

Are You Now, Or Have You Ever Been, An Impermissivist?
A conversation among friends and enemies of epistemic freedom

Sophie Horowitz
Sinan Dogramaci
Miriam Schoenfield
(version of **Oct 7**, 2022)

I. Sophie & Sinan: Introductions and Stage-Setting

Do the *permissivists*, like you, Miriam, have it right? Is it possible for two perfectly rational people with the same total evidence to have different conflicting views?

Or do *impermissivists* like us, Sophie and Sinan, have it right? Does a rational person's total evidence always determine a unique set of beliefs that they must hold? In other words, is *uniqueness* true?

The literature has become massive. Over 80 papers about permissivism have come out in the years since our teacher, Roger White, published “Epistemic Permissiveness” in 2005.¹ Since then, philosophers have come up with so many smart arguments and counter-arguments, new puzzles and new applications. We admire many of these papers (we won't name names), but we won't attempt to engage here with all of the important arguments others have made. Instead, we will focus right away on an issue that the three of us agree is at the heart of the dispute: the pursuit of *accuracy*.

What do we mean by “accuracy”? When it comes to belief and disbelief, being accurate amounts to believing the true and disbelieving the false. But sometimes it's more apt to talk in terms of *degrees of belief*—or “credences” as we also call them—and this requires a more complex theory of accuracy. For example, suppose that it's raining in Singapore today (call this proposition R), and that Sonia, as we'd put it colloquially, is “80% sure” of R, or as we can also put it, Sonia has a credence of .8 in R. It doesn't make much sense to say that her credence of 0.8 in R is *true* or that it is *false*. However, given that it *is* raining in Singapore, we can say things like this: Sonia's 0.8 credence towards R is *more accurate* than a 0.7 credence in R, and *less* accurate than a 0.9 credence in R. In general, when it comes to credences, you're more accurate the higher your credences are in truths, the lower they are in falsehoods. In epistemic utility theory, the accuracy of credences is measured by an epistemic value function (or “a scoring rule”) which uses

¹ See <https://philpapers.org/browse/epistemic-permissivism> and White (2005).

numbers between zero and one to represent how accurate a given credence is at a world.² Accuracy, then, is a concept that can apply both to binary belief and to degrees of belief.

The three of us agree: rationality, in its epistemic sense (the sense we'll always intend), has some kind of deep connection to accuracy, and the value that comes from being a rational believer is largely or even wholly derivative from the value of accuracy. As we'll put it, we have an accuracy-oriented (or "accuracy-first") perspective when it comes to explaining the nature of rationality.³

That much all three of us agree on. But we disagree about the implications for permissivism. The two of us, Sophie and Sinan, think this accuracy-oriented perspective leads to impermissivism, and it leads there in more than one way. It leads there via certain complicated arguments: as we've argued before, we think that if we understand our practice of making rationality ascriptions as aiming to cultivate accurate and trustworthy testimony, *impermissivism* is part of the best explanation of that practice.⁴ We can elaborate on that more later.

We also think the accuracy-oriented perspective leads to impermissivism almost immediately, without any complicated arguments. It leads there via the natural thought that we serve our goal of accuracy by being guided by the evidence: we proportion our beliefs or credences to what our evidence makes *likely*. The connection between *what is likely* and *what it is rational to think* also comes out when we consider the first-personal perspective of the rational deliberator. When we are wondering what we should think about P, we are just wondering how *likely* it is that P is true, given the evidence. This way of thinking fits naturally with impermissivism, and fits awkwardly with permissivism: the question, "how likely is P?" naturally presupposes a single answer. It's extremely odd to imagine that there could be two or more distinct likelihoods that are all correct.⁵ Imagine a weather forecaster reporting that there is both a 40% and 60% chance of showers this afternoon. It seems nonsensical. So, for reasons like these, we thought that, whatever

² For discussion of this approach to accuracy see Joyce (2009) and Pettigrew (2016). For an alternative approach to thinking about the accuracy of credences, see Horowitz (forthcoming). Sometimes we'll talk about the "expected accuracy" of a credence. For example, if Sonia's credence is 0.8, we might ask, from her perspective, how a 0.9 credence looks accuracy-wise. Expected accuracy is calculated just like expected value: if we're thinking about the accuracy that Sonia expects from a credence of 0.9, we take the average of the accuracy scores a 0.9 credence would get in different worlds, and weight that average by the credences Sonia assigns to those worlds obtaining. In general, we'll assume that the epistemic value functions we use are "immodest" or "strictly proper," which means that these functions have the feature that every probability function maximizes expected epistemic value relative to itself. See Oddie (1997), Greaves and Wallace (2006), Joyce (2009), Horowitz (2014) and Pettigrew (2016) for motivations for this assumption.

³ For an extensive discussion see Pettigrew (2016).

⁴ See Dogramaci & Horowitz (2016).

⁵ See White (2013, pp.313-4) and Horowitz (2014, p.46) for versions of this point. See Kelly (2013), Meacham (2013), and Schoenfield (2014) for counterarguments.

other merits permissivism may have, it was at least obvious that *truth* or *accuracy* considerations were on our side.⁶

But a couple of years ago, Miriam argued that permissivism is in fact in a *better* position to deliver an accuracy-connection than impermissivism. Her argument went something like this: if we think of our rationality judgments as endorsements of those belief-forming rules that we *actually* regard as most accurate, and ruling out belief-forming rules that we regard as inaccurate, we'll find that we just can't rule out all that much. We'll be left with multiple, incompatible options for how to respond to our evidence, none of which seem worse, accuracy-wise, than any other. Therefore, we'll attach the word "rational" to multiple, incompatible, belief-forming methods – and in turn, multiple, incompatible doxastic states held in response to a single body of evidence. Miriam posed this as a challenge to the impermissivist: how could we *impermissivists* possibly justify our claim of *uniqueness* about the rational credences (given the evidence) when there is no unique response (to the evidence) that we actually regard as most accuracy-conducive? We couldn't, Miriam argued – at least not unless we believed in ghosts.⁷

This is the question we are prepared to debate today. We all want accuracy, and we all think rationality is about becoming accurate in some way or another. Given our commitment to accuracy, should we be permissivists, or impermissivists? Our goal is to defend impermissivism. Miriam's goal is, well... we'll just let her speak for herself.

II. Miriam: A Permissive Notion of Rationality

I want to convince you that there is no deep disagreement between accuracy-oriented permissivists and impermissivists. (There is, potentially, some disagreement about how to use language, and certain empirical facts, but, I will argue, the disagreement doesn't run philosophically deeper than that). If we think about our shared starting point, this shouldn't be that surprising. Suppose a bunch of archers agree that there is a close connection between good archery and marksmanship. Then the question arises - is there more than one way to shoot well? We can easily imagine contexts in which it makes sense to say the answer is yes, and contexts in which it makes sense to say the answer is no. What it makes sense to say will depend on the interests we have in some particular context (what's being held fixed, what the standards of success are, whether we're writing a guidebook for archers, or judging an archery competition, etc).

⁶ See also Horowitz (2019) for another accuracy-based argument for impermissivism. For reasons of space, neither we nor Miriam will get into those arguments here.

⁷ See Schoenfield (2019).

I'm going to suggest that something similar is true when we ask "is there more than one way to believe well?" But that will take some work to convince you of.

The first step will be to argue that there is at least one fruitful way of using the word 'rational' according to which even you, Sophie and Sinan, should be endorsing permissivism. It will follow from this that if you insist on endorsing impermissivism, you must be using rationality-talk in some other way. If your way of using the word 'rational' supports impermissivism, and my way supports permissivism, then the dispute between us, I will argue, is a verbal one.

The aim of this section will be the first step: to argue that a certain fruitful way of using epistemically normative language supports permissivism. But before beginning, it will be helpful to put forth a metaepistemological ambition of mine that is in the background of some of these arguments: the ambition is that we be able to give an account of our thought and talk about epistemic rationality in naturalistic terms.⁸ I won't try to say anything here about what "naturalistic" means. Rather, I'll just point to the much more comprehensively hashed out debate in metaethics and say: many of the same considerations that motivate views like relativism, error theory, expressivism, and quasi-realism in metaethics, motivate analogous views about epistemic rationality.

The position I'm inclined to reject is what I'll call "hardcore realism" - a view committed to the conjunction of the following three claims: (1) There is some probability function or set of probability functions that are privileged - they are the rational *ur*-prior(s). (2) The property of being rational (or 'epistemically probable') is brute and irreducible to properties that can be described using naturalistic terms.⁹ And (3) Our thought and talk about epistemic rationality is also irreducible: to account for it, we must make reference to properties that are brute and irreducible.

I propose we reject hardcore realism for the following reasons (again, these are similar to the reasons that hardcore realism is rejected by some metaethicists): First, we don't need it. I'll suggest that we can perfectly well make sense of rationality thought-and-talk - or, at very least, the elements of it that are worth thinking and talking about - without being hardcore realists. And second, if making sense of our thought and talk about rationality really does require an appeal to brute and irreducible properties of

⁸ In fact, for my purposes, I only need the weaker claim that we should be able to give an account of our thought-and-talk about what's rational without appealing to *epistemically-normative* terms. If you wanted to appeal to non-naturalistic things like ghosts, qualia, God, or even other normative notions (like the all-things-considered-ought, the practical ought, morality), that wouldn't have any effect on the arguments that follow. The key point I want to press is that I see no reason for us to take *epistemically-normative* language as primitive. For terminological ease, I'll maintain the focus on naturalism.

⁹ As in note 7, 'naturalistic' could be replaced with 'non-epistemically-normative'.

probability functions, it's quite mysterious why this property is relevant to anything we care about.¹⁰

Books have been written about how to develop alternatives to hardcore realism in metaethics, so I won't pretend to offer anything like a comprehensive alternative account to hardcore realism about epistemic rationality. (I do hope to convince those qualified to do so to get to work on this though!) Instead, I'll make some sketchy remarks about one fruitful way to use rationality-language that doesn't require us to make reference to brute irreducible properties. To be clear: I don't want to make any universal claims about what attitudes people who use the phrase "epistemically rational" or "likely given the evidence" are expressing. My suspicion is that this language gets used in all sorts of messy ways that are not easy to systematize. The remarks below are a proposal for a *fruitful* way to use this language - and one that, as it happens, doesn't require us to appeal to irreducible epistemically-normative properties. In fact, I think the proposal also captures the spirit of much ordinary usage - but I don't intend to argue for that.

First remark: There are at least some contexts in which, when we're deliberating about what to believe, what we're interested in is ending up with beliefs that are accurate.

Second remark: We often have views about which ways of believing will promote our aim of accuracy and which ones won't. (Induction, yes. Counterinduction, no).

Third remark: Sometimes these views are "high level" (I expect more accuracy from credal assignments that assign greater credence to simpler theories), and other times they're quite specific (I expect more accuracy from credal assignments that assign a low credence to the proposition that I'll win the lottery).

Fourth remark: Sometimes we're uncertain which credence functions have the properties we regard as accuracy conducive (I want to assign a credence to the proposition that the poker player to my left has a straight based on an indifferent distribution across hands dealt, but I'm not sure which credence that is).

Final Remark: It can be useful to have some language to express our views about which ways of forming beliefs we regard as accuracy conducive. I'll even go out on a limb and say that many of the valuable activities people engage in under the heading of "discussing what is epistemically rational" or "what's likely given the evidence" are well thought of as deliberation about what to believe when the aim is to be accurate.

¹⁰ See e.g. Dreier (2015) and Dasgupta (2017) for discussion of this point as applied to hardcore moral realism.

Next, I want to argue that *if* we use epistemic-rationality language to do roughly the work I sketched above, none of us should be endorsing impermissivism. (This leaves open the possibility that there are other ways of using language which vindicate an endorsement of impermissivism. We'll talk more about that later.)

The Challenge for the Impermissivist

Consider the proposition that it will rain in Honolulu 100 years from this moment. Call this proposition "H." Impermissivists like Sophie and Sinan will say the following sort of thing concerning H: "there is a uniquely rational credence to have in H given my evidence." But which attitude concerning the promotion of accuracy should I take them to be expressing when they make this claim? Perhaps it's something like this: "there is exactly one credal attitude towards H that I expect most accuracy from." But in fact, I will suggest, there is no such attitude. So if they're using rationality-talk in roughly the way I proposed - a way that requires that claims about rationality be expressions of our views about which attitudes we expect to be accurate, they wouldn't go around saying things like "there is a uniquely rational credence to have in H."

To see why, it will be helpful to draw a distinction between two ways of specifying a credal state. First, one can specify a credal state *numerically*. When you specify a credal state numerically, you describe it using an assignment of numbers to propositions (e.g. "the credal state that assigns 0.9 to rain this afternoon, 0.2 to snow this afternoon..." etc). Alternatively, one can specify a credal state *descriptively*. (e.g. "Joe's credences", "the rational credences given my evidence", "the omniscient credences", "the credences that result from conforming to the Principal Principle", ...). Somebody who has beliefs about a credence function specified *numerically* has no uncertainty about which numerical values different propositions are assigned, whereas somebody who has beliefs about a credal state specified *descriptively* may well have uncertainty about which numerical values different propositions are assigned.¹¹

With this terminology in hand, I'm going to defend the following two claims.

Claim 1: There is no *numerical specification* of a credal state, such that my opponents regard it as uniquely rational.

Claim 2: There is no *descriptive specification* – using exclusively naturalistic language – of a credal state such that my opponents regard it as uniquely rational.

¹¹ Thanks to Kevin Dorst for discussion on this point. This distinction also shows up in Schoenfield (2015) and Dorst (2019).

Defense of Claim 1:

Consider Sinan, and the attitude he has towards H: the proposition that it will rain in Honolulu next New Year's. Sinan's credence in H, I claim, is an *imprecise* credence, rather than a *precise* credence. In other (and more colorful) words it's a *mushy* credence rather than a *sharp* credence. Let me explain what I mean by that a bit.

When you have a precise credence in a proposition P, your attitude towards this proposition is well represented by a real number between 0 and 1 (for example, when somebody is, as we would say, "80% confident that P", their attitude towards P is represented by the number 0.8). A highly idealized agent will be well represented by a probability function, which assigns a precise number to every proposition the agent has an attitude towards. In reality though, none of us humans are well represented by such a function. Our doxastic attitudes are simply not fine grained enough. We're better (even if not perfectly) represented by a *set* of probability functions, called a "representor."¹² When we're representing an agent with a set of probability functions, we will often represent attitudes towards particular propositions by an interval rather than by a number. For example, if I talk about Sinan as having an imprecise credence of $[0.4, 0.7]$ in P, that's code for: the numbers r in $[0.4, 0.7]$ are all and only the numbers such that for some credence function f in Sinan's representor, $f(P) = r$.

Back to the main thread: Sinan, I claim, has imprecise credences. I don't make this claim because of some special feature of *Sinan's* psychology. Rather, given what it takes to be best represented by a precise credence function, I think *every* human being has imprecise credences in at least some propositions (see the references in note 12 for more on this point). Furthermore, if Sinan lacks a precise credence, there is no precise credence (specified numerically) that he expects to be most accurate.¹³ If there is no

¹² For example, agents who are "insensitive to evidential sweetening" will be better represented by a set of functions (Schoenfield (2012)) as will those with incomplete comparative confidence orderings (Konek (2019), Builes et al (2022)). See also Rinard (2017).

¹³ If there were such a credence, that would just be *his* credence, assuming the use of a strictly proper scoring rule. Let me explain what I mean by that: Strictly proper scoring rules are those according to which every probability function maximizes expected accuracy relative to itself. (This is one reason for their appeal: there is a thought that rational beliefs states should recommend themselves from an accuracy perspective.) Expected accuracy is always calculated relative to a probability function. So if Sinan were to regard a precise credence c as maximizing expected accuracy, this would require that his own credence function assigns c more expected accuracy than any alternative. Given strict propriety, this will only be true if his credence is c .

precise credence that Sinan expects most accuracy from, then, assuming that rationality talk expresses attitudes about promoting accuracy, he shouldn't say of any credence (specified numerically) that it is the uniquely rational one.

But perhaps the rationally required attitude is imprecise? Why can't Sinan regard his own *imprecise* credence as the uniquely rational one? The answer: given some plausible constraints on accuracy measures, *imprecise credal states don't uniquely self-recommend from an accuracy perspective*. A number of results in the literature show that if you're in an imprecise state, say one represented by $[0.4, 0.7]$ ¹⁴, there will be at least one doxastic state distinct from $[0.4, 0.7]$ that you expect to be no less accurate than your own.¹⁵ This means that insofar as Sinan's claims about what's rational express his views about what conduces to accuracy, he should not claim that $[0.4, 0.7]$, or any other imprecise state (numerically specified) for that matter, is uniquely rational given his evidence. Since any imprecise state leaves open multiple options from the point of view of accuracy, there can be no accuracy-based reason for an agent occupying that state to stay in it, as opposed to switching to one of the other alternatives, even if their evidence remains unchanged.

Defense of Claim 2:

Suppose Sophie comes along and says the following: "I'm not claiming that my own imprecise attitude towards H is uniquely rational - far from it. Indeed, I think that the uniquely rational doxastic state given any body of evidence is precise. After all, if it were imprecise then rationality and accuracy would come apart: Accuracy considerations would permit shifting states but it would be irrational to do so. However, I have no idea *which* precise credence I should adopt in the proposition that it will rain in Honolulu next New Year's. But so what? I'm not perfectly rational. I also have no idea how many gumballs are in that machine - surely you don't think that means that I can't believe that there *is* some precise number of gumballs in the machine!"

¹⁴ If we say that an agent's credence in P is $[x,y]$ what that means is that for all r in the interval $[x,y]$ there is some credence function in the set of credence functions that represent the agent that assigns r to P, and for no r' outside the interval $[x,y]$ is there a credence function in the set that assigns r' to P.

¹⁵ Seidenfeld et al (2012), Mayo-Wilson and Wheeler (2016), Schoenfield (2017), Berger and Das (2020) Builes et al (2022). Seidenfeld (2012) and Konek (forthcoming) describe accuracy measures for self-recommending imprecise states but these measures violate what I take to be extremely plausible constraints on accuracy measures (see the papers above for details). It's also worth noting that Konek's framework is itself motivated by a permissive thought: namely that there should be rational leeway when it comes to how we balance the values of 'pinning down the truth' and 'avoiding error.' His way of thinking about imprecise credences would sit awkwardly within an impermissivist picture.

Certainly I have no problem with Sophie thinking that there is a precise number of gumballs in the machine even though she can't specify that number "numerically." However, if we're using rationality-talk to express our attitudes about what conduces to accuracy, then it's much more difficult to make sense of Sophie's (purported) uncertainty about which credence it is rational to adopt in H than it is to make sense of her uncertainty about the number of gumballs.

To get a sense of why this is, let's contrast the case of Sophie's uncertainty about which credence it is rational to adopt in H with one of the examples I mentioned earlier. Suppose Sophie tells me that there is a uniquely rational credence to adopt in the proposition that I was dealt a straight, but she doesn't know what it is. In this case I can easily make sense of her uncertainty (by filling in some plausible assumptions about her). Sophie, I imagine, regards as accuracy-conducive an assignment of equal credence to each of the possible hands I might have been dealt (see "Third Remark" above). She's not sure, however, which credence in the proposition that I have a straight falls out of this distribution.

More generally, there are a variety of principles (specifiable without using normative language) that we might endorse concerning how to form beliefs when our aim is accuracy. We might, for example, have a preference for simpler theories, or inductive inferences, or priors that match our credences to our expectations of the objective chances. And we also might be uncertain about which credences conform to these general principles. However, given that Sophie is a finite human being, I don't see any reason to think that there is a set of such principles that Sophie endorses that is fine-grained enough to pin down a unique precise credence in H. (I provide further support for this claim in Schoenfield (2019)). It is for these reasons that I think that not only is there no *numerical* specification of a precise credal state that Sophie expects to be most accurate, there is also no *descriptive specification* – *using exclusively naturalistic language* – of a credal state that she regards as most accurate.

Putting Claims 1 and 2 together we get that there is no specification *of any kind* of a credal state, storable naturalistically, such that folks like Sophie and Sinan regard the adoption of that state as uniquely accuracy-conducive. Why is this important?

I want to be able to make sense, in naturalistic terms, of the attitude Sophie and Sinan are expressing when they claim 'rationality is impermissive'. But if we think of claims about rationality as expressions of our views about what promotes accuracy, and their attitudes, when described naturalistically, leave open multiple options from an accuracy perspective, it's quite difficult to make sense of what view about accuracy

promotion they are expressing with the sentence ‘rationality is impermissive.’ Given their own (naturalistically describable) views about what promotes accuracy we’d expect them to say instead that rationality is permissive.

Let’s take a step back: Ultimately, my goal is to convince you that there’s no deep disagreement between the accuracy oriented permissivists and impermissivists. I argued here that if we use rationality-talk to do a certain kind of job, we should all be endorsing permissivism. Sophie and Sinan, will, in the next section give us their own account of the role of rationality-talk. After hearing from them, we can investigate how deep the disagreement runs.

III. Sophie and Sinan: We aren’t mushy permissivists, and moreover, we shouldn’t be.

Miriam said that her main goal is to convince us that the disagreement between permissivists and impermissivists is not substantive. But she has started with a narrower argument: she has argued that, given a certain way of using our language about epistemic rationality, we should defend permissivism, not impermissivism. And given that this argument depends on psychological facts about us (as well as about you!), and what those features of our psychology commit us to, it has the consequence that we are *already* permissivists, albeit confused ones. We are intrigued by this possibility – when we said we agree on many things, we weren’t thinking that *permissivism* would be one of those things!

By way of recap, here is the challenge exactly as Miriam leveled it against us (quoted from above):

(The Narrow Challenge) Which attitude concerning the promotion of accuracy should I take them [Sophie and Sinan] to be expressing when they make this claim [of uniqueness]? Perhaps it’s something like this: “there is exactly one credal attitude towards H that I expect most accuracy from.” But in fact, I will suggest, there is no such attitude. So if they’re using rationality-talk in roughly the way I proposed - a way that requires that claims about rationality are just expressions of our views about which attitudes we expect to be accurate, they wouldn’t go around saying things like “there is a uniquely rational credence to have in H.”

And to recap the argument for this thesis: Miriam argues that since we have mushy credences, there is no single *numerically-specified* credence function which we take to be most accurate. (This is because from the perspective of a mushy state, there are some precise states which are no worse, accuracy-wise, than the mushy state itself; hence, there is no unique state that we'll regard as best in terms of accuracy.)

And she argues that there is no *descriptively-specified* credence function which we take to be most accurate – at least, if we limit ourselves to naturalistically kosher language. So, since we can't pin down a unique credence function that best aims at accuracy, we must be permissivists.

We begin our reply on a concessive note: yes, if we used rationality-talk just to express views about what we expect to be most accurate from our current perspective, then we couldn't and wouldn't say some credence in H is the uniquely rational credence, in *that* sense of "rational". But that concession includes a pretty big "*if*"!

What we want to do now is to elaborate our reply to Miriam by explaining how we use rationality-talk such that we can and should stand by our endorsement of uniqueness. (Though maybe this will play into her hands: she said earlier that she agrees there are impermissivist-friendly ways to use rationality-talk.) Miriam will hopefully be pleased with our reply, because the usage of rationality-talk that we're going to outline meets all of her criteria: (i) it functions to promote accuracy, (ii) we can explain in naturalistic (or at least non-epistemic) terms what attitudes our rationality-talk expresses in speakers and produces in hearers, and how our rationality-talk thereby promotes accuracy, and (iii) our rationality-talk promotes accuracy even when used by ordinary people who lack a precise credence and instead have only a mushy credence for many propositions.

So, in effect, while we can't answer what we labeled Miriam's "Narrow" Challenge, which has built into it that rationality-talk expresses judgments about expected accuracy, we can answer what we will label the "Broad" Challenge:

(The Broad Challenge) If we use rationality-talk to promote our aim of having accurate beliefs, and we want to give a naturalistic specification of how our rationality-talk promotes that aim – and promotes it even if we're mushy – how can we claim to be impermissivists?

The next subsection tries to answer this challenge.

After that, in the subsequent subsection, we'll try to turn the tables on Miriam and argue that her brand of permissivism faces some serious problems.

1. How is impermissivism compatible with naturalism?

Miriam wants us to explain, in naturalistic terms, how we understand our impermissive rationality-talk to promote the goal of accuracy. We cannot do it in the way she considered—we cannot *specify* what the uniquely rational credence is for each proposition. When our credences are sharp, perhaps then we can specify a uniquely rational credence (the one we expect to be most accurate, like Miriam proposes). But when we are mushy, we cannot specify a unique credence that we expect to be most accurate. In these cases, we will say that we are *ignorant* about what the (unique) rational credence is. This is in stark contrast to Miriam's suggestion that, in such mushy cases, we should call multiple credences rational.

(A side note, but an important one: Miriam argues we mushy people cannot even give a descriptive specification of the uniquely rational credence (in terms of epistemic principles that would pin it down, like some principle of indifference), and so long as she's right about that it will mean that we won't be able to know what the uniquely rational credence is in such cases—we'll remain ignorant of it. But we're not so pessimistic. Maybe we can't numerically specify that uniquely rational credence right now, and maybe we can't even descriptively specify the principles that pin it down right now, but we are optimistic that we might discover those principles and those uniquely rational credences eventually. Epistemology makes slow progress, and some significant principles took a long time to discover and still aren't satisfactorily formulated—eg, probabilism, the Principal Principle, the idea of induction as IBE, and so on. As this progress continues, we'll be able to better specify the uniquely rational credences in more and more cases.)

Okay, so, we can't pull off either of the strategies Miriam offered for giving a naturalistic, accuracy-oriented vindication of impermissivism. (That is: we can't explain how we use "rational" in terms of either a numerical or a descriptive specification of all rational credences.) But we have a different strategy. It is not a naturalistic specification of the uniquely rational credence function, but it is a naturalistic explanation of how our impermissive rationality-talk promotes accuracy. We presented it more thoroughly in our past work, but here we'll give it a fresh, new, explicitly expressivist presentation, in order to more directly meet Miriam's challenge while meeting her "anti-(hardcore-)realist" demands.

What attitude does rationality-talk express? Of course it expresses our beliefs about rationality, but we want (in addition) another kind of answer, the kind an expressivist would give. Like some expressivists in metaethics say, we say the language

of epistemic rationality expresses mental states of *planning* and *recommending*.¹⁶ When you praise as rational a belief that's based on certain evidence, you're both planning to form that belief if you should have that evidence, and you're recommending that belief to others if they should have that evidence (we always mean your *total* evidence). When you criticize as irrational a belief that's based on certain evidence, you're planning to not form it, and recommending against forming it, on that evidence. And when you endorse impermissivism, you're doing something more complicated, the details of which will turn on how we solve the Frege-Geach problem, but one rough idea is that you're ruling out having permissive plans (which may just be ruling out the incoherent state having two alternative plans), and you're ruling out giving permissive recommendations, for what to believe (on any evidence).

Given this expressivist model, our epistemically evaluative practice serves a clear function. On this model, our epistemic evaluations promote a kind of coordination in our belief-forming methods: our evaluations serve to recommend that others form the same beliefs that you plan to adopt yourself, given the evidence. The beneficial effect of this is that you are recommending beliefs that will lead to testimony that is as trustworthy as your own beliefs are. This is tremendously valuable from an accuracy-oriented perspective. Our epistemically evaluative practice helps reinforce an efficient way of searching for the truth, and avoiding falsity, as a community of inquirers. We need to rely on the collection and assessment of evidence, and the resulting testimony from other people, in order to have any kind of informed opinion on most topics. But how can you safely trust the testimony of other people, especially when they usually cannot store and share all their evidence with you, and we do not know most of these people very well? The answer is that you can safely trust testimony from people who form beliefs the same way you do. This is an efficient way to establish safety; it doesn't require that people retain and share all their evidence with each other, and it spares us the costs of vetting our testifiers' track records through difficult or impossible investigations into what they've said in the past and what we've independently confirmed the truth of. Thus, by using epistemic language to plan for and to recommend the same beliefs (given the evidence), we help promote and reinforce an efficient way of making testimony safe to trust. It's a defeasible trust of course, easily and often broken, but it's still the default status we grant even to strangers on most topics.

This is our naturalistic, or expressivist, explanation of how we use epistemically evaluative language and attitudes to serve our goal of accuracy. We haven't specified the

¹⁶ The idea that evaluations express recommendation, or approval, is a traditional expressivist idea; see van Roojen (2018). Gibbard (2003) develops the planning idea. Miriam herself is a big advocate of the planning view; see Schoenfield (2018). We like the presentation of the idea that Greco & Hedden (2016) gave, which partly builds on the proposal of Dogramaci (2012).

rational credence in all, or even many, circumstances. But the whole story *requires* that impermissivism be true. This is for two reasons, one to do with the planning aspect and one to do with the recommending aspect of our evaluative practice. First, as Greco and Hedden (2016) argued, you can't coherently hold both of two conflicting plans about what to believe given certain evidence. But that seems to be what the permissivist would be committed to, at least if we're right that to call a belief rational is to express a plan to hold that belief. Second, as we (2016) and Greco and Hedden also argued, if you recommend two different beliefs on the same evidence, you're not going to be generally promoting the most accurate beliefs and the most trustworthy testimony. The permissivist who recommends forecasting a 40% chance of rain and *also* recommends forecasting a 60% chance of rain, given the same evidence, cannot be making the most helpful recommendation if our goal is accuracy.

Now, we imagine Miriam will remind us at this point: we are all mushy! When it comes to most things, including many of our predictions about rain, we mushy people won't view any particular forecast as uniquely best in terms of accuracy. For some topics, like rain in Honolulu in 100 years, one forecast might look no less accurate, from our ignorant mushy perspective, than a very different forecast. So, why should we have any reservations about either adopting a 0.4 credence in rain or adopting a 0.6 credence in rain (at least if both probabilities were in the interval of our mushy representor)?

Admittedly, we cannot plan to do both, but we can leave both options open in our plans, and perhaps we can happily recommend either one, whether we're giving recommendations to other people or to our future selves. We mushy people ought to relax and give ourselves permission to believe in any one of these ways, Miriam might say.

But we are not convinced that it's wise to recommend multiple beliefs, or permit ourselves to adopt any one of multiple beliefs, even when we are in a mushy state. We admit that sometimes—okay, a lot of the time—we have a mushy credence in something or other because we are ignorant about which credence is uniquely rational. Life is like one big non-stop probability exam: you often know your given evidence determines some unique precise credence as the right one to assign to some event, but, under the time pressure and maybe the psychological pressure, you can't immediately say what it is. But there's no need to panic. As we'll explain in the next section, we mushy impermissivists are comfortable being ignorant, and thus staying silent, about what's rational.

We do take those credences we recommend as rational to maximize accuracy, *but* we do not recommend credences just because we can't identify anything better, accuracy-wise—not if we're mushy and so there are multiple credences that meet this criterion. (We'll return to these issues again below.)

2. Some puzzles for Mushy Permissivism

So far we've argued we can think and talk about rationality in a way that supports impermissivism, while also respecting the demands of naturalism and accuracy promotion. In this section we will more directly address Miriam's positive proposal: her suggestion for how we should think and talk about rationality, which would commit anyone who's mushy to permissivism. We think that this form of permissivism has some strange and interesting consequences, which we take to be a mark *against* thinking and talking about rationality in the way Miriam suggests.

Recall Miriam's argument for Mushy Permissivism. An important premise is that when you are mushy about P, you won't regard any doxastic state about P – mushy or precise – as uniquely maximizing expected accuracy. Instead, given one popular view about accuracy, you'll regard your mushy state as “on a par” with all of the precise credences in its representor: none of these are, from your perspective, better or worse than one another in terms of accuracy. (On some other views, there is instead a *single* precise credence that you regard as better or worse than your own, but we will focus on the first alternative.¹⁷) Then, stemming from Miriam's suggestion that we use rationality-talk to make claims about accuracy, it follows that you will call all of these states (the mushy state and the precise members) rational, or rationally permissible. It's worth noting that on this form of permissivism, permission comes cheap. You don't need a specific positive endorsement of these attitudes; rather, they just need not to be ruled out.

On this view, since you regard all of these attitudes as permissible – and since by “permissible” you just mean something like “no worse than any of my other options, in terms of accuracy” – you have no accuracy-based reason not to switch from one to another. In particular, you'll regard it as just fine to switch from your mushy state to one of the precise states in your representor.

You might want to resist this consequence, but since we are getting personal here, we'll note that Miriam and Sophie have both endorsed it in print;¹⁸ so, we'll go along with it for present purposes. The idea that “sharpening” or “contracting” (that is, switching from a mushy credence to one of the precise credences in the mushy state's representor) can be okay has some points in its favor.

If your credence in P is mushy, and you regard the precise members of your state as no less accurate than the mushy state, then in an important sense you don't disagree with them. When you sharpen, you (arguably) aren't “changing your mind” in a

¹⁷ For further discussion, see Seidenfeld et al (2012), Mayo-Wilson and Wheeler (2016), Schoenfield (2017), and Berger and Das (2020).

¹⁸ See Schoenfield (forthcoming) and Builes, Horowitz and Schoenfield (2022).

problematic way, adopting different beliefs which you once thought were *wrong*. You didn't think the beliefs you're adopting were wrong *about P*, anyway. In that way, your change of mind about P doesn't seem problematic in the way that other paradigmatic, arbitrary changes of belief might seem problematic.

But in this case, since we are (if we follow Miriam's suggestion) using rationality-talk to express our accuracy judgments, sharpening carries with it some additional odd consequences. In particular, sharpening means that not only can you go from mushy to precise *about P* with no new evidence – but you can also go *from permissivist to impermissivist*. Here is how.

Capricious permissivism

Step 1: Focus on some proposition P, about which you have a mushy credence represented by, for example, the interval [0.4, 0.7]. Notice that this is a permissive case: your mushy state, as well as all of the precise credences in P within its representor, are all permissible.

Step 2: Sharpen your doxastic state, adopting credence 0.6 in P, just for fun.

Step 3: Again consider the question of whether you're in a permissive situation. Notice that just one doxastic state maximizes expected accuracy, given your evidence – namely, your own credence, 0.6.

Step 4: Conclude that you are, in fact, in an impermissive case.

This seems odd to us. You are able to permissibly change your belief state about P *and about permissivism* without acquiring any new evidence, or even making any changes in your belief state which you regard as an improvement. This means that the mushy permissivism supported by Miriam's argument is an unstable sort of permissivism. It's something that we might happen to endorse because we are mushy, but not something we are committed to.

Before you go through these steps, you can see that they are all permissible. How should you regard your future self, as you consider whether to make these changes in your beliefs? As we said before, it seems that you should *not* think that you have a disagreement with your future self regarding P. (After all, you regard your future self's attitude towards P as neither more nor less accurate than your own.) But you *do* disagree with her regarding permissivism – she has an attitude about permissivism that you now

regard as incorrect, or inaccurate. Does this give you any reason to avoid moving from Step 1 to Step 2? If so, how?

Here the Mushy Permissivist might reply: it's fine to go through those four steps and become an impermissivist, but that's because your evidence changes along the way. Your judgments about permissivism regarding P, given E, flow from your attitudes towards P itself. And so of course if those attitudes change, your views about permissivism might change as well. Compare this situation to the following:

Capricious self-reflection

Step 1: Focus on some proposition P, about which you have a mushy credence represented by the interval $[0.4, 0.7]$. Notice that you have a mushy credence.

Step 2: Sharpen your doxastic state, adopting credence 0.6 in P, just for fun.

Step 3: Reflect on your doxastic state and conclude that you have a sharp credence.

Presumably, if there is no problem moving to Step 2, then there is no problem moving to Step 3. And Step 3 doesn't seem to raise any new problems. Why, the Mushy Permissivist might ask us, should the analogous situation seem any worse if someone changes her mind about permissivism by capriciously sharpening?

Or the Mushy Permissivist might also bring up an example like the following:

Capricious tastes

Step 1: Using your own food preferences to judge, notice that broccoli is disgusting.

Step 2: Cultivate a love of broccoli over several months, experimenting with different recipes and habituating yourself to its taste.

Step 3: Reflect once again and notice that now broccoli is tasty.

There is nothing wrong with this case either. So, the Mushy Permissivist could say, we should think of capricious rationality judgments as akin to capricious self-reflection, or capricious tastes – cases in which we (arguably) do not expect stability.

We think this is probably the right way to think of rationality, for the Mushy Permissivist. But a practice of rationality ascriptions is more useful if it is more stable, if

we see it as aiming for a fixed standard rather than something liable to change on a whim. To illustrate this further, let's look more at how Mushy Permissivism fares over time. Return to the idea that the purpose of rationality ascriptions is to cultivate reliable testifiers, worthy of deference. If we can permissibly sharpen our credences (and at the same time, change our views about permissivism and impermissivism), it seems to us that rationality ascriptions become much worse at serving this purpose.

Consider the following story. The purpose of this story will be to compare two attitudes towards rationality that we might have when we have mushy credences: Mushy Permissivism, following Miriam's argument, and impermissivism (filled out in a certain natural way):

A Tale of Two Rational Advisors: On two snowy mountaintops there are two rational advisors. The advisors are much alike, with similar early life experiences, similar libraries, and similar accuracy oriented values, but one of them is a permissivist and the other one is an impermissivist. Each has their own group of devoted disciples, who wander the valley searching for evidence and forming opinions, which they record in special notebooks. Every so often they come back to report on their progress to the advisors. The advisors check their disciples' notebooks and judge whether the responses are rational. They reward rational credences with a gold star, and mark irrational responses with a red X.

In some cases the advisors' behavior will be the same: they will both hand out a gold star to disciples who defer to objective chance, perhaps. But in other cases the responses will differ. Specifically, they will differ in cases where the advisors themselves have mushy credences. When the *permissivist* advisor is mushy, she will judge many responses to be rational: her own mushy state, as well as any of its precise members. She will hand out many gold stars. But when the *impermissivist* advisor is mushy – given what we consider a plausible interpretation of mushiness and rationality judgements – she will be *uncertain* as to which responses are rational.¹⁹ She will not hand out any gold stars in such cases, though she may still hand out many red Xs for responses she knows to be irrational. Sometimes – when her disciples have sharp credences which fall within the advisor's representor – she will simply leave the ledger blank. (What will she do if both she *and* the disciple are mushy? We'll let Miriam raise that question in her next response.)

¹⁹ We haven't argued directly for this position, but see Carr (2020) for an extended discussion of a very similar view. Carr argues that situations in which some epistemologists are inclined to argue that mushy credences are rationally required are better thought of as situations in which we ought to be uncertain about which (precise) state is rational.

So far these look like two coherent possibilities. But since Mushy Permissivism permits sharpening, we can see that the permissive practice begins to look stranger as time goes by:

A Tale of Two Advisors, Continued: After years of study, one of the *impermissivist's* former disciples returns to her. "Impermissivist Advisor," he says, "under your tutelage, I have honed my belief-forming methods. I have shunned those misguided methods that earned me a red X. I've learned to guide my thinking by those methods that earned a gold star from you. And when you were silent, as in those moments when I sensed a mushiness within you, Advisor, I took no lesson to heart about which belief forming methods to avoid or embrace." The Advisor smiles with pride at her disciple. The disciple continues: "Since I completed my training and left our mountaintop, I have been collecting evidence on the possibility of life on other planets. I gathered all the scientific evidence there is, which there was so much of, and I adjusted my beliefs on its basis. But, Advisor, I fear I will disappoint you: there was so much evidence to consider that I have now forgotten almost all of it." The Advisor gives a forgiving nod to the disciple and says, "Don't worry a bit about that. Just tell me," she says with burning curiosity, "Do you think it's likely there's life on other planets?"

Over on the permissivist's mountain, a similar scene is playing out, though with an important difference. One of the permissivist's former disciples, who honed his belief forming methods to embrace those that earned gold stars and to abandon those that earned red Xs, has returned to the permissivist mountaintop. "Permissivist Advisor," he says, "I have been collecting evidence on the possibility of life on other planets. I have gathered all the scientific evidence there is, and adjusted my beliefs on its basis along the way. Each time, I used belief forming methods that you gave me a gold star for using in the past. Would you like to know what I have discovered?"

The permissivist advisor is proud of her disciple (though again, she must look past the disciple's inability to recall the actual evidence concerning life on other planets) and she is about to ask for his testimony... but then she remembers. In the past, she was mushy on those topics that involved the kind of interpretation of noisy data that the search for extraterrestrial life requires. Suppose this disciple trained with her at that time and asked for her blessing to adopt sharp members of the Advisor's representors. She would have given him encouraging gold stars. And what if, subsequently, *she* had gone on to adopt *different* sharp credences on these matters? Indeed, this is what has happened to her: her credences that are based on

interpretations of noisy astronomical data are *now* sharp... . Although the disciple's transcript is full of gold stars granted by the permissivist advisor herself, and though she trusts he has collected his information only using those methods which earned gold stars, she cannot now trust her disciple's testimony.

Storytime over. What have we learned? It looks like there is a real difference here between the impermissivist advisor and the permissivist advisor. For the impermissivist, a gold star is rarer and worth more – it means that the method that produced that belief was deemed reliable. (In Bayesian terms, it means that the gold starred credence in P accurately reflects the probability of P given the evidence the disciple then had.) For the impermissivist, gold stars cultivate trustworthy beliefs. For the permissivist, though, trusting a gold starred belief can end up being a dangerous gamble. The permissivist advisor can't trust her own disciples, *even if her own evidence has not improved and the disciple has gained new evidence*.

This gives an answer to the worry that remained at the end of the previous subsection, after we summarized our general argument that impermissivists use epistemic evaluations to efficiently cultivate trustworthy testifiers. We were left there with a worry about how impermissive evaluations would help serve this purpose any better than permissive evaluations in a case where you are mushy. Our story here shows how the impermissive advisor emerges as the wiser one for giving out no gold stars when mushy – that is, for not praising any specific attitude as rational. The permissive advisor, who recommends multiple options when mushy, is in the oddly precarious position of potentially transforming into an impermissivist regarding some matter, and then regretting the advice she gave when she was permissive. While in a mushy state, she was happy to recommend several different precise states, and indeed adopt any of them – at the time there seemed to be no reason not to. But once we are precise, our opinions change, along with our recommendations. This result looks odd, and it renders the permissive practice of rationality evaluation considerably less useful for the purposes of promoting trustworthy testimony.

We started off this conversation with Miriam accusing us of being secret permissivists. But now we wonder whether it isn't the other way around. If a permissivist is truly willing to switch arbitrarily from a mushy credence to a precise one – and from permissivism to impermissivism – what kind of permissivism is that? Moreover, we have argued that this kind of permissivism has serious problems.

IV. Miriam: Is the Disagreement Between us Substantive?

1. Verbal Debates in Epistemology

I claimed that if we use rationality-talk to do a certain kind of work (express our attitudes about which ways of forming beliefs conduce to accuracy), Sophie and Sinan shouldn't claim to be impermissivists. Sophie and Sinan then offered us a different way of thinking about the role of rationality-talk: rationality-talk, on their view, serves the role of helping us coordinate our belief formation, at a social level (Dogramaci (2012), Dogramaci and Horowitz (2016)). More specifically, the idea is that claims about what is and isn't rational will encourage others to form beliefs in ways that will make them accurate testifiers, thereby resulting in more accuracy for everybody. Furthermore, they argue, if rationality talk has this coordinative function, endorsing impermissivism makes a lot of sense. In endorsing impermissivism, Sophie and Sinan are (perhaps amongst other things) expressing a commitment to avoid recommending multiple doxastic options in response to a given body of evidence. They argue that if we want accurate testifiers in our community, we *should* avoid recommending multiple doxastic options in response to a single body of evidence.

The first point I want to make is this: if it turns out that one way of using rationality-talk supports permissivism, and another supports impermissivism, that is a happy result. The disagreement between us, it turns out, lies somewhere on the spectrum between merely verbal (we're not really disagreeing at all) and metalinguistic (we're having a substantive disagreement about the best way to use certain bits of language). If the debate is merely verbal we can stop disagreeing (a happy result because: philosophical progress!). If it is metalinguistic, well, it's not clear that it's *our* job - as epistemologists - to figure out which ways of using language are best (a happy result because: we hand off the work to somebody else and go on vacation!).

To make this thought more concrete, let's consider an example in which the question of whether permissivism is true might come up in deliberation about what to do or think. I'll argue that as long as we stay clear-eyed about how we're using language, we will all settle our deliberations in the same way. Consider:

The Move: You currently have an imprecise credence in P of [0.4, 0.7]. Your aim is to end up accurate with respect to P. Much to your surprise, you learn that the weather has a significant impact on people's opinions concerning P. Living in colder climates, you've been told, makes people more confident in P than $\sim P$ (but

never more than 0.7 confident), whereas living in warmer climates makes people less confident in P than in $\sim P$ (but never less than 0.4 confident). You currently live in Seattle but are considering moving to either Massachusetts or Texas. How worried should you be about the anticipated change to your belief? Does your concern with accuracy give you any reason to stay in Seattle?

I don't think so. Builes et al (2022) and Schoenfield (forthcoming) argue that in cases such as this one we needn't worry about certain arbitrary shifts in our belief (assuming our aim is accuracy), because in many cases we start out imprecise.²⁰ If we're imprecise, certain arbitrary shifts will look harmless from an accuracy perspective. So if the rain has got you down, and you're craving New England charm or superb tacos, you might as well move. If I were in this situation, I might express my views by saying "there are multiple rational attitudes to take towards P." In saying those words all I'd be doing is giving voice to the idea that from an accuracy perspective, several doxastic options are on a par.

Now suppose that Sophie and Sinan were in this situation. They probably wouldn't express their attitudes about the multiple options left open to them accuracy-wise with respect to P using the sentence "there are multiple rational attitudes to take towards P." Whatever they end up *doing*, they wouldn't end up *saying* things like the ones I did. But let's focus on the doing for now. (The difference in saying, might, as I suggested earlier, end up being a merely verbal matter). Would Sophie or Sinan have any accuracy-related qualms about leaving Seattle? If their concern is their own accuracy with respect to P - I don't see how they could. If the case is set up right, they won't expect that moving to Texas, for example, will make them any less accurate about P.

Furthermore, given that permissivists and impermissivists (may well) have the same credences in first order matters, it's unclear why they would settle *any* accuracy-aimed deliberative questions differently from one another. (Note that this also holds true for deliberative questions about who to trust or defer to. If a permissivist and an impermissivist have the same first order views, they'll presumably have the same views about who is likely to be testifying truly.²¹) I think that for this reason, when it

²⁰ For similar reasons, it's also true that in many cases, we needn't worry if we learn that some of our beliefs were formed *in the past* as a result of arbitrary influences.

²¹ They also won't have different views about *the extent to which* to deliberate. To see why, suppose a permissivist and impermissivist both have an attitude well represented by [0.4, 0.7] towards the proposition that it will rain in Honolulu next New Year's (call this "H"). The impermissivist might say "I acknowledge that my imprecise credence is irrational, but I'm not just going to give up and leave it at that, or let my belief move around arbitrarily! Instead I'm going to think really hard about H, talk about H with my wise friends who share the same evidence, and so forth. Hopefully by engaging in these processes I'll end up with a (more?) rational opinion about H." Will a commitment to further reflection on H along these lines distinguish the impermissivist from the permissivist? No. For if the impermissivist accepts a close connection between rationality and truth, then the impermissivist must be hoping that engaging in this process will not only result in their opinion being more rational, but also in it being more *accurate*. Perhaps, for example, they think that upon reflection they might discover some new consideration that is relevant to

comes to what we actually *do* or *think* (about first order matters) as opposed to what we *say* (about rationality), we shouldn't expect much in the way of difference between the permissivist and the impermissivist.

The fact that we have no reason to expect the permissivist and impermissivist to have different first order opinions is arguably some evidence that the debate is verbal. Daniel Greco (2015), for example, suggests that worries about mere-verbality arise in epistemology when differences in opinions concerning what it's rational to think about P, don't bear on our opinions concerning P itself.²² This is in contrast to debates that Greco regards as clearly substantive: The skeptic and the anti-skeptic, will (unless they're akratic) have different opinions about the existence of hands, the steadfast and conciliationist will have different credences about political matters, those convinced by fine-tuning arguments and their opponents will have different views about the existence of a designer, halfers and thirders will have different views about coin tosses etc... But, so far, we have been given no reason to think the permissivist and impermissivist disagree on anything in the first order domain.

Greco talks about debates being “merely verbal” - where “merely” is clearly being used in a dismissive fashion. But as I mentioned at the beginning of this section, we could also have substantive debates about the best way to use bits of language, and Sophie and Sinan's arguments could be seen as advocating a practice of using rationality-talk for a different job from the one I proposed. They say: “a practice of rationality ascriptions is more useful if it is more stable, if we see it as aiming for a fixed standard rather than something liable to change on a whim.” And to argue for this they appeal to rationality-ascriptions' role in deference and testimony. So let's turn to this next.

2. A Response to S&S's Argument for the Value of Using Rationality-Talk Impermissively

H that would narrow the credal range, in much the same way that discovering some new piece of evidence concerning H might narrow the credal range. But, again, insofar as the permissivist and impermissivist share first order opinions, they'll have the same opinions about whether or not the reflective process would result in a change of credence, and whether the process of reflection is accuracy conducive. Any hope that the impermissivist has of narrowing their opinions through reflection can be shared by the permissivist. But if I'm right about the psychology of the impermissivist, then, for at least some propositions P, there isn't *any* amount of reflection that could (in a non-arbitrary way) lead them to adopt a precise credence in P given their evidence (see Schoenfield (2019)). The impermissivist who agrees with me about their own psychology will agree with me on this point. There will be no difference, then, in terms of how much reflection the permissivists and impermissivists will choose to undertake.

²² Although Greco raises considerations in favor of this view, and says he is sympathetic to it, he also says that he does not take himself to be offering necessary and sufficient conditions for when a debate is merely verbal. He does, however, think that showing that a debate about the rationality of believing P *does* bear on P (for some first order proposition P) is sufficient to show that the debate is not merely verbal.

Let me first make a concessive remark. “Rational” is a word in the English language, and for all I’ve said it may be that, given the history of its use, the world would be a better and more accurate place if we used this particular word impermissively. I insist that it’s useful to have *some* word that expresses our attitudes about which belief-forming methods conduce to accuracy, but I have no special attachment to it being the word “rational.” If Sophie and Sinan are right, that using “rational” impermissively leads to better results, I am (in the spirit of Carnapian explication) happy to use “rationality” to do the work that I think will bear on deliberative questions like the one about the move from Seattle, and use “rational” to do something else.²³

But on a less concessive note, I’m not entirely convinced by Sophie and Sinan’s arguments that an impermissive practice will be more socially effective than a permissive one. To see this, let’s consider their story about the permissivist and the impermissivist advisors. Here is how the impermissivist advisor operates:

- When the impermissivist advisor is precise, she gives out a gold star to disciples who share her precise credence. (Gold star = ascribing “rational” to the disciple’s attitude).
- When the impermissivist advisor’s attitude is best represented by the interval $[a,b]$, she gives out a red X to any disciples who have a credal attitude that includes credences that are *outside* of the interval $[a,b]$. (Red x = ascribing “irrational” to the disciple’s attitude).²⁴
- When the impermissivist advisor’s attitude is best represented by the interval $[a,b]$, she is silent when approached with a disciple whose credal attitude is a subinterval of $[a,b]$. (silent = ascribing neither “rational” nor “irrational.”).

So the impermissivist advisor has three different signals on offer: gold star (= ascribe “rational”), red x (= ascribe “irrational”), silence (= ascribe nothing).

The first point I want to make about this is that the permissivist also has three signals on offer, and she could use them in a way that corresponds exactly to the gold stars, red x’s, and silence *if she wished to*. Here is the conversion scheme:

- Whenever the impermissivist awards a gold star (ascribes “rational”), the permissivist also awards a gold star. The phrase she uses that is associated with gold stars is “rationally required.”

²³ Indeed, this is precisely what I do in Schoenfield (2019) so as to stay away from verbal disputes.

²⁴ This is intended to also cover a case where the advisor’s attitude is precise, in which case their credence c will be represented by the “interval” $[c,c]$. In this case, they will give out red X-s, to any credal attitude other than a precise credence of c .

- Whenever the impermissivist doles out a red X (ascribes “irrational”), the permissivist also doles out a red X. The phrase she uses that is associated with red Xs is “irrational.”
- When the impermissivist reacts with silence, the permissivist reacts with the phrase “rationally permitted but not required.”

The success of the impermissivists’ society’s practice depends on how people react to the three possible ways the advisor reacts. But however the impermissivist thinks the disciples will react to the stars, X-s, and silences, can, in principle, be mimicked in the permissivist society. There’s no reason why these societies have to operate any differently.

To be clear, it’s not obvious to me that the best practice for promoting accuracy amongst the mushy *does* involve three signals (required/forbidden/silence) rather than two (permissible/impermissible). For in the Tale of Two Advisors, when the impermissivist advisor is imprecise, they still (if they’re pursuing accuracy) might end up contracting arbitrarily to a precise state. (Recall- we don’t expect any differences in first order opinions between the permissivist and impermissivist advisors.) The impermissivist will differ from the permissivist in that the impermissivist will not regard their own contraction as *rational*, but that doesn’t mean we shouldn’t expect them to do it. After all, they *also* don’t regard their own mushy state as rational. Since neither their views about rationality nor their views about what conduces to accuracy motivate remaining imprecise rather than contracting, I see no reason why impermissivist advisors would be any less prone to contraction than permissivist ones. The impermissivist disciple (having received no guidance because the impermissivist advisor initially responded with silence) may also do some contracting. (Once again - staying mushy won’t look any better to the disciple than contracting.) And, as in the permissivist society, there is no guarantee that the advisor and disciple will contract in the exact same fashion. This means that the situation that Sophie and Sinan regard as problematic (the disciple contracting to a different state than the advisor) could happen in the impermissivist society as well as the permissivist one. Nonetheless, the main point I want to emphasize is just that *if* it’s true that a three-signal practice is best, those three signals are available regardless of whether we speak permissively or impermissively.

There is, however, one feature of the impermissivist’s practice that puzzles me. The impermissivist under discussion thinks that *all* imprecise credences are irrational. (Recall: imprecise credences don’t uniquely recommend themselves from an accuracy perspective, so they can’t be uniquely rational). So if the impermissivist advisor’s credence is imprecise, then, when their disciple adopts the very same imprecise credence

that they have, it's unclear why they react with silence, rather than calling the attitude "irrational." But if the practice were revised so that *all* imprecise credences were given red-Xs (both those that overlap with the advisor's and those that don't), then it's unclear to me that there would be much use in this practice. With the exception of propositions about coins and lotteries, and the lucky instances when our phone apps give us precise numbers concerning the probability of rain, nearly *all* of our credences are imprecise. So while the permissivist advisor will distinguish between imprecise credences she regards as no less accurate than her own, and imprecise credences she regards as less accurate than her own (the former she will say are rationally permissible, the latter she will say are irrational), it seems as if the impermissivist is going to be disapproving of nearly every attitude any of us have, and that doesn't seem, on its face, like a practice that is discriminating enough to promote accuracy in a useful way.

3. Deliberation and Evaluation

One thing that I hope has come out of the discussion so far is that talk about rationality can play a variety of different roles. Sometimes discussions of what's rational (or reasonable, or likely) takes place in deliberative contexts: in deciding how to go about forming and managing beliefs, collecting evidence, deferring to experts and so forth. Other times, discussions about what's rational (or reasonable, or likely) take place in evaluative contexts: we're looking at the beliefs of others and making claims about how (along some dimension) "good" or "bad" they are.

When it comes to deliberative contexts, there is no reason, as far as I can tell, for Sophie, Sinan or I to go about things differently when our aim is to end up accurate. We might *say* different things about what's rational, but no such differences will cash out in differences in thought and behavior about first order matters.

When it comes to evaluative contexts, it's arguable that part of what we're interested in is, precisely, *what to say*. I've raised some doubts concerning the claim that it's better, in these contexts, to speak impermissively: First, because the permissivist and impermissivist both have three responses available: (either: rational, irrational, silence, or: rationally required, irrational, (merely) rationally permitted). There is no reason they can't, in principle, distribute these responses to train their disciples in exactly the same way. My second concern was that it seems like the impermissivist should in fact be calling nearly all of our attitudes irrational (since they are imprecise) and this doesn't seem like a good way to promote accuracy.

Given all this, here's what we may end up still disagreeing about:

- Which of two different roles we want the word “rational” to play, if these roles turn out to come apart (Do we want to use the word primarily in deliberative contexts - where it will be a tool to express our attitudes about which ways of forming beliefs we expect to be accurate? Or do we want to use it primarily in evaluative contexts - where it will be a tool to exert social pressure on others to believe in ways that will make them accurate testifiers?)
- Which of two linguistic practices will better serve certain aims

The first, I think, is a merely verbal matter that we shouldn't put too much time into. The second is not merely verbal - but metalinguistic, in the sense that it's a debate about which ways of speaking will lead to better results. Empirical considerations may play a significant role in settling this question, so it's not clear how productive it is to continue debating from the armchair.

V. Sophie and Sinan: A final plea for impermissivism

Let's step back and see where we are. Miriam first introduced a strategy for using our thought and talk about rationality, and argued that given our own mushy credences, it supports permissivism. We then argued for an alternative, which supports impermissivism, and also argued more directly against the version of permissivism that Miriam's present suggestion leads to. Now Miriam raises two challenges for us:

- (1) Miriam predicts that, no matter how we use our language, we impermissivists will act and deliberate exactly the same as we would if we were permissivists. So the difference between the views is, arguably, merely verbal.
- (2) Our objection to permissivism focused on a kind of case in which the impermissivist and permissivist speak differently: one where the advisor (that is, the person doling out the rationality attributions) has mushy credences, and is evaluating sharp credences within her representor. We argued that the impermissivist would be silent in these cases, but the permissivist would award a gold star. This would (we argued) devalue permissive rationality. But Miriam argues that the permissivist could make more fine-grained attributions, differentiating between rational permission and requirement, thereby maintaining the high value of a gold star. Miriam also raises a question about what the impermissivist will do when evaluating mushy states.

We are happy to (mostly) grant Miriam's first point; deliberation and action will look mostly the same for our impermissivist and her permissivist. The differences will mostly show up in our recommendations, and in how we describe and understand our changes in belief. If an impermissivist contracts from a mushy state to a precise one, she won't (like the permissivist) describe herself as doing something rational and give her blessing to others to do the same. Instead, she'll just shrug and say something like, "I don't know if this is right, but I can't think of anything better to do." Of course, though, all of that is verbal behavior.

As Miriam says, this debate is *merely* verbal if language use doesn't matter. But maybe it does. So let's now turn to looking at the language in more detail.

We had argued that the permissivist practice, with its profligate awarding of gold stars, ends up with a notion of rationality that's not worth caring much about. But in the last section, Miriam argued that the permissivist practice need not devalue rationality. The reason, she argues, is that the permissivist has the resources to effectively mimic the impermissivist: where our impermissivist advisor uses "rational", she uses "rationally required", and where our impermissivist advisor is ignorant and makes no evaluation at all, she uses "rationally permitted". In that way, the permissivist can make the same distinctions as the impermissivist.

Miriam argues that this possible mimicry provides more evidence that the dispute is merely verbal. But we think that insofar as the permissivist *wants* to mimic us, this is actually an important concession: it shows the importance of having an impermissive epistemically evaluative concept. (Imitation is the sincerest form of flattery, etc.) Now the question remains: is it *also* important to have a permissive mode of evaluation? That is: is there any benefit to the advisor saying to her disciple, "all of these belief states are rationally permissible" rather than "I do not know which one of these belief states is (the one that's uniquely) rationally required"? Does this permissive norm serve any special role or purpose? We don't see one. We think it makes more sense to have just one normative notion which serves a real purpose (accompanied by admissions of ignorance in certain cases) than two – one of which serves a real purpose and the other of which does not. The full case for permissivism should advance a positive role for the permissive mode of normative evaluation alongside the impermissive one.

We'll close by considering one last important question Miriam raised: how does our view pan out in cases where both the advisor (that is, the person making rationality judgments) and the disciple (the person she is judging) are both mushy? We said that when the advisor is mushy, we should think of her as *uncertain* which precise state is rational. And when she is uncertain what's rational, she won't give out any gold stars. But Miriam points out: shouldn't she at least give out a red X to anyone else who is mushy,

including disciples who share her own mushy state? And if she does, wouldn't this mean giving out too many red X's? (We said that the impermissive advisor would be silent when she is mushy; but if the advisor considers all mushy states to be irrational, Miriam points out, she should not be silent.)

Sticking to our strategy – using rationality judgments to promote trustworthy testimony – there is an interesting empirical question about which strategy would serve our purposes better. According to this strategy, a gold star serves to *encourage* a certain response to evidence, and a red X serves to *discourage* it. And Miriam's challenge opens up some interesting questions about how *irrationality* judgments might be more complicated than we had at first let on. If we want to influence someone else's behavior, it only makes sense to tell them what *not* to do if we think that their following this advice wouldn't result in them doing something *worse*. (You shouldn't outlaw alcohol if everyone will just end up poisoning themselves with moonshine.) And arguably, if we are telling someone *not* to do X, we should also be able to provide a better alternative. (Telling a toddler to pet the cat gently is better than telling her not to hit the cat – her second choice, after hitting, might be pulling the cat's tail.)

This line of thought suggests that if the advisor is unsure what's rational, perhaps she shouldn't be giving out any red X's. If she gives a red X to someone with a mushy credence, when she thinks the uniquely rational precise credence is within the disciple's representor, she runs the risk that the disciple will adopt a credence *outside* the range, which is arguably worse. We'd have to fill in a lot of details to say how or why it would be worse; we will not pursue that line of questioning here. There are a lot of other, related questions we won't pursue either – how different, more fine-grained social practices might encourage people to revise their beliefs in this or that way, how different audiences might react differently to evaluative language, and so forth.

Having opened up this can of worms, then, we will now set it aside. (Readers, please feel free to take a worm.) Let's instead pursue the more simple answer, which is that the impermissivist advisor will give out a red X to any mushy credences, including those instances in which she has the very same mushy credence as her disciple. Can we accept this answer? Maybe it's not so bad. After all, we impermissivists are famously hard to please. We don't expect rationality to come along very often. So we should expect to give out a lot of red X's. We give them out to precise credences that are just a smidge too high or too low, which many permissivists are already upset by. What's the additional harm in giving them out to all mushy credences?

Miriam worries that if there are too many red X's, the evaluative practice won't be very useful. It will be all X and no star. But remember that there is nothing stopping the advisor from adding something to her red X, if she thinks it's useful to say more. She

might say, “you’re irrational, but I think you’re close,” for example. Or she might say, “you’re irrational, but I don’t know what’s rational; sorry I can’t be of more help!” This seems to us to be not far off from how we act in ordinary life. And it’s not as if there are *no* situations in which an impermissivist can offer specific advice; ordinary people are often irrational in dramatic and obvious ways. (Consider sports fans who believe that their team, out of the whole league, will make the playoffs; or parents who think that their child, out of the whole middle school baseball team, is not being given enough time on first base. These people’s main problem is not mushiness.) For these reasons, we think that the impermissivist can avoid the worst consequences of a red-X-heavy practice (first, that it would backfire and move people farther, not closer, to rationality; second, that it’s too coarse-grained to be useful) and also avoid proliferating extra, unnecessary normative concepts.

In sum: we think the impermissivist advisor – and by extension, an impermissive evaluative practice more generally – does okay when it comes to evaluating mushy credences. In such cases, we think it is more useful to have a practice where we discuss our *uncertainty*, rather than deploying a separate evaluative concept.

VI. Miriam: Some Final Thoughts

I’ve been suggesting that (accuracy-oriented) permissivists and impermissivists may disagree about less than is commonly thought. Indeed, I take it that the three of us agree that when it comes to our first order opinions and actions, permissivists and impermissivists will proceed in much the same way.

But that doesn’t mean there is no disagreement to be had. Sophie and Sinan helpfully point to some of the ways in which permissivists and impermissivists will diverge in *verbal* behavior. Furthermore, they claim, the impermissivist way of speaking is superior to the permissivist one. I remain somewhat skeptical of this: I see no reason to expect better results from saying to one’s disciple “I don’t know which credence is permissible” and “they all are”. But since the question of which verbal practice will be more successful is at least in part an empirical one, I don’t expect that we’ll be able to settle the matter here.

I do however want to end by responding to a challenge that Sophie and Sinan pose: Insofar as I’m leaving open the possibility that the *impermissive* practice has social value “is it *also* important to have a permissive mode of evaluation?” In my opinion: Absolutely! For recall: the notion of rationality that is permissive is the one that tracks our attitudes about which ways of forming belief conduce to accuracy. Because accuracy

is important, it is, I think, valuable for us to have ways of thinking and talking about which belief forming methods support this aim. Since permissive rationality just *falls out* of this practice, to think that there is no value to a permissive notion of rationality is to think that there is no value in such a practice.

One final thought: Not only do I think that there is value in a practice that yields a permissive notion of rationality there is, I think, also value in *acknowledging* openly to ourselves and others the fact that the permissivist wants to draw our attention to. Speaking permissively is a way of expressing acceptance (dare I say celebration?) of the fact that our concern with truth doesn't always force our hand. The methods of forming beliefs we endorse for the purpose of believing truly, leave us with an abundant buffet of doxastic options. This means that the fact that our beliefs are sometimes impacted by arbitrary, evolutionary, prudential or moral reasons is not always an epistemic problem. Nor should we always (even when the sole concern is truth!) aim to prevent such influences on belief. The permissivist is drawing our attention to a fact that is epistemically liberating. I hope, then, that even if you are convinced by Sophie and Sinan that it is dangerous (for testimonial reasons) to *speak* as if permissivism is true, you can, in your private reflective moments, feel the tremendous relief (which I have personally experienced and highly recommend) that comes from acknowledging the sense in which epistemic rationality is indeed permissive.

VII. Miriam, Sophie and Sinan: Permissivism and Metaepistemology

Let's end on a note of agreement: all three of us think that the permissivism debate straddles in an interesting way the border between first order epistemology and metaepistemology. While the question of permissivism might be thought of as a question of first-order epistemology (does the evidence support one position or many?), it seems to us that how one answers it depends a great deal on one's metaepistemological views. This thought stems from our shared conviction that there must be some important connection between rationality and truth. Not just as professional epistemologists, but as ordinary civilians who live with and rely on family, friends and strangers too, we are interested in rationality primarily because we are interested in gaining true beliefs and avoiding false beliefs about important topics. It's famously difficult to say exactly how rationality is connected to the truth. But we agree that it would be a problem for any view if there is an argument that the view could never make good sense of the value of rationality in terms of the value of truth or accuracy. This is why, in our debate (not only here but in our past papers too), we have been focused on questions concerning which

practices of thinking and talking about rationality support this aim, and whether such practices are permissive or not.

Acknowledgements

For helpful feedback on the material in this paper, we would like to thank audiences at the Epistemic Utility for Imprecise Probability Conference at Bristol University, the Fine Tuning Conference at Baylor University, MIT, Princeton University, UMass Amherst, Washington University at St Louis, and Roy Sorensen's Spring 2021 seminar at UT Austin.

References

- Berger, D. and Das, N. (2020). "Accuracy and Credal Imprecision." *Noûs* 54 (3):666-703
- Builes, D., Horowitz, S. and Schoenfield, M. "Dilating and Contracting Arbitrarily." (2022). *Noûs* 56(1): 3-20
- Carr, J. (2020) "Imprecise Evidence without Imprecise Credences." *Philosophical Studies* 177, 2735–2758
- Dasgupta, S. (2017). "Normative Non-Naturalism and the Problem of Authority." *Proceedings of the Aristotelian Society* 117 (3):297-319
- Dogramaci, S. (2012). "Reverse Engineering Epistemic Evaluations." *Philosophy and Phenomenological Research* 84 (3):513-530
- Dogramaci, S. and Horowitz, S. (2016). "An Argument for Uniqueness about Evidential Support." *Philosophical Issues* 26(1): 130-147.
- Dorst, K. (2019). "Higher Order Uncertainty." in M. Skipper & A. S. Petersen (eds.). *Higher Order Evidence: New Essays*. Oxford University Press.
- Dreier, J. (2015). "Can Reasons Fundamentalism Answer the Normative Question?" In G. Björnsson, C. Strandberg, R.F. Olinder, J. Eriksson & F. Björklund (eds.), *Motivational Internalism*. Oxford University Press.

- Gibbard, A. (2003). *Thinking How to Live*. Harvard University Press.
- Greaves, H. and Wallace, D. 'Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility'. *Mind* 115(459): 607-632.
- Greco, D. (2014). "A Puzzle about Epistemic Akrasia." *Philosophical Studies* 167(2):201-219.
- Greco, D. (2015). "Verbal Debates in Epistemology." *American Philosophical Quarterly* 52(1):41-55.
- Greco, D. and B. Hedden. (2016). "Metaepistemology and Uniqueness", *The Journal of Philosophy* 113(8): 365-95.
- Horowitz, S. (2014). "Immoderately Rational." *Philosophical Studies* 167(1): 130-147.
- Horowitz, S. (2019). "The Truth Problem for Permissivism." *Journal of Philosophy* 116 (5):237-262
- Horowitz, C. (Forthcoming). "Accuracy and Educated Guesses." *Oxford Studies in Epistemology*.
- Joyce, J. (2009). "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief" in F. Huber & C. Schmidt-Petri (eds.), *Degrees of Belief*. Synthese Library.
- Kelly, T. (2013). "Evidence Can Be Permissive". In M. Steup, J. Turri, and E. Sosa (eds.), *Contemporary Debates in Epistemology 2nd Edition*, Wiley-Blackwell.
- Konek, J. (2019). "Comparative Confidence" in R. Pettigrew and J. Weisberg (eds.) *The Open Handbook of Formal Epistemology* 267 - 348. PhilPapers Foundation.
- Konek, J. (forthcoming). "Epistemic Conservativity and Imprecise Credence." *Philosophy and Phenomenological Research*
- Mayo-Wilson, C. & Wheeler, G. (2016). "Scoring Imprecise Credences: A Mildly Immodest Proposal." *Philosophy and Phenomenological Research* 93: 55–78.
- Meacham, C. (2013). "Impermissive Bayesianism", *Erkenntnis* 79(Supp 6): 1185-1217.
- Oddie, J. (1997). "Conditionalization, Cogency and Cognitive Value" *British Journal of the Philosophy of Science* 48: 533-541.

- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press.
- Ridge, M. (2006). "Ecumenical Expressivism: Finessing Frege." *Ethics* 116(2): 302-226.
- Ridge, M. (2007). "Expressivism and Epistemology: Epistemology for Ecumenical Expressivists." *Proceedings of the Aristotelian Society* 81: 83-108.
- Rinard, S. (2017). "Imprecise Probability and Higher Order Vagueness. *Res Philosophica* 94 (2):257-273.
- Schoenfield, M. (2012). "Chilling Out on Epistemic Rationality: A Defense of Imprecise Credences." *Philosophical Studies* 158 (2):197-219
- Schoenfield, M. (2014). "Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief." *Noûs* 48 (2):193-218.
- Schoenfield, M. (2015). "Bridging Rationality and Accuracy." *Journal of Philosophy* 112(12):633-657.
- Schoenfield, M. (2017). "The Accuracy and Rationality of Imprecise Credences." *Noûs* 51 (4):667–685.
- Schoenfield, M. (2018). "An Accuracy Based Approach to Higher Order Evidence" *Philosophy and Phenomenological Research* 96 (3):690-715.
- Schoenfield, M. (2019). "Permissivism and The Value of Rationality: A Challenge to the Uniqueness Thesis." *Philosophy and Phenomenological Research* 99 (2):286-297.
- Schoenfield, M. (forthcoming). "Meditations on Beliefs Formed Arbitrarily." *Oxford Studies in Epistemology* 6. Oxford University Press.
- Seidenfeld, T., Schervish, M.L, & Kadane, J.B. (2012). "Forecasting with Imprecise Probabilities." *International Journal of Approximate Reasoning* 53: 1248–61.
- van Roojen, Mark, (2018). "Moral Cognitivism vs. Non-Cognitivism", *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2018/entries/moral-cognitivism/>.
- White, R. (2005). "Epistemic Permissiveness", *Philosophical Perspectives*, 19 (1):445-459.

White, R. (2013). "Evidence Cannot Be Permissive". In M. Steup, J. Turri, and E. Sosa (eds.), *Contemporary Debates in Epistemology 2nd Edition*, Wiley-Blackwell.