

Epistemic Unprincipledness

REDACTED

[Draft work: please drop me an email to discuss or see the latest version]

Abstract

Epistemic consequentialist and deontological views have battled in recent debates about the foundations of epistemic normativity. In this paper I'll argue that a view I'll call epistemic unprincipledness provides an attractive alternative.

1 Introduction

Which logically valid inferences are epistemically permissible to make immediately? Many people share the intuition that modus ponens inferences require no further justification, while inferring Fermat's Last Theorem (FLT) from the ZFC axioms does. Despite both being logically valid, beliefs formed through the latter are not considered justified. But why?

This kind of puzzle about permissible inferences has inspired a view I'll call epistemic unprincipledness, on which aspects of our epistemically normative concepts (like justification and adequate argument) reflect contingent human psychology rather than principled intrinsic distinctions. In this paper, I will relate epistemic unprincipledness to recent foundational debates, and argue (partly by expanding on prior work) that it provides an attractive alternative to epistemic consequentialism and deontology.

2 What is epistemic unprincipledness?

2.1 Principled Answers to Foundational Questions?

In addition to familiar questions about the extensions of concepts like knowledge and justification¹, it seems we can ask further foundational questions about epistemically normative concepts. Even if we had a perfectly correct extensional theory of which beliefs are justified, we could still ask why *these* beliefs are justified but not others. In [4] Boghossian expresses this kind of foundational question (focusing on the case of warranted deduction) as follows, “I am asking by virtue of what facts a deductive inference transfers warrant, and not just under what conditions it does so” [4]. Similarly, we might want to know not just *which* factors are relevant to a belief counting as justified but *why* those factors are relevant.

Philosophers often attempt to answer such foundational questions about our epistemic concepts by appealing to intrinsic features like accuracy promotion or respect for accuracy[22, 1]. For example, in recent debates between epistemic consequentialists and deontologists, both parties share the assumption that the boundaries of justification and our other epistemically normative concepts can be fully explained by appeal to some such principled features².

In particular, both parties agree that the fundamental value relevant to epistemology is accuracy and hope to explain the boundaries of our concept of justification (and other such epistemically normative concepts) by appeal to this value, as per the following doctrine.

Valuing Thesis: The fundamental normative explanation of why jus-

¹c.f. the quest for informative necessary and sufficient conditions for knowledge pursued in the Gettier literature[15]

²I take David Enoch to be getting at a similar foundational explanatory project in [11] when he expresses hopes of “vindicating” our basic reasoning methods like IBE by “drawing a principled distinction between [the basic reasoning methods we use and take to be legitimate] and methods we are not justified in employing as basic, a distinction that presents them in a positive light”

tified beliefs are justified is that they manifest certain ways of valuing fundamental epistemic value.[22]

The disagreement between these two groups arises when we ask how to cash out the above notion of valuing accuracy. Consequentialists understand valuing accuracy in terms of accuracy *maximization*. So they answer foundational questions about why certain belief forming methods are justified (and the like) by considering what maximizes accuracy. Deontologists understand valuing accuracy in terms of *respect* rather than maximization³ ⁴. So they would explain why certain kinds of beliefs are justified in those terms.

In contrast to both views, the epistemic unprincipledness thesis I'll be advocating rejects the above shared assumption of principledness (i.e., that foundational questions about our epistemically normative concepts can be answered entirely by appealed to principled intrinsic features of the things being classified, like accuracy promotion or respect). Instead, it posits that some distinctions in our concept of justification arise from contingent features of human psychology, such as which inferences we find compelling⁵.

³Doing this is (among other things) supposed to help explain how believing a dictator's outlandish pet theory can be unjustified though doing so will secure life and funding to acquire many other true beliefs, by appealing to a 'separateness of propositions' [1], analogous to the separateness of persons which forbids killing the one to save five in trolley cases. The deontologist holds that, "[V]alue generates all demands, [but] these aren't exhausted by promotion. [For example] friendship's value demands that I not betray my friend Mike even to cause several new friendships to form." [22].

⁴Although advocating a deontological approach, Sylvan says rather little about how to analyze the relevant notion of respect, beyond associating it with conforming your beliefs to the evidence. He identifies respecting accuracy with conforming your beliefs to the evidence and responding to reasons understood in terms of what beliefs/evidence/facts 'indicate to be true' without endorsing any substantive analysis of the latter indication relation. He says merely, "I leave open how we are to analyze these different indication relations. But one natural view would propose that the truth-indication relation is a special case of the probabilification relation. On this view, the three relations correspond to three different notions of probability: personal, evidence relative, and objective in some sense relevant to epistemology (for example, Keynes (1921)'s sense or the sense in play in some contemporary objective Bayesian views (for example, Williamson (2010)'s)." [22].

⁵In this way, epistemic unprincipledness presents a kind of Humean alternative to the Kantian and Millian approaches to foundational questions above. Hume famously depicted human minds applying moral concepts as, "gilding or staining all natural objects with the colors, borrowed from internal sentiment." And Humean moral sentimentalists hold that moral terms draw distinctions in ways that ultimately reflect mere contingent psychological facts about

Epistemic Unprincipledness: Some important aspects of our concept of justification (and accordingly also our concepts of knowledge, adequate evidence, good argument proof etc.) reflect contingent features of human psychology rather than drawing an intrinsically principled distinction.

In particular, our epistemically normative concepts partly rigidify/project facts about which accuracy promoting reasoning methods actual humans find compelling.

So, for example, an advocate of epistemic unprincipledness might say that at each metaphysically possible world, a subject (regardless of what inferences they find compelling) has defeasible warrant for making only those kinds logically valid deductions⁶ which are psychologically compelling to humans *in the actual world*.

For example Martians who found different logically valid inferences immediately compelling would not count as justified in drawing these extra inferences. However, this fact does not reflect anything metaphysically special about us, or suggest that human psychology correctly matches some independently interesting distinction among different logically valid inferences (in a way that might, e.g., give rise to access worries). Rather, such Martians could have their own equally principled notion of justification*, such that beliefs formed via ‘proofs’ using the extra logically valid inferences Martians find compelling count as jus-

how actual humans are disposed to react, rather than tracking intrinsically principled, natural kind properties of the acts/agents being evaluated. Analogously, my epistemic unprincipledness thesis claims that the boundaries of our concepts like justification and knowledge partly reflect mere contingencies of our psychology (e.g., which of the many logically valid inferences are humans disposed to find immediately compelling?) – rather than anything more principled (like facts about what maximizes accuracy or shows respect for the value of true belief).

⁶Talk of ‘kinds of reasoning’ obviously brings up the generality problem[14]. I won’t try to solve it here, but note that we might want a psychologistic element as well: what kind of things we group together psychologically naturally, so that e.g., stopping finding one inference in the set immediately compelling also disinclines you to make others. Consider how rejecting one instance of affirming the consequent can somewhat automatically make people more wary of others.

tified* but not justified⁷. After all, it's unsurprising that humans would find it useful to have a concept of justification which (among other things) tracks ability to justify ones beliefs via the kind of logically valid inferences *we humans* find immediately compelling.

2.2 Clarifications

With this idea of epistemic unprincipledness now on the table, let me clarify two points.

First, accepting epistemic unprincipledness does not commit us to any kind expressivism or truthvalue anti-realism about epistemic concepts. For (presumably) accepting that the boundaries of our concept of 'handful' reflect contingent facts about the typical size of human hands does not commit one to saying claims like 'there are two handfulls of nuts on the table' fail to be true or false in a completely ordinary and straightforward sense⁸⁹. And my epistemic unprincipledness thesis just makes an analogous claim about some of our epistemically normative concepts.

Second, accepting epistemic unprincipledness (as I will understand it) only requires maintaining that *some* aspects of our epistemically normative concepts reflect contingent features of human psychology rather than something deeply principled and joint carving. So the unprincipledness theorist is free to say

⁷See [9, 10] for some recent discussions of pluralism along these lines in other areas like ethics and mathematics.

⁸What I'm calling epistemic unprincipledness can be closely mimiced by a form of appraiser (as opposed to agent) relativism, which technically allows concepts like justification and knowledge to be entirely principled, by taking them to have a hidden extra parameter. Someone who accepts this kind of relativism might say that Martians (in the example above) often have justification relative to human inference practices, but not relative to Martian ones. And they might say that normally context fills in this extra parameter, such that when we talk about whether such Martians would be 'justified' we mean justified relative to human practices. I take such proposals to be extremely similar to epistemic unprincipledness in spirit, and won't argue against them here.

⁹Epistemic unprincipledness differs from views like Chrisman's epistemic expressivism[8, 7], by allowing one to take taking epistemic normativity claims to be straightforwardly true or false and not including any claim that "epistemic judgments have, at least in part, a desire-like direction of fit with the world".

that an ideal answer to foundational questions about epistemically normative concepts would combine appeal to principled facts (about accuracy promotion or respect etc.) with appeal to contingent facts about actual human psychology¹⁰.

Given this clarification/concession, some readers may feel that it's odd to classify epistemic unprincipledness as an *alternative* to epistemic consequentialism or deontology. They may feel that accepting a pinch of arbitrariness doesn't change our fundamental picture of justification very much. I disagree. For example, note how epistemic unprincipledness suggests that we can sometimes classify someone's belief as being unjustified *without* taking this belief to be bad in any intrinsic or principled sense that the believer in question intuitively should care about. This suggests a poignant sense in which we might turn out not to live in a 'shared space of reasons' with all other thinkers. However, none of the arguments that follow will depend on this view.

3 Motivating Epistemic Unprincipledness

With this characterization of unprincipledness in mind, I now will now review two puzzles which have already been used to argue for epistemic unprincipledness (though not under that name). In the next section I try to buttress the -hitherto less discussed- argument from the second puzzle, by considering and highlighting problems for alternative approaches to handling this puzzle.

¹⁰For example, a fan of epistemic unprincipledness (as I shall understand it) could say that that beliefs formed by deduction from known premises have defeasible justification iff relevant deductions are both logically valid and the kind of inferences actual humans are suitably disposed to make. And, more generally, an advocate of epistemic unprincipledness could say that the acceptable basic inference methods are those which are **both** sufficiently accuracy promoting/respecting and sufficiently attractive to actual world human psychology. This would put them in a position to mirror many appealing consequentialist/deontologist explanations for why certain things *aren't* justified/rational.

3.1 Logical Inferences

The first puzzle concerns logically valid inferences. It appears that one can form justified beliefs by making certain logically valid inferences (like modus ponens) without further argument, but not others –like inferring FLT from ZFC in a single step. But why is this, given that, e.g., both arguments are equally logically valid and hence necessarily truth-preserving? Initial appearances and a history of failure to find a principled response to the above problem¹¹ suggests the following thought.

Perhaps there’s no deep or interesting feature that distinguishes the logically valid inferences which can be made immediately from other logically valid inferences – and no substantive sense in which creatures who were lucky enough to find additional logically valid inferences immediately compelling would be worse off than us. So perhaps corresponding aspects of our concept of justified belief (specifically, facts about when beliefs got from known premises by logically valid inferences count as justified) reflect mere contingencies of human psychology, like the fact that actual humans are disposed to find certain logically valid inferences and not others immediately compelling.

On this view, Martians who found different logically valid inferences immediately compelling would not count as justified in drawing these extra inferences. However, this fact does not reflect anything metaphysically special about us, or suggest that human psychology correctly matches some independently interesting distinction among different logically valid inferences (in a way that might, e.g., give rise to access worries). Rather, such Martians could have their own equally principled (or equally unprincipled) notion of justification*, such that beliefs formed via ‘proofs’ using the extra logically valid inferences Martians

¹¹Schechter’s paper [20] nicely summarizes known problems for a number of existing attempts to provide a principled explanation for which logically valid inferences can vs. can’t be made immediately.

find compelling count as justified* but not justified ¹².

And on this view, contrary to epistemic consequentialism and deontology, the ‘problem’ with a Martian who believes FLT on the basis of a single step deduction from ZFC axioms can’t be cashed out in terms of how any intrinsic features of her situation relate to things we epistemically value like accuracy. We can’t say her deductive methods fail to promote accuracy (plausibly they promote accuracy better than ours do). Nor need she fail to respect accuracy in any obvious way. For example, we can imagine that she finds this deduction immediately compelling, takes such logical intuitions at face value (unless provided positive reason for doubt) like the rest of us, doesn’t use methods she regards as less accuracy conducive than alternatives she knows how to deploy etc.

Our problem with such a reasoner is merely that what she’s doing isn’t helpful to *us* in ways that human mathematical reasoning often is. We don’t classify her ‘proofs’ as adequate mathematical arguments or providing justification because they don’t proceed by the kinds of steps that we humans find a priori compelling (or sufficiently help competent human readers create an expanded proof which does proceed by such steps).

3.2 Priors

A second puzzle arises about when it is epistemically permissible to assign certain metaphysically contingent truth high confidence a priori¹³, and why.

It is often useful to model scientific reasoning in Bayesian terms¹⁴, by supposing a thinker starts with a certain assignment of probabilities (satisfying the

¹²See [9, 10] for some recent discussions of pluralism along these lines in other areas like ethics and mathematics.

¹³Here I present further development of an argument briefly suggested in [2].

¹⁴For concreteness, I will state this puzzle in Bayesian terms. However I expect my arguments can be fairly straightforwardly rewritten any plausible alternative frameworks for thinking about a priori scientific theory choice.

probability axioms) in advance of all experience — and then updates their belief state by conditionalizing on various evidence propositions which they learn from their senses. When we do this, facts about which priors it is epistemically permissible to have will play an important role in determining both which conclusions a person's total body of evidence permits them to accept/assign high probability to and (as a result) what they can be justified in believing.

However, it is possible to raise a puzzle about acceptable priors which parallels the puzzle about basic logical inferences above (and strengthens the case for epistemic unprincipledness) as follows.

Prima facie it seems to be epistemically permissible to assign high a priori probability to certain metaphysically contingent physically necessary truths (e.g., claims that the future resembles the past in certain ways) but not to others (e.g., the facts summarized in the periodic table of elements). And plausibly people can form justified beliefs in this way. But what explains this different epistemic standing of different claims (physically necessary truths)? Just as we can ask 'why is it OK to make some logically valid inferences but not others?' we can ask 'why is it OK to assign very high a priori probability to some truths/physical laws but not others?'

Historically philosophers of science have put a great deal of effort into providing a principled and motivating characterization of what assignment(s) of priors are epistemically permissible. But this project faced significant difficulties. For example, consider Carnap's general difficulties formulating a logic of induction and the specific problem of finding a principled motivation for a specific choice

of Carnapian learning parameters¹⁵ and Bertrand's paradox¹⁶.

And in [2] Berry notes that (for reasons I will expand on below) that our actual priors don't seem to be uniquely, or even maximally, actual world accuracy conducive. So it's not obvious that appeals to the actual world reliability of our methods can be used to find a principled feature that distinguishes epistemically permissible priors.

Epistemic unprincipledness resolves the above puzzle (about which propositions can be assigned high probability a priori) by regarding facts about which priors are acceptable as merely reflecting contingent human psychological tendencies that were evolutionarily useful and remain reasonably accuracy promoting in the actual world.

For example, in [2] Berry uses the above points about priors to motivate the following kind of psychologism, (which I take to be a version of epistemic unprincipledness).

[H]uman beings are inclined to substantially agree in how they assign priors. The particular priors which which human beings are actually inclined [to] 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 rea-

¹⁵When Carnap modified his theory of the logical foundations of probability [6] 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 without any 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. Whether or not you 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.

¹⁶Bertrand's paradox points out that 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 a cube. Suppose a cube is known to have side-length between 0 and 4 meters (and therefore volume between 0 and 64 cubic meters). What probability should we assign to it having side-length ≤ 1 meter (and therefore volume ≤ 1 cubic meter)? If we go by side lengths (assigning equal probability to side lengths of 0-1 meters, 1-2 meters, 2-3 meters and 3-4 meters), we will say 1/4. But if we go by volume, we will say 1/64. This helpful formulation of Bertrand's paradox follows [12]).

sonably useful and easy to physically realize in the human brain in the context of evolution. We think that using these priors and doing conditionalization is truth-conducive and reliable to a certain (fairly significant) degree when in the actual world. 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.

4 Alternative Approaches to the Puzzle About Priors

I will now attempt to enhance the case for epistemic unprincipledness, by developing and expanding on the discussion of priors above. Specifically, in the first two subsections below I will argue there’s little hope of accounting for ordinary distinctions between justified vs unjustified scientific beliefs (reflecting judgments about acceptable vs. unacceptable priors) by appeal to epistemic consequentialists and deontologists’ favored notions of *accuracy conduciveness* and *respect for accuracy*. In §4.3 I will argue that a radical permissive approach (regarded as a live option in prior work like [2]) has little power to let us resist the general epistemic unprincipledness I’ve advocated above.

4.1 Accuracy Promotion

Let's start with the epistemic consequentialists' appeal to accuracy promotion.

I claim that one cannot answer the above question about acceptable priors by appeal to facts about which priors are most accuracy promoting *at the actual world/given actual physical laws*.

For we don't take our priors to be uniquely good (or even optimal) in terms of objective accuracy promotion. Given any reasonable way of cashing out the notion of objective accuracy promotion¹⁷ we can imagine priors that are better at accuracy conduciveness than our own, yet intuitively unjustified. To explain

¹⁷By talking about the objective accuracy promotingness of priors here, I mean something like the expected accuracy (calculated using objective physical probability not subjective probability) of credences got by via Bayesian updating on these priors in the actual world.

We might cash this out by considering the expectation - relative to objective physical probability - for the accuracy (perhaps measured by a Brier score, as described below) of a robot equipped with certain sensory faculties that gets plonked somewhere random in the actual world and does Bayesian updating on its stream of experiences (c.f. the footnote about this above). Note that a suitable notion of objective probability (as needed to calculate an objectively expected accuracy score) can be well-defined and non-trivial even if physics is completely deterministic[21].

There are clearly many choice points faced when filling in this notion. We face choices about how to weight accuracy about different propositions for calculating a Brier score, and how to cash out a robot being equally likely to appear anywhere (what metric on space and time to we assume), and how to weight current accuracy vs. accuracy after making some number of observations.

I won't take a stance on any of these questions. Instead I will try to make arguments that work however we imagine epistemic consequentialists (and other defenders of epistemic principledness) cashing out claims about objective accuracy promotion.

In [5] Brier proposed the following basic formula for the overall inaccuracy of a forecaster who assigns probability to N different events.

$$\frac{1}{N} \sum_{t=1}^N (f_t - o_t)^2$$

where

- f_t is the probability assigned to the claim
- o_t captures whether that event actually occurred (so it is 1 if the event occurred and 0 otherwise).

Note that Brier scores measure inaccuracy – in the sense that a lower score corresponds to better overall accuracy. This basic formula can be naturally generalized to assess forecasters that make infinitely many propositions, given a suitable weighting function $w(i)$, where the weights assigned to individual propositions sum to one.

We can then quantify overall objective accuracy conduciveness, by looking at the expected value of the Brier score for an agent starting with certain priors (and perhaps updating via certain kinds of observations) – while using objective physical probability (not subjective probability) to calculate expected value.

what I have in mind, imagine creatures who

- assign probabilities by Bayesian updating and obey all the standard probability axioms.
- assign high probability a priori to certain exceptionless scientific laws that we are not inclined to accept a priori, (e.g., the facts summarized in the periodic table of elements),
- otherwise have priors that largely resemble our priors

Such creatures would, from our point of view/intuitively err by being dogmatically confident in certain scientific claims, which happen to be physically necessary truths a priori. But there's not any obvious way in which their priors and reasoning methods need be less accuracy promoting than ours. For example, they will (even by our lights) tend to do *better* than us at quickly forming true beliefs/assigning high probability to true propositions in the actual world .

Thus we seem to have an interesting kind of modesty about our priors. We don't take our priors to be uniquely good (or even optimal) at actual world truth-conduciveness, or accuracy promotion. Our notion of justified belief seems to reflect a distinction (between truths which can vs. can't be assigned high probability a priori) that can't be explained solely by appeal to objective accuracy promotion alone¹⁸.

Accordingly it epistemic consequentialists cannot answer foundational questions about why certain priors are the right ones to have (e.g. why certain contingent truths can be assigned high probability a priori, while others cannot) by saying that our correct priors are uniquely objectively accuracy conducive at the actual world.

¹⁸Perhaps saying priors have to (somehow) track natural kinds has some power to block these arguments that almost anything could be assigned high probability a priori. However, it's far from clear that this requirement would eliminate enough intuitively unreasonable priors which are just as accuracy promoting as our own to block the objection from excessive permissiveness.

Could epistemic consequentialists avoid the above problem by instead characterizing accuracy promotion in a way that (somehow) considers how priors perform in all metaphysically possible scenarios (rather than just at the actual world)? Specifically, one might think that correct/permissible priors¹⁹ should be good (i.e. accuracy promoting) in a way that's independent of what possible world happens to be actual. So, perhaps epistemic consequentialists should say that good priors are distinguished by having optimal *a priori expected accuracy* – or some other notion which incorporates the accuracy promotingness of (baysena updating on) priors at *all metaphysically possible worlds*.

However I don't think this approach is viable. For, to acquire such a notion of a priori/general accuracy promotion (rather than accuracy promotingness given objective physical probabilities in the actual world), we'd need to specify some way of *weighting* different metaphysically possible scenarios when calculating expected accuracy. And it's hard to see how to do this without appealing to (our actual or ideally rational) subjective probability assignments for the weighting.

For example, we can certainly say that the rationally correct priors should have optimal expected accuracy *as calculated using rationally correct priors*, rather than objective probabilities in the actual world. One can sharpen this proposal as claiming that rationally correct priors have one of the following two properties

- high expected accuracy as calculated relative to the rationally correct priors
- high expected accuracy as calculated according to themselves i.e., Lewisian immodesty (they don't assign any alternative way of assigning priors a higher expected accuracy)²⁰.

¹⁹Because it is irrelevant to my arguments in this paper, I will bracket debates about whether there's a unique epistemically mandatory way of assigning priors or a small range of slightly different priors that are all epistemically permissible.

²⁰Readers familiar with David Lewis comments about the immodesty of priors in [17] may

However there are many quite different ways of assigning priors which share the Lewisian immodesty property of assigning themselves optimal expected accuracy. So saying the rationally correct priors must have Lewisian immodesty won't let us pick out a uniquely correct set of priors, or distinguish intuitively acceptable from many intuitively unacceptable ones. And neither version of the proposal above seems capable of answering foundational questions about why it's OK to assign high probability to some contingent truths a priori but not others.

One could instead say that the correct priors are only special in having high expected accuracy relative to probability assignments *actual humans find appealing*. That is, one might admit that the only specially good thing about epistemically correct/missible priors is that they have optimal expected accuracy calculated using the kinds of priors humans psychologically tend to have. But this seems to concede the game to advocates of epistemic unprincipledness.

4.2 Accuracy Respect

Now let's turn to the epistemic deontologist. Can we do any better with their idea that we can answer foundational questions about justification (and hence

wonder if there's a conflict. So let me quickly clarify how the modesty I'm claiming is compatible with the immodesty Lewis notes.

Imagine that you've eliminated all but two hypotheses about the fundamental physical laws of the world (e.g., H_1 the world is physically necessarily made of atoms and H_2 the world is physically necessarily made of gunk), and you assign 50% probability to each of these. In this situation, your current assignment of probabilities will likely combine the following features.

- Lewisian immodesty: Your priors assign themselves optimal expected accuracy, in the following sense. There are no alternative priors which your priors assign higher expected accuracy (e.g. a lower expected Breir score) than themselves.
- objective/actual-world modesty: In contrast, your priors do *not* regard themselves as optimal in terms of objective/ actual world accuracy promotion For you are confident that either the world is physically necessarily made of atoms (in which case priors that are dogmatically confident in H_1 will do better) or it is physically necessarily made of gunk (in which case priors start out confident in H_2 will do better). So you are highly confident that *some one* of these alternative way of assigning priors is more objectively truth conducive than your current way of assigning priors. But of course (as per Lewisian immodesty) you don't know *which* of these ways of assigning priors does better.

presumably also about correct priors) by appeal to facts about *respect* for the value of accuracy?

Perhaps I can imagine such an epistemic deontologist arguing that respect for accuracy requires satisfying *structural* requirements, like obeying the probability axioms (and so not being dutch-bookable). And there are many other conceivable ways this appeal to ‘respect’ could be cashed out. So I can’t hope to explicitly consider them all.

But, it’s quite hard to imagine how considerations of respect for accuracy could explain more substantive facts about which metaphysically contingent claims can be assigned high probability a priori. Note that our intuitions about epistemically acceptable priors (and hence also practices of distinguishing good from bad scientific arguments, adequate from inadequate evidence etc. when assessing which beliefs are justified) don’t just require satisfying structural principles like the probability axioms. They also take a stand on substantive matters of how one should rate the relative probability of various contingent empirical hypotheses a priori. They opine on things like how antecedently plausible we should find the claims that world is made of atoms vs of gunk, or the claim that all causation is local, that all space is Euclidean etc. But it’s hard to see how one could get any such substantive positions out of notions like accuracy promotion or respect for accuracy – unless you appeal to further facts about objective accuracy promotion (or actual human psychology as per epistemic unprincipledness). And appeals to objective accuracy promotion actually don’t seem capable of doing this job. For, as noted above, our priors don’t look uniquely good – or even optimal – from the point of view of objective accuracy promotion.

To put this point another way, it’s hard to see how the variant (intuitively bad but equally or more objectively accuracy promoting) priors discussed in §4.1 could be accused of failure to respecting the value of accuracy. For we saw

that these priors can plausibly be cooked up to satisfy all obvious structural requirements (like the probability axioms), while massively erring in more substantive ways, by dogmatically assigning high probability to true physical laws a priori.

And if we imagine meeting creatures who straightforwardly reasoned in accordance with these wonky priors, I don't think we'd intuitively be tempted to accuse them of any kind of lack of *respect for the value* of accuracy or truth. We'd just regard them as having irrational (and perhaps wrong) high confidence that certain particular strange things are true²¹. So it's hard to imagine how a foundational explanation citing respect for accuracy could rule out (or explain what's wrong with) assigning priors in these deviant ways.

4.3 Permissivism

Let me end this section by considering a radical alternative strategy for resisting epistemic unprincipledness: adopting a radically permissive view of how epistemically normative concepts apply.

Above I argued that the epistemic unprincipledness view was better situated than epistemic consequentialism or deontology to explain why our priors might be justified while certain intuitively unjustified - but equally objectively accuracy promoting - priors are not.

What I will call the permissivist strategy avoids the above explanatory challenge by rejecting the premise that there's any legitimate difference in epistemic standing between our priors and these alternative priors to be explained. It aims to defend principledness about justification (and similar epistemically normative concepts), by taking them to apply much more broadly than is normally

²¹The case might be different if they showed ambivalence, self correcting dispositions, or patterns of belief and desire suggestive of wishful thinking or some kind of motivated reasoning.

expected.

For example, we saw above how our concept of justified belief seems to draw an unprincipled distinction between logically valid deductions which can vs. can't be made immediately. A permissivist might resist this conclusion by saying that each thinker is (defeasibly) justified in making whichever (kinds of) logically valid inferences *they* find immediately compelling. Accordingly, they will say that both we and the Martians imagined above can gain justified true belief by making all the logically valid inferences we're disposed to find immediately compelling²².

Similarly a permissivist might respond to the above puzzle about priors, by saying something like the following.

It's epistemically permissible for each thinker *t* to assign priors in any way they find psychologically compelling (even if such a prior would strike us as deeply irrational) — provided these priors obey the probability axioms and are sufficiently objectively truth conducive (e.g., at least as objectively truth conducive as we take ours to be) when deployed in it's possible world.

As regards both the puzzles used to motivate epistemic unprincipledness above, permissivists can agree with advocates of epistemic unprincipledness that there's nothing intrinsically special about the set of logically valid inferences which humans are willing to make immediately/true propositions humans are inclined to assign high probability to a priori. But they avoid the conclusion that our epistemically normative concepts draw unprincipled distinctions, by taking these concepts to apply more broadly than one might have initially expected.

²²The above permissive idea can be spelled out in different ways. For example, an extreme version might say that all specific token inferences that are logically valid have defeasible warrant (regardless of the reasoner's overall inference dispositions). A more moderate permissivist might say that each thinker (only) has defeasible warrant to make a particular token inference that's logically valid when *they* are robustly disposed to find all inferences of this kind (e.g., all substitution instances of this inference) immediately compelling.

However, I will now argue that taking such a permissive approach to priors generates implausibly radical consequences about what individuals are justified in believing.

First this epistemic permissivism seems to conflict with our ordinary practices of epistemically assessing and criticizing people. As we saw in §4.1 above, one can cook up priors that are intuitively unjustified but more objectively accuracy promoting than our own by raising the prior probability assigned to some surprising truths about the world which humans don't assign high probability to a priori (e.g., the facts summarized in the periodic table of elements).

But if the above strategy for boosting the accuracy of our actual priors works, it seems likely that it can also be used to modify many alternative priors in a way that compensates for their initial inaccuracy. Thus we can likely use it to create variant prior which assign high probability to some odd false and intuitively unjustified claims (e.g. 'the moon is made of blue cheese') while being equal to our favored priors in overall accuracy promotingness.

And more generally it seems likely that for many ordinary collections of evidence E and false claims F (not intuitively justified by E) one can cook up probability axiom satisfying priors which assign high probability to F conditional on E but are (overall) just as accuracy promoting as our own, via the following to tricks:

- Compensate for accuracy loss resulting from assign high probability to F with high probabilities assigned to 'surprising but truths' $S_1, S_2, \dots, S_n$ (e.g., the periodic table or Newtonian mechanics) that are very accuracy promoting accept, and humans are disinclined to accept a priori – providing some such surprising informative generalizations can be found which are compatible with F).
- Compensate for accuracy loss from assigning the target claim F high prob-

ability a prior/ conditional on the target evidence E , by assigning surprising general truths about the world $S_1, S_2, \dots, S_n$ high probability (and F low probability if needed) conditional on all actual-world likely extended evidence streams E' extending E . [In this way initial inaccuracy associated with assigning high probability to F given evidence E can be offset by greater accuracy once one has got a chance to update to get to one the extending streams of evidence E' .]

So the permissivist faces significant pressure to allow that there will likely be many false and intuitively unjustified claims F which could be accepted by someone who shared your total evidence and obeyed all the probability axioms and updating rules – provided they had priors which compensated for this by assigning higher probability to some actual physically necessary laws.

But the latter claim is quite unintuitive and conflicts with ordinary practices of describing people as justified or unjustified. For objective accuracy promotion is a holistic matter. And we don't intuitively allow unusually high objective accuracy condusiveness in one domain to compensate for low accuracy condusiveness in an unrelated one. For example, imagine someone whose priors seem to be unusually objectively accuracy promoting about physics and unusually objectively accuracy harming about psychology. We would not say this person's psychological beliefs formed using her wonky priors regarding human behavior can count as justified just because her good sensibilities about theoretical physics make her overall way of assigning priors just as objectively truth conducive as our own.

Furthermore, we are generally willing to criticize someone's individual beliefs as irrational without investigating, or having an opinion about, whether their overall way of assigning priors might be more accuracy promoting than our own (because of compensatory good hunches about unrelated topics).

Still perhaps a permissivist can bite the bullet and accept that everyone is justified in using whatever priors they find *prima facie* appealing -or at least that everyone with sufficiently sufficiently objectively accuracy promoting (in the actual world? in S' world?) priors is justified in using those priors.

But when we turn from claims about what an individual is or would be justified in believing/assigning high probability to more impersonal epistemic evaluations, such bullet biting becomes more costly. For permissivists about priors also face a question about how to make sense of impersonal evaluations of things as *adequate evidence* or *good scientific arguments*. But the points above suggest the permissivist must accept that, for many commonplace bodies of experience E, many false and intuitively unjustified by E propositions F could be rationally accepted/ assigned high probability *by someone* with sufficiently odd priors. So it's hard to see how they could rescue the ordinary and useful distinctions we draw between adequate vs. inadequate scientific evidence/arguments without allowing that at least these notions sort things in a way that reflects unprincipled contingencies of human psychology.

So, to sum up, I don't think permissivism provides a plausible escape from epistemic unprincipledness. For when we think through the question about acceptable priors above, the costs of permissivism about priors are too high. It's one thing to allow that psychologically alien beings could be justified in making additional logically valid inferences. But accepting that, for many evidence sets E, a wide range of intuitively unjustified and unrelated false claims F can be rationally accepted conditional on E (and good scientific arguments can be made for accepting F on the basis of E), is a far less palatable bullet to bite. Thus we seem to need to accept some unprincipledness in our concept of justified beliefs (or at least our concept of adequate scientific evidence/arguments), stemming from unprincipledness in our concept of acceptable priors ²³.

²³And (perhaps) once we've taken this step it's more elegant to give an analogously un-

5 Worries about Circularity and First Person Deliberation

Let me end this paper by briefly considering two possible concerns about the epistemic unprincipledness thesis I've advocated.

First, there might be a vague worry about circularity or explanatory/grounding regress. Epistemic unprincipledness maintains that some aspects of our epistemically normative concepts, like "justification," reflect contingent facts about which accuracy-promoting methods humans find psychologically compelling. So an unprincipledness theorist might say that the meaning of our term "justification" is partly determined by human reasoning dispositions (e.g. reflecting facts about which logically valid inferences we are inclined to make immediately). Yet (in at least some contexts) we can explain the fact that someone believes a certain claim, by noting that this person accepts certain premises from which this claim can be quickly and justifiably inferred.

However I claim there is no problematic circularity here, as can be seen by comparison with uncontroversial facts about our use of color words. Clearly the meaning of a term like "red" is shaped by which shades people are inclined to call "red" under stable conditions. Yet, we can still explain why someone calls a particular object red by noting that the object possesses properties corresponding to the term's meaning. And it is still possible to be wrong about particular judgments about what's red²⁴

principledness based answer to the puzzle about which basic logical inferences can be made immediately above.

It is, of course, in principle possible (and perfectly consistent) to combine permissivism about which logically valid inferences are OK to make with unprincipledness about priors. However, arguably, doing this gives the worst of both worlds, requiring us to bite the bullet of saying epistemically normative concepts are deeply unprincipled and also the bullet of saying the ZFC to FLT proof confers knowledge.

²⁴One could also try to dispel the above worry more directly by noting there's no conflict between the following claims

1. **Metasemantic Fact:** Our basic inference dispositions determine that our term "justification" expresses a concept of justification fitting our dispositions to make basic

A separate (but, perhaps, related) concern might be raised by David Enoch’s interesting arguments about moral phenomenology. In [11] David Enoch highlights a phenomenological difference between the experience of serious first person practical deliberation (as when deciding between graduate study of law and philosophy) and that of ‘arbitrarily picking’ (as when choosing one of two nearly identical cereal boxes from the shelf). And he uses this contrast to argue that we can’t avoid a certain degree of moral realism, in the following sense. He argues (in a somewhat Korsgaardian vein[16]) that we psychologically can’t avoid serious first person deliberation, and when we deliberate we can’t help committing ourselves to the existence of “irreducibly normative facts” (about what we have most reason to do).

Specifically, Enoch emphasizes the (felt) non-arbitrariness of moral deliberation as follows. He writes that “The phenomenology of arbitrary picking is very different from that of deliberation, of trying to make the right decision” and “Deliberation -unlike mere picking- is an attempt to eliminate arbitrariness by discovering normative reasons, and it is impossible in a believed absence of such

inferences etc (rather than other variant concepts of justification* corresponding to variant, equally accuracy promoting practices).

2. The above metasemantic fact (together with our tendency to use “justified” in a way that aligns with the inferences we actually make, as per the theory of motivationally grasped concepts in [3]) helps answer access worries and explain our reliable success in using those reliable inference methods which our tokens of “justified” refers to.
3. **Human Tendency to Justified Reasoning:** People tend to make the kind of objectively accuracy conducive inferences which are justified.
 - And any apparent access worries raised by this striking match between our actual dispositions and justification facts can be answered by (telling a satisfying story about our tendency to use *some* sufficiently accuracy promoting inference methods and then) citing the metasemantic fact above – while taking the philosophical stance[13] of refraining from all claims that our concept of justification is somehow more joint carving or otherwise intrinsically special than the variant concept justification* corresponding to the practice of Martians who favor different but equally accuracy promoting inferences.
4. In explanatory contexts where the above tendency to justified reasoning (an informative general regularity) is accepted background knowledge, one can often explain why a human believes a proposition by noting this proposition quickly follows from others they believe by justified inferences.

reasons to be discovered”. And he emphasizes the way that serious practical deliberation involves a sense of trying to get something right which you could in principle get wrong. “When deliberating, you also try to get things right, to decide as - independently of how you end up deciding - it makes most sense for you to decide”.

Accordingly one might worry that my proposal that our concept of justification reflects contingent and arbitrary features of human psychology (which actual truths/valid inferences we happen to be attracted to) faces a similar challenge to the one Enoch presses against the deniers of all-things-considered ought facts. Specifically, one might note that the experience of serious first person epistemic deliberation (Do I really know that P? Ought I assign more probability to M than N?) involves a sense of effort and possible failure -in contrast to mere plumping- analogous to what Enoch notes in the case of practical deliberation. And one might worry that taking our epistemically normative concepts to draw unprincipled (and in a sense arbitrary) distinctions in the ways I’ve suggested is somehow incompatible with the attitude of serious first person epistemic deliberation.

However, I think that (whether or not Enoch’s argument works in the moral case) this worry is not serious. To address it, I want to point out three things.

First note that typically, epistemic deliberators are concerned with certain first-order truths—e.g., whether a hypothesis is true or whether an inference is valid. These obviously provide one serious subject matter which it is possible to be wrong about, and one sense in which one’s efforts to figure out what to believe could fail. For example, someone who is wondering what they *know* about whether a certain strategy will destroy a looming asteroid without significant casualties will typically also be wondering about about relevant non-epistemic facts about the subject matter at hand (e.g. will pushing the button destroy

the asteroid without significant casualties?),

Second note that even cases of serious deliberation where the *truth* of relevant non-epistemically normative belief is not in question, (e.g., first person deliberations asking "Was I justified in believing the factory we bombed was unoccupied?"), the possibility of failure and felt need to engage in serious consideration and not mere plumping remains. For (as we can see clearly by returning to our analogy with color terms) accepting the epistemic unprincipledness thesis leaves plenty of room for people to be wrong in both their reasoning and their claims about justification—e.g., in cases where their immediate judgments disagree with what they'd accept under reflective equilibrium.

Specifically, my unprincipledness claim about epistemically normative concepts is closely closely analogous to the following hypothesis about our color concepts. Psychologistic hypothesis: the boundaries of some color terms don't reflect chemical natural kinds but rather group certain shades together because of how they relate to distinct quirks of human vision. Anyone who has tried color sorting activities online (e.g. classifying sample patches as green or blue) will remember the feeling of serious deliberation and a possibility of error. And surely accepting the above arbitrariness/psychologism thesis about color concepts is compatible with this experience of serious deliberation (though it might increase your willingness to regard certain classification of borderline cases as having indeterminate truthvalues)²⁵.

Fourth, accepting the unprincipledness thesis above is entirely compatible with allowing facts about justification etc to be independent of speakers' dispo-

²⁵And I see no conflict between this phenomenological seriousness/recognition of the possibility of error allowing that our dispositions to classify colors under ideal access to non-color facts and opportunities for more observation would have to be right in the limit under ideal reflective equilibrium (as per the hypothesis above). Allowing that some color shades will be borderline cases (as perhaps the unprincipledness thesis above somewhat supports) and use of color terms determines meaning doesn't rule out feelings of uncertainty and trepidation, or instances of genuine error. Accepting an unprincipledness hypothesis about color (as above) doesn't in any way conflict with allowing that facts about color classification can be difficult to determine, and it is possible for hasty judgments about such questions to be wrong.

sitions in the following straightforward sense. Although the epistemic unprincipledness view I’ve sketched above suggests that facts about *what concept your word “justified” expresses* are dependent on your dispositions to use language, it does not imply any claim that facts about *what things are justified* are dependent on how you are disposed to infer. (If I’d been disposed to make the ZFC to FLT inference, that inference would still have been unjustified, just as if “tail” meant “leg” horses would still have only one tail).

Thus advocates of epistemic unprincipledness can entirely accept appearances in the context of first person deliberation that facts about whether it would be justified for me to have a certain belief are independent of whether I actually form that belief -and even counterfactually independent of whether I’d be disposed to retain that belief in the limit of reflective equilibrium. In this way they can easily accommodate (the relevant version of) Enoch’s remark that “When deliberating, you also try to get things right, to decide as - independently of how you end up deciding - it makes most sense for you to decide”.

6 Conclusion

In this paper I’ve tried to clarify and defend an emerging epistemic unprincipledness thesis, on which certain aspects of our concepts of justified belief reflect (rigidify and project) mere contingencies of human psychology.

In the first half of the paper I’ve related epistemic unprincipledness to some contemporary foundational debates about epistemic normativity. I’ve argued that it provides an attractive alternative to epistemic consequentialist and deontologist views (while still potentially letting us mirror many valuable ideas from these research programs).

In the second half of the paper, I’ve argued that we should favor epistemic unprincipledness over a rival permissive approach to epistemic normativity, which

seems to address puzzles about why certain logically valid inferences can be made immediately equally well. Specifically I've argued that we need epistemic unprincipledness to account for seeming arbitrariness in which truths can permissibly be assigned high probability a priori ²⁶.

Let me end by contrasting my arguments in this paper with somewhat related prior work in epistemology.

First, note that the arguments above suggest there's something unprincipled about our concepts of *justified assignment of high probability*, not just about our concepts of justified belief and knowledge. In this way I've defended a much more radical form of unprincipledness/context dependence than well known appraiser-relative contextualist theories of knowledge (e.g., views on which the correctness of my attributing someone knowledge that p depends on facts about my current context, like whether p is relevant to any high stakes decision making or what alternative skeptical hypotheses are currently salient)[19]. For these contextualist accounts of *knowledge* (and justified belief) are compatible with taking facts about *justified high confidence* to be entirely principled and context-independent.

²⁶An anonymous referee suggested the following interesting question. Might it be that epistemic unprincipledness truly describes our *actual* current concepts of justification, knowledge etc, but should start using some more principled concepts of justification*, so as to eliminate the relevant element of arbitrariness.

I won't attempt to completely answer this question here, but I strongly suspect the answer is no. For consider what less arbitrary alternatives to our current concept of justification might be like. We could eliminate the kinds of arbitrariness discussed in section 3 by going agent relativist and saying that a person's credences are justified iff Bayesian updating from their priors (i.e., whatever priors best capture their sense of a priori theory plausibility) yields these credences. But requiring speakers to assess such deep psychological questions (attribute certain priors) before classifying others' beliefs as justified or not would be very inconvenient. The point of classifying someone as (say) tending to be irrational on a given topic might be just to flag to myself that I shouldn't expect to model their reasoning in certain ways, that I can't trust their conclusions to be as reliable (and checkable by me) as reliable on this topic as I otherwise would.

Alternatively, we could eliminate arbitrariness by letting a thousand flowers bloom in a very strong sense: saying that justified beliefs given some evidence only have to be justifiable relative to some conceivable priors that are moderately truth conducive. But (as we saw in §4.3), this would mean classifying almost all contingent claims as justified. So switching to a more principled permissive concept of justification would leave us with something less practically useful.

Second, the epistemic unprincipledness thesis I've defended claims that certain aspects of our distinction between justified and unjustified belief reflect/rigidly project mere contingencies of actual human psychology. It thereby requires something more than theories of agent (as opposed to appraiser) relativity, like [18]'s suggestion that whether a subject S has epistemic obligations to engage (politically or cognitively) with arguments against a proposition they believe can depend on the practical feasibility and costs of S so engaging. For the latter proposal only suggests that practical features of an agent's situation can make a difference to whether that agent counts as having a justified belief. It does not suggest that our concept of justification draws any unprincipled distinctions.

Third, in arguing for epistemic unprincipledness, I've claimed that hypothetical beings who found different accuracy-promoting priors and inference methods psychologically compelling could deploy variants on our concepts of justification that are equally good (in the sense of, e.g., being equally metaphysically joint carving and equally helpful in promoting true belief and accuracy for relevant communities). This claim is quite independent from empirical arguments that actual human cultures have different but equally good knowledge or justification-like concepts (e.g., variant concepts of knowledge that handle Gettier cases differently).

References

- [1] Selim Berker. Epistemic Teleology and the Separateness of Propositions. *Philosophical Review*, 122(3):337–393, 2013.
- [2] Sharon Berry. External World Skepticism, Confidence and Psychologism about the Problem of Priors. *The Southern Journal of Philosophy*, 57(3), 2019.

- [3] Sharon Berry. Metaethical Deflationism, Access Worries and Motivationally Grasped Oughts. *Ethical Theory and Moral Practice*, 27(3):301–318, July 2024.
- [4] Paul Boghossian. Blind Reasoning. *Aristotelian Society Supplementary Volume*, 77(1):225–248, 2003.
- [5] Glenn W. Brier. Verification Of Forecasts Expressed In Terms Of Probability. *Monthly Weather Review*, 78(1):1–3, January 1950.
- [6] Rudolf Carnap. *Logical Foundations of Probability*. University of Chicago Press, Chicago, 1950.
- [7] Matthew Chrisman. From Epistemic Contextualism to Epistemic Expressivism. *Philosophical Studies*, 135(2):225–254, September 2007.
- [8] Matthew Chrisman. Epistemic Expressivism. *Philosophy Compass*, 7(2):118–126, 2012.
- [9] Justin Clarke-Doane. *Morality and Mathematics*. Oxford University Press, 2020.
- [10] Matti Eklund. *Choosing Normative Concepts*. Oxford University Press, 2017.
- [11] David Enoch. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press UK, 2011.
- [12] Bas C. Fraassen. *Laws and Symmetry*. Oxford University Press, 1989.
- [13] Bas C. van Fraassen. *The Empirical Stance*. Yale University Press, 2002.
- [14] Alvin Goldman and Bob Beddor. Reliabilist Epistemology. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2021 edition, 2021.

- [15] Jonathan Jenkins Ichikawa and Matthias Steup. The Analysis of Knowledge. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2018 edition, 2018.
- [16] Christine M. Korsgaard. *The Sources of Normativity*. Cambridge University Press, 1996.
- [17] David Lewis. Immodest Inductive Methods. *Philosophy of Science*, 38(1):54–63, 1971.
- [18] Robin McKenna. *Non-Ideal Epistemology*. Oxford University Press, Oxford, New York, July 2023.
- [19] Patrick Rysiew. Epistemic Contextualism. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2023 edition, 2023.
- [20] Joshua Schechter. Small Steps and Great Leaps in Thought: The Epistemology of Basic Deductive Rules. In *Reasoning: New Essays on Theoretical and Practical Thinking*. Oxford University Press, Oxford, 2019.
- [21] Michael Stevens. Probability Out Of Determinism. In Claus Beisbart and Stephan Hartmann, editors, *Probabilities in Physics*. Oxford University Press, September 2011.
- [22] Kurt L. Sylvan. An Epistemic Nonconsequentialism. *The Philosophical Review*, 129(1):1–51, January 2020.