

Against Epistemic Conservatism*

Tim Willenken

Tmw17@pitt.edu

Abstract: Jim Pryor’s “The Skeptic and the Dogmatist” showed the importance of the question: can someone acquire justification via some belief-forming method (perception, induction, etc.) while lacking justification to believe that skeptical scenarios which undermine that method are false? Those who answer *yes* are called dogmatists or liberals, while those who answer *no* are called conservatives.

Recently a number of philosophers have been impressed by an argument that liberalism is inconsistent with Bayesian norms of coherence. This alleged inconsistency has led some authors to reject liberalism (White, Schiffer) and others to urge a rejection of Bayesianism (Pryor, Weatherson).

This paper has two parts. The first part addresses the argument that liberalism is inconsistent with Bayesianism. I identify three auxiliary theses—separate from Bayesianism itself—that are required to generate the inconsistency. I argue that one of these auxiliary principles is false. I conclude that liberalism and Bayesianism are consistent.

The second part offers an argument against conservatism. The crux of the argument is a dilemma about how conservatives understand *propositional justification*. On a robust understanding of propositional justification, conservatism turns out to be false, while on the denial of the robust understanding conservatism turns out to be toothless. In the conclusion I discuss some parallels between the dilemma and the problem with the Bayesian argument for conservatism: in each case, conservatism stumbles in a similar way.

I. Ways for reasoning to fail

Forming new beliefs is perilous. We must—on pain of having unjustified beliefs—avoid making bad inferences: affirming the consequent, committing the gambler’s fallacy, and so on. In the same vein there are non-inferential moves we have to avoid: wishful thinking, believing the putative deliverances of ESP, etc. In addition, each of us must make these moves only when their **enabling conditions** obtain: that is, someone only make an inference when he has justified beliefs in the premises, and only form perceptual beliefs when he is having the relevant

* This is an abridged version of the paper that emphasizes formal epistemology. For the unabridged version, please contact the author.

perceptual experience, and so on. Of course avoiding these bad moves is not enough to ensure that one is a model of epistemic virtue—*that* requires making good moves, not just avoiding bad ones—but we might hope that it is enough to ensure that one’s patterns of reasoning are not defective.

Alas, even that may be too much to hope for. There are bits of reasoning where every step looks fine, yet the reasoning as a whole looks fishy at best.

Consider:

Premise: (Car-1) I intend to walk to the parking lot and get in my car

So: (Car-2) I will walk to the parking lot and get in my car

Therefore: (Car-3) My car has not been stolen or towed from the parking lot

In general, if I intend to Φ , I have justification to believe that I will Φ . The transition from Car-1 to Car-2 is just an instance of this pattern. Car-3 follows by simple deductive inference from Car-2. Each step is kosher, but this is just a crazy way of coming to believe that my car hasn’t been stolen or towed from the parking lot. The same phenomenon is evident in arguments with higher philosophical stakes:

Premise: (Induction-1) In the past, observed regularities have tended to continue into the future

So: (Induction-2) From now on, observed regularities will tend to continue into the future

Therefore: (Induction-3) Induction will continue to be reliable in the future.

Or:

Premise: (MPP-1) Truth tables show that Q is never false when P, $P \rightarrow Q$ are true

true

Premise: (MPP-2) If truth tables show that Q is never false when P, $P \rightarrow Q$ are true, then modus ponens is truth-preserving.

Therefore: (MPP-3) Modus ponens is truth-preserving

Induction-2 follows from Induction-1 by (of course) induction, and Induction-3 follows from Induction-2 by deductive inference. Yet this bit of reasoning is famously suspect. Ditto for MPP 1-3: this argument is deductively valid—in fact, it is sound—yet it does not seem capable of rendering belief in its conclusion justified. The same fishiness arises in bits of non-inferential belief formation. Consider this notorious bit of reasoning:

Enabling condition: (Moore-1) [an experience as of a hand]

So: (Moore-2) I have a hand

Therefore: (Moore-3) I am not a handless brain in a vat

Each step looks unobjectionable on its own—including the non-inferential move from perceptual experience in Moore-1 to a perceptual belief in Moore-2—yet few are convinced by this line of reasoning that we all must have justification to believe that we are not handless brains in vats.

What goes wrong in these bits of reasoning? There are some familiar ways of responding to arguments like these – of course, we are free to offer different diagnoses of the different arguments:

1. **Skepticism.** We might deny that we have justification to believe the second step in each of these arguments on the basis of the first step. For instance, skeptics deny that someone has justification to believe that he has a hand when he has a visual experience as of a hand at the end of his arm.
2. **Mooreanism.** We can accept that an argument can give justification, for the first time, for believing their conclusions. Some fans of Moorean replies might deny that there's anything fishy about these arguments. Others will grant that they fall short in important ways, but still affirm that the reasoning they outline can result in the acquisition of justification for their conclusions.
3. **Closure denial.** *Closure* is the thesis that when you have a justified belief that P, and you know that P entails Q, and you form a belief that Q based on P, the belief that Q must be justified. We can avoid skepticism and Mooreanism—that is, say that we have justification to believe the premises of the above arguments while denying that the arguments need give any new reason to believe their conclusions—if we say that these arguments represent counterexamples to closure. Closure deniers can say that someone can have a justified belief that he has hands, and know that if he has hands then he is not a handless brain in a vat, and yet not be in a position to have a justified belief that he is not a handless brain in a vat.

None of these options looks especially tempting; can't we do better? Fortunately, this list is not exhaustive. The diagnosis I want to focus on for now begins with the observation that the source of the fishiness looks the same in all cases: the transition from the first step in the argument to the second seems, very roughly speaking, to tacitly assume that the conclusion is true. That is, Induction-1 is evidence for

Induction-2 only if we (roughly) assume that induction will continue to be reliable, MPP-1 and MPP-2 give reason to believe MPP-3 only if we (roughly) assume that modus ponens is truth-preserving, my intention to get in my car only bears on whether I will get in my car if we (roughly) assume my car hasn't been stolen from the parking lot, and my experience of a hand only gives reason to believe Moore-2 if I (roughly) assume that I am not a deceived, handless brain in a vat. At the very least, someone who learned that induction was wildly unreliable should not think that Induction-1 is a good reason to believe Induction-2, and so on. We are on the path to a diagnosis of the problem, though it is still an open question how serious it is.

Before proceeding, we should make this diagnosis a bit more precise. Suppose Samantha has some evidence, E , for a proposition p ; E is strong enough evidence to lend Samantha justification to believe p ¹. I'll use the shorthand $\text{Support}(E, p)$ to mean the rational support that E lends to p . A **defeater** for p is a proposition such that, when Samantha comes to reasonably believe that it is true, she no longer has justification to believe p . One type of defeater is an **outweighing defeater**: a proposition that provides evidence for the falsity of p and thus renders belief in p inappropriate in spite of the reasons given by E for counting p true. An outweighing defeater leaves $\text{Support}(E, p)$ unchanged. Another type of defeater is an **underminer**. Learning that an underminer is true eliminates $\text{Support}(E, p)$. For instance, suppose I see a happy looking crowd emerging from Shea stadium one summer night. That evidence gives me justification to believe that the Mets just won a game. Suppose I read in the paper the next day that the Mets lost: my justification to believe that they won has been *outweighed*. Suppose instead that I ask someone in the crowd and they tell me that they came from a rock concert and the Mets played an away game: my justification to believe that the Mets just won has been *undermined*². There may be other kinds of defeaters, but these two will suffice here. I'll call a proposition an **anti-underminer** if its negation is an underminer.

¹ I won't assume that all evidence is propositional; as I'll use the term, sense experience can be evidence.

² Many propositions diminish $\text{Support}(E, p)$ to some degree. I will use the term *underminer*, though, only to refer to propositions that eliminate $\text{Support}(E, p)$ entirely.

The four fishy arguments above each have a conclusion which is an anti-underminer for the support generated by the first step for the second; that is the feature which accounts for the appearance of circularity. Just how epistemically feeble these arguments are turns on a question about anti-underminers: in order to acquire justification to believe p from E , need one have independent (of p) justification to believe all of the anti-underminers for $\text{Support}(E,p)$? Plausible answers to this question include “yes, always” and “sometimes”; nobody should believe that we never need independent justification for an anti-underminer for $\text{Support}(E,p)$ in order to acquire justification to believe p from E . I’ll say that one treats an anti-underminer for the $\text{Support}(E,p)$ **liberally** if one denies that we need independent justification to believe it in order to acquire justification to believe p on the basis of E ; otherwise, one treats it **conservatively**. I’ll call the position that all underminers must be treated conservatively **across-the-board conservatism** (or just **conservatism**, for short) and the denial of conservatism **liberalism**³. A little more carefully:

Conservatism: S can acquire additional rational support to believe that p on the basis of E only if S has p -independent justification to believe that each of the underminers for $\text{Support}(E,p)$ is false.

A note on the sense of **independent** at issue here: though it is easy to gesture at the rough sense in which *independent justification* is used here, it is very difficult to define independence precisely. I will not attempt a precise definition here, though I will have more to say about the issue later on in this paper. As a rough definition, though, we can say that someone has p -independent justification to believe a proposition if he can reason his way to a justified belief in that proposition without going through p as an intermediate step. Note that liberalism is consistent with

³This terminology follows Pryor (2004). Many people use the term “dogmatism” to refer to the position I am calling liberalism, though here too I will follow Pryor and reserve the term *dogmatism* for a very specific variety of liberalism.

treating some anti-underminers conservatively, but conservatism is not consistent with treating any anti-underminers liberally⁴.

Conservatism yields a compelling analysis of what goes wrong in the fishy arguments. Conservatives can embrace closure and thus say that anyone who bases a belief in, say, Car-3 on justified belief in Car-2 must have a justified belief in Car-3. However, conservatives deny a thesis called **transmission**: if you have a justified belief that P, and you know that P entails Q, it must be possible for you to acquire additional justification to believe that Q—or acquire justification to believe that for Q for the first time—by forming the belief that Q based on P. Put another way, when you acquire some reason to believe P, and you know that P entails Q, that reason transmits through the argument to become reason to believe Q. Conservatives think that our four fishy arguments are counterexamples to transmission. Car-3 is an anti-underminer for acquiring justification for Car-2 on the basis of Car-1. So, you can only have justification to believe Car-2 when you have independent justification to believe the conclusion (or some other justification for Car-2 other than Car-1 that isn't undermined by ~Car-3; I'll ignore this possibility). So you can't acquire justification to believe Car-3 via the above argument⁵. So while Car 1-3 is not tautological—that is, its conclusion does not appear among its premises—it is epistemically useless. The same diagnosis applies to the other arguments.

Of course, liberals can agree with this diagnosis in some cases: a liberal can treat Car-3 conservatively relative to Car-1 and Car-2 and agree that the above argument is no way to acquire additional justification to believe Car-3. No example

⁴ The definition of conservatism that I give here is correct only in contexts where beliefs are simply justified or not – we need a slightly more complex statement of conservatism to deal with degrees of justification. The more general statement of conservatism is: S can acquire additional rational support to believe that p on the basis of E only if S has justification to believe that all of the underminers for Support(E,p) are false, *and this justification is independent of whatever support E lends to p*. The difference between this careful formulation and the simpler one in the main text is that in the careful formulation, justification to believe the anti-underminers can depend in part on p, but not on any additional support lent to p by E. I will stick to the simpler all-or-nothing statement in the text – it will not make any difference for our purposes.

⁵ Most conservatives would, I think, endorse not just the claim that if the conclusion of some bit of reasoning is justified, the anti-underminers for that bit of reasoning must be independently justified, but the stronger claim that if the conclusion of a bit of reasoning is justified *to some degree*, the anti-underminers must be independently justified *to at least that degree*. I'll generally ignore questions about degree of justification in this paper, just for the sake of expository simplicity.

of a fishy argument need embarrass liberals as such. In addition, liberals who think that some particular argument is a way of acquiring justification need not deny that it has some genuine deficiencies: for instance, a liberal can say that while Moore 1-4 is a way of acquiring justification to believe that one is not a handless BIV, it is not a way of rationally convincing a skeptic that he is not a handless BIV. Even liberals can have serious reservations about Induction 1-3 or MPP 1-3, and liberals can say that these arguments are instances of transmission failure. Conservatives, though, must say this.

Liberalism and conservatism only give different verdicts in cases where someone is rationally agnostic about at least one relevant anti-underminer. All parties agree that when someone ought to believe that an underminer for $\text{Support}(E,p)$ is true, he cannot acquire justification to believe that p on the basis of E . All parties also agree that when someone ought to believe that all of the anti-underminers for $\text{Support}(E,p)$ are true there is no obstacle for him acquiring justification for p on the basis of E . The disagreement between liberals and conservatives centers on the rational force of agnosticism: is being agnostic about an underminer for $\text{Support}(E,p)$ enough to threaten the support that E lends to p ?

On the face of it, the conservative explanation of what goes wrong in the above arguments is the most plausible. Yet, in this paper I will argue that conservatism is false: it is possible to acquire justification to believe that p on the basis of some evidence E and lack independent (of p) justification to believe in anti-underminers for $\text{Support}(E,p)$. In part II, I will discuss what is at stake in the debate over conservatism – as we will see, figuring out what to say about some perennial skeptical problems turns on getting straight about whether conservatism is true. Parts III and IV will discuss a formal argument purporting to show that Bayesianism is incompatible with liberalism. Part III will explain the argument and Part IV will say where it goes wrong. Finally in Part V I will argue against conservatism itself by posing a dilemma concerning how we should understand the term “justification to believe” which figures in the statement of conservatism.

II. Why does it matter whether conservatism is true?

Whether conservatism is the best story about what is wrong with a particular argument might not seem like a very important question; certainly not much turns on our diagnosis of Car 1-3. But on the larger questions of whether to accept or reject conservatism, the stakes are high. First, and most famously, what we say about conservatism constrains what we say about evil demon-style skeptical scenarios. Jim Pryor has argued that a very common way of presenting such arguments is not at all compelling⁶. Consider:

Premise: (Simple-1) I do not have justification to believe that I am not being deceived by an evil demon

Premise: (Simple-2) If I do not have justification to believe that I am not being deceived by an evil demon then I do not have justification to believe that I have hands

Therefore: (Simple-3) I do not have justification to believe that I have hands

It is not clear that anyone should lose any sleep over this argument. The first premise is itself a skeptical conclusion, and hardly something that most people accept pre-philosophically. It turns out to be quite difficult to argue for Simple-1. So it looks like we are free to reject it. Yet without Simple-1, the skeptic cannot make any mischief with this style of argument.

However, skeptics can produce a more powerful argument that does not rely on anything so contentious as Simple-1. The really worrying skeptical argument is:

Premise: (Nasty-1) Either I do not have justification to believe that I am being deceived by an evil demon *or* I do have justification to believe it, but my justification is based (at least in part) on beliefs justified by perception

Premise: (Nasty-2) If I have justification to believe any proposition on the basis of perception, then I must have independent (of perception) justification to believe that I am not being deceived by an evil demon

So: (Nasty-3) I do not have justification to believe that I am not being deceived by an evil demon

Premise: (Nasty-4) If I do not have justification to believe that I am not being deceived by an evil demon then I do not have justification to believe that I have hands

⁶ See Pryor (2000).

Therefore: (Nasty-5) I do not have justification to believe that I have hands

The third step of this argument follows from the first two because Nasty-1 says that any justification I have to believe that I am not being deceived by an evil demon must depend upon perception, while Nasty-2 says that justification to believe *anything* on the basis of perception requires perception-independent justification to believe that I am not being deceived by an evil demon. More simply, Nasty-1 says that justification to believe that I am not being deceived could *only* come from perception, while Nasty-2 says that justification to believe that I am not being deceived *cannot* come from perception. Nasty-3 follows because there is nowhere left for justification for the anti-skeptical claim to come from.

Nasty-1 is much weaker than Simple-1. More importantly, Nasty-1, unlike Simple-1, is not a statement that most anti-skeptics can flatly reject. So this argument looks sound—or in any case, the range of options of responding to skepticism is narrow—if Nasty-2 is true. But conservatism entails Nasty-2, since *I am being deceived by an evil demon* is an underminer for the acquisition of justification to believe any proposition on the basis of perception. So the really worrying skeptical argument looks difficult indeed to resist if we accept conservatism.

A second area where it matters very much whether conservatism is true—one that has received less attention—concerns the response to skepticism over whether any belief-forming methods generate justified beliefs. People can and do form beliefs in myriad ways: induction, gazing into crystal balls, perception, consulting sacred texts, and so on. We ordinarily think that some of these methods are **correct**—they can generate justified beliefs, at least when certain enabling conditions obtain—whereas others represent pathological ways of conducting epistemic business. What separates the correct methods from the incorrect ones? In some cases, the answer is easy: I have good inductive justification to believe that *The New York Times* is pretty reliable and that *Weekly World News* is quite unreliable, so the method of believing what I read in the *Times* is correct while the method of believing what I read in the *Weekly World News* is not. However, we can't

account for the correctness of every method in this way. At least some methods must be **basic**: their ability to generate justified beliefs does not depend upon using some other method to establish their reliability. So the really difficult question is what separates the correct basic methods from the incorrect ones?

A tempting answer is: nothing. This answer is tempting because of a skeptical argument that purports to show that there could not be any correct basic belief-forming methods. The argument is simple:

Premise: (Regress-1) If any belief-forming method M is correct, it must be possible to acquire justification to believe that M is reliable.

Premise: (Regress-2) If the justification to believe that M is reliable involves the use of another method, N, then M is not a basic method.

Premise: (Regress-3) The justification to believe that M is reliable cannot come from M.

Premise: (Regress-4): Justification to believe that a method M is reliable must either come from M or from some other method.

Therefore: (Regress-5) So, no basic methods are correct.

This conclusion is unacceptable. The second premise looks difficult to reject—it follows from the definition of *basic*—so the place for anti-skeptics to take their stand is to reject Regress-1 or Regress-3. Yet, conservatism entails both Regress-1 and Regress-3. *M is not reliable* is an underminer for any bit of belief formation that uses M: since conservatism says an instance of belief-formation can generate a justified belief only if one has justification to believe the relevant anti-underminers, it says that M can generate justified beliefs only for someone who has justification to believe that M is reliable. Hence, conservatism entails Regress-1. Conservatism entails Regress-3 since it says that M generates justified beliefs for S only if S has *independent* (of any beliefs justified by M) justification to believe that M is reliable. So any attempt to use M to justify an anti-underminer for the use of M is not kosher, according to conservatism. So if conservatism is true, this argument is sound.

Conservatism thus plays a crucial—if often tacit—role in two perennially vexing skeptical problems. If we are going to get a handle on responding to these skeptical arguments, we must figure out whether conservatism is true. Of course, philosophers who embrace conservatism need not be skeptics. Conservatives can

avoid skepticism by embracing **rationalism** and saying that we have apriori justification to believe in a host of important anti-underminers, including *I am not a handless brain in a vat*, *induction is reliable*, and so on. The most compelling kind of rationalism says that apriori justification to believe in anti-underminers does not rely on the use of *any* belief forming method; rationalists can thus deny Nasty-1 and Regress-4. Discussion of the merits and drawbacks of rationalism is a large topic, well beyond the scope of this paper; my aim here is to show that conservatism is false so we are not forced to choose between rationalism and skepticism⁷.

III. An argument for conservatism

So far I have alluded to some general considerations that render conservatism plausible—it explains what is wrong with some fishy arguments—but I have not presented any systematic argument for conservatism. In this section, I will present what I take to be the most powerful argument for conservatism. The crux of the argument is the claim that liberalism is inconsistent with Bayesianism. This argument is well regarded by both liberals and conservatives: the latter see it as a reason to abandon liberalism, the former as a reason to abandon orthodox Bayesianism⁸: that is, to give up the view that epistemic probability, or rational credence, is governed synchronically by the norms of probabilistic coherence and diachronically by the claim the $Pr(p)$ upon learning that E should equal one's prior $Pr(p|E)$. If this argument is sound, we face a paradox: a sound argument for conservatism and a counterexample to conservatism. Fortunately, both sides have got it wrong by overstating what epistemic conclusions we can read off of the Bayesian formalism. In Part IV, I will show that key premise in the argument from Bayesianism to conservatism—the premise that allows us to infer conclusions about propositional justification from premises about formal results in epistemic

⁷ It is possible to be both a rationalist and a liberal. Silins defends just this combination of views (although he defines liberalism somewhat differently than I do).

⁸ The most fully-developed version of the Bayesian argument against liberalism appears in White; a quicker version appears in Schiffer. Pryor (forthcoming a) and Weatherston both acknowledge the force of the argument and accordingly urge revision of Bayesianism to accommodate liberalism. Silins also cites a similar argument which he uses to develop a view that mixes liberalism with rationalism.

probability—is subject to counterexamples that closely resemble the counterexamples that make trouble for conservatism itself.

To help introduce the argument for conservatism, let's look at a specific case⁹. Say that I am having a bit of perceptual experience as of some hands on the end of my arms. Call the proposition that I have hands Q , my token perceptual experience as of hands η (note that this is not a belief), the proposition that I am having experience η E , the proposition that I that Q on the basis of η PERC, and the proposition that I am a deceived brain-in-a-vat having experience η —as well as other perceptual experiences like those of an ordinary person— U . PERC entails both Q and $\sim U$: it is inconsistent with my not having hands or being a deceived BIV. Some liberals, including dogmatists, say that I do not need independent or independent justification for $\sim U$ in order to acquire justification to believe either PERC or Q when I have η . Some results in probability statics seem to make trouble for this view.

First result: $Pr(\text{PERC}) \leq Pr(\sim U)$: that is, the epistemic probability that I am perceiving a hand when I believe that I am having a perceptual experience as of a hand is less than the epistemic probability of the hypothesis that my perceptual experiences are not all misleading. This has to be true because PERC entails $\sim U$: deceived BIVs do not veridically perceive hands.

Second result: $Pr(\sim U|E) \leq Pr(\sim U)$: having experiences like those of a normal person makes it *more* epistemically probable that I am a BIV having experiences like those of a normal person. Coming to believe E lowers the epistemic probability of $\sim U$. So if I didn't have justification to believe $\sim U$ prior to coming to believe E , I do not have it once I believe E .

These two results together generate a challenge to liberalism, because the first says that the epistemic status of PERC at any given time can only be as good as the epistemic status of $\sim U$ at that time, and the second says that the epistemic status of $\sim U$ is worse given E . So, according to the first result, $Pr(\text{PERC}|E) \leq Pr(\sim U|E)$ and according to the second result $Pr(\sim U|E) \leq Pr(\sim U)$. Together: the epistemic status of PERC given E is at best as good as the epistemic status of $\sim U$ prior to updating on E , or $Pr(\text{PERC}|E) \leq Pr(\sim U)$. It seems to follow that if PERC is going to be justified after

⁹ The terminology of this example follows Pryor (forthcoming a); the presentation of the argument and the three bridge principles upon which it relies is also indebted to that paper.

having experience E , $\sim U$ had better be independently justified. I cannot acquire justification to believe that I have hands unless I have independent justification to believe that I am not a deceived BIV. So liberalism about perception is false.

Of course, the above argument cannot be read off of the claims about epistemic probability alone. In order to get from the two results to claims about independent justification we need three bridge principles; I'll state the principle as it occurs in the above argument first and then put it in more general terms.

1. **From experience to epistemic probability:** in the above argument I assume that whatever changes in epistemic probability occur as a result of having experience η all result from conditionalizing on E . More generally, the bridge principle is that all rational change in epistemic probability that comes from some non-doxastic experience involves conditionalization on the proposition that one is having that experience. Call this the **Auxiliary Thesis about Experience**.
2. **From epistemic probability to justification:** the argument assumes that updating on E cannot justify $\sim U$ since it raises the probability of U . More generally, a bit of updating cannot justify a hypothesis if its effect is to lower the epistemic probability of that hypothesis. Call this assumption the **Auxiliary Thesis about Credence Lowering**.
3. **From kinematics to order of rational support:** finally, the above argument assumes that since E can raise the epistemic probability of PERC only to the level of the epistemic probability of $\sim U$ prior to updating on E , E can justify Q only if $\sim U$ is independently justified. More generally: if someone believes that X on the basis of E , and the probability of Y prior to learning that E is an upper bound on the probability of X after learning that E , justification to believe that Y must be independent of justification to believe that X . Call this the **Auxiliary Thesis about Independent Justification**.

None of these principles is part of the Bayesian framework itself: a Bayesian is free to reject any of these three additional assumptions. First, consider (1). Bayesianism says that whenever someone learns that X , the effect of X on the epistemic probability of Y is to update according to $Pr(Y|X)$. However, Bayesianism need not say that *all* updating of epistemic probabilities involves conditionalization on other propositions. In particular, non-doxastic inputs may cause changes in epistemic

probability that cannot go via conditionalization on one's new evidence. For instance, in the above story we assumed that after having experience η I could update on E. But the support that η lends to E cannot have come via conditionalization on η . I could not have had a value for $Pr(E|\eta)$, since η is not a proposition. So Bayesianism itself is silent on the epistemic effects of having η . The above argument assumes that the *only* effect of having η is making it appropriate to conditionalize on E; while it is plausible that η makes it rational to conditionalize on E, I see no reason to assume that is the only effect of η .

The story about acquiring the belief that I have hands via perception changes significantly if we only make the weaker assumption about the rational response to η . We can suppose that experience η makes it rational to conditionalize on E, but there is no impediment in the formalism to η making it rational to conditionalize on other beliefs as well, including $\sim U$. If it does, then even if the Auxiliary Thesis about Independent Justification is true there is no obstacle to acquiring justification for PERC without having some sort of standing justification to believe $\sim U$. The picture is that having η directly justifies—or, more formally, makes it rational to conditionalize on—both E and $\sim U$, and then I can infer PERC from E since I already have justification to believe $\sim U$. This picture accords with all of the principles that conservatives insist upon and yet looks a lot like a form of liberalism: no justification for $\sim U$ independent of having some perceptual experience is required in order acquire justification to believe the content of that experience. So we can easily reconcile Bayesianism and liberalism¹⁰.

So the Auxiliary Thesis about Experience is optional, and looks too strong; but without that bridge principle it is much harder to gin up tension between Bayesianism and liberalism. However, we cannot just end our discussion here because of two difficulties. First, not all liberals will be comfortable getting around

¹⁰ Perhaps some extremely hard-line Bayesians will want to reject talk of experience making it rational to conditionalize on particular beliefs and say instead that any updating by conditionalization is rational. That seems like an odd view, but in any case it is consistent with in being rational to update by conditionalizing on $\sim U$ whenever one updates by conditionalizing on E.

Also, in the story above I ignore Jeffrey conditionalization. Many Bayesians will want to say, with Jeffrey, that in ordinary cases of perceptual experience I should assign the proposition that I am having the experience a likelihood of less than 1 and update accordingly. I've ignored this sort of concern for simplicity; it doesn't make much difference for my purposes here.

the tension with Bayesianism by saying that η directly probabilifies $\sim U$, or makes it appropriate to conditionalize on $\sim U$: it seems a little far-fetched to say that some perceptual experience directly justifies the claim that I am not a BIV, even though nothing in the Bayesian formalism rules it out.

Second, I would ultimately like to make the world safe for liberalism about inferential methods. The same problem arises for inferential methods. Say R is the proposition that it will be cold in Pittsburgh next January, F is the proposition that it has been cold in Pittsburgh every previous January, V is the proposition that induction is reliable until December 2009 but unreliable thereafter, and IND is the proposition that I know R on the basis of inductive inference from F . Some analogous results obtain:

Third result: $Pr(IND) \leq Pr(\sim V)$, since IND entails $\sim V$ if we assume—plausibly—that I can't acquire inductive knowledge about 2009 if induction is wildly unreliable in 2009.

Fourth result: $Pr(\sim V|F) \leq Pr(\sim V)$, since V entails F , learning that F should not lower my confidence in V .

So, the epistemic probability of IND is only as good as the prior epistemic probability of $\sim V$. If the second and third auxiliary theses are correct, I can acquire justification to believe IND only if I have independent justification to believe $\sim V$. Here, relaxing the Auxiliary Thesis about Experience will not help liberals at all. So, while I suspect that we should reject the Auxiliary Thesis about Experience, doing so is not a panacea. We had better dig deeper.

I am going to assume that the Auxiliary Thesis about Credence Lowering is correct, though I am not at all certain that it is. I think that the serious culprit in making difficulties for liberalism is the Auxiliary Thesis about Independent Justification. It is important to recognize the *prima facie* appeal of the Auxiliary Thesis about Independent Justification. That principle says that that if my credence in Y before acquiring evidence E is an upper bound on my credence in X after learning that E (where X is based on E), then if X is justified on the basis of E then I must have independent justification to believe that Y . The thought is that if X is

justified on the basis of E, $Pr(X|E)$ must be high enough to make believing that X appropriate upon updating on E. Since $Pr(Y)$ is an upper bound on $Pr(X|E)$, the prior probability of Y is an upper bound on the probability of X upon learning E. So the prior probability of Y must be high enough to render belief in Y justified. Since Y had this degree of probability before X was justified, Y having this degree of probability cannot possibly depend on X being justified. So, justification to believe that Y must be independent of justification to believe that X.

In spite of its considerable plausibility, the Auxiliary Thesis about Independent Justification is false.

IV. Against the Auxiliary Thesis about Independent Justification

In order to construct a counterexample we need a case where $Pr(X|E) \leq Pr(Y)$, where E is a part of the evidential base of X yet where one does not need independent justification to believe Y in order to have justification to believe that X. Here is such a case:

Etch-a-Sketch with Conditional Pills: Stella discovers that someone has slipped a pill into her drink: this pill has the effect of rendering her incapable of forming cogent beliefs in conditionals—it also makes spurious claims about conditionals pop into her head and appear plausible—but leaves the rest of her cognitive abilities intact. She knows all this about this pill. She has excellent evidence that all of her other capacities are just fine—she has scored well on tests and so forth while on the pill—and she has learned not to trust the thoughts about conditionals that pop into her head.

While knowingly under the effects of the pills, Stella is shown a contraption that she's never seen before: for the sake of concreteness—and nostalgia—make it an Etch-a-Sketch. She watches as several people draw on it, and then shake it, and notes that whenever someone shakes it the picture disappears. She gets an opportunity to draw on it, and she shakes it, and lo, the picture disappears. Then someone else draws on it, and turns it around so that Stella cannot see the screen, and shakes it. Call all of this evidence E. Stella concludes—on the basis of inductive inference from E—that the picture has disappeared (call this p). Her belief that p is impeccably justified.

The proposition that threatens to make trouble for the Auxiliary Thesis about Independent Justification is $E \rightarrow p$. Two conditions need to hold in order for this case to be a counterexample:

1. In the example, Stella must lack justification to believe the conditional $E \rightarrow p$
2. The Auxiliary Thesis about Independent Justification must say that that someone can only acquire justification to believe p on the basis of E if she has independent justification to believe $E \rightarrow p$, since $Pr(p|E) \leq Pr(E \rightarrow p)$.

On the first point: in the example Stella lacks independent justification to believe $E \rightarrow p$ since she lacks justification to believe any conditional at all. Someone with her evidence, which includes the fact that she has taken the incapacitating pills, cannot form any justified beliefs in conditionals—it would be wildly irresponsible for her to do so, given her knowledge that she has taken the pill—so she lacks propositional justification to believe that $E \rightarrow p$ ¹¹.

The second point can seem obvious: it can look like a plain matter of definition that someone’s rational confidence in the consequent of a conditional upon learning the antecedent cannot exceed her prior confidence in the conditional. Matters are not so simple, though: many of these apparent connections between conditional probabilities and the probabilities of conditionals turn out to be false. However, in this case we can prove that $Pr(p|E) \leq Pr(E \rightarrow p)$ without relying on the general thought that probabilities of conditionals need to be upper bounds on conditional probabilities:

1. $Pr(E \rightarrow p|E) \leq Pr(E \rightarrow p)$, since learning that the antecedent of a conditional is true can never raise the epistemic probability of that conditional.
2. p entails $E \rightarrow p$, since the consequent of a conditional entails the conditional as a whole. So $Pr(p) \leq Pr(E \rightarrow p)$.
3. Also, since p entails $E \rightarrow p$, $Pr(p|E) \leq Pr(E \rightarrow p|E)$.
4. So, $Pr(p|E) \leq Pr(E \rightarrow p|E) \leq Pr(E \rightarrow p)$
5. So, $Pr(p|E) \leq Pr(E \rightarrow p)$

¹¹ Worry: can conservatives say that Stella has *propositional justification* to believe $E \rightarrow p$ even though nobody with her evidence can form a justified belief in $E \rightarrow p$? No: watering down the meaning of *propositional justification* in this way will generate a version of conservatism too weak to be of much interest. For more on this point, see section V. For *much* more on this point, see the unabridged version of this paper.

So, if Auxiliary Thesis about Independent Justification is correct, someone can only have justification to believe p if he has independent justification to believe $E \rightarrow p$. Yet in Etch-a-Sketch with Conditional Pills, Stella has justification to believe that p and she lacks propositional justification to believe that $E \rightarrow p$; since she lacks justification to believe that $E \rightarrow p$, she must lack *independent* justification to believe that $E \rightarrow p$. We have a counterexample.

It is worth stepping back and asking what went wrong here. The problem is that Bayesianism says that Stella must assign a high credence to a conditional even though she lacks justification to believe that conditional. How should we resolve this tension? There are three options, none of them mutually exclusive. The first, and most cautious, is to be very careful about equating high epistemic probability with propositional justification; in particular, to be wary of any thesis that takes us from premises about epistemic probability to conclusions about propositional justification. This first way of resolving the tension is my official suggestion; it is, of course, perfectly consistent with Bayesianism, even if it does limit our ability to read epistemic conclusions off of claims about epistemic probability.

Two more radical responses are possible. One type of radical responses aims to develop a formal apparatus that better represents rational agnosticism by allowing that in some situations rational credences can take the value of a certain (perhaps very wide) interval. This sort of view would allow us to avoid Stella assigning any particular credence to $E \rightarrow p$: she could just assign it a (very wide) interval value. This solution faces some serious difficulties, though – very roughly, once we allow rational agents to have *some* interval valued credences, the norms of probabilistic consistency will require them to have interval-valued credences in cases where, intuitively, they should not¹².

A third way around the tension, and perhaps the most radical, is to deny that rational thinkers ought to have probabilistically coherent credences at all. According to this view, the problem with Bayesian idea that we ought to assign credences in a

¹² White (forthcoming) makes this argument quite powerfully.

probabilistically consistent way is not that it is *unrealistic*—Bayesianism is foremost a prescriptive, rather than descriptive view, so the fact that its demands are hard to meet in practice is no great objection—but rather that the advice to have ones credences cohere probabilistically is *pathological*. According to this sort of view, people like Stella ought to have inconsistent credences. David Christensen defends such a view. As Christensen points out, tension arises whenever consistency demands a very high credence in a proposition, but evidence of our own limitations demands assigning a lower credence in that proposition. For instance, people who have evidence that they are not very good at logic have inductive reason to assign not-too-high credence to theorems that they prove (and cannot verify from another source), yet probabilistic coherence requires assigning these theorems probability 1. Stella’s case is just an extreme version of this phenomenon – her inductive evidence demands that she assign low credence to logical consequences of her beliefs, while coherence demands that she assign them high credence. If we jettison the demand for synchronic coherence, the tension goes away, though of course the price of doing so is high.

V. A dilemma for conservatism*

The problems facing the Auxiliary Thesis about Independent Justification are quite similar to a problem facing conservatism itself. Recall that conservatism says:

Conservatism: S can acquire additional rational support to believe that p on the basis of E only if S has p-independent justification to believe that each of the underminers for Support(E,p) is false.

What does “justification to believe” mean here? To state some possible answers, we need some terminology. A statement such as “p is justified for Stuart” is ambiguous. One sense of that phrase is roughly: Stuart has evidence and/or whatever else is needed to provide rational support for the proposition that p. That is **propositional justification**, or justification to believe that p. The other sense is: Stuart has an

* I have condensed the argument in this section in this version of the paper. The unabridged version of the paper discusses the issues I treat here in much greater depth.

appropriate belief that p. This means that, first, Stuart believes that p; second, Stuart has justification to believe p; and third, Stuart's belief that p is **based** on some appropriate grounds: psychologically dependent upon some evidence in virtue of which Stuart has justification to believe that p. This third condition ensures that Stuart does not just have a belief which he is in a position to hold appropriately, but that he really does hold it appropriately. When these three conditions are met, Stuart's belief that p is **doxastically justified**. A doxastically justified belief is an appropriately held belief: someone has it for the right reasons.

Conservatism makes a claim about propositional justification. But note that on one natural way of understanding the relationship between propositional and doxastic justification—where S has propositional justification for p iff S is in a position to form a doxastically justified belief that p—conservatism is clearly false. Plugging this definition of propositional justification we get:

Very Robust Conservatism: S can acquire additional rational support to believe that p on the basis of E only if S is in a position to form p-independent doxastically justified beliefs that each of the underminers for Support(E,p) is false.

To see the problem with this view, consider:

The Naïve Reasoner: Sarah is a normal child of age 4. One day Sarah sees that it is raining outside and she forms the belief that the ground is wet. Sarah does not have any attitude about the underminer *there is a giant tarp covering the ground outside*. Indeed, she lacks the concept TARP and so she cannot form any beliefs about the presence or absence of tarps.

This case makes trouble because Sarah's belief that the ground is wet is justified. The proposition *there is no giant tarp covering the ground outside* is an anti-underminer for Sarah acquiring justification to believe that the ground is wet based on the evidence that it is raining. So, conservatism says that Sarah must have independent (of her belief that the ground is wet) justification to believe that *there is no giant tarp covering the ground outside* if she can acquire justification to believe that the ground is wet given her evidence that it is raining. But she is in no position to form a doxastically justified belief in that anti-underminer: she cannot so much as entertain it, since it involves the concept TARP which she lacks. So Very Robust Conservatism is no good.

A natural reply to the case of the Naïve Reasoner to say that Sarah does have propositional justification, in the relevant sense, to believe that there is no giant tarp covering the ground: after all, she's been outside many times, and always seen bare ground. She can't form the belief that there's no tarp covering the ground because of her undeveloped conceptual capacities, but her evidence strongly supports the claim that there's no tarp covering the ground. A more modest way of understanding conservatism, then, focuses not on what people are in a position to believe in a doxastically justified way, but on what beliefs it is possible to form in a doxastically justified way, via good reasoning. Call the view that S has propositional justification to believe that p iff there is some sound bit of reasoning from S's evidence which concludes in a doxastically justified belief that p the *Robust View* of propositional justification. If we plug this view into the definition of conservatism, we get:

Robust Conservatism: S can acquire additional rational support to believe that p on the basis of E only if there is some p-independent sound bit of reasoning from S's evidence which concludes in a doxastically justified beliefs in all of the anti-underminers for Support(E,p).

By *sound bit of reasoning* I do not mean a logically sound argument—that is much too narrow—but rather any reasoning process where each intermediate step is doxastically justified. The doxastic justification can come from perception, induction, or whatever: it can, but need not, come from deductive reasoning. A sound bit of reasoning is p-independent if it need not have forming a doxastically justified belief that p as an intermediate step. Robust Conservatism has no problem dealing with the Naïve Reasoner: there is a perfectly sound bit of reasoning from Sarah's evidence to the conclusion that there is no giant tarp on the ground—say, an inductive inference from *there has never been a tarp on the ground before* to *there is no tarp on the ground*—so there is no obstacle to her acquiring additional support to believe that the ground is wet based on the evidence that it is raining.

However, Robust Conservatism also faces counterexamples. Here are four:

Etch-a-sketch with evidence pills: Samantha knows that someone has slipped a pill into her drink. The pill has the effect of rendering her incapable of forming cogent beliefs of the form *X is good evidence for Y*—and indeed it makes all sorts of

bogus claims about evidence pop into her head and seem plausible—but leaves the rest of her cognitive abilities intact. She knows all this about this pill. She has been under its affects for some time, and has noticed that her cognitive functioning is as good as ever - indeed, she takes the GRE, plays a game of chess, conducts a scientific experiment, and performs several other challenging tasks while under the effects of the pill and discovers that she do them just as well as when she was not suffering from the pill's effects. She just cannot form beliefs about what is evidence for what, though she has learned not to trust the bogus beliefs about evidence that pop into her head.

While under the effects of the pill she sees an Etch-a-Sketch for the first time, and acquires evidence E (where E is the same as in the previous etch-a-sketch example). Samantha concludes that p (also the same as in the previous example) on the basis of inductive inference from E. Her belief that p is impeccably justified.

The Difficult Class: Steve, an engineering major, decided on a lark to take a class of the epistemology of evidence with the demanding Professor Werkmeister. Steve miserably failed the first few assignments and then dropped the class; he now knows that a lot of what seemed to him like common sense about what is evidence for what is wrong, but beyond that he got nothing out of the class. Steve acquires evidence E, which is in fact excellent evidence for p, and then forms a justified belief in p while being rationally agnostic about *E is good evidence for p*.

The Persuasive Lecture: Seth becomes convinced by a persuasive lecture by a respected authority on ethics who claims that while we should trust our ordinary judgments about the rightness or wrongness of individual acts, it is impossible to tell whether our faculties for making moral judgments are reliable (since the only yardstick we can use to measure their accuracy is our own judgments). Seth can't really follow the details of the arguments, but since the lecturer is respected, he concludes that it must be true that we can trust our ordinary moral judgments and that nobody can say whether his own faculties for making moral judgments are reliable. Seth sees some children setting a cat on fire, and forms a justified belief that they are acting wrongly, while being rationally agnostic about *my faculties for making moral judgments are reliable*.

The Mediocre Mathematician: Stan knows himself to be mediocre at math. One day a friend asks him if in general, assuming that a proposition is false, and then deriving a contradiction, is sufficient to establish that the proposition is true. Stan, quite appropriately, says he's stumped: he just has no idea how to assess whether a method of proof is kosher. He is rationally agnostic about whether *reductio ad absurdam* is a valid method of argument. Later Stan is wondering whether there are any even prime numbers besides 2. He thinks to himself: *suppose n is an even prime number (and not 2), then, since n is even, it is divisible by 2, but since it is prime (and not 2), it is not divisible by 2*. He concludes that there are no even primes besides 2 – and this belief is justified.

This list could go on. While any one example might be open to objection, the pattern is clear: someone can have reasoning-directed evidence which lets him form a doxastically justified belief in p on the basis of E but prevents anyone with his evidence from forming a doxastically justified belief in some anti-underminer for $\text{Support}(E,p)$. Robust Conservatism is no good.

The problem for conservatism is that any disambiguation of conservatism that is weaker than Robust Conservatism is too weak to be at all interesting. Consider any view of the form:

Anemic Conservatism: S can acquire additional rational support to believe that p on the basis of E only if S has p -independent “justification to believe” all of the anti-underminers for $\text{Support}(E,p)$, where this does **not** entail that there is a sound deliberative route from S 's evidence which concludes in a doxastically justified belief that p .

We'd need to say more about just what justification to believe does mean in order to have a bona fide view on the table, but the details of the positive proposal won't matter here. The problem with Anemic Conservatism is that it is consistent with Mooreanism. Imagine someone in a position with the following features (where E is evidence for p , U is an underminer for $\text{Support}(E,p)$, and p entails $\sim U$):

- a. He has p -independent “justification to believe” that $\sim U$ in the anemic sense
- b. There is no p -independent sound reasoning from his evidence which concludes in a doxastically justified belief that $\sim U$; that is, he lacks p -independent justification to believe that $\sim U$ in the robust sense.
- c. He can form a doxastically justified belief that $\sim U$ via reasoning from p , perhaps along with other premises. That is, he can have p -dependent justification to believe $\sim U$ in the robust sense.

Prior to acquiring E , he cannot form a justified belief that $\sim U$ (that's just condition (b)). However, once he does acquire evidence E , he can form a justified belief that p (that follows from condition (a) and the definition of Anemic Conservatism). Once he has a justified belief that p , he form a justified belief that $\sim U$ via inference from p , since p entails $\sim U$. So he can form a justified belief in $\sim U$ for the first time using

Moorean reasoning, and indeed such reasoning is the *only* way he can form justified belief in $\sim U$. Here's a concrete, if far-fetched, example:

Etch-a-Sketch with Evidence Pills Redux: As before, Samantha has taken some pills that interfere with her ability to form beliefs of the form *X is evidence for Y* but leave her cognitive abilities otherwise intact. However, while the pills do interfere with her ability to consider hypotheticals and come up with evidence claims in the abstract, one ability is intact: when she forms a justified belief Y based on evidence X, she can form the belief *X is good evidence for Y* based on Y and her belief that Y is based on X. Samantha knows all about the effects of the pills.

Before observing evidence E, Samantha is asked whether E is good evidence for p. She remains agnostic—and rationally so, given her knowledge of the effect of the pills. Then she observes E, and forms the justified belief that p. She is again asked whether E is good evidence for p. She reasons *I believe that p, and that belief is based on E, so E must be good evidence for p.*

So while some form of Anemic Conservatism may be true, it is too weak to do any interesting epistemic work. In particular, it is too weak to do the work that conservatism is supposed to do in the skeptical arguments discussed in section II: it doesn't entail Nasty-2 or Regress-3. Anemic Conservatism is consistent with using perception to acquire a justified belief, for the first time, that I am not a handless BIV – indeed, it is consistent with such Moorean reasoning being the *only* way to acquire such justification. Similarly, Anemic Conservatism is consistent with using induction to acquire a justified belief, for the very first time, that induction will be reliable in the future. So conservatives face a dilemma: Robust Conservatism is too strong, and runs into counterexamples, while Anemic Conservatism is too weak to be of much interest. There is no true principle that can do the work that conservatism is supposed to do.

VI. Conclusion

The formal argument for conservatism, and conservatism itself, get hung up on similar points: in each case, the problem lies strengthening a fairly harmless claim about prior epistemic probability or “propositional justification” in some anemic sense to get an interesting claim about which claims must be antecedent in the order of sound reasoning. This suggests that some weak claims in the neighborhood

of conservatism are true—for instance, the Bayesian considerations suggest that someone can only acquire additional support for p on the basis of E if the anti-underminers for $\text{Support}(E,p)$ have high prior epistemic probability—though conservatism itself is not. This problem is only apparent when we are careful about distinguishing between different epistemic statuses that a proposition can have – for instance, between being propositionally justified, in the sense that there is a good deliberative route from some evidence to a doxastically justified belief in that proposition, and having a high epistemic probability. It’s often tempting to just conflate these different statuses: cases where they come apart may be fairly rare. Once we do clearly make these distinctions, though, conservatism loses a lot of its plausibility.

Sources:

Christensen, David (2007) "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals," *Oxford Studies in Epistemology* 2: 3 – 31.

Pryor, James (forthcoming a) "Uncertainty and Undermining"

Pryor, James (forthcoming b) "When Warrant Transmits"

Pryor, James (2000) "The Skeptic and the Dogmatist" *Nous* 34, 517-549.

Pryor, James (2004) "What's Wrong with Moore's Argument" *Philosophical Issues* 14, 349-78.

Schiffer, Stephen (2004) "Skepticism and the Vagaries of Justified Belief" *Philosophical Studies* 119, 161-184.

Silins, Nicholas (forthcoming) "Basic Justification and the Moorean Response to the Skeptic".

Weatherson, Brian (2007) "The Bayesian and the Dogmatist" *Proceedings of the Aristotelian Society*, 107, 169-85.

White, Roger (2006) "Problems for Dogmatism" *Philosophical Studies* 131, 525-557.

White, Roger (forthcoming) "Evidential Symmetry and Mushy Credences"

Wright, Crispin (forthcoming) "The Perils of Dogmatism"

Wright, Crispin (2002) "(Anti-) Sceptics Simple and Subtle: Moore and McDowell" *Philosophy and Phenomenological Research* 65, 330-48.

Wright, Crispin (2004) "Warrant for Nothing (and Foundations for Free)?" *Aristotelian Society Supplement* 78 167-212.