

EPISTEMIC UNPRINCIPLEDNESS

SHARON BERRY

Indiana University, Bloomington

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

Not all logically valid inferences seem equally permissible to make without further justification. For example, many people share the intuition that modus ponens inferences require no further justification, while inferring Fermat's Last Theorem (FLT) from the standard ZFC axioms of set theory would – despite both inferences being logically valid. But why?

This kind of puzzle has inspired a view I'll call epistemic unprincipledness, on which the boundaries of our epistemically normative concepts (like justification and adequate argument) partly reflect contingent features of humans rather than drawing entirely principled and intrinsically motivated distinctions. In this paper, I will relate epistemic unprincipledness to recent foundational debates, and argue (partly by expanding on prior work by Dogrammici Dogramaci (2012; 2015) and Berry Berry (2013; 2019)) that it provides an attractive alternative to epistemic consequentialism and deontology¹.

Contact: Sharon Berry <seberry@invariant.org>

1. In a recent paper Dandelet (2024), Dandelet challenges a principledness thesis she calls *veritism*—the view that “what makes epistemically rational belief epistemically rational, and what makes it an achievement of (genuine or perhaps merely domain-relative) significance, is its fitting response to the value of truth.” She presses veritists to explain why it is unjustified to believe P when one assigns P only marginally higher probability than $\neg P$. I am sympathetic to her critique, but here I pursue a more radical form of unprincipledness: the idea that even our judgments about which *probability assignments* are justified lack a principled foundation, not merely our move from justified probability assignments to justified full belief—a step many already treat as pragmatically contextual.

In §2, I will introduce and clarify the epistemic unprincipledness thesis, situating it in contrast with epistemic consequentialism and deontology. In §3 I will motivate unprincipledness by appeal to two familiar but puzzling phenomena: our intuitive distinctions between acceptable and unacceptable basic logical inferences (§3.1), and our intuitive distinctions between propositions which can and can't be assigned high probability a priori (§3.2). In §4, I will argue that neither accuracy-based nor respect-based foundational theories can adequately explain these distinctions—and that permissivist responses, while technically available, collapse important epistemic categories. Finally, in §5, I will consider and respond to a worry that accepting the epistemic unprincipledness thesis somehow conflicts with the experience of first-person epistemic deliberation.

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 justification², 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 Boghossian (2003) 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” Boghossian (2003). 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 Berker (2013); Sylvan (2020). For example, in recent debates between epistemic consequentialists and deontologists (as characterized by Sylvan Sylvan (2020)), 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 features³.

Specifically, both parties agree that the fundamental value relevant to epistemology is accuracy, and hope to answer all foundational questions about epistem-

2. c.f. the quest for informative necessary and sufficient conditions for knowledge pursued in the Gettier literature Ichikawa and Steup (2018)

3. I take David Enoch to be getting at a similar foundational explanatory project in Enoch (2011) 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”

ically normative concepts like justification by appeal to it, as per the following doctrine.

Valuing Thesis: The fundamental normative explanation of why justified beliefs are justified is that they manifest certain ways of valuing fundamental epistemic value. Sylvan (2020)⁴

But disagreement arises when we ask how to cash out the talk 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 maximization⁵ ⁶. 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 assumption that foundational questions about epistemically normative concepts can always be answered by invoking principled intrinsic features of the things being classified, like accuracy promotion or respect. Instead, it posits that some aspects of our epistemically normative concepts draw distinctions that merely reflect contingent features of human psychology (or culture, history, environment etc)⁷.

4. C.f. Dandelet (2024) Dandelet on veritism as a shared commitment of epistemic consequentialists, deontologists and virtue epistemologists.

5. Doing 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' Berker (2013), 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." Sylvan (2020).

6. Although 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)." Sylvan (2020).

7. In 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 how

Epistemic Unprincipledness: Some important aspects of some central epistemically normative concepts rigidify/project contingent features of actual humans (e.g., what accuracy promoting reasoning methods we happen to find compelling) rather than drawing an intrinsically principled distinction.

So, for example, an advocate of epistemic unprincipledness might say that at each metaphysically possible world, a subject has defeasible warrant for making exactly those logically valid deductions⁸ which are psychologically compelling to humans *in the actual world*.

On the resulting view, Martians who found additional logically valid inferences immediately compelling would not count as justified in making these inferences. However, this fact does not reflect anything metaphysically special about us, or suggest that human psychology matches some independently interesting distinction between logically valid inferences in a way that might 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 additional logically valid inferences Martians find compelling count as justified* but not justified⁹. After all, it’s unsurprising that humans would find it useful to have a justification-like-concept which (among other things) tracks ability to support beliefs via the kind of logically valid inferences *we humans* find immediately compelling.

On this view, contrary to epistemic consequentialism and deontology, the ‘problem’ with a Martian who makes additional logically valid inferences immediately can’t be cashed out in terms of her relationship to things we ultimately 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. We can imagine that she finds her extra logically valid deductions 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

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).

8. Talk of ‘kinds of reasoning’ obviously brings up the generality problem Goldman and Beddor (2021). I won’t try to say anything interesting about that problem here.

9. See Berry (2018; 2024); Clarke-Doane (2020); Eklund (2017) for some recent discussions of pluralism along these lines in other areas like ethics and mathematics.

to *us* in ways that human mathematical reasoning often is. Her ‘proofs’ fail to justify because they don’t proceed by the kind of steps humans find compelling (or sufficiently help us produce arguments which which do proceed by such steps)¹⁰.

2.2. Clarifications

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

First, what I have called epistemic unprincipledness does not *require* accepting any kind of expressivism or truthvalue anti-realism about epistemic concepts, though epistemic expressivists of various types¹¹ are likely to accept unprincipledness. 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 handfuls of nuts on the table’ fail to be true or false in a completely ordinary and straightforward sense¹². And my epistemic unprincipledness thesis would be true if some core epistemically normative concepts worked analogously.

Second, accepting epistemic unprincipledness (as I will understand it) naturally motivates, but does not require, giving a partly metasemantic answer to access worries about knowledge of epistemic normativity. For suppose you think some aspect of the boundaries of our normative concepts rigidifies contingent features of actual humans as above. Then you will likely want to answer access worries about our knowledge of this aspect of epistemic normativity partly metasemantically, by saying: if we’d have variant practices in this regard, we’d have expressed a slightly different concept by our words for epistemically normative concepts like justification, but still spoken the truth.

Third, the epistemic unprincipledness thesis (as I will understand it) only requires that *some* aspects of *some* of our central epistemically normative concepts draw unprincipled distinctions. So, an advocate of unprincipledness is free to answer foundational questions about these concepts by combining appeal to principled facts (about accuracy promotion or respect etc.) with appeal to contingent facts about actual human psychology/inference dispositions¹³. For example, the

10. c.f. Dogramaci’s remarks on the usefulness of people who share basic inference methods for epistemic ‘parallel processing’ Dogramaci (2015) and the discussion of this view in §3.

11. c.f. Chrisman’s epistemic expressivism Chrisman (2007; 2012) and Hartry Field’s Field (2009) expressivism-like view.

12. What about forms of appraiser (as opposed to agent) relativism, which technically could allow concepts like justification and knowledge to be entirely principled, by taking them to have a hidden extra parameter? I take such proposals to be extremely similar to (some versions of) epistemic unprincipledness in spirit, and won’t argue against them here.

13. The advocate of epistemic unprincipledness can also allow that some epistemically normative notions (e.g., a notion of validity definable in terms of logical possibility or model

advocate of unprincipledness could say 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. This can let them mirror many appealing consequentialist/deontologist explanations for why certain things *aren't* justified/rational.

Given this clarification, some readers might find it 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 epistemic normativity very much. I happen to disagree. For example, taking the concept of justification to be unprincipled (in the way relevant to the unprincipledness thesis) suggests that we can sometimes classify someone's belief as being unjustified *without* taking this belief to be bad in any intrinsic sense the believer in question should intuitively 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 will now review two styles of motivation for epistemic unprincipledness.

3.1. Logical Inferences

One way to motivate epistemic unprincipledness is via the puzzle about basic logical deduction which I have already mentioned. Why is it epistemically permissible to make certain logically valid inferences immediately, but not others? Initial appearances and a history of failure to find a principled response to the above problem¹⁴ naturally motivates something like the epistemic unprincipledness thesis.

Most prior work motivating epistemic unprincipledness has invoked such considerations about basic logical inference. For example, Sharon Berry (2013) uses the above puzzle to argue that we should adopt *either* epistemic unprincipledness or a permissive view I'll criticize in 4.3. And Dogramaci Dogramaci (2012; 2015) uses a similar puzzle about logical deduction to motivate a doctrine of 'epistemic communism' which (in effect) combines epistemic unprincipledness with a further claim that the function of epistemically normative vocabulary is to promote

theory) do draw entirely principled distinctions.

14. Schechter's paper Schechter (2019) 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.

agreement in reasoning methods, which allows for a kind of useful parallel processing within a community.

Specifically, Dogramaci writes that “With uncoordinated outsiders [who use different logically valid basic inference methods]... it is either dangerous or useless to accept their testimony, because we either cannot confirm their reports, or we have made our own discovery already. But, in a community where we work to coordinate epistemic rules, spreading beliefs through testimony becomes both safe and valuable.” And he characterizes epistemic communism as the view that the function of epistemically normative vocabulary is to promote such useful “parallel processing” by getting people to use the same basic reasoning methods.¹⁵

I take epistemic communism to include commitment to the epistemic unprincipledness thesis I advocate in this paper. And I think Dogramaci is right to identify a kind of epistemic outsourcing—using judgments from people we take to be good reasoners to speed up our own truth tracking—as central to the usefulness of many epistemically normative concepts. However I will not attempt to defend epistemic communism because I question Dogramaci’s suggestion that the function of epistemic evaluation is to promote coordination by *producing agreement* in basic inference methods.

For one thing, I doubt that epistemic praise and blame have much power to change what basic logical inferences people accept. Perhaps they instead promote coordination by helping us *classify* others, i.e., helping us recognize people whose reasoning is already sufficiently similar for epistemic outsourcing to be fruitful¹⁶.

I also think Dogramaci may overstate the extent of actual agreement in basic inference methods. Even in communities which cooperate in broadly the way he imagines, I’m not sure how common it is for people to accept exactly the same basic inferences¹⁷ as opposed to merely compatible ones—in the sense that,

15. He writes, “epistemic communism presupposes that epistemic evaluations influence others’ behavior, including aspects of how they reason, but it doesn’t implausibly presuppose any metaphysically necessary connection here. The connection may be an innate, evolved feature of human psychology. Darwin (1871/74, chapter 5): ‘[A]nother and much more powerful stimulus to the development of the social virtues [than self-interested calculation], is afforded by the praise and the blame of our fellow-men. To the instinct of sympathy, as we have already seen, it is primarily due, that we habitually bestow both praise and blame on others, whilst we love the former and dread the latter when applied to ourselves; and this instinct no doubt was originally acquired, like all the other social instincts, through natural selection.’” Dogramaci (2015)

16. This classificatory function can be valuable even when no possibility of influence exists. For example, an astronaut receiving delayed transmissions from a planet many light years away would still benefit from being able to distinguish between journals or blogs that present arguments in steps she finds compelling—and thus could learn from—and those that do not.

17. Dogramaci allows that trained mathematicians often give step-skipping proofs, but seems to treat them as still implicitly relying on the same basic inference rules, with larger

for any two people x and y in the community, x can generally reconstruct y 's proofs into arguments composed of basic inferences that x finds compelling. The latter mutual interpretability seems sufficient for the kind of useful epistemic outsourcing referenced above¹⁸.

3.2. Priors

A second puzzle which can be used to motivate epistemic unprincipledness concerns the question: which metaphysically contingent truths can/should be assigned high a priori probability and why¹⁹?

It is often useful to model scientific reasoning in Baysean terms²⁰, supposing a thinker starts with a certain assignment of probabilities (satisfying the 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 which conclusions a person's total body of evidence permits them to assign high probability to – and hence what they can be justified in believing.

However, considering this notion of acceptable priors raises a puzzle analogous to the one about basic logical inferences above.

For, *prima facie*, it is epistemically permissible to assign high prior probability to certain metaphysically contingent physically necessary truths²¹ (e.g., claims

steps justified via decomposition into those basics. I don't think that assumption is obviously correct. It may be that mathematical training leads people to find new inferences immediately compelling in a way that is no longer traceable to their earlier acceptance of the more basic forms.

18. It would be natural (but not inevitable) for advocates of epistemic unprincipledness to allow slight variations in the extensions of the predicates expressed by different speakers' terms like 'justification'. One might fear that allowing this kind of variation would make disagreements over epistemically normative concepts too much like disagreements over taste.

However, I don't think it's a problem to say that people can mean slightly different things by 'justified' and accordingly disagreements about certain aspects of epistemic normativity (corresponding to disagreement about the unprincipled/psychologistic aspects of our epistemically normative concepts) can function more like disagreements about taste— where there is only a rational defeasible presumption that other speakers' terms are close that interlocutors praise terms will apply or both fail to apply so as to vindicate some particular claims under dispute. See Berry (2024) for the analog to this proposal as applied to the case of epistemic normativity.

19. Here I develop an argument sketched in Berry (2019).

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

21. One could perhaps make an analogous argument by replacing questions about which physically necessary truths can be assigned high probability a priori with questions about which *non*-physically necessary but (we think) actual world largely reliable claims like "my

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? And answering the latter question in a way that seems entirely principled looks equally daunting.

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 the problem of choosing learning parameters in a principled way²² and Bertrand’s paradox²³.

And in Berry (2019) 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 lets us resolve 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 and/or historically useful and remain reasonably accuracy promoting in the actual world.

As an example of how an advocate of epistemic unprincipledness might approach the problem of priors, consider the ‘psychologism’ about justification (which I take to qualify as version of epistemic unprincipledness) advocated in

perceptual and cognitive faculties are broadly reliable” can be assigned high probability a priori.

22. When Carnap modified his theory of the logical foundations of probability Carnap (1950) 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.

23. Bertrand’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 Fraassen (1989).

Berry (2019), as follows.

[H]uman beings are inclined to substantially agree in how they assign priors. The particular priors 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 reasonably useful and easy to physically realize in the human brain in the context of evolution...

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 ...the kind of prior judgments about theoretical elegance which normal human beings find compelling.

4. Alternative Approaches to the Puzzle About Priors

I will now try to strengthen the above case for epistemic unprincipledness (as an appealing alternative to epistemic consequentialism and deontology) by doing two things. First I'll argue there's especially little hope of accounting for common/intuitive verdicts about acceptable priors entirely by appeal to epistemic consequentialists and deontologists' favored notions of *accuracy conduciveness* and *respect for accuracy*. Second I'll argue against a radical permissive approach to acceptable priors (allowed by Berry Berry (2019) as a legitimate alternative) which would (technically) let epistemic consequentialism and deontologists sidestep these criticisms.

4.1. Limits of Accuracy-Promotion Accounts

Let's start with epistemic consequentialists' appeal to accuracy promotion. I will argue by cases, considering different options for what we mean by accuracy promotion.

First, I claim, we can't answer foundational questions about acceptable priors (why is it OK to assign some claims high probability a priori but not others?) 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 promotion²⁴ we can imagine priors that are better at

24. By talking about the objective accuracy promotingness of priors here, I mean something like the expected accuracy (calculated using objective physical probability not subjective

accuracy conduciveness than our own, yet intuitively unjustified. To explain 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

probability) of credences got 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. 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 Stevens (2011).

There are clearly many choice points faced when filling in this notion. For example, we face choices about how to weight accuracy about different propositions for calculating a Brier score, how to cash out the idea of our robot being equally likely to appear anywhere (what metric on space and time do we assume?), and how to weight immediate 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 Brier (1950) 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 Breier 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.

necessary truths a priori. But there's not any obvious way in which their priors and reasoning methods need be less objectively/actual-world accuracy promoting than ours. For example, they will (even by our lights) tend to do *better* than us at quickly forming true beliefs or 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 alone²⁵. So it appears epistemic consequentialists can't answer foundational questions about why certain priors are acceptable (assuming they accept ordinary verdicts on this), by invoking actual world objective accuracy conduciveness.

Could they 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)? One might think that correct/missible priors²⁶ should be good in a way that's independent of what possible world happens to be actual. So, epistemic consequentialists might instead say correct priors are distinguished by having optimal *a priori expected accuracy* – or some other notion which incorporates the accuracy promotingness of (Bayesian 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 actual world accuracy promotion), 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* over all possible worlds. More specifically, we could note that the rationally correct priors have

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

25. Perhaps 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.

26. Because 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.

higher expected accuracy)²⁷.

- high expected accuracy as calculated relative to the rationally correct priors

However many quite different priors 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 correct priors are only special in having high expected accuracy as calculated relative to probability assignments *actual humans find a priori appealing*. But this seems to concede the game to advocates of epistemic unprincipledness.

4.2. Limits of Accuracy-Respect Accounts

Now let's turn to the epistemic deontologist. Can they do any better at answering foundational questions about which priors are correct, by appealing to considerations of *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.

27. Readers familiar with David Lewis comments about the immodesty of priors in Lewis (1971) may 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.

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. For our intuitions about epistemically acceptable priors don't just require satisfying structural principles like the probability axioms. They also take a stand on substantive matters of what priors to assign contingent hypotheses like 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. And 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 understand this to require something beyond internal coherence like objective accuracy promotion (or fit with actual human psychology/inference dispositions as per epistemic unprincipledness). And we've just seen above that appeals to objective accuracy promotion ultimately don't seem capable of drawing the distinctions we need to answer foundational questions about acceptable priors.

To put this point another way, it's hard to imagine how creatures instinctively inclined towards the variant (intuitively bad but equally or more objectively accuracy promoting) priors discussed in §4.1 could be accused of failing to respect the value of accuracy. For we saw above that priors can plausibly be cooked up which 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 (e.g., those summarized in the periodic table of elements) a priori.

And if we imagine meeting creatures who straightforwardly reasoned in accordance with these wonky priors, I don't think we'd intuitively accuse them of any kind of lack of *respect for the value* of accuracy or truth. We'd just regard them as having unjustified and dogmatic high confidence in specific weird scientific hypotheses which then (luckily for them) happened to be true²⁸. So it's hard to see how the epistemic deontologist could use considerations of respecting accuracy to fully explain facts about correct priors (assuming they accept ordinary verdicts about which priors are acceptable)²⁹.

28. The 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.

29. A respect theorist might reply that a direct deduction from ZFC to Fermat's Last Theorem (FLT) is rational only for agents who can perform it *competently*—that is, by manifesting a reliable deductive disposition. This suggestion forces a dilemma.

(i) **Global reading.** If the competence condition bars *any* rational agent from treating the ZFC to FLT as inference basic, imagine "*mathoids*": engineered reasoners who accept every inference humans find self-evident *plus* the ZFC to FLT step, while not being disposed to make any further invalid basic inferences Berry (2013). Such agents meet the reliability requirement, so the global reading cannot explain why the inference is impermissible for us.

(ii) **Anthropocentric reading.** Alternatively, the competence condition might be used to restrict justified reasoning based on what deductive abilities *humans* can acquire. Then either

4.3. *The Cost of Radical Permissivism*

This brings us to my final topic in this section. One could, in principle, defend epistemic consequentialism or deontology by rejecting the some of the intuitions about cases invoked above – taking a radically permissive approach instead.

This permissive approach aims to save the principledness of our concepts like ‘justified’ and ‘adequate scientific argument’, by taking them to apply much more broadly than is normally expected.

For example, such a permissivist might resist logical-deduction based arguments for epistemic unprincipledness 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 compelling³⁰. And they might address the puzzle about acceptable priors taking the following permissive view.

It’s epistemically permissible for each subject *s* 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 when deployed in *s*’s possible world (e.g., at least as objectively truth conducive as we take ours to be when deployed in the actual world).

So perhaps an epistemic consequentialist could defend the view that all foundational questions about epistemically normative concepts can be answered

-
1. we let different species—or even differently trained humans—have different basic inferences, reviving permissivism and the practical worries from §4.3 (e.g., we do not, and should not, wait to discover which reliable inferential habits people *could* train before judging their arguments); **or**
 2. we privilege the human inferential package, which is just epistemic unprincipledness in another guise.

Hence the competence manoeuvre does not rescue the respect approach from the original puzzle about ZFC to FLT inference. Or, in a more conciliatory vein, one might say that respect theorists can solve this puzzle if they can reformulate their respect thesis to allow room for an element of epistemic unprincipledness, as per the final version of the competence manoeuvre mentioned above.

30. The 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.

by appeal to accuracy promotion, in this way³¹. As noted above, Berry's papers make no effort to block this alternative.

However I will now defend unprincipledness by arguing against permissivism – or at least against the very general permissivism required to avoid epistemic unprincipledness.

First, I claim that adopting epistemic permissivism about justification conflicts with useful ordinary practices of assessing and criticizing people. We saw in §4.1, it's possible to construct priors that strike us as intuitively unjustified but are, in fact, more objectively accuracy-promoting than our own — by increasing the probability assigned to certain physically necessary truths to which we don't, a priori, assign high probability (e.g., the facts summarized in the periodic table of elements).

But if this strategy for boosting accuracy works, we can also use it to lightly modify intuitively bad and accuracy harming priors, in a way that *compensates for* their initial inaccuracy. In this way, we can create variant priors which assign high probability to some odd, false and intuitively unjustified claims (e.g. 'the moon is made of blue cheese') but are no less accuracy promoting than our priors overall³².

So the permissivist will have to allow that many such odd, false and intuitively unjustified claims could be rationally accepted by someone who shared your total evidence but had suitably suitable deviant priors. But allowing this is unintuitive and conflicts with ordinary practices of talking about justification.

Overall, the problem is that objective accuracy promotion is a holistic matter, while our practices of classifying beliefs as justified or not are not. We generally don't allow unusually high objective accuracy conduciveness in one domain to compensate for low accuracy conduciveness in an unrelated one.

For example, imagine someone with unusually objectively accuracy promoting priors about physics and unusually inaccuracy promoting priors about psychology. We'd not say this person's social beliefs formed using her bad priors regarding psychology can count as justified just because her unusually 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

31. They might face some questions about how actual world truth conducive acceptable priors had to be.

32. Similarly, we can cook up priors which are just as objectively accuracy-conducive as our own overall, but assign high probability to something like B: 'the moon is made of blue cheese' conditional on our actual course of experience E. We simply make up for the inaccuracy flowing from assigning B a high probability given E, by assigning it a low probability (and other truths a high probability) conditional on all the actual world likely courses of experience E' extending E.

overall way of assigning priors might be more accuracy promoting than our own (because of compensatory good hunches about unrelated topics).

Now perhaps some readers will be willing to bite the bullet by accepting unconventionally permissive or holistic ideas about which individual beliefs and credences are justified in the cases imagined above³³.

However such bullet biting becomes more costly when we turn our attention to impersonal epistemic evaluations. For what can the permissivist say about impersonal evaluations of things as *adequate evidence* or *good scientific arguments*? The points above suggest the permissivist must accept that, for many commonplace bodies of experience, many false and intuitively unjustified claims could (in principle) be rationally accepted/assigned high probability *by someone* who shares our evidence but has sufficiently odd priors. So going permissive about epistemic evaluations of these notions threatens to collapse the ordinary and useful distinctions we draw between adequate vs. inadequate scientific evidence/arguments³⁴.

For these reasons, I don't think permissivism provides a plausible escape from priors-based arguments for epistemic unprincipledness.

5. Unprincipledness And First Person Deliberation

Let me end by considering a worry that the epistemic unprincipledness thesis somehow conflicts with the phenomenology of first person epistemic deliberation, inspired by David Enoch's interesting metaethical arguments in Enoch (2011).

In Enoch (2011), David Enoch highlights a phenomenological difference between the experience of serious first person practical deliberation (as when de-

33. C.f. Subjective Bayesian views (on which all probabilistically coherent priors are acceptable) Lin (2024). Note that defenses of subjective Bayesianism via arguments that people with all accepted priors would *eventually*/in the limit of updating on evidence agree in their verdicts don't seem to solve the problems that permissivism raises about justification in the here and now. They don't let us vindicate our (seemingly useful) inclination to judge that people who currently judge the moon is made of blue cheese are unjustified, and take their verdicts with a grain of salt accordingly.

34. Could permissivists replace criticisms of scientific arguments from different priors as unjustified with claims about these arguments being *unreliable* (in the sense of not sufficiently actual world truth conducive)? I don't think claims about unreliability would be equally useful for the following reason. We can often usefully criticize arguments using methods we don't recognize as reliable, without positively knowing these methods are *unreliable*. For example, consider an argument which assumes some purported deep scientific law whose truth feels completely non-obvious to me (and which I cannot easily prove from principles I do find obvious). People reasoning in this way may be deploying reasoning methods which are perfectly reliable (if that claim happens to be true). But their reasoning is still useless to me (for the purposes of the kind of epistemic outsourcing/parallel processing discussed in §3), since I can't easily turn their/arguments into ones I find convincing. Thus it is useful for me to criticize/flag this argument as bad/less useful to me, while suspending judgment on whether it is reliable. (Compare this to how one can't replace talk of justification with talk of truth).

ciding 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 (in a somewhat Korsgaardian vein Korsgaard (1996)) that we can't avoid a certain degree of moral realism, in the following sense. 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 highlights the (felt) non-arbitrariness of moral deliberation by noting 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 reasons to be discovered". And he notes how serious practical deliberation involves a sense of possible failure, in the sense 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".

One might worry that advocates of epistemic unprincipledness face a phenomenology based worry, analogous to the one Enoch raises for metaethical deflationists.

For the experience of serious first person epistemic deliberation (Do I really know that P? Ought I assign more probability to M than N?) plausibly involves a sense of effort and possible failure that's very similar to what Enoch highlights in the case serious first person practical deliberation. And one might worry that taking accepting epistemic unprincipledness somehow conflicts with these aspects of the phenomenology of epistemic deliberation.

However, I think that (whether or not Enoch's argument works in the moral case) we can address this worry for epistemic unprincipledness by pointing out three things.

First, typically epistemic deliberators who ask 'but what should I believe?' are concerned with some underlying subject matter facts that don't concern epistemic normativity. For example, someone who is wondering what they *know* about whether a certain strategy will destroy looming asteroid will typically also be wondering about about non-epistemic facts about the subject matter at hand (will this strategy destroy the asteroid?). These non-epistemic facts provide *one* serious subject matter which it is possible for the deliberator to be wrong about, and one important sense in which their cognitive efforts could fail.

Second, even cases where such underlying non-epistemic subject matter facts are not in question (e.g., someone asking "But was I justified in believing that he could legally and safely drive?" after news of inebriated crash), a possibility of failure and possible need for effort (rather than arbitrary picking what to say) remains. For even if we combine the epistemic unprincipledness thesis with

very powerful metasemantic claims about epistemically normative concepts just rigidifying facts about how actual humans are disposed to reason, this still leaves plenty of room for individual judgments and reasoning to err - e.g. in cases where these judgments disagree with what would be accepted under reflective equilibrium.

To dramatize this point, compare the epistemic unprincipledness thesis to the following psychologistic hypothesis about our color concepts: 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 (e.g., classifying sample patches as turquoise or not) will remember the feeling of serious deliberation and a possibility of error. And surely there is no conflict between feeling this effort (and recognizing the possibility of error) and accepting the above arbitrariness/psychologism about color boundaries.

Third, (because of the rigidification involved) accepting epistemic unprincipledness is entirely compatible with allowing facts about justification etc. to be independent of speakers' dispositions in the following straightforward sense. Although the epistemic unprincipledness view suggests that facts about *what concept your word "justified" expresses* are dependent on your dispositions to make certain actual world accuracy-promoting inferences but not others, 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, overall, there are a number of ways that advocates of epistemic unprincipledness can agree that first person judgments about epistemically normative concepts aim to get things right in a sense independent from what we actually wind up deciding. So they can accept (the analog to) 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"³⁵

35. A lingering worry might remain: doesn't endorsing epistemic unprincipledness make *justification* feel less worth caring about than accuracy or truth? If it does, I do not see this as a defect. Philosophical reflection can—though need not—reshape our evaluative priorities, and a theory is not discredited simply because it might prompt such a reorientation.

A different reply concedes no motivational shift. Perhaps we have always cared, ultimately, about accuracy alone, treating facts about justification and good argument merely as psychologically convenient guides to truth. On this picture, the shared consequentialist-deontologist claim that accuracy is the fundamental epistemic good is correct, but the expectation that every foundational question about justification must be answerable *solely* in terms of relations to accuracy is an error.

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 contingent features of actual human inference dispositions rather than anything deeply principled.

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. I've also raised and tried to address worries that accepting epistemic unprincipledness somehow conflicts with the phenomenology of first person epistemic deliberation.^{36,37}

36. An anonymous referee suggested the following interesting question. Might it be that epistemic unprincipledness truly describes our *actual* current concepts of justification, knowledge etc, but we should start using some more principled concept of justification*, which eliminates the 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.

37. 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. Thus it suggests more radical unprincipledness than well known

References

- Berker, S. (2013). Epistemic Teleology and the Separateness of Propositions. *Philosophical Review*, 122(3):337–393.
- Berry, S. (2013). Default Reasonableness and the Mathoids. *Synthese*, 190(17):3695–3713.
- Berry, S. (2018). (Probably) Not companions in guilt. *Philosophical Studies*, 175(9):2285–2308.
- Berry, S. (2019). External World Skepticism, Confidence and Psychologism about the Problem of Priors. *The Southern Journal of Philosophy*, 57(3).
- Berry, S. (2024). Metaethical Deflationism, Access Worries and Motivationally Grasped Oughts. *Ethical Theory and Moral Practice*, 27(3):301–318.
- Boghossian, P. (2003). Blind Reasoning. *Aristotelian Society Supplementary Volume*, 77(1):225–248.
- Brier, G. W. (1950). Verification Of Forecasts Expressed In Terms Of Probability. *Monthly Weather Review*, 78(1):1–3.
- Carnap, R. (1950). *Logical Foundations of Probability*. University of Chicago Press, Chicago.
- Chrisman, M. (2007). From Epistemic Contextualism to Epistemic Expressivism. *Philosophical Studies*, 135(2):225–254.
- Chrisman, M. (2012). Epistemic Expressivism. *Philosophy Compass*, 7(2):118–126.
- Clarke-Doane, J. (2020). *Morality and Mathematics*. Oxford University Press.
- Dandeleit, S. (2024). Epistemic Rationality and the Value of Truth. *The Philosophical Review*, 133(4):329–365.
- Dogramaci, S. (2012). Reverse Engineering Epistemic Evaluations. *Philosophy and*

appraiser contextualist theories (on which, e.g., the correctness of my attributing someone knowledge/justified belief 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) Rysiew (2023). For the latter contextualist theories are compatible with taking facts about *justified high confidence* to be entirely principled and context-independent.

Second, the epistemic unprincipledness thesis I'm advocating differs from theories of agent relativity, like McKenna (2023)'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 contingent features of their situation like the practical feasibility and costs of S so engaging. For the latter view only requires that practical features of an agent's situation impact the stringency of epistemic requirements on that agent (in a principled way). It does not suggest anything unprincipled in our overall ways of classifying people as justified or not.

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 variant knowledge and justification like concepts (e.g. corresponding to different verdicts about Gettier cases).

- Phenomenological Research*, 84(3):513–530.
- Dogramaci, S. (2015). Communist Conventions for Deductive Reasoning. *Noûs*, 49(4):776–799.
- Eklund, M. (2017). *Choosing Normative Concepts*. Oxford University Press.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press UK.
- Field, H. (2009). Epistemology without metaphysics. *Philosophical Studies*, 143(2):249–290.
- Fraassen, B. C. (1989). *Laws and Symmetry*. Oxford University Press.
- Goldman, A. and Beddor, B. (2021). Reliabilist Epistemology. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2021 edition.
- Ichikawa, J. J. and Steup, M. (2018). The Analysis of Knowledge. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2018 edition.
- Korsgaard, C. M. (1996). *The Sources of Normativity*. Cambridge University Press.
- Lewis, D. (1971). Immodest Inductive Methods. *Philosophy of Science*, 38(1):54–63.
- Lin, H. (2024). Bayesian Epistemology. In Zalta, E. N. and Nodelman, U., editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2024 edition.
- McKenna, R. (2023). *Non-Ideal Epistemology*. Oxford University Press, Oxford, New York.
- Rysiew, P. (2023). Epistemic Contextualism. In Zalta, E. N. and Nodelman, U., editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2023 edition.
- Schechter, J. (2019). 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.
- Stevens, M. (2011). Probability Out Of Determinism. In Claus Beisbart and Stephan Hartmann, editors, *Probabilities in Physics*. Oxford University Press.
- Sylvan, K. L. (2020). An Epistemic Nonconsequentialism. *The Philosophical Review*, 129(1):1–51.