

Epistemic Sentimentalism

Sharon Berry

Abstract

In this paper, I'll argue that considerations about basic a priori knowledge suggest a problem for both consequentialist and deontological views in recent debates about the foundations of epistemic normativity. I'll advocate an alternative approach, which I'll call Epistemic Sentimentalism, and contrast this view with its closest relative in the literature (Dogramaci's epistemic communism).

1 Introduction

Which logically valid inferences is it OK to make immediately? Many people have the following intuition. It's OK to make modus ponens inferences immediately, without need for appeal to any further evidence or justification¹. In contrast, it's not OK to infer Fermat's Last Theorem (FLT) from the standard ZFC (Zermelo-Fraenkel with choice) axioms of set theory immediately – even though this inference is just as logically valid (and hence necessarily truth-preserving) as modus ponens. Someone who does the latter does not thereby acquire a justified belief. And one cannot gain knowledge of FLT from this too quick 'proof'. But why is this so? What accounts for the different epistemic status of these inferences? And more generally, why is it OK to make some logically valid inferences immediately, but not others?

Works like [2, 8] have used difficulties in answering the above questions in a principled way to (in effect) move us towards a view I'll call epistemic sentimentalism.

¹Rejecting this threatens to generate an infinite regress of justifications as per 'What the Tortoise Said to Achilles'.

Epistemic sentimentalism: Some aspects of our concept of justification (knowledge, adequate evidence, good argument proof etc.) reflect contingent features of human psychology.

In particular

- Berry [2] argues there's no deep difference between us and creatures fortunate enough to find additional logically valid inferences compelling. She (in effect) uses the puzzle about basic logical inferences above to motivate a weaker claim that we should *either* embrace epistemic sentimentalism or take a permissive 'let 1000 flowers bloom' approach, which counts all logically valid inferences as equally justified.
- Dogramaci[8] uses this puzzle to argue for a view called epistemic communism, which (I'll suggest) includes epistemic sentimentalism but adds the further controversial claim that the main function of epistemic evaluation is to coordinate reasoning methods.

In §2 I'll situate epistemic sentimentalism in the literature. First I'll review some motivations for epistemic sentimentalism, and give an example of what an epistemic sentimentalist theory of justification might look like. Then I'll clarify how I take epistemic sentimentalism to parallel Humean moral sentimentalism, and provide an alternative to both the epistemic consequentialist and deontological views currently slugging it out in philosophy journals.

Then, I'll devote the bulk of the paper to advocating epistemic sentimentalism simpliciter as the best response to the above puzzle about basic logical inference (contra previous work on this topic), in the following sense.

In §3 I'll argue that the permissive option Berry [2] leaves open as an alternative to epistemic sentimentalism should be rejected. I'll note that a variant on the arbitrariness puzzle above about basic logical inferences can be formulated, regarding the problem of priors. And I'll argue that (unlike epistemic

sentimentalism) permissivism yields implausible consequences when applied to this question about priors.

In §4 I'll attack Dogramaci's epistemic communism[8], by attacking the further function claim it combines with epistemic sentimentalism. I'll argue that a certain kind of spontaneity of our reasoning methods (which Dogramaci himself emphasizes in other contexts), makes it implausible that justification talk can have the function Dogramaci imagines. And I'll briefly suggest some alternative purposes for justification talk (compatible with epistemic sentimentalism).

2 What is Epistemic Sentimentalism?

2.1 Motivation and A Sample Epistemic Sentimentalist View

So let's begin with existing motivations for epistemic sentimentalism. As noted above, the question 'Which logically valid inferences is it OK to make immediately?' can seem to reveal a puzzle about our common concepts of justification and knowledge.

Many people feel it's OK to make modus ponens inferences immediately, but not OK to infer Fermat's Last Theorem (FLT) from the ZFC axioms of set theory immediately. Why is this, given that 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 problem² suggests the following thought. Perhaps there's no deep or interesting sense in which creatures who were lucky enough to find additional logically valid inferences compelling would be worse off than us. So perhaps facts

²Schechter's [14] 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.

about justification (in particular, about logically valid inferences it's OK to make immediately) reflect mere contingencies of human psychology. Perhaps the only reason why (we have and use a concept of justification on which) it's OK to make modus ponens inferences immediately, but not the FLT inference, is that humans are disposed to find the former and not the latter immediately compelling³. And perhaps creatures who found different logically valid inferences immediately compelling would have their own equally principled notion of justification*, such that 'proofs' using the additional logically valid inferences which they find compelling would count as providing justification*). 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.

Accordingly, considering the question 'which logically valid inferences is it OK to make immediately?', can motivate the following view.

Epistemic sentimentalism: Some aspects of our concept of justification (knowledge, adequate evidence, good argument proof etc.) reflect contingent features of human psychology.

2.2 Example

The above idea can be fleshed out in different ways. But a simple and natural example of an epistemic sentimentalist account of justification (motivated by the considerations above) might say the following. There are certain moderately truth-conducive⁴ kinds of reasoning which we are inclined to find immediately compelling, once exposed to the relevant concepts. Modus ponens is an exam-

³Human beings find the standard combination of introduction and elimination rules for the material conditional (which include modus ponens) immediately acceptable (and so can learn and accept this concept by immersion), but they do not take that attitude towards concepts whose combined introduction and elimination rules let you quickly go from the ZFC axioms to FLT, like the schnumber concept of [2]

⁴C.f. the discussion of (actual world) truth conduciveness below.

ple of such reasoning. The short ‘proof’ of Fermat’s Last theorem considered above is not. Our epistemic concepts like justification (and adequate argument and evidence) **rigidify/project** these facts about which kind of (moderately reliable) reasoning which “we” are thus inclined to find immediately compelling, in roughly the following sense.

At every metaphysically possible world w , a creature will only count as gaining justification from an inference, if can break this inference down into the kind of steps⁵ which are psychologically compelling to humans *in the actual world*.

Thus, for example, a community of aliens who are lucky enough to find the ZFC to FLT inference immediately compelling would not count as justified in drawing this inference⁶ – while we are justified in making modus ponens inferences immediately (and the same goes for all other logically valid inferences we’re inclined to find immediately compelling). However, this fact does not reflect anything metaphysically special about us, or suggest that human psychology correctly reflects some kind of independently interesting distinction among different logically valid inferences (in a way that might, e.g., create an access problem)⁷.

On this version of epistemic sentimentalism notions like ‘justification’ and ‘rationality’ reflect contingencies of human psychology (e.g., what kinds of reasoning humans are inclined to find immediately compelling), like the boundaries of our concept of ‘edibility’ reflect biological facts about human digestion and

⁵Talk of ‘kinds of reasoning’ obviously brings up the generality problem. 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.

⁶And nor would humans in an alternate possible world who did the same.

⁷These aliens might have their own concept, justification*, which reflected their different (but equally truth conducive) methods of first order logical deduction and standards for “adequate” mathematical proof. And we need not take the notion of justification to be more intrinsically special or joint carving than that of justification*.

those of ‘a handful’ reflect the size of human hands.

Note that in saying there’s a handful of fruit, I’m not advocating having hands the size of human hands; I’m just using a concept whose boundaries it would be foolish to try to derive and explain from general principles without reference to human hands. And (of course) in adopting the above theory about the boundaries of our concepts ‘handful’ and ‘edible’, one need not suppose that claims made using these concepts have any kind of special ‘world to mind’ direction of fit – or deny that such claims can be true or false in the most ordinary and straightforward sense of the term.

2.3 Foundational Epistemology

Epistemic sentimentalism promises to add a third option to currently lively debates between consequentialism and deontology in foundational epistemology.

In addition to familiar questions about the extensions of concepts like justification and knowledge (e.g., the quest for informative necessary and sufficient conditions pursued in the Gettier literature[10]), it seems we can also ask a further, more foundational, type of question. Even if we had a perfectly correct extensional theory of which beliefs/inference methods are justified, we could still ask why *these* beliefs/inferences are justified but not others. Boghossian evokes this kind of further 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]. We might want to know not just *which* factors are relevant to some belief being justified but *why* those factors are relevant.

Many philosophers have hoped to answer such foundational questions by finding some *principled and intrinsic good feature*, which distinguishes justified beliefs, warranted inferences etc. In this way, one could hope to both explain

and justify our drawing the epistemic distinctions we do.

In particular, epistemic consequentialists and deontologists [15, 1] share the project of attempting to find such a principled and intrinsic good feature. Indeed both sides agree that the fundamental value relevant to epistemology is accuracy, and hope to somehow entirely explain the boundaries of our concept of justification by appeal to this value, as per the following claim.

“Valuing Thesis: The fundamental normative explanation of why justified beliefs are justified is that they manifest certain ways of valuing fundamental epistemic value.” [15]

These two dueling schools of philosophers just disagree about how to provide relevant explanations. The consequentialist contingent understands ‘valuing accuracy’ in terms of accuracy *maximization*. So they say that the reason why certain ways of reasoning are justified (in the sense of the foundational question above) is that they maximize accuracy. And the epistemic deontologist differs by understanding ‘valuing accuracy’ in terms of *respect* rather than maximization^{8,9}.

In contrast, the epistemic sentimentalist rejects the assumption shared by both epistemic consequentialists and deontologists, that we can hope to explain

⁸Doing 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.” [15].

⁹Although advocating a deontological approach, Sylvan says rather little 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).” [15].

the boundaries of our concept of justification by citing some principled good feature. For (as motivated by the puzzle about which logically inferences are OK to make immediately, at the beginning of this paper) the epistemic sentimentalist maintains that some aspects of our concept of justification reflect (project and rigidify¹⁰) mere contingencies of human psychology. Accordingly, (they will say) we shouldn't expect to be able to explain all facts about which kinds of beliefs are justified by citing some principled good feature that distinguishes justified from unjustified beliefs.

Epistemic sentimentalism: Some aspects of our concept of justification (knowledge, adequate evidence, good argument proof etc.) reflect contingent features of human psychology.

In this way, I think Epistemic Sentimentalism presents a kind of motivated Humean alternative to the Kantian and Millian approaches to foundational questions about epistemic notions referenced above. For 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 actual humans are disposed to react, rather than tracking intrinsically principled, natural kind properties of the acts/agents being evaluated. Analogously, my epistemic sentimentalist 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).

However, we should note that epistemic sentimentalism (as characterized above) only requires that *some* aspects of our concept of justification aren't

¹⁰See §2.2 for more detail.

deeply principled. The epistemic sentimentalist can allow that a correct conceptual analysis of justification could make some use of principled notions of accuracy promotion or respect (if such can be found). What they deny is that such a conceptual analysis could be given using only such principled notions. So, for example, my epistemic sentimentalist might say that acceptable basic inference methods must be **both** accuracy promoting/respectful of accuracy and bear a certain relationship to actual world human psychology (as per the simple story sketched above)¹¹. Accordingly, she may be able to take on board many epistemic deontological or consequentialist explanations for why certain things *aren't* justified/rational.

In view of this clarification, some readers may feel that it's odd for me to classify epistemic sentimentalism an *alternative* to consequentialism or deontology. For (they would say) accepting the pinch of arbitrariness epistemic sentimentalism requires doesn't change our fundamental picture of knowledge very much,

I happen to disagree. I think allowing even a small amount of arbitrariness/psychologism in core epistemic notions like rationality and justification makes a big philosophical difference. For example (as we saw the case of the martians who find additional logically valid inferences compelling), it suggests that one can correctly classify someone's belief as being unjustified and irrational *without* committing oneself to this belief being bad in any sense that the person in question should intuitively care about. It thereby suggests a poignant sense in which we might turn out not to live in a 'shared space of reasons'.

¹¹Including this kind of truth-conduciveness/accuracy promotion requirement for justified basic inference methods can help the Epistemic Sentimentalist explain how that fallacies (i.e., psychologically tempting but not justified types of reasoning) are possible. However, there are other possible ways of making room for fallacies. For example, a more internalist epistemic sentimentalist might maintain that facts about a priori justification reflect facts about which assignments of priors we are not disposed to give up upon further reflection. That is, we are justified in assigning priors in those ways which we find initially attractive and wouldn't reject on further consideration.

However, none of the arguments that follow will be affected by this mostly presentational question. So I invite readers who feel differently (about how much of a difference embracing epistemic sentimentalism makes) to just think of this paper as arguing that a certain caveat should be added to extant formulations of consequentialist and deontological views.

Also note that epistemic sentimentalism differs from views like Chrisman’s epistemic expressivism[7, 6], by being compatible with taking epistemic normativity claims to be straightforwardly true or false and denying that they have a particular function of exhortation.¹²

3 Against Permissivism

With this clarification of epistemic sentimentalism (and a simple version of this view I find attractive) in mind, I will turn to the task of arguing for it.

In [2] Berry suggests that puzzles about which logically valid inferences can be made without further argument motivate adopting either (what I am calling) epistemic sentimentalism or a let 1000 flowers bloom pluralist alternative (which I will call the permissive approach). To explain and motivate the latter permissive approach, recall that the epistemic sentimentalist holds that our concept of justification is unprincipled in a way that reflects (and rigidly projects[11]) mere contingencies of human psychology. So, for example, an epistemic sentimentalist might say that the logically valid inference methods which thinkers at all possible worlds have basic warrant for using are just those ones which *we*

¹²Both views hope to treat epistemic and moral vocabulary alike, and are inspired by Hume’s famous lines about our minds gilding and staining the world around us. However, the epistemic expressivist takes metaethical expressivism as their model, and maintains that “epistemic judgments have, at least in part, a desire-like direction of fit with the world”. Either (as per classic moral expressivism) justification and rationality talk resembles cries of ‘boo’ or ‘hurray’ in not making claims that can be true or false, or (as per sophisticated contemporary moral expressivism) justification and rationality talk can assert truths– but only in a secondary sense that requires appeals to special minimalist theories of truth and assertion to be understood.

in the actual world are disposed to find sufficiently compelling. Accordingly, we are justified while Martians who reason via different but equally logically valid inference methods (ones that ‘skip steps’ from our point of view) would not be.

In contrast, the permissive view holds that a thinker is (defeasibly) justified in immediately making whatever of logically valid inferences *they* find immediately compelling. Accordingly, (on this view) both we and the Martians can gain justified true belief and knowledge from reasoning in the way we do¹³. Advocates of the permissive approach can allow there’s nothing intrinsically special/better about modus ponens compared to the inference from some ZFC axioms to FLT. But they will reject the epistemic sentimentalist conclusion that our concept of justification is unprincipled and categorizes the world in a way that reflects mere contingencies of human psychology.

3.1 Arbitrariness Worry About Priors

In this section, I will criticize permissivism, by arguing that it yields unattractive consequences regarding a certain variant on the puzzle about logically valid inferences at the beginning of this paper. This variant puzzle concerns the ‘problem of priors’, so let me begin with some stage setting about that.

It is often useful to model scientific reasoning in Bayesian terms, by supposing a thinker starts with a certain assignment of probabilities (satisfying the probability axioms) in advance of all experience — and then updates their beliefs by conditionalizing on various evidence propositions which they learn from their senses. When we do this, facts about which priors it is OK to have will play an important role in determining which conclusions a person’s total body

¹³The 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.

of evidence permits them to accept/assign high probability to.¹⁴¹⁵.

In [3] Berry argues that there's an unprincipled aspect to our concept of epistemically acceptable priors (much like the unprincipled aspect to our notion of acceptable basic logical inference noted above).

To initially motivate this idea, recall how philosophers of science have put much effort into providing a principled and motivating characterization of what assignment(s) of priors are epistemically permissible, however significant difficulties have arisen. Consider, for example, Carnap's general difficulties formulating a logic of induction and specific issues concerning Carnapian learning parameters¹⁶ and Bertrand's paradox¹⁷

I will now argue that, for any reasonable way of cashing out the notion of objective/actual world truth conduciveness¹⁸, you can always increase actual

¹⁴I will follow [3] in stating the arbitrariness problem in terms of this framework. Some readers may prefer a different framework for thinking about scientific reasoning, on which we somehow don't assign probabilities to any contingent propositions *a priori*, but nonetheless favor some empirically adequate hypotheses over others once we get evidence, performing something like inference to the best explanation. It will become clear that an arbitrariness argument exactly analogous to the one I describe concerning priors arises regarding these dispositions to favor some hypotheses compatible with our total body of evidence over others *a posteriori*.

¹⁵The considerations presented in this subsection will most directly suggest that there's arbitrariness/psychologism in our notion of acceptable priors. However, as will become clearer, it thereby indirectly suggests corresponding arbitrariness in our notions of what probability assignment(s) are permitted to someone with a total body of empirical evidence, and what constitutes an adequate scientific argument.

¹⁶When Carnap modified his theory of the logical foundations of probability [5] 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. Yet, even if one doesn't find Carnap's theory persuasive, one must either abandon learning from experience or pick some number of observations after which such a probability assignment is epistemically permissible.

¹⁷Bertrand'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 and influential formulation of Bertrand's paradox comes from [9]).

¹⁸By actual world truth-conduciveness here, I mean something like the physically objective probability of producing true beliefs/objectively expected accuracy of belief states via Bayesian updating on these priors in the actual world. We might cash this out by considering

world truth conduciveness by tweaking our actual priors to assign high probability to some claims that are (physically necessary) informative scientific laws about the actual world, which we assign low probability to a priori. In this sense we think our priors are fairly good (objectively truth/accuracy conducive), but not maximally or uniquely so ¹⁹. Specifically consider priors which resemble ours but assign high probability (in advance) 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). Creatures who assign probabilities by Bayesian updating on these priors would, from our point of view/intuitively err by dogmatically assuming certain things that happen to be true a priori. However, even by our own lights, such creatures will tend to do better at quickly forming true beliefs/accurate credences than creatures with our correct priors.

In this way, we don't take our priors to be uniquely good (or even optimal) as regards actual world truth conducive. Thus we have a kind of modesty about our priors that interestingly coexists with the kind of immodesty studied by David Lewis [12](i.e., the fact that our priors assign themselves optimal *a priori* expected accuracy). To highlight the difference, imagine that you have

the objective expected (in)accuracy (perhaps measured by a Breyer score) of a robot equipped with certain sensory faculties that gets plonked somewhere random in the actual world and does Bayesian updating on the stream of experiences it would have. Note that this notion of objectively expected accuracy/truth conduciveness can be well-defined and non-trivial even if physics is completely deterministic. There are clearly many choice points faced when filling in this notion. For example, in addition to a way of choosing and weighting the importance of relevant propositions for calculating a Breyer score, the above picture appeals to a favored notion of spatial and temporal distance/volume (to cash out the idea of a robot being equally likely to appear anywhere) and a weighting of options about many observations the robot is allowed to conditionalize on before we assess its accuracy with the Breyer score. I'm not sure that there's a uniquely preferred and principled way of filling in all these parameters – or even just the weighting of propositions needed to specify accuracy via Breyer score (which could be a problem for epistemic consequentialists and deontologists who take accuracy to be an entirely principled notion which can be used to give a precise and principled explanation for all other epistemically normative notions). Instead (much like people making accuracy dominance arguments for conditionalization, and/or obeying the probability axioms) I want to suggest that, for each reasonably attractive way of filling these parameters in, there is some implementation of the strategies described below (e.g., for cooking alternative priors which are more objectively truth conducive than our actual ones) that works.

¹⁹That is, there should be many ways of assigning priors that qualify as more objectively truth conducive than the priors we actually use.

eliminated all but two hypotheses H_1 and H_2 about the fundamental physical laws of the world (e.g. the world is physically necessarily made of atoms and the world is physically necessarily made of gunk), and you assign 50% probability to each of these. In this situation, your probability assignment might take itself to have maximum rationally expected accuracy (e.g., it doesn't expect H_1 or H_2 to be more accurate than itself). In contrast, you will assign high probability to your way of assigning probability not being maximally objectively truth-conducive. 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 will be confident in the existential claim that *some* alternative way of assigning priors is more objectively truth-conducive to your own – though you won't know which one.

In fact, this modesty goes a step further, which will be important to keep in mind for the arguments below. I claim that, for almost any metaphysically contingent claim, however false or absurd, there's an alternate assignment of priors which gives this claim high probability²⁰, but is (overall) just as actual world truth-conducive as our way of assigning priors. For example, suppose we want to cook up an assignment of priors that is just as actual-world truth-conducive as our ordinary way of assigning priors, but assigns probability .99 to 'the moon is made of blue cheese'. Just modifying out actual assignment of priors to include this probability will reduce their objective truth-conduciveness. However one can (in most cases) compensate for this by the trick above, of further modifying these priors to assign high probability to true scientific laws (in unrelated areas), which ordinary humans don't assign high probability to a priori²¹.

²⁰The same goes for an assignment of priors that assigns high probability to that claim given some specific, narrowly constrained, type of evidence.

²¹Admittedly, this strategy may not be applicable in cases where the false proposition

So we get the following variant on our initial puzzle about which logically valid inferences can be made immediately (which I will argue poses special difficulties for the permissivist): why is it OK to assign high a priori probability to certain physical truths (e.g., that the future is like the past in certain ways) but not to others (e.g., the facts summarized in the periodic table of elements)? Our judgments about justification and what constitutes an adequate scientific argument for a claim reflect this distinction. However, it's hard to find any principled intrinsically special nature which distinguishes the kinds of contingent scientific truths which can be assigned high probability a priori from those which can't.

The difficulty of providing a principled answer to the above puzzle can suggest something like the following picture (which is a form of Epistemic sentimentalism). Human beings are inclined to substantially agree in how they assign priors. The probability assignments we actually favor involve a mix of symmetry intuitions, preference for simplicity and permission to learn from experience. But there's nothing special about this mix: it just happened to be reasonably useful and easy to physically realize in the human brain in the context of evolution.

We think reasoning from these priors is objectively truth-conducive to a certain (fairly significant) degree in the actual world²². And, given that we are psychologically inclined to assign priors in this way, it is useful for us to have a concept of justification which tracks what beliefs/probability assignments can be gotten via using evidence to update the (moderately objective-accuracy

chosen is extremely powerful and theoretically central. For, such powerful false beliefs might be impossible to logically coherently combine with compensatory significant true beliefs about other matters. However, even some of these cases might be salvageable by choosing priors which assign high probability to the relevant powerful contingent falsehood *a priori*, but low probability to it *conditional on having certain experiences actual world epistemic agents are very likely to have*.

²²For example, one might crudely cash this out in terms of the objective physical probability of this method producing true beliefs – given some measure on the space of beliefs and locations where one might find oneself in the total history of the actual world.

promoting) ways of assigning priors that humans (in the actual world) favor. However, there is nothing uniquely intrinsically special about this exact choice of priors which distinguishes them from others that are equally objectively truth conducive (and obey the probability axioms).

3.2 How this Creates Trouble for Permissivism

Now let's consider how the permissive approach handle this puzzle about priors. As noted above the epistemic sentimentalist can say that our concept of justification requires (from thinkers at all possible worlds) something like adherence to the actual, moderately objectively truth conducive, way humans actually do assign priors. In contrast, it seems the permissivist should say that it's epistemically permissible for a thinker to assign high probability to everything they find a priori attractive (even if such a prior would strike us as deeply irrational) — provided these priors obey the probability axioms and are sufficiently objectively truth conducive (i.e., as objectively truth conducive as we take ours to be).

However, I claim the latter application of permissivism to the problem of priors yields radical and counter-intuitive consequences. Specifically, it threatens to trivialize talk of propositional justification (i.e., impersonal talk about what propositions one can justifiedly believe/assign high probability to, given a certain total body of evidence). For, given almost any such proposition ϕ (that's not a metaphysically necessary falsehood) and evidence set E , one can use the trick above (i.e., compensating for dogmatically assuming a falsehood by dogmatically assuming some unrelated physically necessary truths) to produce priors which assign high probability to ϕ after updating on evidence E , and meet the requirements above (i.e., obeys the probability axioms and is comparable to our own priors in overall objective truth conduciveness). Thus, extending the

permissive approach to our new puzzle about probability thretans to imply that almost any contingent claim C can be assigned high probability conditional on whatever evidence E you currently have.

Second, arguably the permissive approach to priors fits poorly with our actual current practices and behavior towards people who seem to be reasoning differently from us. When we say someone else's belief is unjustified (given some shared body of evidence), we commonly (intuitively) don't commit ourselves to them having similar priors to us concerning other matters, or otherwise try to rule out the possibility that they are starting from different but equally objectively truth conducive priors. The hypothesis 'maybe this person favors the epistemically worse (unjustified) explanation for this data because they're starting from different but equally objectively truth conducive priors' seems like a coherent possibility (if unlikely). But according to the permissive view it should not be; for (on this view) everyone can form justified beliefs just by correctly updating on whatever sufficiently truth-conducive priors they happen to have.

Thus, although the permissive approach can be tempting when we just consider questions about permissible logically valid inferences, it implies implausible consequences when applied to the (hitherto less emphasized) question of epistemically permissible assignments of priors.

4 Against Epistemic Communism

In this final section, I'll contrast epistemic sentimentalism with one of the most similar views in the literature: Dogramaci's epistemic communism [8]. Although Dogramaci's work on basic logical inferences is a significant influence and motivator for my current proposal, Dogramaci's epistemic communism goes beyond epistemic sentimentalism by adding a further function claim, that I'll try to call

into doubt.

But first let's start with some points of agreement. In formulating epistemic communism, Dogamici says, and I agree, that there are a range of comparably truth conducive (and in all intuitive senses intrinsically equally good), variants of our concepts of justification. So I take it that epistemic communism counts as a form of epistemic sentimentalism.

However, Dogramaci's epistemic communism goes beyond epistemic sentimentalism (as I've defined it) by including a further claim that I want to dispute. He proposes that epistemic evaluations serve the following function: we use these evaluations to try to get others to use the same epistemic rules (e.g., basic reasoning methods and constraints on acceptable assignments of priors) that we do, which we take to be good rules. This promotes coordination in our epistemic community, and in the long run helps us to make testimony more trustworthy.

In contrast, the epistemic sentimentalism I've defined and advocated doesn't make any claim about the function of epistemically normative vocabulary. In this section, I will argue that we should resist (or at least question) Dogramaci's further claim that *the* core function of epistemic concepts like justification is to exhort others to change their basic reasoning methods²³, and briefly suggest some alternative functions.

4.1 Background on Epistemic Communism

So let me begin with a bit more detail about Dogramaci's epistemic communism. As noted above, epistemic communism maintains that, "the function (in the sense of point or purpose)" of concepts like justification is to promote coordination and division of epistemic labor, in roughly the following sense. If

²³One might say the topic of debate here specifically concerns the function of our a concept of justification that distinguishes between different kinds of valid arguments or objectively truth conducive priors in the way mine does.

everyone agrees to use the same basic reasoning methods in their proofs, then we can trust their testimony because they have given us the kind of proofs we can confirm. The main function of concepts like justification and rationality (according to Dogramaci) is to exhort others to change their basic reasoning methods so as to better coordinate.

In [8] Dogramaci provides some motivation for his function claim by attacking a common alternative view (which I agree should be rejected) — that the main function of concepts like justification is to guide first-person deliberation and inquiry. On this view, “[serious inquiry] involves a thinker’s asking herself whether she would be justified (or rational) in believing this or that” [8], in such a way that (contra epistemic communism) the concepts of justification and rationality play an essential role in guiding the first person deliberation and belief formation.

Dogramaci criticizes this view by noting that most beliefs are formed spontaneously, with no appeal to concepts of rationality and justification made or needed. He allows thoughts about rationality and justification can play more of a role in drawn-out conscious first-person deliberation. However, he suggests that thoughts about notions like rationality and justification are generally dispensable in process, as they replaceable with, e.g., considerations of reliability.

A thinker may use the concept of rationality in the course of her deliberation, but it would be totally inessential to the deliberative process. Imagine deliberating about something, say, whether there is alien life. How do you deliberate about whether there is alien life? You ask what the data is: what do we see through telescopes, what do radio signals from space indicate, how many other planets are as hospitable as ours? You ask whether that data provides a reliable indication of the existence of alien life. By answering those

questions, you may settle your deliberations about whether there is alien life. Do you, or should you, also ask whether the rational thing to believe, given that data, is that there is alien life? No, that is not normally relevant. Or imagine, to return to the deductive side of reasoning, deliberating about whether Peirce's Law, $((p \supset q) \supset p) \supset p$, really is a theorem (as you might have just recently done, if you had not encountered it before reading this paper). You've settled the matter when you've settled whether there are any truth-value assignments that invalidate the claim. Do you also take up the question of whether it is also rational to think there is such a truth-value assignment? Again, that is inessential, and not normally relevant. For the most part, an ideal deliberator may proceed perfectly properly without giving any consideration to questions of rationality

Whether or not appeals to epistemically normative notions (like justification and rationality) are entirely dispensable from first person deliberation, Dogramaci is surely right to emphasize the spontaneity of normal belief formation (how forming new beliefs does not generally require considering any claims about epistemic normativity). For example, it would be absurd to suppose (on the model of belief-desire psychology for action) that forming a new belief that P requires some kind of combination of a desire to be rational with a recognition that believing P is the indispensably necessary means to being rational²⁴. Indeed I'll suggest below that Dogramaci's epistemic communism may not go far enough in acknowledging the spontaneousness of applying our basic reasoning methods in deliberation²⁵.

However, even if we grant that justification thoughts are entirely dispens-

²⁴Among other things, this picture would seem to generate a regress, where inferring that P requires first forming the belief that rationality requires believing P, but forming the latter belief itself requires previously forming a belief that rationality requires forming this belief.

²⁵By this I mean making the kinds of inferences we're inclined to make without further argument.

able from first-person deliberation, this leaves plenty of room to doubt epistemic communist claims that exhortation to coordinate basic reasoning methods is the main function of our concept of justification. For note that appeals to epistemically normative concepts seem to be just as dispensable in such attempts to promote coordination, as they are in first person deliberation. Couldn't we exhort others to change their basic method of reasoning by simply saying things like 'please reason differently!' instead? And, below I will argue that considerations of about the spontaneity of reasoning Dogramaci emphasizes make it unclear whether justification talk can ever (or often) succeed in promoting coordination in basic reasoning methods.

4.2 Basic Reasoning Methods Mostly Beyond Voluntary Control?

In this section, I'll question Dogramaci's claim that the main function of epistemically normative concepts is the exhortation of third parties to coordinate with us by changing their basic reasoning methods (e.g., which logically valid inferences they'll immediately accept), by raising a worry about feasibility.

I'll argue that success at this project is sufficiently unlikely to make it implausible as a candidate for the 'main function' of concepts like justification and rationality (especially given the range of important and more feasible rival functions which I'll mention below)²⁶. In particular, I'll argue that it's unclear whether people can often or ever consciously change their basic sense of a priori theory plausibility, to better coordinate with others – even when doing so would clearly be useful in the ways Dogramaci imagines.

Obviously, we do often criticize peoples' *arguments*, in hope of getting people

²⁶Of course, the mere fact that Fs rarely succeed at ϕ ing doesn't automatically ensure that the main function of Fs isn't ϕ ing. For example, the main function of a matchmaking app might be to make introductions that lead to great marriages, even if it has little power to do this and only succeeds in .01% of cases.

to change their minds. But it's not clear that, in such cases, we hope to change peoples' basic reasoning methods. We might instead simply hope the people we're addressing have failed to fully and carefully apply their own current basic reasoning methods (e.g., that they have failed to realize dispositions they already have to make certain logically valid inferences immediately but not others), and hope to change their minds by inviting them to do some double checking. So, I claim that normal uses of justification and rationality talk for third-person persuasion don't require that there's any possibility of changing interlocutors' basic reasoning methods exhortation²⁷ .

And, in fact, I think there are good reasons to doubt the feasibility of inspiring people to change their basic reasoning methods by exhortation (and thus also to doubt Dogramaci's suggestion that producing such change is the main function of concepts like justification). To motivate this point, imagine the following scenario. You learn that some mathematical claim, which doesn't immediately strike you as obviously true or false, is actually true by some method other than a priori argument, (e.g., by testimony, or some kind of physical experiment hypercomputers²⁸). Then you learn that humans have got into contact with a larger galactic community, and that most members of this community find the relevant mathematical fact a priori obvious. In this case, there would seem to be strong coordination benefits of the kind Dogramaci cites to starting to find this principle immediately compelling (and so, e.g., accepting 'proofs' which take this fact as an unargued premise as providing a priori knowledge of

²⁷One might compare criticizing some step in a scientific argument as unjustified to suggesting a dish might be improved by substituting one spice for another. In making such claims about cooking, you don't usually hope to change anyone's palate (basic reasoning methods) by exhortation. One doesn't hope to shame, peer pressure or argue the cook into finding different things delicious. Instead, you simply think a change of spices would be an improvement from the point of view of your palate. And you hope that the cook's palate is sufficiently similar to yours, that you have good hope such a change will be an improvement to for palate as well (so merely inspiring further reflection and experiment could get them to change their approach).

²⁸See Malament's [13] involving throwing a computer into a black hole and waiting to hear back while it checked infinitely many cases in what was (for you) finitely much time.

their conclusions²⁹). But it seems plausible (to me at least) that, however clear and strong this prudential argument was, one wouldn't be able to do this. And an analogous case could be made involving different ways of judging a priori theory plausibility.

Admittedly, we do seem willing to change our basic reasoning methods in some ways. For example, contemplating actual (or concretely imaginable) cases where some inference method leads from truth to falsehood can cause you to stop finding that method compelling. And repeated experiences of going through a certain kind of multiple-stage argument could lead you to mentally chunk the process together, acquiring a disposition to go from the beginning to the end of arguments of the relevant type immediately.

But none of these types of change involve finding different premises/inferences immediately compelling in response to mere exhortation or conviction that reasoning differently would better serve the practical goals of social coordination Dogramaci mentions.³⁰³¹

Accordingly, I claim that appeals to justification (and other such epistemically normative vocabulary) have little power to get others to better coordinate by changing their basic reasoning methods.

²⁹Of course, such proofs could give one a posteriori knowledge of the same claims, which justification partly depending on your initial testimonial or scientific evidence for believing the initial non-obvious mathematical truth mentioned in this thought experience.

³⁰One might claim that some interpersonal influence on reasoning methods exists in the following sense. Community members tend to imitate one another, so that the fact that everyone around you starts to accept some claim or reasoning method suddenly makes you start to find it plausible. But such influence via imitation does not involve deployment of a justification concept, and hence cannot be used to defend the claim that reasoning-coordination is the main function of concepts like justification.

³¹Perhaps a version of epistemic communism which maintains the only function of justification talk is to influence *your own* future reasoning would survive the worries above. However, this is clearly more individualistic than Dogramaci has in mind. Also, I'm not sure whether this position can be reconciled with Dogramaci's claims about the dispensability of justification talk (the thought that we can always rationally reconstruct first-person deliberations without using that notion) cited above.

4.3 Other Functions of Epistemic Vocabulary

If the main function of justification talk isn't to exhort others to change their basic reasoning methods (in hopes of promoting better coordination), what else could the main function of such talk be? Is there a plausible answer to this question that's compatible with epistemic sentimentalism?

To address this challenge, I'll end this section by briefly reviewing some alternative functions for our epistemically normative concepts (compatible with Epistemic Sentimentalism), which are plausibly as, or more, important than Dogrammici's claimed main function³².

The points above already suggest one way epistemically normative notions like justification can be useful: in prompting others to recheck their work. If it's a well-known empirical fact that people tend to approximately agree in which basic inferences they find compelling³³, then (in cases where we share sufficiently much relevant evidence) the fact my basic reasoning methods lead me to reject a claim that you accept, can provide some evidence that you have failed to apply your own methods in this case (performance has fallen short of competence) and inspire you to recheck your work. So talk of justification can perform the useful function of helping us help our neighbors to better apply their own favored basic reasoning methods.

It also seems that a completely solitary person could use the concept of justification (understood upon epistemic sentimentalist lines) to gather and classify resources for their own attempts to form true beliefs about the world, as follows.

Consider a case where I'm trying to determine whether P, and thus looking for

³²I won't argue that justification plays an indispensable role in these functions. For (as argued above) I don't think epistemically normative terms like are indispensable to Dogrammici's purported main function of coordinating belief formation. And I don't take indispensability to be conceptually necessary for a main function claim (there is no conflict between holding that the main function of brooms is to clean floors and that one can clean floors by vacuum cleaner instead).

³³More pedantically perhaps (since mathematical experience can teach people to skip steps), they tend to agree in all accepting arguments that they are capable of justifying from certain shared premises.

certain kinds of arguments and experiences that will help with that project. What kinds of arguments/experiences am I looking for? Crudely put I'm looking for things like

- a priori arguments for P (or not P) – which employ the kind of basic truth-preserving steps that I'm disposed to find immediately convincing.
- sensory evidence that would warrant high confidence in P (or not P) — relative to my priors (as updated according to my current evidence).

Given our reoccurring interest in finding these particular kinds of things (experiences and arguments that relate to our psychology re: priors and convincingness in this useful way), it makes sense that we'd have concepts that classify the world accordingly. So it seems that even a causally impotent individual (with no hope to change others reasoning methods) could usefully employ concepts like 'proof' and 'adequate evidence' for the purposes of describing and keeping track of the kinds of experiences and arguments they are looking for.

Similarly, it can be useful for me to classify others as sources of information (assess their beliefs as having or lacking adequate evidence or argument - i.e., the kind of truth-conducive evidence or argument that would be able to settle the question for me – and rate their rationality according to their disposition to be guided by such arguments) regardless of whether I have any hope or desire to change how other people reason. For example, a spaceman in a different solar system might facilitate her own learning from news by classifying politicians and new anchors as rational or irrational (and disposed to make good arguments or not), even if this news reaches her centuries too late to influence anyone with these classifications.

Thus I think an epistemic sentimentalist can acknowledge various possible functions of epistemically normative concepts, which are plausibly more im-

portant and central than what Dogramaci claims is the main function of such concepts (exhortation to coordinate basic reasoning methods).

5 Conclusion

In this paper I have tried to develop and defend epistemic sentimentalism — a view of fundamental epistemic normativity on which certain aspects of our concepts of justification, knowledge, proof, adequate evidence and the like reflect (rigidify and project) mere contingencies of human psychology.

I've sketched an example of what a simple form of epistemic sentimentalism might look like. And I've argued for both favoring epistemic sentimentalism over a more permissive approach and resisting the additional claims the main function of justification talk included in Dogramaci's epistemic communism.³⁴

³⁴An anonymous referee suggested the following interesting question. Might it be that epistemic sentimentalism 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 2.1 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 §3), this would mean classifying almost all contingent claims as justified. So switching to a more principled conception of justification which meant (in effect) 'justified relative to some moderately truth conducive assignment of priors' would mean switching to a concept that's much less practically useful.

Finally, one might consider replacing our actual concept of propositional justification with a variant that's more principled because it reflects relationships to some kind of optional psychology (as regards the promotion or respect of fundamental epistemic value like truth or accuracy), where our concept of justification reflects messy contingencies of actual human psychology. For example, creatures with this optimal psychology would presumably be disposed to believe all logical and mathematical necessary truths a priori immediately. And they might assign priors to contingent statements in a way that's optimally truth conducive given the facts about the actual world (relative to some weighting of the questions they are likely to ask and the worldly locations they might find themselves in). It's not obvious that there

References

- [1] Selim Berker. Epistemic Teleology and the Separateness of Propositions. *Philosophical Review*, 122(3):337–393, 2013.
- [2] Sharon Berry. Default Reasonableness and the Mathoids. *Synthese*, 190(17):3695–3713, 2013.
- [3] Sharon Berry. External World Skepticism, Confidence and Psychologism about the Problem of Priors. *The Southern Journal of Philosophy*, 57(3), 2019.
- [4] Paul Boghossian. Blind Reasoning. *Aristotelian Society Supplementary Volume*, 77(1):225–248, 2003.
- [5] Rudolf Carnap. *Logical Foundations of Probability*. University of Chicago Press, Chicago, 1950.
- [6] Matthew Chrisman. From Epistemic Contextualism to Epistemic Expressivism. *Philosophical Studies*, 135(2):225–254, September 2007.
- [7] Matthew Chrisman. Epistemic Expressivism. *Philosophy Compass*, 7(2):118–126, 2012.
- [8] Sinan Dogramaci. Communist Conventions for Deductive Reasoning. *Noûs*, 49(4):776–799, 2015.

would be a single psychology that's optimal in this way. But if there were, perhaps one could have a concept of justification* that's entirely principled, because it reflects relationships to this unique optimal psychology in the way that (the epistemic sentimentalists maintain) the notion of justification relates to human psychology.

Would it be practically better (or otherwise more illuminating or interesting) for us to transition to using this notion of justification* rather than our current notion of justification? It's hard to imagine how it could be. For example, how could a notion of proof that broke things down into steps you didn't find viscerally convincing (counting every mathematical proof as provable in a single step) be as useful as our current notion? How could a notion of scientific justification/inference to the best explanation that didn't relate to the kinds of explanations you find most plausible and are disposed to believe? Classifying arguments and experiences based on their relationship to such imaginary psychologically optimal creatures seems both practically difficult and largely not very useful or illuminating.

- [9] Bas C. Fraassen. *Laws and Symmetry*. Oxford University Press, 1989.
- [10] 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.
- [11] Joseph LaPorte. Rigid Designators. In Edward N. Zalta and Uri Nodelman, editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2022 edition, 2022.
- [12] David Lewis. Immodest Inductive Methods. *Philosophy of Science*, 38(1):54–63, 1971.
- [13] David Malament. Science Without Numbers by Hartry H. Field. *Journal of Philosophy*, 79(9):523–534, 1982.
- [14] 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.
- [15] Kurt L. Sylvan. An Epistemic Nonconsequentialism. *The Philosophical Review*, 129(1):1–51, January 2020.