infinitely many of these ur-priors for a Bayesian subject is that any subject will only ever receive a proper subset of all the evidence he *could* receive, and any proposition that the subject doesn't ever conditionalize on adds infinitely many degrees of freedom to the ur-prior, by varying the ur-prior's conditional probabilities on that proposition. If we fix the subject's conditional probabilities not only on the propositions he actually receives as evidence, but on every possible proposition, then his ur-prior will be unique. This uniqueness is still relative to the subject, for there are infinitely many different ways of abiding by **Probabilism** and **Conditionalization** (or **Jeffrey-Conditionalization**). This subjectivity is seen by many (including myself) as a problem for Bayesianism. On the other hand, the other extreme of supposing that only *one* of the possible *ur-priors* is the uniquely rational one has seem to many others to be just as problematic as assuming that any of them will do—many, for instance, see the failure of Carnap's project of determining the unique prior syntactically as showing that there is no such ur-prior. I think that the problems for the uniqueness of the ur-prior are more a result of drawing the wrong consequences from some features of the model than deep philosophical reasons for being weary of uniqueness. But this issue need not detain us here. Whether we think there is only one or infinitely many rational ur-priors, they allow us to present an alternative to both **Conditionalization** and **Jeffrey-Conditionalization**, namely:

Ur-prior Conditionalization: If a subject's evidence at t is E , then $Cr_t(P) = Cr_u(P|E)$.

Ur-prior Conditionalization is like **Conditionalization** in that it treats the impact of experience on our evidence propositionally, but it is like **Jeffrey-Conditionalization** in that it is not cumulative. Consider the above case again. When the subject sees a red ball go into the urn, it is part of his evidence that the urn contains a red ball, and so he assigns credence 1 to this proposition. But once he performs the 10,000 samples with replacement

and does not obtain a red ball from them, it is no longer part of his evidence that there is a red ball in the urn. Rather, his evidence now consists of the results of those samples. And the ur-prior conditional probability that there is a red ball in the urn given the results of the sampling (if his ur-priors are sane at all) is very low indeed. Similarly, when I open my eyes and have an experience as of a red wall in front of me, it is part of my evidence (I say) that the wall is red. But later, when I learn that the lights are red, that proposition is no longer part of my evidence. Rather, my evidence now consists of the proposition that the lights are red, and the conditional probability that the wall is red given that the lights are red is, according to my ur-prior (which, I assure you, is the correct one), approximately 0.5. By appealing to ur-priors, then, we can model the solution I offered to the generalized version of Weisberg's puzzle in the previous section.⁸

7. Conclusion: Epistemological Particularism

The appeal to **Ur-prior Conditionalization** has a pretty obvious downside, one which to some extent or another it shares with other forms of Bayesianism. The downside is that some important epistemic work is done "behind the scenes," in that it is not modeled in the formalism. What is not modeled is the very fact that evidence is defeasible. For we cannot model the principles in virtue of which when you learn that the lights are red now you no longer have the proposition that the wall is red as part of your evidence (or, for that matter, the principles in virtue of which when you perform the 10,000 sample with replacement you no longer have the proposition that there is a red ball in the urn as part of your evidence)—the best we can do is to record those facts by saying that now you should conditionalize your ur-prior on that newly acquired evidence alone, not in conjunction with the old evidence. All forms of Bayesianism relegate some interesting epistemic work to behind the scenes—in particular, both Traditional and Liberal Bayesianism treat the impact of experience on our doxastic life as a given, exogenous input. But Weisberg's puzzle highlights that the defeasibility of evidence is another aspect which must be treated as an exogenous input.

However, what I want to suggest is that just as we should not treat Weisberg's puzzle as a problem only for formal epistemology, but rather as showing that possessing evidence can generate new evidential relations, so too we should not view the treatment of the defeasibility of evidence as an exogenous input as peculiar to formal epistemology either. Rather, the impossibility of modeling the defeasibility of evidence is a symptom of a more general fact: that particularism holds in epistemology. By "particularism" in this context I do not mean the view, opposed to "methodism," according to which judgments about particular cases have priority over judgments about general principles. Rather (and roughly speaking), I mean the view, opposed to "generalism," according to which the role of

⁸ Others have presented different formal treatments of Weisberg's puzzle—see, for instance, Gallow (2014) and Schwarz (2018). To my mind, the solution in terms of ur-priors has the advantage of being up-front about what is left out of the model, an issue discussed in the next section.

principles in epistemology is necessarily limited. The question on which epistemological particularists and generalists divide is whether epistemic relations are codifiable. Can we state exceptionless and informative principles regulating when a given doxastic attitude towards a given proposition is justified? Generalists say "Yes," particularists say "No."

Now, at the level of evidential relations between propositions which are subject-independent maybe generalism is true. At the very least, neither Weisberg's puzzle nor my generalization are relevant to this issue. But the fact that possessing evidence creates new evidential relations means that particularism is true at the level of propositional and doxastic justification. To see why, let us go back to Pollock's account of reasons and defeaters. In Pollock's semantics, initial nodes (which are given by experience) can never be defeated. Given this commitment to the indefeasibility of basic evidence, one can give a recursive account of *prima facie* justification in Pollock's framework. A belief is *prima facie* justified if and only if it is justified by an experience of the right kind or it is justified by other beliefs which are themselves *prima facie* justified. We can then see (the rest of) the semantics as dealing with defeaters, and we can define a justified belief as a belief which is *prima facie* justified and undefeated. Crucial to the recursive definition of *prima facie* justification is the fact that the basic evidence itself cannot be defeated. Experiences always provide *prima facie* justification to the corresponding beliefs, and as long as this justification is not defeated, those beliefs are all-things-considered justified. That defeat, crucially, can only come from other beliefs which must themselves be all-things-considered justified. Those beliefs might be inferentially justified by other beliefs, but given this foundationalist picture that justification must eventually trace back to some other experiences. The experiences themselves can never be defeated as providers of *prima facie* evidence. We can generalize Pollock's picture and talk of basic evidence instead of experiences. A belief is justified if and only if it is backed by a justificatory chain which goes back to some basic evidence and it is not defeated. Defeat, in turn, can only come from other justified beliefs. Basic evidence itself can never be defeated.

Given that basic evidence can be defeated, we must abandon this Pollockian picture of the structure of epistemic justification. What can we put in its place? If what we are looking for are general principles of the sort that allow for an informative (perhaps recursive) characterization of the justification relation, then perhaps the answer is: nothing. The generalization of Weisberg's puzzle suggests that, because even basic evidence is subject to defeat, that kind of characterization is not available. Pollock's framework was supposed to give us a liberalized version of foundationalism. Traditional foundationalism had it that the foundations were provided by beliefs about our own mental states, and these beliefs were indefeasible because our own mental states are transparent to us. From this foundation, the rest of our beliefs were to be justified by infallible inferences. Pollock's liberalization was two-fold: in the first place, the foundation was provided not by infallible beliefs about our mental states, but by our mental states themselves, and justification was allowed to be defeasible. Pollock's identification of reasons with mental states obscured the difference between the reasons there are for subjects to believe and the reasons those subjects have. Experiences are either had or not, and so there is no place in the Pollockian framework for unpossessed basic evidence. This obfuscation of the distinction between the reasons there are and the reasons we have in turn obscured the extent to which the Pollockian

framework retains the indefeasibility of traditional foundationalism. True, in the Pollockian framework our beliefs are themselves always defeasible, but that is because the foundational evidence is not doxastic—and this non-doxastic basic evidence is itself indefeasible. If, as we should, we give up on the indefeasibility of basic evidence, then, I am suggesting, we have to also give up on another feature of traditional foundationalism: the possibility of characterizing justification on the basis of general principles.⁹

References

- Casullo, Albert. 2018. "Pollock and Sturgeon on Defeaters." *Synthese* 195: 2897–2906.
- Comesaña, Juan. 2020. *Being Rational and Being Right*. Oxford University Press.
- Gallow, Dmitri. 2014. "How to Learn from Theory-Dependent Evidence; or Commutativity and Holism: A Solution for Conditionalizers." *British Journal for the Philosophy of Science* 65 (3): 493–519.
- Jeffrey, Richard. 1965. *The Logic of Decision*. 2nd. ed. McGraw Hill.
- Knowledge and Justification*. 1974. Princeton University Press.
- Kotzen, Matthew. 2019. "A Formal Account of Epistemic Defeat." In *Themes from Klein*, edited by Cherie Braden Branden Fitelson Rodrigo Borges. Springer.
- Pryor, James. 2000. "The Skeptic and the Dogmatist." *Noûs* 34: 517–49.
- . 2013. "Problems for Credulism." In *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford University Press.
- Schwarz, Wolfgang. 2018. "Imaginary Foundations." *Ergo* 5 (29).
- Sturgeon, Scott. 2014. "Pollock on Defeasible Reasoning." *Philosophical Studies* 169 (1): 105–18.
- Weisberg, Jonathan. 2009. "Commutativity or Holism? A Dilemma for Conditionalizers." *British Journal for the Philosophy of Science* 60: 793–812.
- . 2015. "Updating, Undermining, and Independence." *British Journal for the Philosophy of Science* 66 (1): 121–59.

⁹ Thanks to Rhys Borchert, Ding, Tim Kearl, Carolina Sartorio, Will Schumacher, Kurt Sylvan, and Jonathan Weisberg for very helpful comments on previous drafts.