

External World Skepticism, Confidence and Psychologism about the Problem of Priors

Sharon Berry

Abstract

In this paper I will draw attention to an important route to external world skepticism, which I will call confidence skepticism. I will argue that we can defang confidence skepticism (though not a meeker ‘argument from might’ which has got some attention in the 20th century literature on external world skepticism) by adopting a partially psychologistic answer to the problem of priors. And I will argue that certain recent work in the epistemology of mathematics and logic provides independent support for such psychologism.

1 Introduction

Consider the following skeptical scenario.

PEASOUP: Everything outside a 5 foot radius is peasoup, which forms up around me in such a way as to mimic the behavior of a persisting physical world obeying uniform laws.¹

Thinking about this kind of outlandish skeptical hypothesis² (and how it appears to be compatible with all our sensory experiences) can motivate worries

¹Note that for the purposes of this paper we can consider the more general possibility that everything outside of a d foot radius (for some value of d greater than the length of our arms) is made of some uniform material that forms up around me to mimic the behavior of a persisting physical world obeying elegant scientific laws. This addresses any concerns that it is the particulars (the distance and the type of material) which guarantee a low probability for the hypothesis.

²Given Putnam [20] and Chalmers [6] arguments suggesting that even if we were a brain in a vat our beliefs about the external world would still be true (but merely about the simulated vat world), one might wonder whether being highly confident skeptical scenarios are false is even necessary to defend our epistemic access to facts about the external world. However, Putnam’s argument only seems to succeed against certain skeptical scenarios. For instance, it seems difficult to swallow that if the world came into existence 5 minutes ago then claims

about our epistemic access to facts about the external world (i.e., about whether we can have knowledge, justified belief and/or justified confidence in claims about the external world).

In this paper I will consider how we can respond to a particularly threatening form of external world skepticism, Confidence Skepticism.

First I will distinguish the Belief Skeptical thesis that we cannot be justified in believing contingent claims about the external world from the stronger³ Confidence Skeptical thesis that we cannot even be justified in assigning high probability to/being very confident in these claims. I will note that Confidence Skepticism is what most people intuitively worry about when considering classic skeptical arguments, even though Belief Skepticism has gotten more explicit attention.

I will then develop what I take to be the intuitive motivations for Confidence Skepticism. Very crudely, I'll suggest that confidence skeptical worries arise from an intuition that it is dogmatic to assign high probability to a contingent proposition like \neg PEASOUP a priori, unless a principled story can be told about what intrinsic features of this proposition license us in doing so.

Finally, I will conclude by sketching a program for answering the Confidence Skeptical challenge. I will develop a positive picture of the foundations of a priori knowledge, on which significant arbitrariness is to be expected. I will also note that this picture is naturally motivated by some recent independent work on the nature of basic a priori knowledge of mathematics and logic.

about past events really only refer to the remembered story. Some people have the same intuition about peasoup scenarios. See [3] for a discussion of a range of other cases where the Putnamian response feels unsatisfying. Thus, the arguments presented here should be thought of as filling a gaps left by Putnam's response.

³Presumably one can't be justified in believing a proposition which one cannot have justified high confidence in.

2 Confidence Skepticism

2.1 On Skepticism and What Responding to it Requires

Let me begin with a quick note about the philosophical challenge posed by skeptical arguments, and the sense in which I aim to respond to this challenge.

Skeptical arguments appear to give us reason to stop taking certain epistemic attitudes (like belief or confidence) towards a certain range of propositions, by arguing from *premises which we find prima facie attractive* to the conclusion that it's epistemically impermissible for us to take these attitudes. The skeptic doesn't just autobiographically mention that she suspends judgment regarding some topics on which we are opinionated. Rather, she attempts to destabilize our attitudes by showing that some of our beliefs about epistemic normativity⁴ (among other things) imply that we must suspend the epistemic attitudes in question.

Accordingly, I take our task in responding to skepticism to consist in finding a way to coherently and plausibly block such arguments. Answering skepticism doesn't require us to attempt the ambitious (and questionably motivated) project of using shared beliefs to convince the skeptic⁵. Instead we must (merely) find a credible way to resist the skeptic's arguments that it is impermissible to have the epistemic attitudes which we do.

2.2 Confidence vs. Belief Skepticism

With these general preliminaries in place, I can now characterize the specific form of skepticism which will be my target in this paper. In this section, I will distinguish two varieties of external world skepticism (belief and confidence

⁴By this I mean beliefs about what it is ok to believe, infer, assign high probability to etc.

⁵I think there are systematic reasons to suspect that success at this project is both difficult and not of great philosophical importance. The bland fact that certain subsets of your total web of beliefs (e.g. your number theoretic beliefs) don't suffice to imply the whole, or that someone could consistently have all the same sensory experiences as you but continue to disagree with you forever doesn't seem very troubling on its own – except insofar as the skeptic can develop this point into a reason for doubting our beliefs about the external world *from our own point of view*.

skepticism), which can be inspired by considering hypotheses like PEASOUP, and explain why Confidence Skepticism will be my target in what follows.

Considering the epistemology of lotteries suggests that one's evidence may justify being very very confident in a claim without justifying believing that claim. For imagine someone who has just bought a ticket in a fair million-ticket lottery. Presumably she would be justified in being very confident that her ticket won't win, and in betting and otherwise acting accordingly (e.g., selling that ticket for a very small sum of money compared to the lottery's payout). However, many philosophers maintain that she would not be justified in believing (or asserting) that the her ticket won't win⁶.

Accordingly, we can distinguish skeptical arguments which call into doubt our justified *high confidence* in claims about the external world from those which merely call into doubt our justified *belief* in these claims. It is possible to question someone's claims to justified belief in some proposition without calling their justification for being highly confident in that proposition into doubt. In what follows, I will use 'Belief Skepticism' to mean the claim that it is epistemically impermissible for us to believe ordinary claims about the external world, and 'Confidence Skepticism' to mean the stronger thesis that it is epistemically impermissible for us to (even) be highly confident in such claims.

Although philosophical treatments of external world skepticism have tended to explicitly focus on Belief Skepticism, I take it that Confidence Skepticism is what captures the popular imagination and is what people tend to have in mind (or care most about) in typical discussions of external world skepticism.

For, when we introduce students to skeptical worries via Descartes' Meditations[8], we don't just raise the (comparatively) bloodless and technical worry that, although it's fine to assign very high probability to claims like 'Mount Everest exists', it would be wrong to *assert or believe* these claims – in the way that it is (arguably) wrong to assert or believe a specific lottery ticket won't win (but

⁶For example, see [14]. Note that if you can justifiably believe that lottery ticket #1 will lose, there's a puzzle about why you can't make the same point about each of the other tickets and then conjoin all these claims to conclude that no ticket will win.

perfectly fine to bet at heavy odds against it and adopt plans that will have fairly severe bad consequences if it does win etc). Instead we raise the more disturbing possibility that we have ‘no reason to believe’ that the skeptical hypothesis is false, in some sense which makes even high confidence in its falsehood unjustified⁷.

Accordingly, I will focus the bulk of this paper on the question of how to understand (and respond to) intuitive motivations for Confidence Skepticism.

3 Motivating Confidence Skepticism

3.1 Blumenfeld and Blumenfeld

With this target in mind, I will now propose a way of spelling out the intuitive motivations for Confidence Skepticism. First we should note that some classic formulations of the external world skeptic’s argument only motivate Belief Skepticism – so that we will need to look elsewhere if we want to flesh out the intuitive motivations for Confidence Skepticism.

Blumenfeld and Blumenfeld’s influential formulation of an argument for external world Belief Skepticism begins with an argument from the idea that our experience doesn’t ‘guarantee’ the falsehood of skeptical hypotheses to the conclusion that we lack knowledge that a skeptical hypothesis (like BIV or PEASOUP) is false, as follows:

- 1 Some PEASOUP scenarios are experientially indistinguishable from my current situation⁸.
- 2 So the qualitative character of my experience does not guarantee that \neg PEASOUP.

⁷There’s been some discussion, e.g., by Williamson [23], of whether we can eliminate questions about knowledge of the external world in favor of questions about justified belief/acceptance “Once Gettier showed in 1963 that justified true belief is insufficient for knowledge, and therefore that knowledge is unnecessary for justified true belief, it became natural to ask: if you can have justified true beliefs, why bother with knowledge”. Here I am trying to draw a difference contrast, between skepticism about justified belief vs. justified high confidence.

⁸Perhaps other motivations for (2) can be found, but I won’t challenge the move from (1) to (2) so I think the exact choice of (1) won’t matter.

- 3** If the qualitative character of my experience does not guarantee that \neg PEASOUP, then I can't know that \neg PEASOUP.

Conclusion So, from 2 and 3, I can't know that \neg PEASOUP.

Plausibly one can transform the above into an argument that it would be unjustified/epistemically impermissible for us to believe \neg PEASOUP⁹. However, the mere fact that we can't know or be justified in believing a proposition doesn't directly imply that we can't be justified in being highly confident in it. For, as noted above, people holding lottery tickets can plausibly have justified high confidence without justified belief or knowledge.

What happens if we replace appeals to knowledge in Blumenfeld and Blumenfeld's argument with appeals to justified confidence? We get the following argument:

- 1** Some PEASOUP scenarios are experientially indistinguishable from my current situation.
- 2** So the qualitative character of my experience does not guarantee that \neg PEASOUP.
- 3'** If the qualitative character of my experience does not guarantee that \neg PEASOUP, then I can't be justified in assigning high probability to \neg PEASOUP.

Conclusion I can't be justified in assigning high probability to \neg PEASOUP.

However, this argument provides little intuitive support for its conclusion because it's not clear why we should accept premise 3'. Certainly the following general inference form is not immediately compelling: if the qualitative character of my experience does not guarantee that ϕ , then I can't be justified in assigning high probability to ϕ . Indeed, I suspect that this general inference form will strike most contemporary readers (who reject Descartes' doctrine that knowledge requires certainty) as clearly fallacious¹⁰.

⁹For example, if the argument above works, then one can plausibly use it to infer that one lacks knowledge of \neg PEASOUP. Plausibly one can then infer from 'I know that I can't know that P' to 'I can't be justified in believing that P'.

¹⁰Indeed, even the external world skeptic will plausibly want to allow that one can sometimes

Perhaps the fact that PEASOUP is compatible with all one's evidence *does* motivate assigning PEASOUP a non-zero probability¹¹. However, this doesn't imply one shouldn't assign a *very low* probability to PEASOUP. Accordingly, appeals to the mere *possibility* of PEASOUP given one's total course of experience don't seem adequate to motivate the claim that one should not be very confident in \neg PEASOUP.

So, to summarize, Blumenfeld and Blumenfeld's formalization of the skeptical argument provides little help with articulating the intuitive motivations for Confidence Skepticism¹².

3.2 The Humean Dilemma

Instead, I propose, we can do a better job of spelling out the intuitive motivations for confidence skepticism by appealing to the idea that we lack some kind of necessary positive support for \neg PEASOUP, rather than from the mere (apparent) compatibility of PEASOUP with all our experiences invoked by Blumenfeld and Blumenfeld's argument. Thus, we can begin to explicate these motivations by posing the following, vaguely Humean, dilemma¹³.

1. One is either justified in assigning high probability to \neg PEASOUP a priori or a posteriori.

be justified in assigning high probability to claims which one's experience does not guarantee to be true. For example, consider claims like 'Either this is not a single ticket from a fair lottery with 10000 tickets, or it won't win.'

¹¹Maybe not though. Plausibly there are propositions compatible with all one's experience which one should assign probability 0 to, such as, 'this stick is exactly π inches long'.

¹²If one accepted the claim that one's evidence for anything must qualify as knowledge (as proposed in Williamson's *Knowledge and its Limits*[23]), then perhaps one could parlay these reasons for doubting that we know/are justified in accepting anything about the external world into an argument that we cannot be a posteriori justified in being highly confident of anything about the external world. But this argument would need much developing. For example, one could in principle say that one's evidence is always sense data/claims like 'I am having an experience as if of...' and combining that with priors that suitably privilege theories that explain sense data by appeal to an external world would still let one have high confidence in claims about the external world, even if the latter claims are never taken to qualify as knowledge and hence cannot themselves be cited as evidence for anything. Even if such an argument could be developed, I think it would require contentious extra premises in a way that makes my own proposed Humean argument a better rational reconstruction of the common Matrix-viewer and/or freshman philosophy student's motivations for accepting confidence skepticism.

¹³The dilemma is Humean in the sense that it resembles his famous argument that our belief in scientific induction is unjustified [16].

2. One can't be justified in assigning high probability to \neg PEASOUP a priori (or to any conditional whose antecedent is something metaphysically compatible with PEASOUP and whose consequent is \neg PEASOUP, like 'If I seem to see a tower 50 feet away then \neg PEASOUP.' or 'If my experiences would be elegantly explained by scientific laws positing uniform matter then \neg PEASOUP.').
3. We can't be justified a posteriori in assigning high probability to \neg PEASOUP. For example, by 2 we can't gain this justification by conditionalization on sensory information which is compatible with PEASOUP.
4. So, if you can learn \neg PEASOUP a posteriori, you could also have learned it a priori.
5. If one is not justified in assigning high probability to \neg PEASOUP, then one is not justified in assigning high confidence to various contingent propositions about the external world like 'Mount Everest exists.' whose truth implies \neg PEASOUP.

Conclusion: Belief in everyday propositions referencing distance objects is unjustified.

As mentioned in the introduction to this paper, I think¹⁴ that we should re-

¹⁴A different way to respond to the Humean challenge would be to resist premise 3. But I find this unattractive for reasons that have been often discussed in the literature. I will briefly review some of my favorites, just to motivate focus on resisting premise 2 above. But I do not claim to absolutely refute a posteriori routes (there being a huge literature and many possible moves and counter-moves to make here).

First, there's a standard probabilistic argument (given in [14] among other places) against the position looking at a ball can let us know 'that ball is red' and hence 'if that ball looks red it is red' (call this proposition NONDECEPTIVE), in cases where we could not know the latter proposition immediately in advance of looking at it. For discovering that the ball looks red should (if anything) *lower* our confidence in NONDECEPTIVE, since it only rules out scenarios where where the ball doesn't look red, i.e., ones in which NONDECEPTIVE is trivially true. Thus it seems strange to suppose that learning that the ball looks red should put you in a position to know NONDECEPTIVE if you were not originally.

Second Cohen develops an appealing variant on this point in his anti-bootstrapping argument in [7]. It would be weird to say that one could acquire justification for being very confident that your senses are reliable (which you otherwise lacked) just by flipping through a slideshow with different colored swatches and saying of each swatch (without any further checking) 'that looks red, so it is red, so my color vision is accurate in this case' and then doing scientific induction from this history of reliability. It seems to me that one can only permissibly 'reason' in this way to the extent that one is already permitted to assume that one's color vision was reliable to a certain degree – and that attractive probabilistic models of

ject premise 2. But to give full credit to skeptical intuitions, I must admit that premise 2 can look prima facie attractive. So let's complete our attempted rational reconstruction of the motivation for Confidence Skepticism by considering why premise 2 is attractive, and what can be said in its favor.

3.3 Concerns About A Priori Confidence in \neg PEASOUP

I will discuss a few ways the skeptic might try to intuitively motivate premise 2, and end with what I take to be the strongest motivation¹⁵.

A very simple way to motivate premise 2 would be to say that it's never permissible to have high a priori confidence in a claim without some further non-circular argument which we lack for \neg PEASOUP. However, insofar as the Confidence Skeptic is not a total Pyrrhonian skeptic (i.e., a skeptic who argues that we should suspend judgment regarding all propositions), even they will probably¹⁶ allow that we can be permissibly confident in some (analytic or logical) propositions in this way¹⁷. So, I don't think the Confidence Skeptic can provide much intuitive motivation for premise 2 by suggesting that it's never permissible to be highly confident of a proposition which you cannot support by providing a non-circular, non-regressive argument or sensory evidence for¹⁸.

scientific induction will not let one increase ones initial confidence in the reliability of ones' senses. But I won't try to argue for the latter conclusion here.

In a similar vein, Cohen also notes that it would be intuitively strange for him to answer a son who wants to buy a red table, but is worried that the one he is looking at in the shop is a white table made to look red by trick lights, by looking at the table, concluding that it is red and then saying 'it is red, so there aren't trick lights'. In contrast it seems much more natural to imagine the man responding (as per the idea that we are antecedently permitted to assign low probability to certain skeptical scenarios), 'But don't you think it's implausible that the furniture store has trick lights? That kind of thing doesn't happen very much/Why would they have them?'

¹⁵Hawthorne points out that while it seems weird to say 'I know a priori that if something looks red it is red' (even though from a Bayesian point of view where we are permitted to assign high confidence to this claim we probably do want to say this). But I think it is much less weird to say (as I am advocating saying here) that it is permissible to be highly confident in this claim a priori.

¹⁶In principle, one can imagine a non-pyrrhonian external world skeptic who maintained that a priori knowledge is impossible.

¹⁷For, as Lewis Carroll points out in in 'What the Tortoise Said to Achilles'[5], it seems that if we are justified in making modus ponens inferences (as presumably we are) it must be permissible for us to make some such inferences immediately, without further supporting argument. And seems (at least) plausible that we can have 'basic knowledge' of - and hence epistemically permissible high confidence in - some analytic or logical truths in the same way, i.e., know them without appeal to further (non-circular and infinite descending chain involving) argument or sensory experience.

¹⁸That is, the skeptic might try to maintain that, along the lines of Unger's Ignorance [22]

The skeptic can avoid this problem by instead appealing to a narrower principle: that one can't be highly confident in any deeply contingent¹⁹ proposition like \neg PEASOUP without further argument. This principle would also forbid assigning high a priori confidence in claims like 'the future is like the past' or 'unobserved parts of the universe are like observed ones' relevant to scientific induction.

However, accepting any such principle would essentially force us to say that no process of assigning priors (in the way that Bayesianism, the most widely accepted model for how to do probabilistic reasoning, demands) is epistemically permissible. For, on this picture, good reasoning involves starting with some probability-axiom-satisfying assignment of 'prior' probabilities to propositions, and then reacting to observations by updating these probabilities by conditionalization. But if *some* assignment of priors is epistemically permissible, we can show that one must assign very high (and very low) prior probability to some deeply contingent propositions. For²⁰ it's not hard to find reasonably large numbers of propositions which are 'metaphysically independent' in the sense that all truth-functional combinations of them are metaphysically possible, e.g., 'The first US President, if they exist, likes peanut butter.', 'The second US President, if they exist, likes peanut butter.', 'The third US President, if they exist, likes peanut butter.' etc. If you have n such propositions there are 2^n such truth functional combinations. The axioms of probability guarantee the probabilities assigned to these truth-functional combinations sum to 1. Since

part 5, "it.. is inconsistent to say, 'He is reasonable in believing that, but he has no reason at all for believing it.'" But as the literature 2000-2010s literature on basic a priori knowledge shows, it's hard to reconcile this idea with an attractive picture of logical knowledge. For, by essentially the argument in 'What the Tortoise said to Achilles'[5] it seems appealing to say that we have (defensible) prima facie warrant for inferences like modus ponens and believing various logical truths without need for further justificatory appeal to anything [2][9]

¹⁹Such a skeptic might want to allow that one can be highly confident in some special contingent claims like Kripke's 'The canonical meter stick in Paris is a meter long'[18] but not propositions like \neg PEASOUP. But they might say that deeply contingent truths, which Hawthorne characterizes as "one[s] for which there is no semantic guarantee that there actually exists some verifying state of affairs" (p. 247) [13], cannot be known a priori. See [1] and [21] for discussion of whether we can have a priori *knowledge* of such deeply contingent truths even if we can't be highly confident of them.

²⁰Assuming we assign priors to statements like 'If there is an n -th US president, he or she like(s) peanut butter.'

each of these truth-functional combinations is incompatible, the probabilities assigned to them must sum to (at most) 1. Thus there must be at least one of these truth functional combinations which has probability $\leq 1/2^n$ (and the negation of this proposition has probability $\geq 1 - 1/2^n$). Since it's not hard to get 20 or 30 such metaphysically independent propositions, some contingent proposition must be assigned a very high probability ²¹.

I propose that intuitive resistance to the idea that we can permissibly assign high confidence in a proposition like \neg PEASOUP a priori is motivated by a kind of non-arbitrariness intuition. For once one allows (contra Pyrrhonian skepticism) that chains of justification can permissibly come to an end somewhere, it's a natural thought that they should end in elegant general principles, not just admissions that certain kinds of theories seem immediately plausible or implausible. It might well seem hard to imagine any deeply philosophically principled reason why (even if \neg PEASOUP is true) this should be one of the contingent facts about the world which it is ok to be highly confident of a priori whereas (say) the periodic table of the elements or Schrödinger's equation are not.

4 Acceptable Priors Might be Somewhat Arbitrary

Now let us turn to the main goal of this paper: providing therapy for the confidence skeptical argument outlined above. Can the above argument for confidence skepticism be answered? Is it really plausible that it would be dogmatic and epistemically impermissible to assign high confidence to a deeply contingent proposition like \neg PEASOUP without some principled explanation of what (about the intrinsic nature of this proposition) makes it OK to assign high prior probability to it? Is it really plausible that no such principled explanation can

²¹Okasha's 'Bayesianism and the Traditional Problem of Induction'[19] makes a version of this point: namely that Bayseanism is incompatible with certain ideas about suspending belief/confidence.

be given? Or can this intuition be somehow dissolved?

As a first stab at motivating the idea that facts about acceptable priors *actually can* be somewhat arbitrary, note that when Carnap modified his theory of the logical foundations of probability [4] to allow learning, he had to include a choice of a factor for how quickly one projects from past experiences. For example, if you start with out any relevant prior information, how many black balls do you have to pull out of an urn before it is ok to assign 60% probability to the claim that they are all black? To say that any particular value for this factor is epistemically correct can seem arbitrary. Yet, even if one doesn't find Carnap's theory persuasive, one must either abandon learning from experience or pick some number of observations after which such a probability assignment is epistemically permissible.

Also recall Bertrand's paradox. One must choose between assigning equal probabilities to 'analogous' options with regard to possible side-lengths, side areas, or volumes when deciding what probability to assign to cube which is known to have side-length between 0 and 4 meters having side-length between 0 and 1 meter ²² For if you assign equal probability to this cube having side lengths between 0 and 1, 1 and 2, etc., then the probability of the cube having side length 1 should be 1/4. But if you assign equal probability to its having volume 0-1, 1-2 etc. then the the probability of the cube having side length 1 should be $1/4^3 = 1/64$.

Considering these points suggests that (regardless of skeptical worries about whether *our* non-skeptical/ scientific induction friendly way of assigning priors is the unique deeply intrinsically special one) there is no unique deeply intrinsically special way of assigning priors.

I will now try to further motivate the idea that facts about what priors are epistemically permissible to have (and more generally what propositions it is OK to assign high confidence to without justificatory appeal to any further sensory evidence or argument) can include some elements which are not deeply

²²This helpful and influential formulation of Bertrand's paradox comes from [10].

principled (like the choice of learning parameter in the example above) in a more positive way, by motivating and sketching a moderately attractive and modestly deflationary family of views - what we might call '(partial) psychologism' - about the nature of epistemic normativity which imply this. Obviously, providing a complete and compelling account of what priors it is epistemically ok to assign is beyond the scope of this paper. Instead I will try to make partial psychologism attractive by (just) sketching a story on which facts about acceptable priors are not deeply principled and then indicating some promising possibilities for how this story could be developed.

4.1 Motivation from Philosophy of Mathematics

The easiest way to motivate partial psychologism about acceptable priors which I want to advocate, is by considering recent work on analogous questions in the philosophy of mathematics and logic.

Human society has a well established and useful (if somewhat vague) practice of classifying some necessarily truth preserving arguments as proofs, and taking reasoning through these arguments (but not other ones which are equally valid and necessarily truth preserving and have the same premises) to let one go from knowledge of the premises to knowledge of their conclusions. We expect genuine proofs to break things down into certain kind of 'simple' steps. So, for example, we don't take a one line derivation of Fermat's Last Theorem (FLT) from the widely accepted standard axioms of PA or ZFC to count as a proof or to deliver adequate justification for believing its conclusion, even though this argument is logically valid. But we do think that a longer proof which walks us through many logically valid inferential steps which nearly all human beings find obvious could confer knowledge. And also we think that a more condensed proof which 'skips' steps can confer knowledge to a trained mathematician who is able to unpack them.

So a question arises about how epistemic normativity facts relate to logico-mathematical facts about validity and necessary truth: what determines *which*

logically valid inferences it is ok make without further argument?

Philosophers have tried to give a deeply principled explanation for why some logically valid arguments qualify as proofs and can confer justification, by saying that proofs need to break things down to premises and inferences which are conceptually necessary, in the sense that anyone who counts as having the the relevant concepts must find them (fairly immediately) attractive. One can then say that the difference between the inferences which figure in an acceptable proof of FLT and the (equally valid but) somehow inappropriate inference which occurs in the one line proof is that only the former are conceptually necessary.

But this idea turns out to be quite hard to develop and defend. The problem is that, as Boghossian[2] has pointed out, theories which say it is default reasonable to accept any genuine conceptual truth give rise to problem-cases with regard to concepts that seem to ‘pack too much in.’ The conceptual truth proposal authorizes us to reason about any coherent concept, not merely those we are justified in believing are coherent. But, this permission to use concepts that are non-obviously coherent can be parlayed into permission to explicitly believe correspondingly non-obvious necessary truths via a short argument. In particular, we can design a concept whose coherence depends on the truth of FLT and use that concept to infer the truth of FLT in the numbers.

To illustrate this point, consider the concept ‘schnumber’ characterized by the following bundle of claims:

1. The schnumbers satisfy the Peano axioms [17] for arithmetic.
2. No proper initial segment of the schnumbers satisfies the Peano Axioms.
3. The schnumbers satisfy FLT. i.e.. There are no schnumbers a, b, c, n with $n > 2, a, b, c > 0$ such that $a^2 + b^2 = c^2$.

Since FLT is true of the numbers, the above claims characterize a coherent concept (one can assign the same extension to the schnumbers as one does to the numbers). Thus, if some creature were assume these statements and thereby

possess the concept of schnumber, they can give a quick 'proof' that Fermat's last theorem holds for the numbers as follows:

Theorem 4.1. *There are no natural numbers a, b, c, n with $n > 2, a, b, c > 0$ such that $a^n + b^n = c^n$.*

Proof. By standard results in mathematical logic, any structure satisfying the Peano axioms has an initial segment isomorphic to the numbers. By 2 that initial segment can't be proper and hence the schnumbers are isomorphic to the numbers. By 3 the schnumbers satisfy FLT and by isomorphism so do the numbers. □

Thus, it would seem that any creatures with the schnumber concept can come to know that FLT is true by way of the short argument above. Yet, intuitively, the argument above is just as inadequate as the one line proof we originally considered²³.

We may attempt to block the proposal above by denying that there's a genuine concept schnumber corresponding to the bundle of inferences above. One might argue that, although coherent, the schnumber concept is bad because it 'packs something extra in.' Thus, one might try to say that there are only genuine concepts corresponding to bundles of inferences such that none of the relevant premises or inference rules are redundant.

One problem with this line of response is that it's not clear that there is any psychologically realistic way to individuate the constitutive premises and inference rules for our mathematical reasoning which satisfies this constraint. The premises and inferences which we find immediately obvious and take as unargued premises in apparently justified mathematical reasoning seem to involve a great deal of redundancy. For instance, the least number principle and the principle of induction both seem obvious and can figure in apparently justified deductions about the integers.

²³Note that we routinely use facts about the numbers to demonstrate various claims hold of certain groups, or facts about the complex numbers to show things about the natural numbers. Thus one can't object to this proof on the ground that it relies on conceptual truths about a concept other than the one the conclusion is about.

More importantly however, it is trivial to modify the schnumbers concept so it's not redundant. Rather than asserting that "The schnumbers satisfy the Peano axioms for arithmetic." we simply modify the first claim to instead assert that "If FLT is true in the schnumbers then the schnumbers satisfy the Peano axioms for arithmetic."²⁴

In view of this problem (and other problems for other principled stories, as discussed further in REDACTED ²⁵), it seems attractive to say that what's really going on is just this. Human beings are lucky enough to find the logical coherence of PA immediately (or almost immediately) plausible. Thus, proofs that derive results from PA are useful to us, because they 'break things down' into premises and inferences which human beings find immediately compelling.

In contrast, human beings do not find the logical coherence of PA+FLT immediately obvious. Anything but! Thus a proof which invokes the concept of schnumber, and appeals to an un-argued presumption that PA+FLT is coherent, is not useful or convincing to human beings. Similarly, human beings are not inclined to find the one step (logically valid) inference from PA to FLT immediately compelling. So the one line argument for FLT does not count as a proof, or provide adequate justification (in the sense of our concept of justification) for believing its conclusion.

So, to summarize, there are reasons to think that our notions of justification, proof and adequate mathematical argument reflect contingent facts about what logical truths we humans are inclined to find immediately obvious. The facts about what logically valid inferences and assumptions it is epistemically permissible to make without further argument (partly) reflect contingent psychological facts about how much insight into logic and coherence human beings are lucky enough to have.

²⁴See REDACTED for a more expanded version of this argument and related ones.

²⁵Also see [24] for an independent defense of this idea that there is nothing intrinsically special about the subset of valid logical deductions which we allow to figure in acceptable mathematical arguments.

4.2 Simple Psychologism about Acceptable Priors

I think the negative points above (about arbitrariness in learning parameters, applying uniformity intuitions and the epistemology of mathematics) and this positive example of the attractiveness of partial psychologism about a different epistemically normative notion ('adequate logico-mathematical proof'), provide some motivation for partial psychologism about acceptable priors. One might also take the history of 20th century philosophical failures to discover a plausible principled answer to the problem of priors to remove much of the sting from 'settling for' a story like mine on which there is nothing intrinsically special about the epistemically correct priors.

However one might worry that this doctrine isn't compatible with any plausible larger picture of the nature of epistemic normativity. So in the next two sections I will discuss a few options for fleshing out partial psychologism about priors which I advocate, and fitting it into such a picture. But remember that I don't pretend to provide a complete true story. Instead, I am merely attempting to dissolve the intuition that there *must* be something intrinsically good and special about those claims which it is epistemically permissible to assign high probability to in advance, by giving some plausible examples of theories of epistemic normativity on which this is not the case.

This conclusion is all I need for the task at hand. For once we block the assumption that facts about acceptable priors must be principled, we block the skeptic's argument that because no deeply principled theory of how to assign priors could let one assign high prior probability to \neg PEASOUP, it is not epistemically permissible to do so.

So, let me begin by introducing what I will call 'Simple (Partial) Psychologism' about priors (to contrast with the more complex variants which I consider in the next section).

This doctrine begins with the idea that human beings are inclined to substantially agree in how they assign priors. The particular priors which which human beings are actually inclined use involve a mix of symmetry intuitions,

preference for simplicity and permission to learn from experience.

There's nothing special about this mix: it just happened to be reasonably useful and easy to physically realize in the human brain in the context of evolution. We think that using these priors and doing Bayesian conditionalization is truth-conducive and reliable to a certain (fairly significant) degree when in the actual world²⁶. But this degree of reliability does not distinguish this way of assigning priors from various other ways of assigning priors.

Because human beings have this kind of large agreement on priors, it would not be surprising if we developed an (approximately) shared notion like 'adequate scientific argument' and 'good reasoning' which distinguishes empirical arguments which establish their conclusion from combining sensory experience with (something like) *the kind of prior judgments about theoretical elegance which normal human beings find compelling* from those which do not.

This psychologistic approach to the problem of acceptable priors may seem strange²⁷. However, I think this approach is supported and motivated by the fact that there are plenty of other concepts which seem to work approximately this way. For we clearly have some notions which behave like this, e.g., 'edible' applies to the kinds of substances which human beings can eat. Couldn't notions like 'justification' or 'adequate argument' behave similarly – applying to exactly those arguments which establish things from the kind of prior judgments of theoretical elegance which human beings happen to find compelling²⁸?

We don't take there to be any deep fact which the edible substances have in common beyond their being substances which can be digested by and nourish normal adult humans. Nor do we take there to be anything intrinsically joint

²⁶For example one might crudely cash this out in terms of the objective physical probability of this method producing true beliefs – given some measure on the space of beliefs and locations where one might find oneself in the total history of the actual world.

²⁷Perhaps Nelson Goodman's idea of what distinguishes projectable predicates like 'blue' from non projectable ones like 'grue' in [11] partly reflects contingent facts about what predicates have been 'entrenched' by the use in making successful inductive generalizations in the past has some kinship to what I am proposing here, and people have definitely found that strange.

²⁸Note that the mention of evolution above is just one way of dramatizing the image of ourselves as creatures who happen to find some arguments compelling and not others. It is not supposed to provide a further *justification* for taking these at face value.

carving about the boundary between more and less than a handful of stuff. Rather, the notions of ‘edible’ and ‘handful’ seem to track a distinction in the world which is interesting to us and useful to pay attention to because it tracks a natural distinction in *how normal humans can relate to* certain objects, rather than because it tracks a deeply natural kind which is independently interesting²⁹. I contend that justification is another such notion.

One might think that accepting this kind of psychologism forces us to accept the unattractive consequence that if human psychology were different then different arguments would be justified. However it seems quite natural and plausible that terms like justified (and even ‘edible’ or ‘handful’) should apply in a rigidified way – that they should, in all possible worlds, apply to the class of substances which people in the actual world can eat. Thus psychologism about the problem of priors is perfectly compatible with strong internalist intuitions about priors. One can say that even thinkers at possible worlds where human psychology is very different would still be justified in assigning the same priors, and making the same kind of inferences from observations which we do (see 4.4 for more detail on how I think we should think about such worlds).

One might also worry that giving this kind of account requires making a very strong and controversial psychological assumption that people exactly agree in how they assign priors. However, this is not so. Even though there is substantial agreement in different people’s judgments of theoretical elegance, it is plausible that they will not all exactly agree. But we can accommodate this possibility by saying that any range of priors which is accepted by sufficiently many people is a permissible starting point, so there will be some cases where two epistemically virtuous agents disagree because they are starting with slightly different popular and permissible assignments of priors. We could also say that each of us uses a slightly different concept of justification (reflecting the kind of priors we are psychologically inclined towards), but converse in a way that equivocates between these very similar concepts is common because it is predictably

²⁹See the large literature on response dependent concepts in ethics.

harmless.

Finally, note that (even the rigidified) version of this view avoids the appearance of coincidence regarding our accuracy about acceptable priors. For, insofar as we do not take there to be anything special and reference magnetic about these priors, we can explain the match between human psychology and acceptable priors by noting the following. If people had been inclined to use slightly different priors, they would have wound up using a different but equally interesting notion of justification' and spoken truly about that notion, rather than falsely about justification³⁰.

4.3 Reliabilist Psychologism about Acceptable Priors

One might worry that on the psychologistic story above *any* immediately attractive method of judging theoretical elegance will be epistemically permissible to use – if not epistemically mandatory. However, it's well known that in deductive reasoning people are subject to a variety of systemic biases, like the gambler's fallacy and 'affirming the consequent'³¹ and it's natural to think that there are similar biases in our judgments of theoretical elegance which don't violate the axioms of probability theory, but such biases shouldn't confer justification.

For example, maybe human beings have a 'looks-essentialist' tendency to assign higher prior probability to hypotheses on which visual differences between people go along with deep and unchangeable psychologically differences, but this is (nonetheless) unjustified. We can accommodate such intuitions by modifying the story above in one of two ways.

First, one can tell an internalist story, on which facts about a priori justification reflect facts about which assignments of priors we are not disposed to give up upon further reflection. That is, we are justified in assigning priors in those ways which we find initially attractive and wouldn't reject on further

³⁰Also, note that my strategy is compatible with saying that believing as we do is epistemically obligatory (not just permissible), and defending this bold claim from skeptical objects (but not, of course, convincing the skeptic that it is true). It might be mandatory for justified belief to assign priors as we do, but aliens with equally (actual world) truth conditions would have justification* for reasoning as they do.

³¹That is, inferring from $P \rightarrow Q$ and Q to P .

consideration.

Second, one can tell a more externalist story, on which we could discover that certain ‘components’³² of our ways of evaluating theories a priori (e.g., the ‘looks-essentialist’ theoretical preference above) reduce our tendency to form true beliefs in the actual world. One might then say that these aspects of our methods of a priori theory choice are ‘bad’ (in being insufficiently actual-world truth conducive), and characterize the acceptable priors as those you would get by using all the psychologically natural mechanisms of a priori theory choice *minus these bad apples*³³.

Thus, partial psychologism about justification, as sketched above, can be spelled out in a way which allows some aspects of the way we currently evaluate the a priori attractiveness of scientific theories to be unjustified.

4.4 On Epistemic Twin Earth and the Intended Deflationary Significance of This Proposal

One can dramatize the intended substantiveness and (moderately) deflationary import of these psychologistic approaches to epistemic normativity by considering Terence Horgan and Mark Timmons’ moral twin earth paper[15] and the Sentimentalist tradition in moral philosophy.

Horgan and Timmons consider a scenario in which reflecting on our moral sentiments would lead to a deontological ethics, but there is a twin earth containing people with slightly different moral sentiments (specifically, a different balance of propensities to guilt vs. sympathy) such that reflection would lead them to a consequentialist ethics. From a Humean/sentimentalist point of view,

³²Note that the literature on the generality problem suggests that many epistemological theories will need to presume something like this (i.e., some preferred or natural way of thinking about our actual reasoning as belonging to preferred general methods and types) like different weighted elements in an a priori plausibility determining function.

³³This is just a very crude sketch intended to give readers the idea of this approach. A serious theory of this form would need to deal with the possibility of interactions between different ‘components’, e.g., maybe just removing component a makes you get to the truth faster, and just removing component b does, but removing both of them is makes you much less reliable. A serious theory would need to say what combination of components corresponds to justified assignment of priors in this case. There will be some amount of vagueness about justification corresponding to things like how unreliable a procedure has to be to count as ‘sufficiently unreliable’. But nearly all natural language terms have the same kind of vagueness.

it seems immediately attractive to say that such people would be using a different concept of morality with equal metaphysical interest to our own, and not necessarily getting anything wrong³⁴. But many people who consider themselves moral realists find this a deeply unattractive conclusion and have the strong intuition that people on this twin earth must be latching on to the same concept of morality as we have, and substantively disagreeing with us about how it applies.

Analogously, we can imagine some analogs to moral twin earth involving various epistemically normative notions like ‘adequate mathematical proof’ or (our topic here) ‘acceptable assignment of priors’. We can imagine a twin earth inhabited by people who find³⁵ ‘If PA_2 then Fermat’s Last Theorem’ immediately obvious, and hence say different things about which arguments qualify as “mathematical proof”. Similarly, (more directly to the point here) we can imagine a twin earth whose inhabitants assign higher prior probability to truths, like the atomic hypothesis, or anti-vitalism, than we do – and thus say different things about “justification”.

What shall we say about these people and their concepts of “mathematical proof” and “adequate scientific justification” ? They don’t provide proofs, or justifications of a kind that would let them share their knowledge with us.

I want to embrace the conclusion that denizens of both moral and epistemic twin earth would be speaking the truth. I recognize that this conclusion may be surprising and strike some readers as initially unpalatable, in much the same way Humean projectivism/sentimentalism/response dependence about morality does. But I think that closer consideration of the merits of this position (along the lines developed in this paper above) gives it similar philosophical attraction.

Let me conclude this subsection by answering an objection which may naturally come to mind, which may be raised by my characterizing partial psychol-

³⁴It also seems attractive to say that such a person would not be missing something deep about the world by lacking our concepts any more than we are by not thinking about their concepts

³⁵If desired one can imagine that they are designed by a benevolent biohacker to find these extra truths obvious in whatever way we could be said to have been designed by evolution to find the logical truths we find obvious obvious.

ogism an epistemic analog to humean sentimentalism about morality.

One might think fear this view leads to some kind of paradoxical and/or self-undermining ‘epistemic relativist’ doctrine that all conceivable variant justification and proof practices (e.g., astronomy vs astrology) are somehow *equally truth conducive*. However, this is not so. For the (partial) psychologism about epistemic normativity I am advocating endorses straightforward realism about all kinds of ‘descriptive’ (in the sense that is traditionally contrasted with normative) facts.

It makes no claim that all possible notions of justification are equally truth conducive. Instead, it says that such comparisons of truth conduciveness aren’t what determine justification facts, and that acknowledging that some other possible shmustification practices (like those of the luckier mathematical and scientific reasoners above) are more truth-conducive (in these various externalist/descriptive senses) than our justification practices doesn’t lead to any kind of problem or paradox concerning our own practices of trying to apportion our beliefs to what is justified.

5 Conclusion

In this paper I have drawn a distinction between two different kinds of external world skepticism: Belief and Confidence skepticism. I have suggested that Confidence Skepticism is what most people care about, and argued that it is largely motivated by the expectation that facts about epistemic normativity must be principled in a certain way. I have argued that taking facts about epistemic normativity to be somewhat unprincipled is unavoidable, and tried to make that conclusion palatable by developing an attractive positive story which illustrates how our practice of talking about justified confidence could function and be useful without being principled in the relevant fashion.

Let me close suggesting that there are actually two deeply different motivations for classic external world Belief Skepticism, which might benefit from

separate treatment in the future. First there's the motivation for Confidence Skepticism which I have discussed in this paper: a feeling that justified high confidence in a proposition like \neg PEASOUP requires being able to point to a certain kind of 'positive reason to believe' this proposition, which we find ourselves unable to provide. Since justified belief presumably requires justified high confidence, this provides one route to Belief Skepticism (a route which I have tried to help close off).

Second, there there's a quite different route to Belief Skepticism which might be called the 'argument from might'. This route exploits the general fact that our failure to assign probability 1 to a proposition can be used to raise doubts about knowledge of that proposition (via natural language reasoning about 'might') – even in cases where the correctness of all one's degrees of confidence is taken entirely for granted. Consider the following argument, which I have adapted from Hajek's 'Most Counterfactuals are False'[12]: 'Quantum mechanics says there's a very small but positive probability that a giraffe will appear in this room in the next 5 seconds. Thus, a giraffe might appear in this room in the next five seconds. So I don't know that a giraffe won't appear in this room in the next 5 seconds'. As PEASOUP doesn't seem to warrant total rejection a priori and none of our a posteriori observations are incompatible with it, this argument suggests we can't know \neg PEASOUP. This is what I take to underly Blumenfeld and Blumenfeld's argument.

Attempting to answer or provide philosophical therapy for this (second) motivation for (belief) skepticism about the external world is beyond the scope of this paper. I would merely like to suggest³⁶ that it forms a specific instance of a general problem about how to fit together the binary vocabulary of belief with the continuous vocabulary of confidence. It may be that claims about justified confidence capture everything you'd want to know in a more nuanced way than belief claims, but moving between the two notions is a source of many puzzles. The argument just described raises issues about the relation of confidence and

³⁶See [14] for a related position .

knowledge but similar issues arise just moving between degrees of confidence and belief. For example, consider the following puzzle³⁷: suppose I assign 80% probability to something and then it happens. Do I count as having already believed that it would happen? How high a probability would I have to assign (or what other contextual factors are relevant to this)?

References

- [1] Yuval Avnur. Hawthorne on the deeply contingent a priori. *Philosophy and Phenomenological Research*, 83(1):174–183, 2011.
- [2] Paul Boghossian. Blind reasoning. *Aristotelian Society Supplementary Volume*, 77(1):225–248, 2003.
- [3] Tim Button. *The Limits of Realism*. Oxford University Press UK, 2015.
- [4] Rudolf Carnap. *Logical Foundations of Probability*. Chicago]University of Chicago Press, 1950.
- [5] Lewis Carroll. What the tortoise said to achilles. *Mind*, 104(416):691, 1995.
- [6] David J. Chalmers. The matrix as metaphysics. In Christopher Grau, editor, *Philosophers Explore the Matrix*, page 132. Oxford University Press, 2005.
- [7] Stewart Cohen. Why basic knowledge is easy knowledge. *Philosophy and Phenomenological Research*, 70(2):417–430, 2005.
- [8] Rene Descartes. *Meditations on First Philosophy*. Oxford University Press, 2008.
- [9] Hartry Field. Recent debates about the a priori. In Tamar S. Gendler and John Hawthorne, editors, *Oxford Studies in Epistemology*. Oxford University Press, 2005.

³⁷I learned about this issue from REDACTED, in conversation.

- [10] Bas C. Fraassen. *Laws and Symmetry*. Oxford University Press, 1989.
- [11] Nelson Goodman. *Fact, Fiction, and Forecast*. Harvard University Press, 1955.
- [12] Alan Hájek. *Most Counterfactuals Are False*. [Http://Philrsss.Anu.Edu.Au/People-Defaults/Alanh/Papers/Mcf.Pdf](http://Philrsss.Anu.Edu.Au/People-Defaults/Alanh/Papers/Mcf.Pdf), Accessed June 12, 2014, 2014.
- [13] John Hawthorne. Deeply contingent a priori knowledge. *Philosophy and Phenomenological Research*, 65(2):247–269, 2002.
- [14] John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, 2003.
- [15] Terence Horgan and Mark Timmons. New wave moral realism meets moral twin earth. *Journal of Philosophical Research*, 16:447–465, 1991.
- [16] David Hume. *A Treatise of Human Nature*. Oxford University Press, 2000 [1739].
- [17] Richard Kaye. *Models of Peano Arithmetic*. Clarendon Press, 1991.
- [18] Saul Kripke. *Naming and Necessity*. Blackwell, 1972.
- [19] Samir Okasha. Bayesianism and the traditional problem of induction. *Croatian Journal of Philosophy*, 5(2):181–194, 2005.
- [20] Hilary Putnam. Brains in a vat. In Sven Bernecker and Fred I. Dretske, editors, *Knowledge: Readings in Contemporary Epistemology*, pages 1–21. Oxford University Press, 2000.
- [21] John Turri. Contingent a priori knowledge. *Philosophy and Phenomenological Research*, 83(2), 2011.
- [22] Peter K. Unger. *Ignorance: A Case for Scepticism*. Oxford University Press, 1975.
- [23] Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, 2000.

- [24] Timothy Williamson. *Absolute provability and safe knowledge of axioms*.
Oxford University Press, 2016.