BOOTSTRAPPING, DOGMATISM, AND THE STRUCTURE OF EPISTEMIC JUSTIFICATION

SHYAM NAIR

Arizona State University

CONSIDER the following quixotic attempt to determine whether your color vision is reliable. You ask your friend to set up a slide show in which every few seconds a colored slide will appear on the screen. Your friend does not tell you what color the slides are or in what order they will appear. You sit down and the slideshow begins. The first slide comes up and you look at it and, on the basis of your visual experience, you come to believe:

(1) The slide is red.

You next notice that you are having a visual experience of the slide being red and come to believe:

(2) The slide looks red.

You then reason from these two beliefs to form the belief:

(3) The slide is red and it looks red.

Having gotten this far, you conclude:

(4) My color vision worked this time.

The next slide comes up. And you go through the same reasoning for this slide. And so on for the whole slideshow. From this stock of beliefs together you conclude:

Contact: Shyam Nair <gsnair@asu.edu>

Track Record: My color vision worked n times.

where n can be as large as we like. Based on *Track Record* you finally conclude:

Reliability: My color vision is reliable.

Despite your efforts, you have not made progress toward figuring out whether your color vision is reliable.¹ If you did not know or were not justified in believing *Reliability* in the first place, this process of reasoning does not provide you with justification or knowledge now.

The bootstrapping problem in epistemology arises for views that seem to be committed to the implausible result that this form of reasoning does generate justified belief in or knowledge of *Reliability*.² Though this problem has been raised for a number of different kinds of theories, in this paper I will concentrate on the problem as it arises for so-called dogmatist theories of epistemic justification.

My argument will be that the only way for the dogmatist to avoid the bootstrapping problem is to claim that epistemic justification fails to have a structural property, which I will describe in detail later, known as *CUT*. This allows the dogmatist to admit that each step in this reasoning considered on its own is acceptable, but when stitched together, these pieces of reasoning are unacceptable (§2). The fact that this is the only plausible solution to the bootstrapping problem is, in one way, bad news. This is because, as I will show, it adds another member to a family of recently uncovered results that show dogmatism is incompatible with certain connections between epistemic justification and probabilities (§3).

1. Here we make a number of simplifications and idealizations for ease of presentation. First, we put aside the idea that you get some evidence that you are reliable simply by *not* having extremely bizarre perceptual experiences. Second, it would do no harm to assume the first two steps occur simultaneously. Third, plausibly the relevant track record propositions should record the ratio of success to failure as well. In our context, of course, you always succeed and never fail. Third, for the inference to *Reliability* from *Track Record* we assume n is large enough to make the inference plausible but not so large that it raises concerns about conjunction introduction (see §3.1 for discussion). Fourth, *Reliability* perhaps should be qualified so it mentions a context. It does nothing to reduce the force of the problem if we simply make the context the specific one the agent is currently in. Thanks to Brad Armendt and Ángel Pinillos for pushing me to clarify these issues.

2. See Cohen (2002) for a seminal discussion of bootstrapping for dogmatism as well as a closure problem. Following Vogel (2007) and Weisberg (2010), I focus only on the bootstrapping problem in this paper. I do not believe the arguments of this paper on their own are sufficient to solve the closure problem (see §C for discussion). Weisberg (2012) is a survey of work on the bootstrapping problem. Other discussions of the bootstrapping or closure problems include Altschul (2012), Becker (2012), Black (2008), Breisen (2013), Brueckner (2013), Brueckner and Buford (2009), Davies (2004), Douven and Kelp (2013), Cohen (2005; 2010), Hawthorne (2004: 73–77), Kalstrup (2012), Kornblith (2009), Markie (2005), Neta (2005), Scheall (2011), Titelbaum (2010), Weisberg (2010), Vahid (2007), van Cleve (1979; 2003), Vogel (2000; 2008), and Zalabardo (2005).

But instead of stopping here and concluding that dogmatism is false, I try to make the best of it on the dogmatist's behalf. I show that within a certain kind of foundationalist framework, we can make good on this idea that epistemic justification fails to satisfy CUT in a way needed to solve the bootstrapping problem (§4–5). But let's begin by more clearly explaining what dogmatism is and what the bootstrapping problem is.

1. Dogmatism and Bootstrapping

Though there are different versions of dogmatism, there are three core theses that I will take to be definitive of this family of views. The first thesis is that that a visual experience, for example, of a table being red, provides defeasible justification for believing, for example, that table is red. The second thesis is that this justification does not require the agent to have any prior knowledge or justified belief that in the present situation appearances are not deceiving. And the third thesis is that this justification does not require the agent to have any prior knowledge or justified belief that her visual system is reliable. Instead, dogmatism allows that an agent is defeasibly justified simply in virtue of having the visual experience (see Huemer 2001; Pollock & Cruz 1999; and Pryor 2000).

I will also take these theses to be commitments of dogmatism whether we read 'justification' as so-called propositional justification or so-called doxastic justification. This is not a universally followed practice. Some theorists identify as dogmatist but only accept these theses understood as claims about doxastic justification; they reject thesis two and three understood as claims about propositional justification (Silins 2008). This paper however will use 'dogmatism' and its cognates to refer to the stronger view that accepts the three theses on both interpretations. I focus on this stronger view because it cleaves more closely to certain motivations for dogmatism related to avoiding so-called deeply contingent *a priori* knowledge that we will discuss in greater detail below (§3.2.4, §C.2) and because solving the bootstrapping problem for these views is especially hard.³

So in what follows, I will not belabor the distinction between these two kinds of justification because my claims apply to both. That said, at times, I will put things in terms that are more congenial to one interpretation than the other for simplicity. I leave it to the reader to do the simple translation required to fit my comments with the other interpretation of justification.

A second clarification is also in order: the justification at issue is *defeasible*. That is, the agent could learn some new information that makes her no longer

3. Those who only accept the weaker dogmatist claims may endorse the neo-Cartesian solutions offered by Cohen (2010) and Wedgwood (2013). See §C.3 for discussion.

justified in believing that the table is red. For example, the agent might learn that her perception is unreliable or that in the current situation appearances are deceiving. But when the agent simply lacks knowledge that her perception is reliable and lacks knowledge that in the current situation appearances are deceiving, she is justified in believing the table is red on the basis of her visual experience as of it being red. And this is true whether or not as a matter of fact her visual perception is reliable or as a matter of fact in the current situation, appearances are deceiving.

The bootstrapping objection to dogmatism is that it entails that the reasoning that I began the paper with provides new justification for believing *Reliability*. In order to present this objection and my solution perspicuously, it will help to adopt a bit of formalism for representing claims about epistemic justification.

Officially, the dogmatist says that having a certain visual experience, for example, of the table being red justifies you in believing, for example, that the table is red. This suggests that a natural way to think of epistemic justification for an agent according to the dogmatist is as a relation that holds between mental states. While this is the natural and most direct way to think of the dogmatist view, it will ultimately prove convenient to formalize this view in a less direct way that allows us to connect our discussion to simple formal properties of relations that are familiar from philosophical logic.

So we will model epistemic justification for an agent, a , at a time, t , using a relation $\models_{a, t}$ and in fact, we will from now on leave reference to the agent and time implicit and just write $\models$. This relation will not be used to relate mental states, but instead will be used, like a logical consequence relation, to relate a set of sentences and a sentence. It will be useful later to use lower case Greek letters to denote sentences (α, β, γ , etc.) and uppercase Greek letters to denote sets of sentences (A, B, Γ , etc.).

So for example, we can write:

$$\{\text{'the table is red'}, \text{'the table is round'}\} \models \text{'the table is round and the table is red'}$$

Or if we wish to be more succinct, we may for convenience omit the braces and write:

$$\text{'the table is red'}, \text{'the table is round'} \models \text{'the table is round and the table is red'}$$

This is intended to represent the claim that an agent who is epistemically justified in believing that the table is red and is epistemically justified in believing that the table is round is epistemically justified in believing that the table is

round and the table is red. So more generally, we read this as saying an agent who is epistemically justified in believing each of the propositions *expressed* by the sentences on the left is epistemically justified in believing the proposition expressed by the sentence on the right.

This smoothly covers cases in which beliefs justify other beliefs. But the dogmatist treats cases of perceptual justification differently than cases of beliefs justifying other beliefs. As we saw, the dogmatist says that merely having a perceptual experience can provide epistemic justification for beliefs. In order to capture this in our formalism in a way that will still allow us below to make use of some simple formal properties of relations, we will have to engage in a bit of harmless equivocation. So we will write:

‘the table looks red’ $\Vdash$ ‘the table is red’

and this will be used to represent the claim that an agent who is in the state *reported* by the sentence on the left is epistemically justified in believing the proposition expressed by the sentence on the right. For those concerned that this equivocation is problematic, Appendix A introduces the complexity needed to show how our discussion can be successfully conducted in a non-equivocating formalism

With this formalism in hand, let’s return to the bootstrapping reasoning and consider which claims about epistemic justification would have to be true for that reasoning to lead to you having a justified belief in *Reliability*. To start, notice that since dogmatism does not require agents to be justified in believing *Reliability* from the start, it entails that there are some agents who are not justified in believing *Reliability* before the bootstrapping reasoning. Imagine you are such an agent who is engaging in the bootstrapping reasoning. It begins with you believing that the slide is red based on your visual experience of the slide being red. If we let LR be ‘the slide looks red’ and R be ‘the slide is red’ the claim about epistemic justification that would have to be true for this step is the following:

Step 1: $LR \Vdash R$

The next step is noticing that the slide looks red to you. For simplicity, we will just assume you are justified in believing that the slide looks red to you from the start and so omit this step.

From R and LR you then come to believe that the slide looks red and the slide is red. If we let $R \wedge LR$ be ‘the slide is red and looks red’, the claim about epistemic justification that corresponds to this step is the following:

Step 2: $R, LR \Vdash R \wedge LR$

From $R \wedge LR$, you conclude that your color vision worked. If we let *Worked* be ‘my color vision worked’, the claim about epistemic justification that corresponds to it is the following:

Step 3: $R \wedge LR \Vdash \textit{Worked}$

Next after reasoning analogously to n instances of something like *Worked* that we will represent as $\textit{Worked}_1, \textit{Worked}_2, \dots, \textit{Worked}_n$, you are able to conclude that your color vision worked n times which is the proposition expressed by *Track Record*:

Step 4: $\{\textit{Worked}_1, \textit{Worked}_2, \dots, \textit{Worked}_n\} \Vdash \textit{Track Record}$

Finally from *Track Record* you conclude that your color vision is reliable which is the proposition expressed by *Reliability*:

Step 5: $\textit{Track Record} \Vdash \textit{Reliability}$

So steps one through five lead to the result that you are epistemically justified in believing *Reliability* based on the bootstrapping reasoning.

The objection to dogmatism is that each of steps two through five is acceptable and yet nonetheless an agent is not justified in believing the proposition expressed by *Reliability* in virtue of the bootstrapping reasoning, so *Step 1* must be incorrect. But *Step 1* is a commitment of dogmatism.

Let’s look at this. To see that *Step 1* is a commitment of dogmatism, we need only note that it is an application of the dogmatist’s general idea that having a visual experience as of p epistemically justifies you in believing p . The remaining steps are quite plausible. *Step 2* and *Step 4* involves performing “conjunction introduction”. While there are contexts in which this is not acceptable that we will discuss in greater detail later, in the present context it is acceptable (§3.1). *Step 3* involves a simple analytic entailment. Finally, *Step 5* is an instance of concluding by enumerative induction that your color vision is reliable based on many instances of it working correctly.

In sum, it is unacceptable to think that the bootstrapping reasoning leads us to be newly justified in believing the proposition expressed *Reliability*. So we must reject the conclusion of this reasoning. But this triggers what Jonathan Vogel (2007) calls “rollback”: We must say which step in the reasoning goes wrong. But since conjunction introduction, the simple analytic entailment, and induction looks overwhelmingly plausible when considered alone, it looks like we roll all the way back to saying the first step in the reasoning is faulty and that’s just a commitment of dogmatism.

2. The Structure of Epistemic Justification

The only way out of this result is to somehow claim that while each step is acceptable on its own, you cannot perform the steps back-to-back.⁴ While our formalism may have been initially unnatural, what makes it worthwhile is that we can now state the idea that you cannot do the steps back-to-back simply and precisely as the idea that $\Vdash$ is not a transitive relation, where we say:

$\Vdash$ is *transitive* just in case if $A \Vdash \beta$ for all $\beta \in B$ and $B \Vdash \gamma$, then $A \Vdash \gamma$

If the dogmatist were to deny this claim, then she could accept each step in the reasoning, but deny that they can be performed one after another.⁵

As it turns out, transitivity follows from the following two properties that are commonly called *MONOTONICITY* and *cumulative transitivity* or for short, *CUT*:

$\Vdash$ satisfies *MONOTONICITY* just in case if $A \Vdash \gamma$ and $A \subseteq B$, then $B \Vdash \gamma$

$\Vdash$ satisfies *CUT* just in case if $A \Vdash \beta$ for all $\beta \in B$ and $A \cup B \Vdash \gamma$, then $A \Vdash \gamma$

So if the dogmatist wishes to solve the bootstrapping problem, she must reject at least one of these claims.⁶ And if this solution is to be anything but *ad hoc* the dogmatist must explain why epistemic justification fails to satisfy either *CUT* or *MONOTONICITY* in a way that solves the problem. The remainder of the paper is dedicated to exploring whether such an explanation can be provided.

In particular, I will explore whether the dogmatist can provide an explanation of why epistemic justification fails to satisfy *CUT* in a way that solves the problem. I start with some bad news for the dogmatist. Recently, Jonathan Weisberg has attempted to use broadly Bayesian considerations to argue that epistemic justification fails to satisfy *CUT* in a way that solves the bootstrapping problem. But I argue that in fact dogmatism is incompatible with this Bayesian picture (§3).

While I am sympathetic to the reaction that this bad news is a *reductio* of

4. Cf. Vogel (2007: footnote 43) who suggests but does not develop such a solution and Vogel (2000: footnote 30) who is unimpressed by a related strategy. But Vogel's (2000) pessimism is directed at an application of it that says *Step 3* cannot be performed after *Step 2* and as we will see, my application of this strategy is different.

5. §5.2 gives an example that shows why epistemic justification is not transitive in general. But the interesting question is whether it is transitive *over the bootstrapping reasoning*.

6. To see that *CUT* and *MONOTONICITY* entail transitivity, assume the antecedent of transitivity. That is, assume (i) $A \Vdash \beta$ for all $\beta \in B$ (ii) $B \Vdash \gamma$. By applying *MONOTONICITY* to (ii), we have (iii) $A \cup B \Vdash \gamma$. Finally by applying *CUT* to (i) and (iii), we have the consequent of transitivity, $A \Vdash \gamma$.

dogmatism, I spend the rest of the paper trying to make the best of it on the dogmatist's behalf. In particular, I will develop an alternative explanation of why epistemic justification fails to satisfy CUT within a certain foundationalist framework and apply it to the bootstrapping problem (§4–5). In doing this, I will be assuming that if the dogmatist can deny that the bootstrapping reasoning goes all the way through to the end of *Step 5*, this will suffice to solve the problems. I however discuss in detail some recent arguments that suggest that there is a problem much earlier in the reasoning in §5.4 and Appendix C.

But before turning to these tasks, it is worth pausing for a moment to consider whether the problem can be solved by claiming that epistemic justification fails to satisfy MONOTONICITY over the bootstrapping reasoning. This is especially important to do because it is a well-known and relatively uncontroversial commitment of dogmatism that epistemic justification fails to satisfy MONOTONICITY in general.

To see this, we need only note that MONOTONICITY essentially says that if the states corresponding to the sentences in A justify you in believing the proposition expressed by γ then no matter what new information you may learn from some external source you will still be justified in believing the proposition expressed by γ . But dogmatism says that while something looking red epistemically justifies you in believing it is red, if you were to find out from some external source that the lighting conditions are bad for distinguishing red objects from other objects, you would no longer be justified in believing it is red. So dogmatism already entails that epistemic justification fails to satisfy MONOTONICITY.

Despite this fact, it is nonetheless unsurprising that theorists have not thought that rejecting MONOTONICITY could be used to solve the bootstrapping problem. As I noted, MONOTONICITY fails because learning some new information from an external source might lead an agent to be no longer justified in believing something that she was originally justified in believing. Now consider that the bootstrapping reasoning begins with the agent learning that the slide is red and that it looks red. It is a commitment of dogmatism that the agent learns the slide is red. And it is obvious that an agent can learn that the slide looks red from introspection.

For the failure of MONOTONICITY to solve the problem at this stage in the reasoning, we would have to implausibly claim that learning that the slide looks red makes you no longer justified in believing that it is red. This way of trying to solve the bootstrapping problem is especially implausible for the dogmatist. She would be saying that noticing the perceptual state that justifies you in believing the slide is red defeats the justification for that belief.⁷

7. Cf. Titelbaum's (2010: 121) discussion of how appeal to defeaters can't help the reliabilist and Barnett's (in press: §4.3) discussion of this defeater for dogmatists.

Having seen why MONOTONICITY failure cannot help us at the start of the bootstrapping reasoning, we can also see why it will not help in later stages in the argument. It won't help because these steps do not involve learning any new information from an external source. Instead, all of the new claims are inferred from previous claims. And this in turn shows us why rejecting CUT looks to be a more initially promising candidate than rejecting MONOTONICITY. Rejecting CUT says that the fact that you inferred each of the sentences in B from A rather than having them both from the start makes a difference to whether you can conclude γ . In other words, it says that how you inferred a claim may matter for what it can justify.

So it is possible that by rejecting CUT we may say that the bootstrapping way of inferring is what makes it so we cannot perform the steps back-to-back. Our question, then, is what could explain why CUT fails in this way.

3. The Bad News

Jonathan Weisberg (2010) has a simple answer to this question. He points out that so-called *Bayesian epistemologists* have a number of different probabilistic models of epistemic justification. And an interesting fact is that all of these different models of epistemic justification entail that epistemic justification does not satisfy CUT.^{8, 9} In this section, I will present the basic idea behind why these Bayesian approaches entail that epistemic justification fails to satisfy CUT and how this is supposed to solve the bootstrapping problem (§3.1). I will then argue that Bayesian approach is incompatible with dogmatism (§3.2). This argument is similar to and is another member of a family of results demonstrating tension between dogmatism and probabilistic models of justification (§3.2.5).

3.1. A Probabilistic Explanation of Failures of CUT

As I mentioned, there are a number of different ways of modeling epistemic justification using probabilities. But in order to introduce the basic idea, I will adopt one specific and simple way of understanding these issues in Bayesian terms.

I will develop the idea with the help of three assumptions. First is the claim that agents have not just all or nothing beliefs but also have degrees of belief

8. More precisely, Weisberg posits a defeater he calls *No Feedback* (2010: 533–534) and for this defeater obtain it must be that epistemic justification fails to satisfy CUT on probabilistic grounds.

9. Probabilistic models are also often used to model a related issue of transmission failure (see Moretti & Piazza 2013; Okasha 2004). For non-probabilistic philosophical discussion of transmission failure, see Wright (2004) for non-dogmatist perspective and Pryor (2012) for a dogmatist perspective.

or credences and the degrees of belief or credence of a fully rational agent are representable by a probability function. Second is the claim that when a (fully rational) agent learns a new claim for certain, her credences evolve in way that can be modeled as updating by conditionalization. That is, per the first claim, the agent's credences before she learns the new claim are represented by some probability function Pr . And when she learns some new claim p for certain, her new credences are representable by a probability function Pr^* where for any claim q , $\text{Pr}^*(q) = \text{Pr}(q|p)$. The third and final claim is that a (fully rational) agent is epistemically justified in believing some claim only if her credence in that claim is above some sufficiently high (but below 1) threshold t .

With these assumption in hand, we can see how epistemic justification can fail to satisfy *CUT* by considering the example of a lottery. The lottery has one hundred tickets and exactly one ticket is a winner. Let A be a set of sentences describing this lottery. Next let B be the set of sentences 'ticket one loses', 'ticket two loses', . . . , 'ticket fifty loses'. And let γ be the conjunction of all the claims in B . Finally suppose the relevant threshold is .9. Now imagine an agent who learns each of the claim expressed by the sentences in A for certain. This agent now would be epistemically justified in believing that ticket one loses because we may suppose given the description of the lottery, the agent prior to learning the propositions expressed by the sentences in A had credences representable by a probability function Pr such that $\text{Pr}(\text{ticket one loses} | \Lambda a) = .99$ where ' Λa ' expresses the proposition expressed by the conjunction of the sentences in A . And by analogous reasoning the agent would be epistemically justified in believing the proposition expressed by each sentence in B . Nonetheless the agent would not be justified in believing the proposition expressed by γ because conditional on the description of the lottery, all of the first fifty tickets losing is much less probable than .9. And this is true despite the fact that $\text{Pr}(c | a \wedge b) = 1$ where ' c ' expresses the proposition expressed by γ and ' $a \wedge b$ ' expresses the proposition expressed by the conjunction of the sentences in A with the sentences in B .

What's happening here? That all fifty tickets lose is epistemically justified for an agent who knows each of the propositions expressed by the sentences in A and B for certain but not justified for an agent who knows the propositions expressed by sentences in A for certain and therefore is justified in believing the propositions expressed by the sentences in B . This leads to epistemic justification failing to satisfy *CUT* because the kinds of inferences you can perform with certainties differ from the kinds of inferences you can perform with uncertainties. In particular, if you are justified in believing the proposition expressed by each of the members of B but are uncertain, when you conjoin all of these claims the uncertainty of the conjunction may be as high as the sum of the uncertainties of each conjunct. Since this means the conjunction may be much more uncertain than any given conjunct, it may be that you are not epistemically justified in

believing it. This cannot happen for certainties because certainties are not uncertain at all so the sum of their uncertainties is still zero and so the conjunction of certainties must be epistemically justified as well.

So two claims explain why epistemic justification fails to satisfy CUT according to Bayesian epistemology. First, some of the things you are justified in believing are certain and others are uncertain. Second, what can be justified by a certainty differs from what can be justified by an uncertainty.

Despite the fact that we have used conjunction introduction to illustrate how epistemic justification fails to satisfy CUT, Weisberg does not apply the Bayesian approach to the bootstrapping problem by saying we cannot perform conjunction introduction after the other steps. Instead, he denies that we can do the last step, *Step 5*, after all the other steps. This is because it is actually not plausible to reject the conjunction introduction steps in the context of bootstrapping.

Consider *Step 2* that proceeds from a belief that the slide looks red and a belief that the slide is red to the belief that the slide looks red and is red. In general, coming to form beliefs about the external world via perception gives us highly certain (though not totally certain) beliefs. And in general coming to have beliefs about our present visual experiences via introspection gives us highly certain (though not totally certain) beliefs. These kinds of beliefs outside of bootstrapping contexts obviously admit of conjunction introduction. If, for example, I form the belief that the table is red based on visual perception and I form the belief that the chair appears blue based on introspection, I am in normal contexts justified in believing the table is red and the chair appears blue.

In other words, in standard contexts, summing the uncertainties associated with our beliefs about the external world that we arrived at through perception and the uncertainties associated with beliefs about our visual experiences arrived at through introspection does not make the uncertainty so high that we are not justified in believing the conjunction. So the general fact that Bayesian epistemology predicts that epistemic justification does not satisfy CUT is of no help to block this step absent some explanation of why this instance of conjunction introduction is relevantly different than other acceptable instances of it (cf. Vogel 2000: footnote 24).

We can make a similar point with regard to the *Step 4* that proceeds from instances of color vision working to the conjunction of those instances. We want to allow in general that agents can perform enumerative induction of the form that we have in *Step 5*—proceeding from a track record of it holding to it holding generally. So for example, we want agents to be able to conclude that swans are white from track record of instances of seeing white swans. But the standard way in which one amasses such a track record is by separately establishing each of the instances. That is, one comes to believe that some swan *a* is white, then believes that swan *b* is white, and so on.

So if our practices of enumerative induction are to be sound, we must be able to perform conjunction introduction of the sort we have in *Step 4*. So again, absent some special story, the general fact that epistemic justification fails to satisfy CUT is of no help for denying the instances of conjunction introduction.¹⁰

For this reason, it is sensible for Weisberg to deny that we can perform *Step 5* after the other steps rather than deny that we can perform the conjunction introduction step. His idea is that while knowing for certain from the start that it looks red and is red justifies you in believing that your color vision is reliable, this does not mean that finding out that the slide is red from the slide looking red justifies you in believing that your color vision is reliable.

Thus, Weisberg's theory tells us that epistemic justification fails to satisfy CUT based on the independently motivated Bayesian grounds for thinking that epistemic justification fails to satisfy CUT. And as can be seen by how Weisberg proposes to use this idea, his theory also tells us exactly which step cannot be performed after the other steps, *Step 5*.

3.2. *The Incompatibility*

Unfortunately, as I will now argue, Weisberg's account of why it is that we cannot perform *Step 5* after the other steps is incompatible with dogmatism. To present this argument in the most approachable way that I know how, I will begin by simplifying our discussion in three respects.

First, I will ignore the fact that in the bootstrapping reasoning, we look at many slides. Instead, I will just assume that we have a single instance of seeing a slide that looks red. And accordingly I will assume for simplicity that just a single instance of my color vision working is strong evidence that my color vision is reliable. Alternatively, we can think of this simplification as a case in which the agent sees the whole series of slides all at once and the agent is so constituted that she can attend to them in the same way that we can attend to a single slide.¹¹

Second, since we have not yet found any reason to be suspicious of the simple analytic entailment from 'the slide looks red and the slide is red' to 'my color vision worked' considered on its own or after the other steps, we will ignore this

10. Cf. Cohen who writes "we can stipulate that the number of conjuncts is small enough that I remain justified in believing the entire conjunction, but long enough to justify inductive inference. This is possible on the assumption that enumerative induction is possible" (2010: 143).

11. An anonymous referee points out still another helpful way to think about why this simplification is harmless. Suppose instead of thinking about an ordinary agent we are considering a case where "someone with excellent (maybe super-human) color vision looks at a paint chip and can tell that it's some very particular shade of blue (out of arbitrarily many possible shades). She can know that it's arbitrarily unlikely that she'd be right about which very specific shade it is by chance, so if she nevertheless knows which shade it is, this seems like as good evidence for reliability as a series of coarse-grained identifications would be."

step in the bootstrapping reasoning, *Step 3*.¹² Finally, since the previous subsection showed that conjunction introduction is plausible even in the context of the bootstrapping reasoning, we will also ignore the steps that feature conjunction introduction, *Step 2* and *Step 4*.

This leaves us then with just *Step 1* and *Step 5*. With all of our simplifications, these steps look like this: In *Step 1*, the agent has a visual experience of a red slide and concludes the slide is red. And as we did earlier, we assume this agent is justified from the start in believing she is having a visual experience of a red slide. Then in *Step 5* the agent concludes her color vision is reliable based on her belief that the slide is red and her belief that she has a visual experience as of a red slide.

With these simplifications in hand, we can start to think about whether the dogmatist can make use of Weisberg's solution. For the dogmatist to be able to make use of the solution, she must allow that her commitments can be usefully modeled probabilistically. So the dogmatist must allow that we can think of an agent's credences as representable by a probability function. And I will assume that the dogmatist in this setting will say that a belief is justified only if the agent's credence in it is above some threshold t .¹³ So let us, in particular, use Pr_0 for the probability function that represents the agent's credences prior to the slide show.

3.2.1. A Constraint on the Prior Probability of Reliability

The first step in showing Weisberg's solution is incompatible with dogmatism is to isolate a constraint on the prior probability of your perception being reliable that the solution entails. If we use ' r ' to express the proposition that the slide is red, ' lr ' to express the proposition that the slide looks red, and '*reliability*' to express the proposition that your perception is reliable, the constraint is this:

$$\text{Pr}_0(\text{reliability}) \geq \text{Pr}_0(r|lr)\text{Pr}_0(\text{reliability}|r \wedge lr) + \text{Pr}_0(\text{reliability} \wedge \neg r|lr)$$

Though this inequality can be an eyesore to look at initially, with a little work it won't be hard to demonstrate why it must hold if Weisberg's solution is to work.

To start, notice that to solve the bootstrapping problem we really need to be able to show that we do not get *any* new justification for believing that our color vision is reliable. The bootstrapping reasoning not only doesn't newly justify us in believing that our color vision is reliable, it does not even give us more reason to believe this (cf. Cohen 2002: 317; Weisberg 2010: 532; White 2006: 543). In Bayesian terms, this means that our probability that our color vision is reliable should not increase at all upon learning that the slide looks red.

12. But see §5.4 and §C for discussion.

13. §3.2.5 considers how our results fair under weaker assumptions.

So if we let $\Pr_1$ be the probability function that results after the agent learns that the slide looks red, for Weisberg's solution to work it must be that:

$$\Pr_0(\text{reliability}) \geq \Pr_1(\text{reliability})$$

Now by definition of $\Pr_1$, we know that $\Pr_1(\text{reliability}) = \Pr_0(\text{reliability} | l r)$. And it is not hard to prove that:

$$\Pr_0(\text{reliability} | l r) = \Pr_0(\text{reliability} \wedge r | l r) + \Pr_0(\text{reliability} \wedge \neg r | l r)^{14}$$

and that:

$$\Pr_0(\text{reliability} \wedge r | l r) = \Pr_0(r | l r) \Pr_0(\text{reliability} | r \wedge l r)^{15}$$

So putting these together we have that:

$$\Pr_1(\text{reliability}) = \Pr_0(r | l r) \Pr_0(\text{reliability} | r \wedge l r) + \Pr_0(\text{reliability} \wedge \neg r | l r)$$

Finally recall that in order for Weisberg's solution to work it must be that $\Pr_0(\text{reliability}) \geq \Pr_1(\text{reliability})$. So this gets us our constraint:

$$\Pr_0(\text{reliability}) \geq \Pr_0(r | l r) \Pr_0(\text{reliability} | r \wedge l r) + \Pr_0(\text{reliability} \wedge \neg r)$$

This gives us a lower bound the prior probability of *reliability*. The next thing to do to pave the way for my argument is to illustrate why this is a demanding constraint. I do this by showing how high this lower bound must be.

3.2.2. This Constraint Is Demanding

Recall that the solution that we are considering accepts each step on its own, but rejects the claim that they can be performed one after the other. So it accepts *Step 1* and *Step 5* on their own. This tells us something about the value of $\Pr_0(r | l r)$ and $\Pr_0(\text{reliability} | r \wedge l r)$ respectively.

To say *Step 1* on its own is acceptable is to say that if you learned that the slide looks red, you would be justified in believing that the slide is red. In the

14. In general, $\Pr(p) = \Pr(p \wedge q) + \Pr(p \wedge \neg q)$, so $\Pr_1(\text{reliability}) = \Pr_1(\text{reliability} \wedge r) + \Pr_1(\text{reliability} \wedge \neg r)$. And so by the definition of $\Pr_1$, $\Pr_0(\text{reliability} | l r) = \Pr_0(\text{reliability} \wedge r | l r) + \Pr_0(\text{reliability} \wedge \neg r | l r)$.

15. This assumes that $\Pr(p | q) = \Pr(p \wedge q) / \Pr(q)$. So $\Pr_0(r | l r) \Pr_0(\text{reliability} | r \wedge l r) = [\Pr_0(r \wedge l r) / \Pr_0(l r)] [\Pr_0(\text{reliability} \wedge r \wedge l r) / \Pr_0(r \wedge l r)]$. Next $[\Pr_0(r \wedge l r) / \Pr_0(l r)] [\Pr_0(\text{reliability} \wedge r \wedge l r) / \Pr_0(r \wedge l r)] = [\Pr_0(r \wedge l r) \Pr_0(\text{reliability} \wedge r \wedge l r)] / [\Pr_0(l r) \Pr_0(r \wedge l r)]$. Then we have that $[\Pr_0(r \wedge l r) \Pr_0(\text{reliability} \wedge r \wedge l r)] / [\Pr_0(l r) \Pr_0(r \wedge l r)] = \Pr_0(\text{reliability} \wedge r \wedge l r) / \Pr_0(l r)$. Finally by definition, $\Pr_0(\text{reliability} \wedge r \wedge l r) / \Pr_0(l r) = \Pr_0(\text{reliability} \wedge r | l r)$. So $\Pr_0(\text{reliability} \wedge r | l r) = \Pr_0(r | l r) \Pr_0(\text{reliability} | r \wedge l r)$.

Bayesian model that we are working with this means that if you learn the slide looks red your probability that it is red must be t or greater¹⁶:

$$\Pr_1(r) = \Pr_0(r|lr) \geq t$$

Next, to say that *Step 5* on its own is acceptable is to say that if you learned from the start that the slide looks red and is red, you would be justified in believing that your color vision is reliable. In the Bayesian model we are working with this means that if you learn that the slide looks red and is red, your probability that your color vision is reliable must be t or greater:

$$\Pr_0(\text{reliability} | r \wedge lr) \geq t$$

This facts tell us something about how high our lower bound must be.¹⁷

Our lower bound recall is this:

$$\Pr_0(\text{reliability}) \geq \Pr_0(r|lr)\Pr_0(\text{reliability} | r \wedge lr) + \Pr_0(\text{reliability} \wedge \neg r | lr)$$

The work we just did allow us to see how high this lower bound must be:

$$\Pr_0(r|lr)\Pr_0(\text{reliability} | r \wedge lr) + \Pr_0(\text{reliability} \wedge \neg r | lr) \geq t^2 + \Pr_0(\text{reliability} \wedge \neg r | lr)$$

This essentially teaches us that your prior probability in *reliability* must in fact be quite high. After all, it is hard to deny that the threshold, t , must at least be as high as .9. So the lower bound on *reliability* must be at least as high as .81 summed with $\Pr_0(\text{reliability} \wedge \neg r | lr)$. Having to have at least .81 credence in *reliability* prior to the slide show is quite a demanding constraint.

16. This assumes that the dogmatist notion of coming to have justified beliefs based on perception can be modeled as conditioning on the claim that you had the perception. This assumption has been rejected by some who discuss dogmatism (see Morreti 2015 and Barnett in press for discussion). I am sympathetic to this idea, so below in §3.2.5 and §B I show how the argument of this section can be adapted so that it does not rely on this assumption.

17. The comments at the end of §3.1 that explain why conjunction introduction is safe in the bootstrapping context are important for seeing why these assumptions are innocuous given the simplifications that I have made. $\Pr_0(\text{reliability} | r \wedge lr) \geq t$ is innocuous because we are treating $r \wedge lr$ as simplifying proxy for *Track Record*. $\Pr_0(r|lr) \geq t$ then must be understood as saying, in effect, that *Track Record* is probable on each perceptual experience. This too is innocuous because we have argued that not only must the belief that slide one is red, the belief that slide two is green, etc., and the belief that slide one looks red, the belief that slide two looks green be above t , they must be sufficiently above t to allow for conjoining them all together.

3.2.3. Possibly Inconsistent

And the fact that your prior probability must be quite high is incompatible with dogmatism. To begin, it may in fact be that the lower bound is straightforwardly inconsistent with dogmatism. For suppose the values of $\Pr_0(r|lr)$ and $\Pr_0(\text{reliability}|r \wedge lr)$ are not right at the threshold but much higher. This has some plausibility: Perception and enumerative induction are not just any old ways of forming justified beliefs but are among the most epistemically important ways of forming beliefs. Arguably, part of their importance is that they give us beliefs that are not just justified but strongly justified.

If that is right, then it may be that the product of $\Pr_0(r|lr)$ and $\Pr_0(\text{reliability}|r \wedge lr)$ summed with $\Pr_0(\text{reliability} \wedge \neg r|lr)$ is greater than or equal to t itself. For example, if t were to be .96 or less and $\Pr_0(r|lr)$ and $\Pr_0(\text{reliability}|r \wedge lr)$ were to be .98 or greater, the product of $\Pr_0(r|lr)$ and $\Pr_0(\text{reliability}|r \wedge lr)$ alone would be greater than the threshold. And so the inequality would entail that your prior probability in *reliability* is greater than t .

This is flatly inconsistent with dogmatism. As we said, the dogmatist (of the strong sort that we are discussing in this paper) says you can form justified beliefs about the color of objects even without being justified in believing that your color vision is reliable in the first place. If we are to explain why you are not justified in Bayesian terms, we would need to require that your prior probability in *reliability* is less than t . But we just saw that if $\Pr_0(r|lr)$ and $\Pr_0(\text{reliability}|r \wedge lr)$ are sufficiently greater than t , Weisberg's solution would require your prior probability in *reliability* to be greater than t .

3.2.4. Incompatible with the Spirit of Dogmatism

But even if $\Pr_0(r|lr)$ and $\Pr_0(\text{reliability}|r \wedge lr)$ do not have values that are sufficiently greater than t to lead to this simple inconsistency, the fact that the inequality requires us to have high, though perhaps less than belief-level, credence in *reliability* alone is incompatible with the spirit of dogmatism. This is because the reasons that tell against needing to be justified in believing that our color vision is reliable in order to form justified beliefs about the color of objects also tell against needing to be justified in having high credence in our color vision being reliable.

Let me explain. The question of whether we need a prior justified belief that our color vision is reliable in order to form justified beliefs about the color of objects is a pressing one in epistemology because of the *problem of the criterion*. Here is Stewart Cohen's tight presentation of that problem:

A natural intuition [. . .] is that [. . .] sense perception can not deliver knowledge unless we know [. . . it] is reliable. But surely our knowledge that sense perception is reliable will be based on knowledge we have

about the working of the world. And surely that knowledge will be acquired, in part, by sense perception. So it looks as if we are in the impossible situation of needing sensory knowledge prior to acquiring it. [. . .] Skepticism threatens. (2002: 309)

Though Cohen puts the issue in terms of knowledge, it is clear that there is a parallel problem concerning justification. The problem can be resolved in one of two ways: reject the assumption that justification for believing that our color vision is reliable requires prior justified beliefs about the color of objects or reject the assumption that we need to be justified in believing our color vision is reliable prior to forming justified beliefs about the color of objects. The dogmatists adopts the second solution.

What supports going in for the second option over the first? The first option would require that we can have a justified belief in our color vision being reliable even without first having any justified beliefs about the color of objects. This would mean that we first are justified in believing that our color vision is reliable.¹⁸ But since it is contingent whether our color vision is reliable, this, it seems, would amount to countenancing a kind of deeply contingent *a priori* knowledge or justified belief.¹⁹ That's puzzling.

There is, of course, much more to say about deeply contingent *a priori* justification before rejecting it. But my goal here is not to evaluate the dogmatist grounds for rejecting it. My goal is only to describe the grounds that support its rejection: the claim is that since it is a contingent whether our color vision is reliable, we are justified in believing that it is reliable on *a posteriori* grounds—only experience can rule out genuine possibilities.

Insofar as we accept these grounds for rejecting deeply contingent *a priori* justified belief, we should also be uncomfortable with our epistemic space being strongly biased toward our color vision being reliable *a priori* given that it is contingent whether our color vision is reliable. Just as only experience can justify us in ruling out genuine possibilities, only experience can justify us in be strongly biased against certain possibilities. This is an especially plausible point in the present setting where we are assuming a belief that *p* is rational only if it is rational to have an extremely strong bias in favor of *p*. But if only experience can justify such an extremely strong bias, it is hard to see why experience would not be needed to justify a strong bias.

18. Here I put aside the coherentist solution that the dogmatist must also reject that says our justification comes “at the same time”.

19. While Kripke may have introduced convincing cases of contingent *a priori* knowledge, this example does not fit that model (e.g., it does not involve any reference fixing descriptions). I, following Hawthorne (2002) (who takes the term from Gareth Evans), flag this difference by calling the proposition *deeply* contingent.

This is what makes Weisberg's solution incompatible with the spirit of dogmatism. Even if it allows that we can form justified beliefs about the color of objects without first having a justified belief that our color vision is reliable, it still requires that our epistemic space be strongly biased toward our color vision being reliable in order to form justified beliefs about the color of objects. Though this bias may not be strong enough to amount to belief, it is still puzzling (from the dogmatist's perspective at least) why we are justified in being biased toward this contingent proposition on *a priori* grounds.²⁰

Finally, we can bring this same argument into still sharper focus if we look at it from a slightly different perspective. Consider what would happen if we tried to avoid being strongly biased and only exhibited a small bias toward our color vision being reliable, say have .65 credence in *reliability*. Our lower bound entails that in order for this to hold it must be that $\Pr_0(r|lr) < .81$ or $\Pr_0(\text{reliability}|r \wedge lr) < .81$. That is, either enumerative induction or perception only leads us to having .81 credence in our conclusion.

This is problematic not just because it is unintuitive but also because it severely limits the role that beliefs formed in this way can play. In particular it would be difficult to conjoin such beliefs. For example, if we had two probabilistically independent such beliefs, our credence in their conjunction would be just below .66. So if we were only slightly biased, this would entail that we cannot usefully combine the information that we gain through induction or perception. But perception and induction are epistemically important belief forming processes. It is unacceptable to relegate them to this impotence with respect to conjunction introduction.

Thus, even if Weisberg's solution is not flatly inconsistent with dogmatism (of the strong sort that we are discussing in this paper), it is incompatible with the spirit of dogmatism. As a dogmatist, we want to allow that we can form justified beliefs about the color of objects without having our epistemic space be strongly biased toward our color vision being reliable. This is not something that Weisberg's Bayesian solution can allow.

3.2.5. Bayesianism Is Incompatible with Dogmatism

This is not a criticism of Weisberg. His idea may be correct as a diagnosis of what is wrong with the bootstrapping reasoning. It is just of no help to the dogmatist. Our result is essentially a new member to a family of results that have been uncovered in recent years that show dogmatism is incompatible with even relatively modest Bayesian assumptions about justification (see Cohen 2005: 424–

20. Cohen (2005: footnote 1) conjectured that the bootstrapping problem would still arise for views that allowed that one had strong *a priori* evidence that falls short of knowledge or fully-justified belief. Our discussion of Weisberg suggests that this conjecture is incorrect but nonetheless, that this should be of no comfort to the dogmatist.

425; Hawthorne 2004: 73–77; Wedgwood 2013: §2; White 2006). What makes our result different from these other results—and especially relevant to this paper—is that it directly shows how probabilistic failures of *cut* that suggest that the bootstrapping reasoning is defective are incompatible with dogmatism. Since I argued (in §2) that dogmatists can solve the bootstrapping problem only if epistemic justification fails to satisfy *cut*, this probabilistic tension is of special importance to the issues discussed in this paper.²¹

Now in initially presenting the argument of this section, I relied on some fairly strong assumptions to fix ideas. But the assumptions that are needed for the argument are in fact rather modest. Let's look at this: The dogmatist says that having a visual experience of the slide being red newly justifies us in believing that the slide is red. We have assumed that this means that upon learning the slide looks red, your credence in the slide being red will be quite high (e.g., .9 or higher). Enumerative induction is a way of being newly justified in believing that your color vision is reliable based on the slide looking red and being red. We have assumed that this means upon learning the slide looks red and is red, your credence that your color vision is reliable will be quite high (e.g., .9 or higher). Finally, the spirit of dogmatism says that we not only fail to be *a priori* justified in believing our color vision is reliable but also fail to be justified in being strongly *a priori* biased toward thinking that our color vision is reliable. We have assumed that this means your credence in your color vision being reliable before the bootstrapping reasoning is moderate and not strongly biased (e.g., not above .8).

Insofar as there is some connection between epistemic justification and credences understood probabilistically, these assumptions about how that connection works out in this particular case are plausible. They are compatible with many different views about the relationship between full belief and partial beliefs and take no stand on what the correct probabilistic measure of evidential support is. But as we have seen, these assumptions are incompatible with one another.

That said, these assumptions do become less natural if we move away from thinking of learning, as we have, in terms of standard conditionalization and turn to thinking of it in terms of Jeffrey conditionalization. Whereas standard conditionalization requires that we assign probability 1 to propositions that we learn, Jeffrey conditionalization does not require this. Thus, one possible dogmatist solution to this problem is to adopt Jeffrey conditionalization and reject the assumptions that I have been working with.

21. Other results focus on dogmatism's tension with other probabilistic constraints. For example, some results focus on showing that the dogmatist idea that a perceptual experience allows us to newly rule out the possibility of a skeptical scenario obtaining is incompatible with certain probabilistic constraints (§C provides a description of these results and discusses their bearing on the issues of this paper).

It is important here to note that all of the probabilistic arguments against dogmatism (as far as I know) are vulnerable to this response.²² But interestingly, the argument given here (unlike the other arguments) can be made directly in terms of Jeffrey conditionalization without making any assumptions about its connection to standard conditionalization. I leave this generalization, however, to Appendix B because it introduces a number of complications that would distract from the main thread of this paper.

Thus, I conclude that dogmatism is incompatible with relatively modest (albeit not irresistible) connections between justification and probabilities.

4. CUT in Foundationalist Epistemology

One reaction to this result is that it is a *reductio* of dogmatism: Bayesianism is a well-understood, powerful, formal theory of rational belief and dogmatism either suffers from the bootstrapping problem or is incompatible with it. So much the worse for dogmatism.

While I'm sympathetic to this reaction, I wish to make the best of it on the dogmatist's behalf and explore an alternative perspective. Indeed, I believe that this alternative perspective is one that is, in any case, a more natural fit with dogmatism. This is because prominent dogmatists such as John Pollock and James Pryor are also critics of Bayesianism.²³

Now since we are rejecting Bayesianism and since Bayesianism is the only model that we have looked at so far that allows epistemic justification to fail to satisfy CUT, the first thing that we must do is develop an alternative picture of how it could be so much as possible that epistemic justification fails to satisfy

22. Those who give probabilistic arguments against dogmatism (e.g., White 2006: 534–535) standardly respond to the objection from Jeffrey conditionalization by arguing that there is no reason why it should work relevantly differently from standard conditionalization: it would be bizarre, they say, for there to be a relevant difference between learning directly to degree n that the slide is red because you have a perceptual experience of it being red (the Jeffrey conditioning model) and learning indirectly the slide is red to degree n based on learning directly for certain that you are having the visual experience (the standard conditionalization model). Whatever the merits of this response (I myself am of two minds about it and think it is at least not dialectically effective), it is an interesting fact that the argument of this section does not need this response to succeed in order for it to be sound (in particular, I do not need to assume that the $\Pr(r|lr)$ is high). Instead, the argument can be made without making any claim about the connection between Jeffrey conditionalization and standard conditionalization; it can be made directly in terms of Jeffrey conditionalization (see §B). Thanks to John Hawthorne and several anonymous referees for pushing me to consider this issue.

23. Pollock and Cruz (1999) and Pryor (2013) are dogmatist critics of Bayesian approaches. Pryor's (2013) is closely related to Christensen (1992) and Weisberg (2009) and (2015). Together they develop, in my view, the most serious dogmatist challenge to using a probabilistic framework to model justification.

CUT.²⁴ While I cannot develop a theory as powerful and mathematically precise as Bayesianism here, what I will do is show how a broadly foundationalist epistemology has the resources to explain why epistemic justification fails to satisfy CUT.

I begin by developing a minimal account of what foundationalism is (§4.1) and then frame the question of whether epistemic justification satisfies CUT within this perspective (§4.2). After this, I will show how it could be that epistemic justification fails to satisfy CUT (§4.3). This discussion of how it is possible for epistemic justification to fail to satisfy CUT closely follows ideas that I have presented in earlier work (Nair 2019: §5). The task of Section 5 will be to explain how this possibility can be applied to the bootstrapping problem.

4.1. Bare Bones Foundationalism

There are two components to the broadly foundationalist framework that I will be employing: one psychological, the other epistemological. The psychological component is that an agent's epistemic states are structured so that at a given time, some of her attitudes (beliefs, perceptual states etc.) count as foundational and others count as non-foundational. And the non-foundational states are based on the foundational states.

Now different theories will have different accounts of what the foundational states are. Dogmatism, for example, is committed to the idea that some perceptual states are foundational states. A Cartesian might say it is indubitable beliefs that are foundational. But the general foundationalist psychological claim is neutral about this.

The epistemological component consists of two theses. The first thesis is that the justification of a non-foundational state is determined by the relation in which these state stand to foundational states. The second thesis is that the justification of non-foundational states is not determined in this way. On traditional approaches, foundational beliefs will somehow be immediately justified and non-foundational beliefs that can be inferred from them by a process of good reasoning will be justified in virtue of this relation. According to dogmatism, a perceptual state is not the kind of thing that can be justified and is a foundational state. And the non-foundational belief that the world is the way it is represented as being by that perceptual state is justified in virtue of you being in that perceptual state.

So foundationalism consists of a claim about the structure of an agent's

24. In fact, as I observe in Nair (2019), all formal theories other than probabilistic theories and certain systems of inheritance reasoning entail that epistemic justification satisfies CUT. Nair (2019) also considers from an abstract perspective whether this convergence among formal theories reflects a deeper truth about justification and argues it does not.

psychology together with a claim about the epistemological significance of this structure. Specific theories will tell us which states are foundational states and what epistemological role these states play. And similarly, they will say which states are non-foundational states and how the epistemic status of these states is dependent on the epistemic status of foundational states.

This, then, is the minimal foundationalist framework that I will be working with.

4.2. Interpreting Structural Properties of Justification

I now want to explain how to think about the different claims about the structure of justification that we have been discussing. In order to do this, it helps to start by looking at the structure of what we might think of as standard cases of learning.

Consider, then, the following:

Begin by focusing on what is on the left. As before, A is a set of sentences and we are meant to understand this diagram as saying the set of mental states corresponding to the sentences in A are the foundational states that the agent has. The $\uparrow$ represents epistemic justification. So it says that the mental states corresponding to A justify the non-foundational mental states corresponding to α and β respectively. As we move to the right, we have *Add γ* and this tells us that the agent comes to be in the state corresponding to γ and this state is a foundational mental state for that agent. We then see that this new collection of foundational states justify not just the non-foundational states corresponding to α and β but also a new non-foundational state corresponding to δ .

This is the structure of standard cases of learning in the sense that typically when we learn some new information we acquire additional new justified non-foundational states. With this structure in mind, let's consider the two properties that I mentioned earlier, *MONOTONICITY* and *CUT*.

As I said earlier, it is relatively uncontroversial that epistemic justification does not satisfy *MONOTONICITY*. In foundationalist terms, this means that we can have a situation like this:

That is, adding a new foundational state can make it so you lose justification for being in the non-foundational state corresponding to β .

The easiest way to understand what *CUT* is about is to compare it to standard cases of learning. The natural way to interpret the diagram of standard cases of learning is so that γ is distinct from the sentences in A as well as α and β . But *CUT* concerns cases where γ is the same as one of these, say the mental state corresponding to β :

If epistemic justification satisfies *CUT*, then the situation depicted by this diagram cannot arise. If it fails to satisfy *CUT* then it can arise.

So far then, we have an interpretation of each of the structural claims about epistemic justification that we are interested in. With this in hand, I want to turn to showing how it might be that epistemic justification fails to satisfy *CUT*.

4.3. How Justification Might Not Satisfy *CUT*

The first thing to notice is that the foundationalist framework is compatible with epistemic justification satisfying *CUT* as well as epistemic justification failing to satisfy *CUT*. Whether epistemic justification satisfies *CUT* turns on the following question: can adding a state corresponding to β to the foundations lead you to have at least one new justified non-foundational state if you were already justified in having a state corresponding to β among your non-foundational states? As far as I know, no major foundationalist has posed or answered this question. Nonetheless, it is a substantive question that is not answered by the bare bones foundationalist framework.

And in fact, it is somewhat unsurprising that foundationalism is compatible with epistemic justification failing to satisfy *CUT*. After all, according to foundationalism, the foundational states are epistemically special in that they get their epistemic status in a different way than non-foundational states. Epistemic justification failing to satisfy *CUT* suggests that they are special not just in how they get their epistemic status but also in what they are capable of justifying. In light of this, it is natural to wonder what we would have to add to the bare bones foundationalist theory in order for it to entail that epistemic justification fails to satisfy *CUT*. I will answer this question in an abstract in this subsection. And then in Section 5 I apply this answer to the bootstrapping problem.

To begin, let's consider an informal way of thinking about defeasible justification. One way to think about the idea that having a visual experience as of

something being red defeasibly justifies you in believing that it is red is as saying that the visual experience justifies you in believing that it is red *unless* you believe that the lighting conditions are bad, believe that your visual system is malfunctioning, etc. In general, foundationalism can be informally thought of as making claims like: the proposition expressed by α justifies you in believing the proposition expressed by β unless such and such conditions holds. And according to foundationalism what goes in these 'unless'-clauses are some claims about what other epistemic states you are in.

With this informal picture in mind, we can notice that the foundationalist psychology does not just claim that agents have epistemic states. It also claims that the epistemic states of an agent are *organized* in a certain way. And this means that the foundationalist could mention these facts about the organization of our beliefs in the 'unless'-clauses as well.

To explain why this might be a promising idea, I want set dogmatism aside for a moment and work through an example in terms that will be most congenial to more traditional forms of foundationalism. In particular, one form of foundationalism says that the foundational states are the beliefs that things seem to you a certain way that are the direct result of perception (rather than the foundational states being the perceptual states themselves). According to this theory, we have foundational beliefs like the belief that *o* seems to have a rough surface. And this belief counts as defeasibly justified because it is the direct result of a perceptual experience. This belief in turn defeasibly justifies the non-foundational belief that *o* has a rough surface. Now suppose also that you are justified in believing that your sense of touch is malfunctioning. Are you justified in believing that *o* has a rough surface or not? One natural (albeit not inevitable) answer is that it depends on which perceptual experience your belief that *o* seems to have a rough surface was the direct result of. If that belief was the direct result of visual perception, then it is plausible to think that you are still justified in believing that *o* has a rough surface. But if it was the result of touching, you are not justified in believing this.

And one plausible (albeit not inevitable) way of getting this result is to allow 'unless'-clauses that make reference to what your belief was the direct result of. In particular, we might say that the belief that *o* seems to have a rough surface and the belief that your sense of touch is malfunctioning justify you in believing that *o* has a rough surface unless your belief that *o* seems to have a rough surface is based on touching.

This case illustrates that how your belief came about might affect what it can justify. There are, of course, alternative ways of treating this case. But my only goal here is to sketch one plausible way of treating this case that involves an 'unless'-clause that mentions not just other epistemic states but how those states came about as well. What's more, I believe that there is some plausibility to the

idea that there are true claims about defeasible justification that have ‘unless’-clauses of this sort.

This is because these kinds of claims occupy a strategic position between two alternative kinds of clauses that we could use to handle these cases. In these cases, I have been suggesting (and here I am simplifying a bit) that an epistemic state corresponding to α justifies an epistemic state corresponding to β unless the epistemic state corresponding to α came about in a certain way. The two main alternatives to this suggestion would be to simply deny that the epistemic state corresponding to α justifies the epistemic state corresponding to β or to claim that the epistemic state corresponding α justifies the epistemic state corresponding β no matter how it came about. The proposal I am sketching strikes a strategic balance between these by not being as conservative as the first option or as liberal as the second.

So the example that I sketched together with the fact that this claim occupies a kind of strategic position argue in favor of thinking that claims about defeasible justification might have ‘unless’-clauses that mention how an epistemic state comes about. Now the example that I gave does not involve a failure of *cut* but it does illustrate the significance of how a belief came about by considering the significance of two different ways the belief that *o*’s surface appears to be rough can come to be a foundational state.²⁵ So we might also consider different ways in which a state can come about. For example, we can consider the difference between a state coming about in the way a foundational state comes about and a state coming about in the way a non-foundational state comes about.

Here is how that might go. Suppose (1) the epistemic states corresponding to *A* defeasibly justify the epistemic state correspond to β and (2) the epistemic states corresponding to *A* together with the epistemic state corresponding to β defeasibly justify the epistemic state corresponding to γ . And suppose that if we unpack (2)’s ‘unless’-clause it would say the epistemic state corresponding to *A* together with the epistemic state corresponding to β justify the epistemic state corresponding to γ unless the epistemic state corresponding to β is solely based on the epistemic states corresponding to *A*.

25. It is also important that this first case applies to a particular form of non-dogmatist foundationalism according to which the causal etiology of our beliefs about appearances may matter. The case to follow (the one that illustrates a failure of *cut*) is consistent with a dogmatist setting in which our foundational states may be perceptual states themselves and the causal etiology of these states is irrelevant. It has been pointed out to me by an anonymous referee that some dogmatists have argued in response to concerns about cognitive penetration that the causal etiology of our states is irrelevant (see Siegal 2012 and 2013 for the concern and, e.g., Huemer 2013 and Tucker 2010: §6.1 for the response). The case that illustrates the failure of *cut* does not assume the causal etiology of our states (in general or our perceptual states in particular) is relevant; it only assumes the current basing structure of our states is relevant (a claim which dogmatist are in any case committed to).

Let's look at what these claims predict. If we add the supposition that the epistemic state corresponding to δ defeasibly justifies the epistemic state corresponding to β , they predict that we can have a situation like this:

Though this case is not what I called a standard case of learning, it shares its structure. In standard cases of learning the new information plays a more or less direct role in justifying the new epistemic state. But in this case, it is the states corresponding to A and the state corresponding to β that justify the new state. What the state corresponding to δ does is make it so the state corresponding to β is no longer solely based on the states corresponding to A and this makes it so the 'unless'-clause of (2) is no longer true.

So in that example the proposition expressed by β 's justificatory role changes by becoming justified by the state corresponding to δ . But the proposition expressed by β 's justificatory role could also change if the state corresponding to β came to be part of the foundations. That is, we could have a situation like this:

And now notice that this situation is one that demonstrates how epistemic justification might fail to satisfy CUT. Thus, once we admit that claims about epistemic justification might have 'unless'-clauses that mention how beliefs come about, it is easy to come up with abstract claims about epistemic justification that would lead to epistemic justification failing to satisfy CUT.

And since I have argued that such claims are plausible, this shows that within a purely qualitative foundationalist epistemology we have the resources to explain why epistemic justification does not generally satisfy CUT (cf. Weisberg's 2010: 536 discussion of evidentialism).

5. The Solution

Of course, the fact that epistemic justification fails to satisfy CUT in general does not by itself show that *Step 5* cannot be performed after the other steps. But it is not hard to see how it can be applied in a way that secures this result.

Consider the claim that *track record* defeasibly justifies an agent in believing *reliability*.²⁶ There are many well-known defeaters for inductive justification. For example, there are defeaters concerning the existence of competing evidence and concerning certain defective ways of gathering data. So the ‘unless’-clause associated with the relevant claim about justification will mention these defeaters. I wish to add another defeater to this ‘unless’-clause: I posit that the ‘unless’-clause says ‘unless the belief that *track record* is solely based on the perceptual state of the slide being red and introspective awareness of this state’.^{27, 28} This solves the bootstrapping problem.

5.1 How the Solution Works

Let’s look at this in detail. As we know from our work in Section 4, this kind of ‘unless’-clause predicts two things. One thing it predicts is that in cases where an agent’s belief that *track record* is based on something other than just the perceptual state of the slide being red and introspective awareness of this state, the belief in *reliability* may be justified. And this is the right result. For example, if we learn the slide is red from testimony from the person who set up the slide show, this would allow us to gain some justification for *reliability*.

Of course, we use perception in learning by testimony. But my solution does not say that this is problematic. The account says justification for *reliability* is defeated if the belief that the slide is red is solely based on the perceptual experience of it being red. But the account does not say justification for *reliability* is defeated if the belief that the slide is red is based solely on other perceptual experiences such as the experience of the person who set up the slide show saying that the slide is red.²⁹ In this way, the solution allows that we can learn *reliability* in the appropriate circumstances.³⁰

26. Whereas ‘*Track Record*’ refers to the sentence given at the beginning of the paper, ‘*track record*’ expresses the proposition expressed by that sentence.

27. Though there are a number of counterexamples to this formulation of the ‘unless’-clause, I will work with it in what follows because it is the simplest way to capture the spirit of my proposal. A formulation designed to avoid these counterexamples would read something like ‘unless the belief that the slide is red that is involved in supporting *track record* is partially essentially based on the perceptual experience of slide being red’. Thanks to John Hawthorne for asking me to consider these examples.

28. I recently learned of the NSC defeater in Butzer (2017: 2090) that is similar to the defeater here. While I do not essentially disagree with Butzer, his work does not situate the defeater in the general context about the structure of justification (with both the problems based on probabilistic results and the prospects based on the informal model developed in the previous section) that is center stage in this paper.

29. This shows how the present account would treat the case of Eliza in Weisberg (2010: 530; cf. White 2006: 530) and thereby avoid a problem that plagues Vogel’s (2007) views.

30. Cf. Weisberg (2010: 538–539). It may also be worth noting that this response assumes (to my mind, harmlessly) that the dogmatist will individuate the reasons that we have from testimony

The solution also predicts that in the bootstrapping reasoning, we are not justified in believing *reliability*. This is because in this case, our belief that *track record* originates just in the perceptual experience of the slide being red and introspective awareness of this state. And the ‘unless’-clause says that this is a case in which *track record* does not justify *reliability*.

To appreciate what is going on here, let’s contrast this solution with merely claiming that the bootstrapping reasoning is bad reasoning.³¹ Recall that the bootstrapping problem relies on a number of claims. It requires the claim that enumerative induction, two instances of conjunction introduction, and a simple analytic entailment are justified inferences considered alone. And it also requires the claim that over the bootstrapping reasoning, epistemic justification satisfies transitivity so that each of these inferences can be performed one after another. If all of these claims are true and *reliability* is not justified, then dogmatism is false.

To only say that bootstrapping reasoning is bad would raise the question of which of the assumptions that the reasoning is based on is false. And this, in turn, would raise the specter of dogmatism being false because the reasoning might “rollback” all the way to the dogmatist’s claim.

My solution, by contrast, says exactly where the reasoning fails and does so in a way that inoculates dogmatism from the bootstrapping objection. My solution says that enumerative induction, two instances of conjunction introduction, and the simple analytic entailment are all justified inferences considered alone but cannot be performed one after another because epistemic justification fails to satisfy transitivity over the bootstrapping reasoning. More precisely, it fails to satisfy *CUT* over the bootstrapping reasoning. And Section 4 developed a picture of why epistemic justification fails to satisfy *CUT* in general: it showed how within a broadly foundationalist epistemology, ‘unless’-clauses that reference how epistemic states come about could lead to epistemic justification failing to satisfy *CUT*.

What I have claimed in this section is that the enumerative induction involved in *Step 5* has an ‘unless’-clause of this type that says ‘unless the belief that *track record* is solely based on a visual experience of the slide looking red and introspective awareness of this experience’. This claim about defeasible justification occupies exactly the strategic position that I identified earlier. The alternatives to it would be to reject *Step 5* even on its own or to accept it even in the context of bootstrapping reasoning. And we have already looked in detail at why both of these options are undesirable.

to believe that the slide is red from the reasons we have from the experience of the slide being red to believe the slide is red. See Black (2008; especially footnote 9) and Vogel (2008: footnote 17 and 527–528) for closely related issues concerning individuation.

31. Cf. Weisberg’s (2010: 537–538) response to a similar worry and Vogel’s (2000: 615–619) objection to the claim that bootstrapping is unreliable.

This, then, is my solution to the bootstrapping problem on behalf of the dogmatist. Let me close by discussing two worries for this solution and a possible alternative to it.

5.2. *The Non-Independence of the Solution*

The first worry is that I have not given an independent argument for this defeater. And that is true.

What I have done instead is argue that the only solution that the dogmatist can give to the bootstrapping problem is one that says epistemic justification fails to satisfy CUT. And after showing that this is problematic on probabilistic grounds, I tried to make the best of this strategy on the dogmatist's behalf. And the result is that this defeater is the dogmatist's way out of the problem. Insofar as we are attracted to dogmatism and find the bootstrapping reasoning unacceptable, my hypothesized defeater is the best explanation of how both of these views could be true.³²

This feature of my solution may be unsatisfying to some. But there are, I believe, two important points that help to put this shortcoming in perspective. I begin with the more general of the two points.

Consider ordinary perceptual defeaters. For example, consider finding out that the lighting conditions are abnormal and what this does for perceptual justification for believing the slide is red. It is widely agreed that this is a defeater for the justification for believing that the slide is red. And this is because it is a natural way to reconcile our theoretical commitments about perceptual justification and our pre-theoretical judgments about this case.

There is considerably less agreement about what *independent* explanation can be given of why this is a defeater. Some (e.g., Pollock & Cruz 1998) believe that it is a basic feature of the concepts involved (e.g., the red concept) that this is a defeater of perceptual justification. And they believe that the way we find out that this concept has this feature is by looking at cases like this and seeing what is required to get the right results about them. Others believe that we can say more, for example, by looking at the nature of perception itself.

32. This underestimates the resources available to independently motivate this defeater. First, this defeater may be a specific instance of a more well-known defeater (e.g., a biased sample defeater). Second, the defeater that I have posited captures the intuitions that support a number of principles that have been offered in the literature such as *The No Self-Support Principle* (Bergmann 2000: 168; Cohen 2002: 319; Fumerton 1995: 180), the *Independence Principle* (Black 2008: 606–609; Cohen 2005: 428; Markie 2005: 414), the *No Rule Circularity Principle* (Vogel 2007: 531), and certain epistemic difference-making principles (Comesaña & Sartorio 2014). Third, Butzer (2017: §4ff.) argues for a similar defeater and defends it against objections. I do not myself wish to endorse (or deny) any of these claims as motivations for the defeater. Instead, in work in progress, I try to motivate this defeater by appeal to a certain general feature of the sort discussed in Footnote 35.

Still others might appeal to probabilities and note that the conditional probability of the slide being red on it looking red and the lighting conditions being abnormal is low. This presupposes that our conditional probabilities have this feature and raises the question of what explains it. And here again we find different answer, some are subjectivists and some think there are more objective claims that can explain this.

In this respect, the bootstrapping problem and my solution to it is no more problematic than the problem of justifying ordinary perceptual defeaters for the dogmatist. There are many possible answers but no particular answer is obviously correct. No doubt this is a hard problem, but it is not obvious that there is any special problem with explaining the defeater that I posit.

The second point is more specific to the kind of defeater that I am posting here. The defeater that I posit says that the structure of one's mental states is important. In particular, it is important whether there is a state other than the visual perception of a red slide that can justify you in believing that the slide is red. According to my development of dogmatism, if you lack such a state, the defeater is triggered. What I wish to point out now is that everyone (dogmatist and non-dogmatist alike) should accept that whether such a state is available is epistemically significant.

In particular if we consider whether a given collection of states *improves* one's epistemic position, all the going theories agree that the availability of an alternative basis for your beliefs about the color of the slides matters. Obviously, no theorist thinks that the bootstrapping reasoning *improves* one's epistemic position; that's just the datum that we began the paper with. This paper develops a dogmatist account of why this is so but non-dogmatists have their own account. They reject a commitment of dogmatism: They typically claim one must be justified in believing that *reliability* prior being justified in believing the slide is red based on visual perception. They add that the bootstrapping reasoning itself adds nothing to your justification for believing that *reliability*.

On the other hand, theorists should agree that if you were somehow directly justified in believing that the slide is red by a state other than the visual perception of a red slide, this would improve your epistemic position with respect to the belief that *reliability*. And theorists should agree if you were somehow indirectly justified in believing the slide is red in a way that did not depend on having a visual perception of a red slide, this too would improve your epistemic position.

For example, testimony that the slide is red from the person running the slide show would improve your epistemic position with respect to the belief that *reliability* (perhaps by increasing your amount of justification or by making one's justification more resilient to defeat). According to some views, testimony that (or hearing that) the slide is red directly justifies you in believing that the slide is

red.³³ According to other views, testimony indirectly justifies you believing the slide is red.³⁴ Either way, your epistemic position improves and this is because of the availability of an alternative basis for your beliefs.

Once we accept that even non-dogmatists take the availability of alternative bases for your belief that the slide is red to be epistemically significant, it strikes me a special pleading to suggest that the dogmatist's own way of accounting for the epistemic significance of this is especially problematic.

Thus, both on the general grounds of how defeaters are typically explained, and on the more specific grounds that the availability of alternative bases for believing the slide is red is epistemically significant, I believe there is no special problem for the dogmatist. Indeed, as we have seen, the particular brand of dogmatism that I have developed on which *cut* fails is one on which the existence of defeaters of this type is unsurprising and natural (notice also here that the two structures that improve your epistemic positions correspond exactly to the abstract structures in §4.3). So it is not surprising that such a defeater might hold the key to the solution to the bootstrapping problem. For this reason, we should not reject my solution because we lack an independent explanation of the defeater that I posit (that said, I briefly sketch how to start to give an explanation for this defeater in a note.)³⁵

33. This is one development of the Reidian non-reductionist approach to testimony that is perhaps best represented by Burge (1993).

34. Lackey (2006) provides a rich introduction to the variety of views about testimony.

35. Begin by noting that one reason why justification is defeasible is that it is ampliative in the sense that a proposition that you justifiably believe, say *p*, can justify belief in another proposition, say *q*, without entailing it. This opens the door for an agent to learn new information *d* that is consistent with *p* but is inconsistent with *q*, makes *q* unlikely, or makes *q* otherwise unreasonable. Thus, the ampliative relations among propositions are one source of defeaters in the sense that allowing for defeaters is the only way such ampliative leaps could be reasonable. But notice that in forming beliefs, we not only take risks related to the propositional relations among our beliefs but also with respect to our belief forming processes themselves. For example, even when we form beliefs based on deductive inference from other things that we believe, the mental process that subserves this transition is not completely reliable. This makes it natural to allow defeaters that are sensitive to this way in which our rational belief formation incurs risk. And once we admit that justification should be sensitive to the propositional relations among beliefs as well as the nature of the belief forming processes themselves, this opens the door for the details of the structures of those processes to influence which things are defeaters. (One way to think of this is as saying how you *have* your evidence is epistemically important; cf. Neta 2013.) It is here that within a foundationalist framework which wishes to ward off *a priori* knowledge of deeply contingent propositions that I believe an explanation for the defeater that I posit can be found. These remarks require further development (of the sort that I give in work in progress), but it should be enough to provide some license for optimism that such an explanation can be provided.

5.3. *How Epistemic Justification Spends*

A second objection can be gleaned from the work of Peter Markie (2005: 415; cf. Cohen 2005: 428). Markie criticizes another attempt to solve the bootstrapping problem for being incompatible with the following intuitive idea about epistemic justification:

If it is reasonable for us to believe p and the truth of p increases the likelihood that another proposition, q , is true, then p is a reason (perhaps defeasible) for us to believe q . Just as money, however gained, still spends the same, so too reasonable beliefs, however gained, still epistemically support the same beliefs. (2005: 415)

Now we have already discussed in some detail how probabilistic considerations are in fact deeply at odds with the dogmatist's theory and decided to put them aside and try to make the best of it. So we will not concern ourselves with the aspects of Markie's remarks that rely on probabilities.

However, Markie's more general idea that "reasonable belief, however gained, still epistemically support the same beliefs" is plausible in its own right but is incompatible with my solution. This is because my solution says that the belief that the slide is red, when gained by visual perception of the slide being red, does not support the same beliefs as the belief that the slide is red gained by testimony from the person who set up the slide show. Nonetheless, I wish to now argue that epistemic justification is not in fact like money in the way Markie suggests.³⁶

To see this, recall that we are pursuing the idea that epistemic justification does not satisfy transitivity because it fails to satisfy CUT. But we also noted that epistemic justification may fail to satisfy transitivity because it fails to satisfy MONOTONICITY. To illustrate, consider a case where a certain body of evidence that can be represented by the set A justifies concluding γ by induction, but A together with, for example, 'the evidence gathered is biased' does not. This leads to a failure of epistemic justification to be transitive as well, for it is plausible that believing the big conjunction of the propositions expressed by the sentences in A and 'the evidence gathered is biased' justifies believing each of the propositions expressed by the sentences in A . So we have it that the big conjunction justifies believing each of the propositions expressed by the sentences in A and the propositions expressed by the sentences in A justifies the propositions expressed by γ , but the big conjunction does not justify the proposition expressed by γ .

36. As a referee pointed out to me, a different challenge to Markie's principle comes from the literature on so-called transmission failure, see Footnote 9.

This shows why Markie's principle is wrong. A, however gained, does not support γ . A gained by deduction from the big conjunction, for example, does not support γ . Now Markie may respond to this by pointing out that he is talking about *defeasible* justification and so A however gained does defeasibly support γ . It is just that the support is defeated in this case. In particular, it is undercut so that there is no actual support for A from γ in this case.

I am happy with this response. But it entails that my solution is in fact compatible with Markie's principle. For I may say that *track record* defeasibly supports *reliability* however gained. It is just that in the bootstrapping reasoning the support is defeated. It is undercut so that there is no actual support for *reliability* from *track record*.

Thus, everyone must reject or qualify Markie's principle. And so qualified, the principle is compatible with my idea. Of course, what defeats the justification in the example that I just gave is another belief whereas in the bootstrapping reasoning the *organization* of your beliefs defeats the justification. But we have developed a general theory on which we can make sense of exactly how and why the organization of your beliefs matters.

5.4. An Alternative Approach?

We close our discussion by considering an alternative solution to our problem that is suggested by the model developed in Section 4. The approach posits that there is a defeater for the inferences involved in *Step 3*; the inference from 'the slide is red and it looks red' to 'my color vision worked this time'.³⁷

The idea is that the claim that the slide looks red and is red defeasibly justifies the claim that my color vision worked this time. But, according to this alternative implementation, the "unless-clause" that mentions the defeaters of this inference reads "unless the belief that the slide looks red and is red is solely based on the perceptual state of the slide being red and introspective awareness of this state" (cf. Black 2008; Weisberg 2010: §4.3).

There is much to recommend this alternative approach: According to the original approach one can safely reason to believing *track record*. But it may seem that believing *track record* is as problematic as believing *reliability*. And it may seem puzzling that you can believe that you have track record of great success but be unsure about whether you are reliable.³⁸

37. Thanks to an anonymous referee and the editor for encouraging me to address this issue and for helpful suggestions about how to address it.

38. For example, White (2006: 538; cf. Cohen 2010: 144) suggests that the best explanation of our color vision working is our color vision being reliable so getting this far is a problem. My approach, of course, would simply posit the same defeater for inference to the best explanation that we posited for enumerative induction.

Though this paper is officially neutral about whether this alternative approach is better than the approach that we have been exploring, I tentatively prefer the original approach. To see why, notice the following difference between the reasoning in *Step 3* and the reasoning the final step, *Step 5*: *Step 3* is *deductive* reasoning while *Step 5* is *inductive* reasoning. This is important because induction is a form of *ampliative* reasoning, reasoning where the truth of the premises does not guarantee truth of the conclusion. For any piece of ampliative reasoning, it is always possible for an agent to acquire some new information compatible with her premises that is inconsistent with the conclusion or makes the conclusion otherwise improbable or unreasonable. When an agent acquires such information, it is no longer reasonable to maintain belief in the conclusion. This is why any ampliative step of reasoning will include certain “unless-clauses”.

But since *Step 3* is a piece of deductive and hence non-ampliative reasoning, these same grounds cannot be given for thinking there is a defeasible relationship between the premise and conclusion.³⁹ More generally, rejecting *Step 3* commits one to rejecting the following so-called *single premise closure* principle:

if S is justified in believing *p* and competently deduces *q* from *p* then S is justified in believing *q*

as one can (we may assume) deduce ‘my color vision worked’ from ‘the slide is red and it looks red’. Many philosophers (including myself) have found this principle hard to deny.⁴⁰ For these philosophers, this is a serious cost of rejecting *Step 3*.⁴¹

There are also important differences between merely accepting *track record* and accepting *reliability* that still makes it worth stopping the reasoning at the final step. In particular, reliability information has broader effects on our epistemic state. If we have justified beliefs about how reliable some belief forming process is, this may affect how much justification we can get from this process. The same does not hold for mere track record information.

And though it may initially seem strange for an agent to have a long track record of success but be unsure about how reliable she is, this is the correct attitude for the agent to have. If the agent prior to beginning the reasoning considers whether amassing a track record through using this reasoning should change

39. Of course, as we have seen, the model in Section 4 allows for defeaters of a distinctively structural sort as well. And this may give us room for positing defeaters for even deductive inference (see Footnote 35). But the point in the text is that there is no independent prior reason to think *Step 3* is a defeasible inference whereas there is for *Step 5*.

40. It is a delicate matter how best to state the principle of single premise closure. We pass over this important issue.

41. Rejecting this principle will also make it difficult for the dogmatist to offer their typical Moorean response to skepticism. To some (such as myself), this is a serious cost as well.

her mind about her reliability, her answer should be ‘no’. This same answer is appropriate after the reasoning.

More generally: There can be an undercutting defeater of an inductive inference even if the defeater does not suggest an alternative better explanation of the evidence that you have. To show that some conclusion *C* does not inductively follow from *E* in a given case is not thereby to show that there some alternative conclusion *C'* that does inductively follow from *E*. So it is not surprising that an agent may not conclude *Reliability* from *Track Record* even in cases where the agent has no an alternative better explanation of the track record that she has amassed.

That said, these brief considerations are far from decisive. And there are still other powerful arguments suggesting that stopping the reasoning at *Step 3* is important. I leave further discussion of these issues to Appendix C. And I leave it to the reader to judge which solution is best.

6. Conclusion

Overall, then, the results of this paper are mixed for the dogmatist. We have seen that the only way for the dogmatist to solve the bootstrapping problem is for her to claim that epistemic justification does not satisfy *CUT* (§1–2, §A). And this is some bad news for the dogmatist because probabilistic considerations suggest that this solution is not compatible with dogmatism (§3, §B). But we then tried to make the best of it on the dogmatist’s behalf by developing an alternative non-probabilistic framework in which to implement our solution. The framework we used was a broadly foundationalist epistemology (§4). We saw that in this framework, we can explain why epistemic justification fails to satisfy *CUT* in a way that solves the bootstrapping problem (§5).

Though this doesn’t solve all of the problems for dogmatism (§5.2, §5.4, §C), it, I submit, is the best the dogmatist can do by way of solving the bootstrapping problem. Whether the dogmatist best is good enough depends on whether sufficiently powerful alternatives to probabilistic frameworks for understanding justification can be developed, whether deeply contingent *a priori* justification is really problematic, and whether alternative theories can capture the attractive feature of dogmatism. For these reasons, it is a question best left for another day.

A. A More Realistic Formalization

In the main text, I simplified things by giving an equivocating formalization of bootstrapping reasoning. Here I give a non-equivocating formalization that

gives a richer albeit more complicated picture of the bootstrapping problem and our solution.⁴²

To begin, we will need two relations $\models$ and $\models$ to give a non-equivocating formalization. These will both be relations between mental states. To see how the first relation works, let $|V(p)|$ refer to the mental state of visually perceiving p (where we, perhaps, abuse language by treating ‘perceiving’ as non-factive) and let $|Bel(p)|$ refer to the mental state of believing p . Then using $\models$ we formalize the dogmatist idea that perceiving justifies you in believing, as accepting each instance of the following scheme for suitably restricted values for p :

$$V(p) \models Bel(p)$$

$\models$, then, represents cases where *merely being* in the state on the left *justifies* you in being in the state on the right.

$\models$, on the other hand, is used to discuss cases in which being *justified* in being in a certain mental state *justifies* you in being in another mental state. So for example

$$Bel(\text{snow is white}) \models Bel(\text{snow is white or grass is green})$$

is to be understood as saying that being justified in believing that snow is white justifies you in believing that grass is green.

With these two relations in hand, we can now give a non-equivocating formalization of the reasoning. To start, we assume that you have a perceptual experience of the slide being red— $V(\text{the slide is red})$ —and an introspective experience as of having a visual perceptual experience of the slide being red— $I(\text{the slide looks red})$. From here, the rest of the reasoning proceeds in the following six steps.

- Step 1 $V(\text{the slide is red}), I(\text{the slide looks red}) \models Bel(\text{the slide is red})$
- Step 2 $V(\text{the slide is red}), I(\text{the slide looks red}) \models Bel(\text{the slide looks red})$
- Step 3 $Bel(\text{the slide is red}), Bel(\text{the slide looks red}) \models Bel(\text{the slide looks red and is red})$
- Step 4 $Bel(\text{the slide looks red and is red}) \models Bel(\text{my color vision just worked})$
- Repeat
- Step 5 $Bel(\text{color vision just worked } 1), \dots, Bel(\text{color vision just worked } n) \models Bel(\text{my color vision worked } n \text{ times})$

42. Thanks to Andrew Bacon for pushing me to discuss the problem that this section addresses.

Step 6 Bel(my color vision worked n times) $\models$ Bel(my color vision is reliable)

Next once we have these six steps, the issue of whether these steps can be performed one after the other turns on whether the two relations satisfy two structural principles. To state these principles in a way that is analogous to the way that we stated them above, let us now use uppercase Greek letters to refer to sets of mental states and use lowercase Greek letters to refer to mental states. With this change of notation, the structural properties are these:

BRIDGE CUT: if $A \models \beta$ for each $\beta \in B$ and $B \models \gamma$, then $A \models \gamma$

PLAIN CUT: CUT for $\models$ as in the main text but adopting the present reading of the Greek letters

In this richer notation, the view that I develop rejects BRIDGE CUT (confirming the informal idea in the text that it is the special role of the foundational states that makes the difference). Thus, though slightly more complex, this formalization involves no equivocation and the remaining discussion in the paper fits smoothly with this more complex formulation with some easy translation that I leave to the interested reader.

B. Generalization for Jeffrey Conditioning

How do our results fare if we don't rely on the assumption that we learn that the slide is red by conditioning on the slide looking red and we don't rely on the assumption that we know for certain the slide looks red? Suppose in particular, instead, we think the natural modelling choice for learning is Jeffrey conditionalization. Indeed, I believe this is the best probabilistic development of the dogmatist view. According to the view, one directly learns that the slide is red, but does not become fully certain that the slide is red because the belief is only defeasibly justified.

Matters become more complicated in this setting. To begin, Jeffrey conditionalization is defined over a partition of an agent's epistemic space. Normally, when we think of learning p by Jeffrey conditioning we think of p as an atom of the partition. If we were to apply this approach to our current situation, our model would be that we start out by conditioning on the slide being red, so the partition would include the slide is red as an atom. And then we condition on the slide looking red, so the partition would include the slide looks red as an atom.

But it is obviously possible that the slide looks red and is red, so these propo-

sitions cannot be atoms of the same partition. This way of modeling the case, then, assumes that we change partitions as we learn. Now it is known that once we allow switching of partitions over learning episodes, there are no constraints at all provided by your prior probabilities on the final result after conditionalization. This is more or less to say that if we think of the situation this way, the probability of *reliability* could be anything from 0 to 1 no matter what your priors are.

This, in principle, is one way out for the dogmatist. It would be one which more or less says that we have sequences of Jeffrey updates involving a change of our partition and that probabilities play no interesting role in our epistemology of these cases. Perhaps this is what the dogmatist should say. But if she wishes to say this, the interesting question is what interpretation she can provide of the epistemic role of a partition and of changes in partitions. This is a question that has, to date, been relatively underexplored in traditional epistemology.⁴³

In light of this issue involved with shifting partition, the next question to consider is whether we can work with a fixed partition for both Jeffrey updates and what if anything that would tell us. Now in doing this, we should remember that the cases we are discussing are ones where we learn things like the slide looks red and the slide is red. And these, as we said, cannot be atoms of the same partition.

The way to interpret, say, learning the slide is red directly is to think of us as Jeffrey updating our partition (whose atoms does not include the slide is red or the slide looks red) so that the sum of the probabilities of the atom at which the slide is red is greater than or equal to the threshold for justification. There are of course many different kinds of partitions we could use. But the most natural choice is the one that is just as fine grained as we need: $\{r \wedge lr, r \wedge \neg lr, \neg r \wedge lr, \neg r \wedge \neg lr\}$.⁴⁴ Then we imagine a first Jeffrey update on r here understood to tell us that the sum of the probabilities in $r \wedge lr$ and $r \wedge \neg lr$ after the update are at least as high as t and then a Jeffrey update on lr understood analogously.

So recall now our setup. Our setup is that your probability that your color vision is reliable shouldn't go up after these two Jeffrey updates. So using Pr_0 for your probabilities before Jeffrey updating at all and using Pr_2 for the results after two Jeffrey updates, we have it as a constraint that:

$$\text{Pr}_0(\text{reliability}) \geq \text{Pr}_2(\text{reliability})$$

43. But see Greco (2017) for a framework that could be applied to the present context to help to make this option more palatable.

44. This follows Jeffrey's advice (1983: §11.6) about how to deal with cases where one simultaneously learns several propositions. Since no harm is done by assuming one simultaneously learns the slide is red and the slide looks red in our puzzle, this way of modeling the case is the modeling choice recommended by Jeffrey. Thanks to Brad Armendt for pointing this out to me.

Next, since $\Pr_2$ is the result of Jeffrey conditioning on the above described partition, we know that:

$$\Pr_2(\text{reliability}) = \Pr_1(\text{reliability} | r \wedge lr) \Pr_2(r \wedge lr) + \Pr_1(\text{reliability} | r \wedge \neg lr) \Pr_2(r \wedge \neg lr) + \Pr_1(\text{reliability} | \neg r \wedge lr) \Pr_2(\neg r \wedge lr) + \Pr_1(\text{reliability} | \neg r \wedge \neg lr) \Pr_2(\neg r \wedge \neg lr)$$

From here, we can notice that:

$$\Pr_1(\text{reliability} | r \wedge lr) \geq t$$

This follows from the following two facts

(1) conditional probabilities on atoms of a fixed partition are rigid over Jeffrey conditioning in the sense that if e_i is an atom of the partition and we have any sequence of Jeffrey updates on that partition $\Pr'(\cdot | e_i) = \Pr''(\cdot | e_i)$ where $\Pr'$ and $\Pr''$ are probability functions from anywhere in that sequence. So applying this fact to the case at hand, we have it that $\Pr_1(\text{reliability} | r \wedge lr) = \Pr_0(\text{reliability} | r \wedge lr)$.

(2) $\Pr_0(\text{reliability} | r \wedge lr) \geq t$ because we are assuming that learning the conjunction directly really would be good evidence.⁴⁵

And we can also notice that:

$$\Pr_2(r \wedge lr) \geq t$$

This is because our discussion in Section 3.1 showed that perception and introspection are not just ways of forming justified beliefs in r and in lr respectively but also ways of forming beliefs in r and in lr that are sufficiently good to allow us to safely conjoin these beliefs. This then allow us to get a similar constraint on the probability of *reliability*:

$$\Pr_2(\text{reliability}) \geq t^2 + \Pr_1(\text{reliability} | r \wedge \neg lr) \Pr_2(r \wedge \neg lr) + \Pr_1(\text{reliability} | \neg r \wedge lr) \Pr_2(\neg r \wedge lr) + \Pr_1(\text{reliability} | \neg r \wedge \neg lr) \Pr_2(\neg r \wedge \neg lr)$$

This is our problematic result again except with more complicated terms summing with t^2 .

45. This represents inductive reasoning being good reasoning in this context. Notice that we make no assumption about $\Pr(r | lr)$ and indeed this probability may even be low. This accords with a dogmatist perspective on which one directly learns the slide is red but one does not accept the inference from the belief the slide looks red to the belief the slide is red. As far as I know, the result here is the only probabilistic objection to dogmatism that can be directed at this kind of dogmatist.

This shows that given the natural choice of a fixed partition, our result can be established for Jeffrey conditioning. The best dogmatist response to this would be to identify some other fixed partition and show that given this partition our result cannot be established. And I do not at this time have a general argument that shows this is impossible.

But much like the dogmatist response that involved switching partitions, the best thing to say at this time is that if the dogmatist wishes to adopt this approach it is incumbent on them to say more. In particular to develop this response, the dogmatist must identify the partition that avoids the results of this appendix and give some grounds for thinking that this partitioning of the space is the correct one.

Thus, I conclude that while matters are more complicated if we consider Jeffrey conditioning, dogmatist views about justification continue to have an uneasy relationship with probabilistic models of justification.

C. The Problem Comes Earlier

We saw that the solution I have developed can be adapted to target *Step 3* rather than *Step 5*. Though the paper is officially neutral about this issue, I prefer to adopt the solution that targets *Step 5*. This appendix considers further arguments for targeting *Step 3* rather than *Step 5* that I believe are inconclusive (though I leave it to the reader to assess the costs and benefits for themselves).⁴⁶

C.1. The Earlier Problem

Recall that the main ground for accepting *Step 3* is the following closure principle:

if *S* is justified in believing *p* and competently deduces *q* from *p* then *S* is justified in believing *q*

On the assumption that simple analytic inferences are the kinds of things that we can competently deduce, after *Step 2* closure allows us to conclude that our color vision worked.

What I wish to consider here are two important arguments that show that if

46. See Cohen (2010: 145, 149), Kallestrup (2012: §4), Titelbaum (2010: 120–121, 128–129), White (2006: §7). Though it is actually not clear exactly which steps each of these philosophers find problematic, I focus on *Step 3* to fix ideas. See also Butzer (2017: §6) for a response to this objection that usefully clarifies some potential misunderstandings and responds to some reason to reject *Step 3* but, in my view, does not adequately address the two arguments below.

one accepts the closure principle and one accepts *Step 3*, then certain problems arise for dogmatism. If these arguments do show there are some deep problems for dogmatism, these are grounds for rejecting *Step 3*.

First there are cases where instead of continuing with the bootstrapping reasoning after concluding that our color vision just worked, we are to imagine an oracle says something like:

if your color vision works n times, then your color vision is reliable

We may imagine that the oracle's testimony justifies you in being fully confident in this claim. So after performing n instances of reasoning to the conclusion that your color vision worked, you may safely conjoin this with the oracle's claim to get:

'my color vision worked n times and if my color vision worked n times, then my color vision is reliable.'

Reliability then would follow from a simple deduction from this claim. Nonetheless, it seems that this reasoning cannot newly justify you in believing your color vision is reliable.⁴⁷

This case cannot be explained by the solution that targets *Step 5* because it does not involve a step of enumerative induction (though I sketch a response using other dogmatist resources in a note).⁴⁸

47. See Cohen (2010: 145, 149) and Titelbaum (2010: 120–121).

48. Though my official response to this objection is the somewhat concessive one given below, the dogmatist may be able to respond to this objection less concessively. I can only sketch the response here.

Begin by considering the case where the oracle tells you that if your color vision works once, then it is reliable. This amounts to telling you that your color vision will either never work or it is reliable. If you were fully justified in being confident in this disjunction and spreading your confidence equally over each disjunct, the dogmatist should say this defeats your perceptual justification for believing that the slide is red. This is just an undercutting defeater like learning that in the present circumstances it is as likely as not that the slide is red if it looks red

This covers the case where n is one. But not all values of n will look like cases of undercutting as this one does. However, we may be able to generalize the strategy to cover these cases. The idea would be that all of them adjust how much justification you get for believing that the slide is red based on it looking red. The amount of justification you get for the slide is red constrains whether you are justified in believing this and whether this belief can be conjoined with other ones. If we hypothesize that the level of justification will never be high enough to admit of conjunction introduction of n conjuncts for n high enough to allow enumerative induction unless the oracle's claim is sufficient on its own to justify us in believing our color vision is reliable, we avoid the objection. (This response twists Cohen's, 2010: 146–148, response to the super reliability related objection to be a response to his own objection to reaching *Step 3*.) Though the details need to be spelled out, the basic idea is that the oracle's testimony in effect gives you information about your reliability and this constrains the bootstrapping reasoning in a way that allows us to avoid this problem.

The second argument is probabilistic. Though we won't go through the details here, the basic idea is this. We can deduce 'it is not the case that the slide looks red and is not red' from 'your color vision worked'. So presumably we are newly justified in believing this as well. But it can be shown with the help of some minimal assumptions that your credence in 'it is not the case that the slide looks red and is not red' in fact *decreases* upon seeing that the slide looks red. So it is hard to see how you could become newly justified in believing this claim.⁴⁹

The solution that targets *Step 5* does not tell us anything that would allow us to resolve this tension between the dogmatist claim about justification and probabilities. Indeed, this second argument is one of the members of the recent family of results that show dogmatism is incompatible with Bayesianism that I alluded to earlier.

These two arguments make vivid that there is a tension between dogmatism and accepting *Step 3* and the principle of single premise closure.

C.2. Two Problems

That said, I believe that this is a different problem for dogmatists.⁵⁰ To illustrate why, it is helpful to know that those who have raised these objections (and also accept single premise closure) think that it shows that dogmatism is wrong because we are in fact *a priori* justified in believing that if the slide looks red, it is red. If this were so, then we could rule out from the start the possibility that the slide looks red while it is not red and thereby avoid the implausible results mentioned by the second argument. And if this were so, then it would be plausible that learning that the slide looks red allows us to learn our color vision worked, and combining this with the oracle's testimony allows us to conclude that our color vision is reliable.

Now we know that dogmatists do not believe that we are *a priori* justified in believing that if the slide looks red, it is red and do not believe that we are *a priori* justified in believing that our color vision is reliable. To the dogmatist, this is a suspicious kind of deeply contingent *a priori* justification.

But notice that the arguments that we have just considered only force us to admit that we are *a priori* justified in believing that if the slide looks red, it is red. They do not force us to admit that we are *a priori* justified in believing that our color vision is reliable. Perhaps, in order to solve these problems, we really do

49. See Cohen (2005: 424–425), Hawthorne (2004: 73–77), White (2006: §1–6). But these results cannot be established in the Jeffrey conditioning setting (cf. §B) at least in the sense elaborated in Footnotes 22 and 45.

50. It was originally presented as a different problem in Cohen (2003). More recent discussion such as Vogel (2007) and Weisberg (2010) also treat it as a different problem.

need to be *a priori* justified in believing that certain specific skeptical scenarios (e.g., appearing you have hands when you don't and it appearing the slide is red when it isn't) do not obtain.

But this fact alone is not a reason to admit that we are also *a priori* justified in believing our color vision is reliable any more than it would be a reason to admit, for example, that our scientific knowledge of the contingent workings of our world is *a priori*. Even if we must come to grips with a particular piece of deeply contingent *a priori* justification, this fact alone is no reason to admit other deeply contingent claims are *a priori* justified unless we can construct similar arguments to show that they are: While it may be impossible to have a plausible epistemology without allowing some contingent propositions to be *a priori* justified, this is no excuse for allowing arbitrary contingent propositions to be *a priori* justified. So the dogmatist who accepts closure on hence *Step 3* would do well then to stop the spread of the contagion of *a priori* justified contingent propositions by seeing if it can be quarantined to a limited domain.

Now one of the results of this paper is that we can give a probabilistic argument for the claim that we are at least *a priori* justified in being strongly biased toward our color vision being reliable. And as I have said this is bad news for the dogmatist. But we have tried to make the best of it for the dogmatist by putting probabilistic considerations aside. We saw that if we do this, we can solve the bootstrapping problem. What we have not yet seen is why on non-probabilistic grounds we should admit that we are *a priori* justified in believing that our color vision is reliable.

For this reason, I admit that there is a problem that arises once we get to the claim that our color vision works, and my solution does not solve this problem. But the problem that we have focused on in this paper is a different problem. And it is one that is solved by the arguments of this paper.

C.3. Backstories

Now this response is bound to be unsatisfying to some so I want to consider one objection to it: The case of *a priori* contingent justified scientific beliefs is not on all fours with the case of the *a priori* contingent justified belief in your color vision being reliable. This is because whatever "backstory" explains why we have *a priori* contingent justified belief in if the slide looks red, then it is red will also explain why we have *a priori* contingent justified belief that *reliability*.⁵¹

Of course, the success of this objection turns on what exactly the backstory is,

51. I only discuss the possibility of generalizing the backstory about why the claim that if the slide is red, then it looks red is *a priori* justified to explain why the claim that your color vision is *a priori* justified. I leave to one side direct attempts to explain why the claim that your color vision is reliable is *a priori* justified such as the one in Wright (2004).

and I know of no systematic way to evaluate all possible backstories that there could be. So to respond to this objection, I will focus on a specific backstory that is due to Stewart Cohen.⁵²

Before I do this, let me be clear that what follows is not a criticism of Cohen. Cohen uses this backstory as part of explaining how to give a non-dogmatist diagnosis of what's wrong with the bootstrapping reasoning. But our purpose is to consider whether this backstory forces the *dogmatist* who does admit *a priori* justified belief in if it looks red, then it is red to also admit *a priori* justified belief in *reliability*. As we will see below, we will get off board with a step in the argument because the dogmatist is not forced to accept it. This does nothing to show that the non-dogmatist shouldn't accept this step.

Cohen's backstory begins by observing that we can engage in a certain practice of suppositional reasoning. We suppose α and reason to some conclusion β on the basis of α and assuming that this reasoning is good, it seems that we are epistemically justified in believing the proposition expressed by $\lfloor \text{if } \alpha, \text{ then } \beta \rfloor$. So, for example, consider the modus ponens rule that says from α and $\lfloor \text{if } \alpha, \text{ then } \beta \rfloor$, you may conclude β . Now I can engage in a bit of suppositional reasoning. I start by supposing α and $\lfloor \text{if } \alpha, \text{ then } \beta \rfloor$ and then reason by modus ponens to the conclusion β . So I am now epistemically justified in believing a conditional corresponding to the modus ponens rule.

And modus ponens is just one example. The same kind of reasoning can be applied to any basic rule of reasoning to show that we are justified in believing a conditional corresponding to this rule. If we assume that it is a basic rule of reasoning that we can conclude 'the slide is red' from 'the slide looks red'. This would entail that we are *a priori* justified in believing that if the slide looks red, then it is red.

Now notice that just saying this much does not get us *a priori* justification in *reliability*. The claim that our color vision is reliable is not a conditional corresponding to any basic rule of justification. Instead, it is a conclusion inferred by enumerative induction. And the rule of enumerative induction that we need for scientific inference involves a track record proposition not the proposition that if the slide looks red, then it is red. Cohen's argument does not say that the track record proposition is *a priori* justified so this does not explain why *reliability* would be *a priori* justified.

But we can get an argument going for being *a priori* justified belief in *reliability* if we engage in the following longer chain of reasoning: First suppose 'the slide looks red' and infer by the perceptual inference rule 'the slide is red'. Next

52. See Cohen (2010: 150–156; cf. Hawthorne 2002: §1.2). Wedgwood (2013: §3, especially n. 14) offers a similar backstory but does not claim that we are *a priori* justified in believing that our color vision is reliable (see Footnote 53 for further discussion).

conclude by conjunction introduction from these two sentences ‘the slide looks red and is red’. Then perform the simple analytic inference that allows you to conclude ‘your color vision worked’. Repeat this process for many different colors and afterward infer *Track Record*. Finally, conclude *Reliability* by enumerative induction.

In other words, Cohen’s idea is to suppose each sentence reporting a perceptual experience and then show *Reliability* can be concluded using the step that we identified at the beginning of the paper. So it appears that we are *a priori* justified in believing that whatever perceptual experience of the color of the slide we have, our color vision is reliable. And that more or less amounts to simply saying that we are *a priori* justified in believing our color vision is reliable.

But the dogmatist is not in fact forced to accept this argument. This is because it assumes more than just that suppositional reasoning allows us to be *a priori* justified in believing a conditional corresponding to each basic rule. It assumes that epistemic justification satisfies transitivity. And we have already seen that the dogmatist must deny this claim and seen how they should do so. We said that the inductive inference to *Reliability* fails because our basis for *Track Record* is just ‘the slide looks red’.^{53, 54}

Thus, Cohen’s backstory would force the dogmatist to say that we are *a priori* justified in believing our color vision is reliable only if epistemic justification is transitive in the way that I have argued the dogmatist must deny. But since dogmatists who accept closure and hence *Step 3* are not eager to add to the stock of deeply contingent *a priori* justified beliefs that they must accept, this is corroborating evidence that the view I have been developing is the best solution for the dogmatist to adopt and that the bootstrapping problem is a separate problem over and above the problem about single premise closure.

53. Though Wedgwood’s (2013) ideas are similar to Cohen’s they differ in important ways. In particular, Wedgwood only claims that we are *a priori* justified in believing that if the slide looks red and my experiences contain no defeaters, then it is red and similarly for all of the conditionals corresponding to rules. This difference leads us to be unable to perform the suppositional reasoning that I sketched on Cohen’s behalf. This is because we cannot coherently add the supposition that no defeaters to enumerative induction are present after inferring the slide is red from the slide looks red because this itself is a defeater of that inductive inference according to the present account. To better understand this point as well as the point that I make about Cohen’s solution, it may be useful to compare how Cohen and Wedgwood would treat the failure of transitivity discussed in Section 5.3.

54. The dogmatist may also claim that supposing a proposition that reports one is having a visual experience only supports believing under the supposition what one is justified in believing when one *believes* one is having a visual experience. But dogmatism has no commitments about this; it only has commitments about what one is justified in believing when one *has* the visual experience.

Acknowledgements

For comments on an early draft thanks to the USC dissertation seminar, several anonymous referees, Andrew Bacon, Kenny Easwaran, John Hawthorne, Ben Lennertz, Jacob Ross, Mark Schroeder, Scott Soames, Gabriel Uzquiano, Ryan Walsh, Ralph Wedgwood, and Jonathan Weisberg. For comments on later drafts, thanks Brad Armendt, Bryan Lietz, Ángel Pinillos, and especially the area editor and anonymous referees at *Ergo*.

References

- Altschul, Jon (2012). Entitlement, Justification, and the Bootstrapping Problem. *Act Analytica*, 27(4), 345–366.
- Barnett, David James (in press). Perceptual Evidence and the Cartesian Theater. In Tamar Szabo Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 6). Oxford University Press.
- Bergmann, Michael (2000). Externalism and Skepticism. *The Philosophical Review*, 109(2), 159–194.
- Becker, Kelly (2012). Basic Knowledge and Easy Understanding. *Act Analytica*, 27(2), 145–161.
- Black, Tim (2008). Solving the Problem of Easy Knowledge. *The Philosophical Quarterly*, 58(233), 597–617.
- Briesen, Jochen (2013). Reliabilism, Bootstrapping, and Epistemic Circularity. *Synthese*, 109(18), 4361–4372.
- Brueckner, Anthony (2013). Bootstrapping, Evidentialist Internalism, and Rule Circularity. *Philosophical Studies*, 164(3), 591–597.
- Brueckner, Anthony and Christopher Buford (2009). Bootstrapping and Knowledge of Reliability. *Philosophical Studies*, 145(3), 407–412.
- Burge, Tyler (1993). Content Preservation. *Philosophical Review*, 102(4), 457–488.
- Butzer, Tim (2017). Bootstrapping and Dogmatism. *Philosophical Studies*, 174(8), 2083–2103.
- Christensen, David (1992). Confirmation Holism and Bayesian Epistemology. *Philosophy of Science*, 59(4), 504–557.
- Cohen, Stewart (2002). Basic Knowledge and the Problem of Easy Knowledge. *Philosophy and Phenomenological Research*, 65(2), 309–329.
- Cohen, Stewart (2005). Why Basic Knowledge Is Easy Knowledge. *Philosophy and Phenomenological Research*, 70(2), 417–430.
- Cohen, Stewart (2010). Bootstrapping, Defeasible Reasoning, and A Priori Justification. *Philosophical Perspectives*, 24(1), 141–159.
- Comesaña, Juan and Carolina Sartorio (2014). Difference-Making in Epistemology. *Nous*, 48(2), 368–387.
- Davies, Martin (2004). Epistemic Entitlement, Warrant Transmission, and Easy Knowledge. *Proceedings of the Aristotelian Society*, 78(1), 213–245.
- Douven, Igor and Christoph Kelp (2013). Proper Bootstrapping. *Synthese*, 190(1), 171–185.

- Greco, Daniel (2017). Cognitive Mobile Homes. *Mind*, 126(501), 93–121.
- Hawthorne, John (2002). Deeply Contingent A Priori Knowledge. *Philosophy and Phenomenological Research*, 65(2), 247–269.
- Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford University Press.
- Huemer, Michael (2001). *Skepticism and the Veil of Perception*. Rowman & Littlefield.
- Huemer, Michael (2013). Epistemological Asymmetries Between Belief and Experience. *Philosophical Studies*, 162(3), 741–748.
- Kallestrup, Jesper (2012). Bootstrap and Rollback. *Synthese*, 189(2), 395–413.
- Kornblith, Hilary (2009). A Reliabilist Solution to the Problem of Promiscuous Bootstrapping. *Analysis*, 69(2), 263–267.
- Lackey, Jennifer (2006). Knowing from Testimony. *Philosophy Compass*, 1(5), 432–448.
- Markie, Peter (2005). Easy Knowledge. *Philosophy and Phenomenological Research*, 70(2), 406–416.
- Moretti, Luca (2015). In Defense of Dogmatism. *Philosophical Studies*, 172(1), 261–282.
- Nair, Shyam. (2019). Must Good Reasoning Satisfy Cumulative Transitivity? *Philosophy and Phenomenological Research*, 91(1), 123–146.
- Neta, Ram (2005). A Contextualist Solution to the Problem of Easy Knowledge. *Grazer Philosophische Studien*, 69(1), 183–205.
- Neta, Ram (2013). Easy Knowledge, Transmission Failure, and Empiricism. In Tamar Szabo Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 4, 166–184). Oxford University Press.
- Pollock, John and Joseph Cruz (1999). *Contemporary Theories of Knowledge* (2nd ed.) Rowman & Littlefield.
- Pryor, James (2000). The Skeptic and the Dogmatist. *Nous*, 34(4), 517–549.
- Pryor, James (2013) Problems for Credulism. In Christopher Tucker (Ed.), *Seemings and Justification* (89–132). Oxford University Press.
- Scheall, Scott (2011). Later Wittgenstein and the Problem of Easy Knowledge. *Philosophical Investigations*, 34(3), 268–286.
- Siegel, Susanna (2012). Cognitive Penetrability and Perceptual Justification. *Nous*, 46(2), 201–222.
- Siegel, Susanna (2013). The Epistemic Impact of the Etiology of Experience. *Philosophical Studies*, 162(3), 697–722.
- Silins, Nicholas (2008). Basic Justification and the Moorean Response to the Skeptic. In Tamar Szabo Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 2, 108–141). Oxford University Press.
- Sosa, Ernest (1997). Reflective Knowledge in the Best Circles. *Journal of Philosophy*, 94(8), 410–430.
- Tucker, Christopher (2010). Why Open-Minded People Should Accept Dogmatism. *Philosophical Perspectives*, 24(1), 529–545.
- Titelbaum, Michael (2010). Tell Me You Love Me. *Philosophical Studies*, 149(1), 119–134.
- Wedgwood, Ralph. (2013). A Priori Bootstrapping. In Albert Casullo and Joshua Thurow (Eds.), *The A Priori in Philosophy* (226–246). Oxford University Press.
- Weisberg, Jonathan (2009). Commutativity or Holism. *British Journal for the Philosophy of Science*, 60(4), 793–812.
- Weisberg, Jonathan (2010). Bootstrapping in General. *Philosophy and Phenomenological Research*, 81(3), 525–548.

- Weisberg, Jonathan (2012). The Bootstrapping Problem. *Philosophy Compass*, 7(9), 597–610.
- Weisberg, Jonathan (2015). Updating, Undermining, and Independence. *British Journal of Philosophy of Science*, 66(1), 121–159.
- White, Roger (2006). Problems for Dogmatism. *Philosophical Studies*, 131(3), 525–557.
- Wright, Crispin (2004). On Epistemic Circularity: Warrant for Nothing (and Foundations for Free)? *Proceedings of the Aristotelian Society*, 78(1), 167–212.
- Vahid, Hamid (2007). Varieties of Easy Knowledge Inference. *Act Analytica*, 22(3), 223–237.
- van Cleve, James (1979). Foundationalism, Epistemic Principles, and the Cartesian Circle. *The Philosophical Review*, 88(1), 55–91.
- van Cleve, James (2003). Is Knowing Easy – Or Impossible?. In Stephen Luper (Ed.), *The Sceptics* (45–59). Ashgate.
- Vogel, Jonathan (2000). Reliabilism Leveled. *The Journal of Philosophy*, 97(11), 602–623.
- Vogel, Jonathan (2008). Epistemic Bootstrapping. *The Journal of Philosophy*, 105(9), 518–539.
- Zalabardo, José (2005). Externalism, Skepticism, and the Problem of Easy Knowledge. *The Philosophical Review*, 114(1), 33–61.