Why Subjectivism?

Chloé de Canson

Abstract. Few philosophical positions are as unpopular as radical subjective Bayesianism. In this paper, I seek, if not to rehabilitate subjectivism, at least to show its critic what is attractive about the position. I argue that what is at stake in the subjectivism/anti-subjectivism debate is not, as is commonly thought, which norms of rationality are true, but rather, the conception of rationality that we adopt: there is an alternative approach to the widespread telic approach to rationality, which I call the poric approach, on which subjectivism is an attractive position. I end by considering the prospects for escaping subjectivism, and I suggest that they are bleak.

0.

One would be hard-pressed to find a self-proclaimed radical subjective Bayesian (henceforth, *subjectivist*)[1] today, and with good reason: the view suffers from serious problems. Critics of the position point out that, to the extent that it mandates particular forms of omniscience, it is too demanding; and to the extent that this is all that it mandates, it is not demanding enough. In light of these problems, subjectivism has been widely rejected.

In this paper, I seek, if not to rehabilitate subjectivism, at least to show its critic what is attractive about the position. But my strategy won't be to engage head-on with the anti-subjectivist. Rather, I will highlight that many anti-subjectivists adopt a particular account of epistemic justification—one which might in principle be rejected. I will then propose an alternative account that one might adopt in its stead, and show not only that the serious problems with subjectivism lose their traction on this new account; but also that a strong argument for subjectivism emerges out of it. Thus I will carve an alternative view of rationality on which—I will argue—subjectivism holds. Although I will not dispute the standing of anti-subjectivism on the original account of rationality, I will end this paper by suggesting that the account I propose is better suited to play the role that the dominant account is supposed to play. This tells in favour of my account, and thus, of subjectivism.

[1] The term *subjectivism* unfortunately has two meanings in the philosophy of probability. It refers both to the epistemic interpretation of probability functions, as opposed to the various ontic interpretations; and to a particular view about which epistemic probabilities are justified. In this paper, I am concerned with the latter.

The plan for the paper is as follows. I begin with a presentation of subjectivism and its two usual objections (§1). Then, I bring out the assumption that much of Bayesian epistemology has made, and I propose an alternative to it (§2). I propose an argument for subjectivism based on this new assumption (§3), and I show that the two usual objections fail if it is accepted (§4–5). I conclude by examining the prospects for escaping subjectivism, and I suggest that they are bleak (§6).

1.

The consideration at the heart of Bayesian epistemology is that beliefs come in degrees: agents may be more or less confident in the truth of any particular proposition.[2] The agent's degree of belief in a proposition is called her *credence* in that proposition, and her total epistemic state can be represented by a credence function $p : \mathcal{A} \to [0, 1]$, which assigns a credence to each element in the set $\mathcal{A}$ of all propositions the agent entertains. The formal structure of $\mathcal{A}$ is as follows. We begin with a non-empty set Ω. We then construct $\mathcal{A}$ to be a Boolean algebra of Ω, which is to say that it is a set of subsets of Ω, such that $\Omega, \emptyset \in \mathcal{A}$, and $\mathcal{A}$ is closed under union and negation. I leave the question of how to precisely interpret these mathematical objects for later on. However I want to draw the attention of the reader to the distinction between the propositions we might call *trivial* (Ω and $\emptyset$), and those we might call *non-trivial* ($A_1, A_2, ...$). This distinction will turn out to be important.

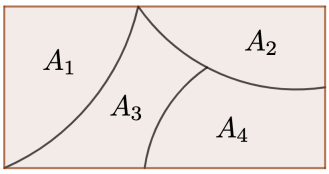

Much of Bayesian philosophical scholarship is organised around the question of which norms, if any, govern an agent's credences on such sets of propositions. Subjectivism is a particular view on the matter:

Subjectivism. There are two (and only two) credal norms of rationality:

Probabilism. An agent's credence function p must be a probability function, that is, such that:

(a) $p(\Omega) = 1$ and $p(\emptyset) = 0$.

(b) $p(A_i \vee A_j) = p(A_i) + p(A_j)$ for all inconsistent propositions A_i, A_j.

Conditionalisation. Upon receiving some evidence E, an agent's new credence function p' must be such that:

(a) $p'(E) = 1$ and $p'(\neg E) = 0$.

(b) $p'(A_i|E) = p(A_i|E)$ for all propositions A_i.

[2] I use the term *Bayesian* in a broad way to pick out any theory or theorist concerned with agents' credences.

This somewhat unorthodox presentation of subjectivism highlights the fact that each of the norms that compose it are made of one sub-norm **(a)** we might call *substantive*, which prescribes particular credences to particular propositions; and of another sub-norm **(b)** we might call *formal*, which makes particular coherence prescriptions on the entire credence function(s).[3] These can be presented in the table below.

	Probabilism	**Conditionalisation**
Substantive	Trivial Omniscience	Evidential Omniscience
Formal	Additivity	Rigidity

In this paper, I shall be concerned solely with the substantive sub-norms of subjectivism, for two reasons. The first is that, as we shall see, they are the ones at stake in the dispute between subjectivists and their opponents. The second is that formal sub-norms require an altogether different treatment, different enough that it can be neatly treated separately. As a result, I leave the question of formal sub-norms for future work.

In light of this restriction, we can consider a restatement of subjectivism.

Subjectivism. An agent is (substantively) rational if and only if she is

Trivially omniscient. $p(\Omega) = 1$ and $p(\emptyset) = 0$.

Evidentially omniscient. $p(E_i) = 1$ and $p(\neg E_i) = 0$, for all evidential propositions E_i.

And this restricted claim may be restated in a different form once again:

Subjectivism.

Necessity claim. Agents ought to be trivially and evidentially omniscient.

Sufficiency claim. There are no other (substantive) requirements of rationality.

With such a statement in hand, we are in a position to see why subjectivism is so widely rejected: there are what seem to be decisive arguments against both the necessity and the sufficiency claims. But before reviewing these arguments, let us first clarify what is meant by trivial and evidential omniscience. Let us begin with trivial omniscience. Without an interpretation of the agent's algebra, we do not know what Ω and $\emptyset$ represent, and so we cannot know what it means to be omniscient with respect to them. Now, the problem with the necessity claim in the literature is often called the *problem of logical omniscience*,[4] and correspondingly, probabilism is typically taken to entail that agents ought to have credence 1 in all

[3] I call these sub-norms *substantive* and *formal* because I think that these labels capture nicely the way in which they differ, but nothing of substance hangs on this. If the reader rejects this characterisation of the distinction, she will (I hope) still accept the restriction of my concern to one of the disjuncts.

[4] In fact, what is usually called "the problem of logical omniscience" actually refers indiscriminately to the problems associated with two distinct requirements: (a) that agents be certain of the truth-value of all logical propositions; but also (b) that agents' credences in different propositions be related in particular ways, for instance that an agent always be more confident in a proposition X's logical consequence than in X. This amounts to treating substantive and formal components of probabilism at once. But as I announced earlier in the paper, I will restrict all my claims to substantive issues; in this case, to (a).

logical truths and credence 0 in all logical falsehoods. This shows that Bayesians typically interpret the formalism in such a way that logical truths and logical falsehoods—tautologies and contradictions—should be modelled by the trivial elements. So, we shall assume that trivial omniscience (at least) entails logical omniscience. (More on this assumption later.)

What about evidential omniscience? Again, without an interpretation of the algebra, we are not in a position to determine what qualifies as an evidential proposition. This is because, in order to determine what evidential omniscience requires of agents, we must adopt a specific conception of evidence. Now, there are two options. On the *externalist* conception of evidence, an agent's evidential propositions express facts about the external world, facts such as those expressed by the sentences "the sun has risen today" and "it is a book I am holding". On the *internalist* conception of evidence by contrast, evidential propositions express facts about the agent's internal world, facts such as those expressed by sentences like "it seems to me that the sun has risen today" and "it seems to me like it is a book that I am holding". For now, it will suffice to remark that most Bayesians (explicitly or implicitly) assume an externalist conception of evidence. Typical examples of evidence in the literature include things like: the coin has landed heads, and there is a goat behind the door.

Let us begin with the necessity claim: what is the problem with mandating logical and evidential omniscience? In order to have credence 1 or 0 in all logical and evidential propositions, the argument goes, one must be in a position to recognise them as such. But real human agents are not in a position to do this. Indeed, not only are there infinitely many logical truths, some of which very complex, but many have been shown not to be determinable even in principle. So many human agents, when faced with a logical proposition, are not in a position to tell that they are. Furthermore, real human agents are not always in a position to recognise what evidence (externalistically conceived) they possess. For instance, an agent may not be able to tell whether the proposition expressing that she has hands is an evidential proposition for her. So many human agents, when faced with an evidential proposition, are not in a position to tell that they are. It follows that agents cannot be omniscient in the ways mandated by subjectivism: this is the over-demandingness objection.

Over-demandingness objection.

1. Agents ought to be trivially and evidentially omniscient. (Necessity Claim)

2. Ought implies can. (Premise)

3. Agents can be trivially and evidentially omniscient. (1, 2)

4. They cannot. (Premise)

5. So, the necessity claim is false.

In light of this objection, most have rejected the necessity claim as interpreted above. Now, because of a widespread reticence to abandoning probabilism (the reason for which I discuss in fn. 11), the main way in which Bayesian scholars have done this is not by rejecting trivial omniscience, but by reinterpreting it in such a way as to undermine premise **4**. For example, Hacking (1967) and more recently Pettigrew

(minga), and in a different way Bradley (2017), have proposed reinterpretations of the trivial/non-trivial distinction on which trivial omniscience no longer entails logical omniscience. In that way, the fact that real human agents cannot recognise every logical truth no longer challenges the mandate of trivial omniscience. An analogous strategy has been attempted to ward off the challenge from evidential omniscience: some have suggested reinterpreting the evidential/non-evidential distinction in an internalist way, so that evidential propositions express *seem*-facts and agents are always in a position to be certain of their evidence.

Now, an in-depth examination of the reinterpretation strategies is beyond the scope of this paper. However, it will be relevant to note that such reinterpretations stand in tension with what is often taken to be the target-subject of Bayesian epistemology; namely, inquiry into the empirical world. By virtue of the formalism, the propositions in which we are interested—in which we want to know what credence to have—are the non-trivial propositions. Furthermore, by virtue of the formalism too, each of these propositions is liable to become evidence. In light of this, the Bayesian enterprise is naturally understood in such a way that non-trivial propositions are *empirical* propositions, that is, propositions about the external world that might come to be known by observation (capaciously understood). But reinterpretations of the algebra which allow non-trivial propositions to be logical propositions or propositions about the agent's mind stand in tension with this. For on these reinterpretations, logical propositions could in principle constitute evidence, and propositions about the external world could not. Clearly, my claims here do not constitute a definitive refutation of reinterpretation strategies. They do however call for an answer to the question of which domain(s) of inquiry we take Bayesian epistemology to be concerned with.

While the concern about treating logical propositions as potential evidence has, to my knowledge, not been discussed in the literature, the internalist reinterpretation of evidence has been. As Jeffrey (1968) puts it: "for excellent reasons, this move is now in low repute" (p. 171). He argues that we should reject evidential omniscience, and replace it with a norm according to which, when agents acquire evidence, they do not adopt credence 1 in E and credence 0 in $\neg E$, but instead come to have (probabilistic) non-trivial credences in these propositions. In that way, agents can come to assign credence .9 to the proposition that they are holding a book, and credence .1 to the proposition that they are not (because, perhaps, they are being deceived by an evil daemon).[5] More recently, rejections of conditionalisation motivated by the failure of evidential omniscience can be found in Bronfman (2014), Schoenfield (2017a), and Gallow (2017, forthcoming).

Let us now move on to the second sub-norm of subjectivism: the sufficiency claim. This claim has been the target of the vast majority of anti-subjectivist criticism, to the point that the subjectivist/anti-subjectivist debate is often presented in those terms only. What the sufficiency claim asserts is that there is no substantive requirement of rationality beyond trivial and evidential omniscience. In other words, as long as agents are omniscient in those ways, they are rational. So, take an arbitrary non-trivial

[5] Jeffrey retains rigidity (the formal sub-norm of conditionalisation), and ends up endorsing a new updating norm called *Jeffrey conditionalisation*.

non-evidential proposition: any credence in this proposition is rational. But surely that cannot be, the argument goes. Surely someone who has an extremely high credence that a coin will land heads, despite knowing nothing about the coin in question, is irrational. So, subjectivism is not sufficiently restrictive: this is the under-demandingness objection.

Under-demandingness objection.

1. All that's rationally required of an agent is for her to be certain in the truth of trivial and evidential propositions. (Sufficiency Claim)

2. A credence of .99 in the proposition that a coin will land heads is rational. (1)

3. It is not. (Premise)

4. So, the sufficiency claim is false.

In response to this objection, many have rejected the sufficiency claim, and proposed variously strong norms to supplement probabilism and conditionalisation. In one tradition, people have attempted to constrain credences by appealing to norms according to which agents' credences are required to match exogenously given probabilities. In another (sometimes overlapping) tradition, people have proposed indifference norms on credences, which prescribe a particular unique credence to each proposition in the algebra in the absence of evidence. In yet another (sometimes also overlapping) tradition, people have appealed to evidentialist considerations by insisting that credences must reflect the evidence. I will discuss these traditions in more detail in §5.

This objection has turned out to be particularly potent in deterring people from subjectivism. Indeed, not only are there, to my knowledge, no contemporary defences of the sufficiency claim in print, but there are also reasons to believe that even the original subjectivists—Ramsey, de Finetti, and Jeffrey—rejected it. Ramsey says: "if we are told that one of these people's names begins with A and that there are 8 of them, it is reasonable to believe to degree $1/8^{\text{th}}$ that any particular one's name begins with A, and this is what we should all do (unless we felt that there was something else relevant)" (1926, pp. 100–101). Jeffrey writes that, as far as the assignment of credences is concerned, his view is "often faulted as uncritical [...]; 'anything goes'. But the adoption of [credences] is a subject-matter dependent *techne*, an art of judgment [...]. Although [the expert agent] is far from knowing how [her mechanism for assigning credences] works, she can know *that* it works, pretty well" (1991, pp. 11-12).[6] And even de Finetti says: "though maintaining the subjectivist idea that no fact can prove or disprove belief, I find no difficulty in admitting that any form of comparison between probability evaluations (of myself, of other people) and actual events may be an element influencing my further judgment, of the same status as any other kind of information." (1962, p. 360).[7]

Given the strength of both the over- and under-demandingness objections, and the attractiveness of many available alternatives, it should come as no surprise that subjectivism enjoys a low level of

[6] So it turns out that Jeffrey rejects both the necessity and the sufficiency claim!

[7] For a history and critical discussion of the early subjectivists' commitments, see Galavotti (2011, 2016, 2018).

popularity. Thus, although subjectivism tends to be viewed as the default position in Bayesianism, this is not because it is taken to be true. Rather, I submit, it is because it is taken to be simple; and so it is discussed, not in the sense of being engaged with, but in the sense of being used to define putatively more plausible positions by contrast. My aim in this paper is to counter this trend, and to provide an argument *for* subjectivism. I will present a way of thinking about justified credence on which subjectivism is an appealing answer to the question of what norms govern epistemic attitude. I will then show that the over-demandingness and the under-demandingness objections fail to have any grip on this alternative way of thinking about justification. Finally, I will take seriously the thought that it is in some sense bad if subjectivism is true, but will suggest that, in light of my argument, the prospects for escape are bleak.

2.

As I remarked in the previous section, Bayesians have organised their field of study around the question of which graded epistemic attitudes are justified. What I want to show in this section is that many Bayesians have been taking a particular interpretive approach to this question; and that there exists an other one might take. For now, I shall call them the *standard approach* and the *alternative approach* respectively.

Let me begin by sketching the standard approach. Bayesians working within this approach are ultimately interested in questions like: What is the best credal state to have for a climate scientist in this (specified) situation? What is the best credal state to have for a person who wonders whether it is about to rain? These questions are interesting to the standard Bayesian because of the two characteristics he takes epistemic states to have. On the one hand, credences aim to represent the world; on the other, they guide rational action. So, the standard Bayesian wants to know whether the scientist should have a high credence that temperatures will rise because he wants to know whether she is correct in her confidence that they will, and because he wants to know what she (and therefore we) should do in response to the climate crisis. So, determining what it is best to believe, and thereby (given fixed desires) what it is best to do, is at the heart of the standard Bayesian project. Given this broad purpose, his methodology tends to look as follows. He starts with a (real or imagined) human agent, such as the climate scientist or the rain-averse pedestrian. He then isolates the normatively relevant aspects of the agent's epistemic state: her credences; which he represents, or *models*, using a credence function.[8] Then, he asks: What shape should this credence function have? Or equivalently, which credences are justified? By this, he means: Which credences would it be best for the agent to have, given the aim of credences?

Here is an alternative approach to Bayesian epistemology. On this approach, the question in which the theorist is ultimately interested is: which justified epistemic states can our means of inquiry afford us? In other words: What justified credences can we achieve on the basis of our means of inquiry?

[8] For an investigation into modelling in Bayesian methodology, see Roussos (ms).

How successful can we be at finding out about the world? In order to address these questions, the alternative Bayesian adopts a different methodology than the one described above. She begins by listing the agent's means of inquiry. These might be reason, observation, testimony, introspection, and so forth. Then, for each proposition that the agent considers, she asks: do the agent's means of inquiry warrant a particular credence in this proposition? For example, has the agent observed that it is true? Has she been told that it is true? And so forth. More generally: on the basis of her means of inquiry, what kind of epistemic attitudes towards these propositions are warranted?

These two approaches differ in at least one salient way: they differ in their account of justification. On the standard approach, a credal state is justified just in case it is the best credal state for the agent to have; but on the alternative approach, a credal state is justified just in case it is warranted by the agent's means of inquiry. Thus we can call the standard approach *telic*, after the Greek τέλος, which refers to an end or aim: on it, to be justified is to be related in a particular way to the end, aim, or goal of epistemic inquiry. By contrast, we can call the alternative approach *poric*, after the Greek πόρος, which refers to the means to one's ends:[9] on it, to be justified is to be related in a particular way (warrant) to the means of inquiry.

> *The poric account.* An epistemic state is justified just in case it is warranted by the agent's means of inquiry.
>
> *The telic account.* An epistemic state is justified just in case it is related in the right way to the aim of inquiry.

On the poric approach, one might say: you ought to have credence 1 in a particular proposition because you have observed that this proposition is true. On a telic approach, one might say: you ought to have credence 1 in a particular proposition because you would be guaranteed to lose money if you did not. We can see from this example that the two approaches can in principle be extensionally equivalent. Indeed, both the telic and the poric theorists rule that the fictional agent above ought to have credence 1 in the proposition in question, albeit for different reasons.

The next three sections will function, among other things, as an illustration and precisification of poric Bayesianism. But in the rest of this section, I will show that the telic approach is widespread in Bayesian epistemology, and I will comment on how the telic/poric distinction relates to epistemic consequentialism and evidentialism. Before I start, a terminological note is in order. I have defined justification on the poric account as consisting of a particular relationship between epistemic states and the means of inquiry, namely, warrant. By contrast, I have left the relationship between justification and the end of inquiry indefinite in my exposition of the telic account. This is because there are many ways of filling in the account; or in other words, there are many possible telic accounts of justification. In what follows, I will sometimes speak correctly and say of a particular account that it is telic, or of a

[9] The word also refers to a passageway, especially a passage over a body of water, such as a bridge; and derivatively, to a journey or crossing.

particular argument that it presupposes a telic account of justification, but I will also sometimes talk of *the* telic account. This will be shorthand for the most common account of rationality in the Bayesian literature, which, as I am about to show, is telic. To insist once again: although there is a single poric account, there are (at least in principle) several telic accounts.

As the reader who has made it this far will know, the most popular types of arguments in Bayesian epistemology are given in the table below,[10] together with the applications of these argument-types to defences of probabilism and conditionalisation.[11]

	Probabilism	**Conditionalisation**
Dutch Book	Ramsey (1926), de Finetti (1937)	Teller (1973), Lewis (1999)
Accuracy Dominance	Joyce (1998)	Briggs and Pettigrew (2020)
Expected Accuracy	Leitgeb and Pettigrew (2010a,b)	Greaves and Wallace (2006)
Decision Theoretic	Ramsey (1926), Savage (1954)	Savage (1954)

As I will show, these arguments share a common structure: they conform to the following recipe. **(1)** Take a putative norm of rationality X; for instance, probabilism. **(2)** Show that, if an agent's credences violate X, there is an alternative, better credence function that the agent could have had, which satisfies X. **(3)** Conclude that an agent's credences must conform to X. It follows that the proponents of these arguments adopt a telic conception of justification: in virtue of their shared structure, they link up justification to what it would be best for the agent to believe. Indeed this is made explicit by Pettigrew, who calls his position "teleological" (p. 11, 2016a).

Let us begin with the Dutch Book argument and the two accuracy arguments. Firstly, let us consider the Dutch Book argument.[12] It was first formulated by Ramsey (1926), and goes as follows. **(1)** Consider the norm of probabilism. **(2)** Show that, if an agent fails to conform to probabilism, she will accept a series of bets such that she is guaranteed to lose money, no matter how the bets are resolved. **(3)** Conclude that credences ought to be probabilistic. It is plain that the Dutch Book argument follows the above recipe. As we shall see, the accuracy do arguments too. The accuracy dominance argument is very closely related to the Dutch Book argument; it differs in its step **(2)**, as it shows that if an agent fails to conform to probabilism, there will be an alternative probabilistic credence function which is guaranteed to be more accurate, no matter how the world turns out to be. The expected accuracy argument is again closely related; it differs in its step **(2)**, which shows that if an agent fails to conform to probabilism, she will fail to minimise the expected inaccuracy of her credences. These three types

[10] For an overview and critical discussion of: Dutch Book arguments, see Pettigrew (forthcominga); accuracy dominance and expected accuracy arguments, see Pettigrew (2016a) and Pettigrew (2019); arguments for conditionalisation, see Pettigrew (2020).

[11] Bayesians have been much more convinced by arguments for probabilism than they have been by the corresponding arguments for conditionalisation. As a result, probabilism is more widely endorsed than conditionalisation. This might explain why people have been quicker to reject evidential than trivial omniscience!

[12] The etymology of the term is unknown, and a better name for it might be the *utility dominance argument.*

of arguments presuppose a telic account of justification: it is because the agent will be (pragmatically or alethically) badly off if she doesn't that her credences must satisfy probabilism. (This presentation of the arguments allows us to give a nice characterisation of the relationship between Dutch Book arguments and accuracy arguments: they disagree about what has final value—Dutch Books: utility; accuracy arguments: accuracy—but they agree on the way in which what it is rational for an agent to do relates to what has final value.)[13]

Let us now turn our attention to decision theoretic arguments. These arguments also originate with Ramsey (1926), who proposed them as an alternative to the Dutch Book argument, which he took to have serious flaws. We shall see that they too follow the three-step recipe laid out above, though in a slightly more complicated way. **(1)** Consider the norm of probabilism. **(2a)** Prove what is known as a *representation theorem*, which shows that an agent's preferences over a set of outcomes satisfy particular constraints Y if and only if that agent can be represented as (among other things) having probabilistic credences.[14] **(2b)** Show that, if the agent does not satisfy constraints Y, things will go badly for her in some sense. For instance, all major representation theorems are such that among constraints Y is the constraint that preferences be *acyclical*; that is, not such that $a \prec b \prec c \prec a$, where $\prec$ represents the preference relation and a, b, c are outcomes. The reason given for such a constraint is the *money pump argument*, which goes like this: if $a \prec b$, the agent will pay to swap a for b; if $b \prec c$, the agent will pay to swap b for c; if $c \prec a$, the agent will pay to swap c for a; and so on—the agent can be made to give out money indefinitely.[15] **(3)** Conclude that the agent ought to satisfy constraints Y—and thus ought to have probabilistic credences. Thus we see that decision theoretic arguments too assume a telic account of justification: it is in virtue of how they relate to (pragmatic) aims that probabilistic credences are rational.

I have shown that the major arguments in the Bayesian literature are all telic.[16] Before examining the relevance of this fact for my purposes in this paper, let me explain how the telic/poric distinction relates to the literature on epistemic consequentialism. Berker (2013a), with whom this literature originates, defines what is now known as *epistemic consequentialism* to be the thesis that agents ought to believe what is most conducive to the aim of inquiry.[17] He writes: "according to this picture, there are certain

[13] Note that many Bayesians believe that credences have dual aims: that of accurately representing the world, and that of adequately guiding actions. This gives them reason to accept both types of arguments.

[14] There are a number of different representation theorems in the decision theory literature. The classic ones which establish that credences must be probabilisitc are: Ramsey (1926), Savage (1954), and Jeffrey (1965) and Bolker (1966); and some more recent influential ones include Joyce (1999), Buchak (2013), and Bradley (2017).

[15] The money pump argument has roots in Ramsey (1926), and is later found in Davidson et al. (1955).

[16] Besides these, there are argument-types in Bayesianism that I have not discussed, including symmetry arguments (van Fraassen, 1989; Zabell, 2016), conservativeness arguments (Dietrich et al., 2016), and calibration arguments (van Fraassen, 1983; Lange, 1999). But I think it's fair to say that the arguments I have discussed are by far the most popular amongst contemporary Bayesians. Showing that they are telic and that there are good reasons for telics to be anti-subjectivists will suffice to account for the widespread rejection of subjectivism.

[17] Berker names this position *epistemic teleology*, but it has come to be known as *epistemic consequentialism*. Thus the two terms are often taken to be synonymous. But the point I am making in this passage is that consequentialism and what I am calling the *telic view* are not synonymous: the former is a special case of the

epistemic *ends* or *goals* that it is epistemically good for us to promote, and the question of what we *should* believe is determined by how well our believing conduces toward the fulfilling of those goals, or the furthering of those ends" (p. 340, emphasis original). Thus it is easy to see that epistemic consequentialism is a telic position: what matters when determining whether a credal state is justified is how it relates to the aim of inquiry.[18] However, there also exist telic positions besides epistemic consequentialism; in other words, not all telics are epistemic consequentialists. Indeed, epistemic consequentialists make two assumptions beyond telism. The first concerns the nature of the aim of epistemic attitudes: epistemic consequentialists assume that this aim is truth or accuracy. But if, as many Bayesians do, we take seriously the fact that credences are also that which rationalises action, we might consider pragmatic goodness to be an aim of credence, too—as indeed the proponents of Dutch Book and decision theoretic arguments do. And the second concerns the nature of the relationship to the aim of inquiry. As Berker notes, "what is distinctive about [epistemic consequentialism] is not just its taking value to be fundamental but moreover its attitude toward the nature of value and how we should respond to it. According to the [epistemic consequentialist], the proper response to value is to bring it about, and the proper response to disvalue is to stop it from being brought about: in short, for the [consequentialist] all value is 'to be promoted,' and all disvalue is 'to be prevented'" (p. 343). But it is possible to imagine a position whereby justification is a matter of relating to the aim of inquiry (a telic position), but where this relation is not one of conduciveness or promoting. For instance, Sylvan (2018, 2020) argues that justification is a matter of *respecting* the truth. Thus epistemic consequentialism is one among several possible telic views.

This is not to say however that the poric/telic disjunction exhausts the space of possible views about epistemic justification. Indeed, both the telic and the poric epistemologist view epistemology as a means-ends endeavour of sorts: there is an epistemic good towards which the agent strives—a goal to be achieved on the basis of one's means. I shall call this view a *means-ends* view of epistemology.[19] But we can imagine accounts of epistemic justification that are not means-ends in this way. Take *evidentialism* for instance, the view originating with Feldman and Conee (1985, 2004) according to which an epistemic state is justified to the extent that it reflects the agent's evidence. Berker (2013b) argues that Conee and Feldman are not consequentalists: he says that "they do occasionally lapse into talk of 'epistemic value' and 'epistemic goals,' but I think this talk can be excised from their program

latter. The terminology might be confusing, but unfortunately there are only so many words. I have tried to mitigate this by using "telic" and not, as he, "telelological".

[18] The accuracy arguments mentioned earlier in the section are widely acknowledged to be epistemically consequentalist, and this generates problems for them. For discussion, see Greaves (2013), Carr (2017), Pettigrew (2018) and Konek and Levinstein (2019).

[19] This label is imperfect. Indeed, it has an instrumental flavour, and evokes a view called *epistemic instrumentalism*, considered by Kelly (2003), according to which epistemic rationality is akin to practical (instrumental) rationality, in the sense that the agent determines her aim. On such an account of epistemic rationality, an agent whose aim it is to have false beliefs would be justified in believing that the earth is flat. But I want to stress that, the aims of means-ends epistemology are not particular to the agent or decidable by her; rather, the epistemic aims are those in relation to which epistemic states are evaluated, whether or not the agent additionally values these aims.

without major loss" (p. 380). If he is right, their account of justification is not only non-consequentialist, but it is also non-telic, and non-means-aims. For them, epistemic justification is not a matter of truth-seeking or any other epistemic striving. I will discuss evidentialism in greater detail in §5, where I will consider various interpretations of evidentialism. But I want here to illustrate the point that there exist alternatives to means-ends epistemology.

Let us now circle back to my argument. As I announced earlier, I am going to argue that subjectivism is an appealing position on the poric approach. But if, as I suggested in this section, all the main Bayesians are telics, then should we not suspect that the widespread arguments against subjectivism presuppose a telic account of rationality? And if that is the case, would it not be natural to suspect that subjectivists and anti-subjectivists are engaged in a verbal dispute? That subjectivists and their critics do not disagree about anything substantive, but merely operate with different conceptions of justification? I will show in §4–5 that the first suspicion is correct: the over- and under-demandingness objections rely on telic assumptions, and crumble if these assumptions are replaced by poric ones. But I will also show that the second reaction would be too hasty. I will argue that we cannot do away with the poric approach. If this is right, it poses a dual problem for the status quo. Firstly, if subjectivism is indeed false on the telic approach (that is, if widespread arguments to this effect are sound), and if we cannot do away with the poric approach on which it's true, we might be faced with the pressure to to abandon telism altogether. And secondly, we will have to grapple with the fact that subjectivism is, in many ways, a terrible predicament.

3.

In this section, I present an argument for subjectivism, on the poric approach. I begin by making two (contentious!) assumptions, which I will discuss later on.

> *Dualism.* Agents have two means of inquiry: reason, and observation.
>
> *Idealisation.* Agents have their means of inquiry perfectly.

So, according to dualism, agents have two distinct ways of inquiring into the world: they have reason, on the one hand, and observation (or experience, or sense-perception) on the other. And according to idealisation, they have these means of inquiry perfectly. What does that mean? Philosophers have long discussed how successful our means of inquiry might be. For instance, epistemologists of perception have discussed whether and how observation could be a source of justified beliefs about the external world; that is, whether and how our capacity to observe could warrant the belief that, for instance, the book I am holding is red, rather than just the belief that the book seems red to me.[20] Similarly, epistemologists of the *a priori* have debated whether and how reason could be a source of justified belief about non-empirical realms, such as the logical realm, the mathematical realm, the moral realm, and so on.[21] The idealisation assumption is designed to circumvent these challenges. It states (boldly!)

[20] This is the problem posed by Descartes (1993). See Lyons (2017) for an overview of the contemporary literature.
[21] See Russell (2014) for an overview of the literature.

that they have been answered, such that Bayesian agents are in a position to determine that the book really is red on the basis of observation, and that 2+2 really does equal 4 on the basis of reason.

It follows that taken together, the dualism and idealisation assumptions entail that Bayesian agents have two properties. Firstly, they are *perfect reasoners*. They are always in a position to recognise any proposition that can in principle be settled by reason, that is, any proposition decidable *a priori*, and always in a position to be certain of such propositions. This may include logical propositions, mathematical propositions, metaphysical propositions, moral propositions, and so on, depending on one's views on those realms of inquiry. And secondly, they are *perfect observers*. They are always in a position to recognise any proposition that can be settled on the basis of their observations, and always in a position to be certain of these propositions. So, for agents with such properties, propositions come in three kinds. The first kind are the propositions determinable *a priori*, those whose truth-value the agent can settle by exercising her capacity to reason. This leaves the propositions not decidable by reason, that is, the propositions determinable *a posteriori*, or empirical propositions. A further distinction between those is salient to our agent. On the one hand, there are the propositions she is in a position to settle on the basis of her observations; we can call those the propositions *about the observed*. And on the other, there are the propositions that she is not in a position to settle on the basis of her observations; we can call those the propositions *about the unobserved*.[22]

The next step is to model this agent with the help of the Bayesian mathematics. How can we model the *a priori*/*a posteriori* distinction? A very natural option is to use the trivial/non-trivial distinction. So, all propositions determinable *a priori* are modelled by Ω or $\emptyset$ (depending on whether they are true or false), and all the propositions not so determinable are represented by the non-trivial elements of $\mathcal{A}$. Among those, we can further distinguish between $E_1, E_2, ...$, the propositions about the observed, and $A_1, A_2, ...$, the other propositions, about the unobserved. We now find ourselves with an interpretation of the Bayesian algebra of entertained propositions. This interpretation has, I think, a distinctly Bayesian flavour.[23] The A_is are propositions that the agent might come to learn by empirical inquiry: propositions that might constitute, or come to constitute evidence—that might become E_is. By contrast, the trivial propositions Ω and $\emptyset$ are those whose truth-value cannot be determined by looking in the world. Note that on this interpretation, trivial omniscience is not mere logical omniscience: it is omniscience about all *a priori* matters.

Now, onto the argument for subjectivism. Remember, subjectivism is the conjunction of the necessity claim, which says that agents ought to be trivially and evidentially omniscient, and the sufficiency claim, which says that there are no other requirements of rationality. Here goes the argument.

[22] Some propositions may be determinable both *a priori* and *a posteriori*. For instance, if I want to know what 2+2 is, I can calculate it mentally, or I can look at the result on my calculator. But this does not threaten my arguments. Indeed, because of the idealisation assumption, any proposition that is knowable *a priori* is already known by the agent. And as will become clearer, the propositions that are of interest to the poric are those whose truth-value the agent's means do not suffice to determine.

[23] Note too that it is a much better candidate than the orthodox metaphysical interpretation of the algebra to account for uncertainty about the necessary *a posteriori*. See Chalmers (2011).

1. An agent is always in a position to be certain of the truth-value of *a priori* propositions. (Dualism and idealisation assumptions.)
2. The *a priori* propositions are Ω and $\emptyset$. (Modelling assumption.)
3. An agent is always in a position to be certain that Ω is true and $\emptyset$ is false. (1, 2)
4. An agent's credences should be such that $p(\Omega) = 1$ and $p(\emptyset) = 0$. (Poric justification, 3.)

5. An agent is always in a position to be certain of the truth-value of propositions about the observed. (Dualism and idealisation assumptions.)
6. The propositions about the observed are $E_1, E_2, \ldots$. (Modelling assumption.)
7. An agent is always in a position to be certain that E_i is true and $\neg E_i$ is false, for all i. (5, 6.)
8. An agent's credences should be such that $p(E_i) = 1$ and $p(E_i) = 0$, for all i. (Poric justification, 7.)

9. *Necessity claim.* Agents ought to be trivially and evidentially omniscient. (4, 8.)

10. The agent has no means of inquiry beyond reason and observation. (Dualism assumption.)
11. *Sufficiency claim.* There are no other requirements of rationality. (Poric justification, 10.)

12. *Subjectivism.* (9, 11.)

So why be a subjectivist? Because trivial and evidential omniscience, but nothing else, is warranted by Bayesian agents' means of inquiry. These agents' perfect capacity to reason and to observe warrant certainty in the *a priori* propositions, and the propositions about the observed. But these capacities do not warrant any particular epistemic attitude towards the third kind of propositions under consideration: the propositions about the unobserved. As such, there is nothing that agents ought to believe about them.

4.

We have made some progress towards subjectivism: we have an argument in favour of the position. But it remains to be shown that its proponent (along the lines above) can address the objections to the necessity and the sufficiency claims. I will begin with the over-demandingness objection in this section, and will move on the under-demandingness objection in the next. For each, I will concede that the telic epistemologist should reject the relevant claim in light of the argument, but argue that the poric epistemologist should not. The first part of this argument—that the telic Bayesian falls prey to these objections—will play a dual dialectical role. On the one hand, and given my argument in §2 that all major Bayesian arguments are telic, it will explain why these objections have been so successful in persuading Bayesians away from subjectivism. On the other hand, it will function as a challenge to the status quo in Bayesian epistemology. For if subjectivism is false on the telic view, but it is true on the poric view with which—I will argue—we cannot do away, we might be forced to reject telism altogether. I discuss this in §6. In the meantime, my aim is not so much to entrench the telic's rift from subjectivism as it is to show that there is no such rift for the poric. Accordingly, the standard

of argument will be higher for the latter goal than for the former: I will do little more than rehearse already widely accepted explanations for the falsehood of subjectivism on the telic approach, but I will seek to actively convince the reader that the objections fail on the poric approach.

With the dialectic made clear, we are ready to begin. Let me first reproduce the over-demandingness objection.

Over-demandingness objection.

1. Agents ought to be trivially and evidentially omniscient. (Necessity Claim)

2. Ought implies can. (Premise)

3. Agents can be trivially and evidentially omniscient. (1, 2)

4. They cannot. (Premise)

5. So, the necessity claim is false.

Let us begin with premise **2**. There are two kinds of telics in the literature. The first kind seeks to issue *guidance* to the agent on what to believe; for instance, Bradley seeks to construct a theory which "would provide guidance on how bounded agents should represent the uncertainty they face, how they should revise their opinions as a result of experience" (p. xiii, 2017). The second kind seeks to evaluate or *appraise* the agent's epistemic state; Pettigrew for instance wants to know "when it is appropriate to criticize an individual for their logical ignorance" (forthcoming).[24] How does premise **2** fare for each type of telic Bayesian?

It is straightforward that Bayesians who seek to guide cannot reject it: were it not the case that ought implies can, judgements of rationality could not serve as advice. (What kind of advice is unfollowable advice?) It is less straightforward, but still true I think, that Bayesians who seek to appraise cannot either. There are two things that one can do under the guise of appraising: (i) comparing an agent's epistemic state to the perfect epistemic state; and (ii) comparing an agent's epistemic state to the best such state that the agent could have been in. If telic Bayesians were in the business of doing (i), then they would be in a position to reject premise **2**. But they would also need to mandate empirical omniscience—a principle according to which agents must be certain of the truth-value of all empirical propositions. This is because the perfect epistemic state is that whereby all truths are believed with certainty and all falsehoods correspondingly disbelieved. Given that they do not mandate empirical omniscience, we can infer that it is (ii) in which they are involved. Thus we should interpret Pettigrew's question of when it is "appropriate" to criticise someone's beliefs, as asking which of the beliefs that the agent could have had (for some sense of "could") was the right one to have had. For the appraising Bayesian too, ought implies can. (I should remark that the sense of "can" plausibly differs in the guiding

[24] Pettigrew steers clear of the guidance aim in response to the recognition that human agents do not in general have control over what they believe: *epistemic voluntarism* is false. While I agree with him that epistemic voluntarism is false in general, I think that the cases in which (telic) Bayesian epistemology is most useful are cases where it is true, and I suspect that those are the cases that Bradley has in mind. For instance, we can easily imagine a scientist genuinely wondering what credences to form in response to a body of evidence.

and appraising projects[25]—but it does not differ in a way that is relevant to us; namely, in whether Bayesian agents *can* be trivially and evidentially omniscient.) Therefore, telic Bayesians of all stripes must accept premise **2**.

Since it is because of the telic's intention to guide and/or appraise that he must accept premise **2**, and since the poric Bayesian need not share that aim, might it be that she can reject premise **2**? No. For her, an epistemic state is justified just in case it is warranted by the agent's means of inquiry. In other words, what it means to say that an agent ought to have a particular epistemic state is that her means of inquiry allow her to have it, warrantedly. So the poric Bayesian too must accept premise **2**: it is constitutive of what she means by rationality that ought implies can (in the sense of "can" relevant to the over-demandingness objection).

Let us now turn our attention to premise **4**. Assuming that we take trivial omniscience to entail logical omniscience, and that we adopt an externalist account of evidence, it is clear, as we saw in §1, that if we take Bayesian agents to be (models of) real human agents, this premise is false. Actual humans are not always in a position to recognise logical and evidential propositions as such. But Bayesians are well known to be engaged in idealised epistemology: epistemology that brushes over humans' various cognitive and perceptual limitations. And if an idealisation assumption can be justified on which Bayesian agents can be trivially and evidentially omniscient, the necessity claim might be protected from the over-demandingness objection. For instance, on the poric position outlined in §3, the theorist is not particularly interested in real human agents. Instead, the agents that she considers are highly idealised. In fact, they are perfect reasoners and perfect observers, who are, by assumption, always in a position to settle any proposition determinable *a priori*, and always in a position to settle any external-world proposition about the observed. It follows that, on this view, agents can be trivially and evidentially omniscient, contra premise **4**. So, since it is with the idealisation assumption, that, perhaps unsurprisingly, the necessity claim stands or falls, let us examine how it fares on both the telic approach and the poric approach.

Can telics justifiably accept the idealisation assumption? It is noteworthy that few think they can: Bayesians have long believed idealisation assumptions to be ultimately unwarranted, and the recent trend towards de-idealisation is gaining momentum. And indeed, this follows from the considerations rehearsed above. If telics seek to apparise and/or guide real, human agents, any idealisation assumption will translate into an unachievable requirements on their agents, and so, ultimately, in the failure of their project. In light of this, the role of idealisation for them is transitional at best—that is, if they do idealise, it is with the intention of subsequently weakening the idealisation assumptions.[26] The goals of telic Bayesianism undermine idealisation assumptions.

But what about poric Bayesianism? We saw in §2 that the key question on the poric approach is that of which epistemic states an agent might achieve on the basis of her means. And we saw in §3

[25] For instance, guiding Bayesians must interpret "can" in such a way that it implies epistemic voluntarism (see fn. 24), whereas appraising Bayesians need not.

[26] For a discussion of transitional idealisation, see Staffel (2017) and de Bona and Staffel (2018).

that a particular domain of inquiry stands out: the empirical unobserved. Indeed, we saw that the dualism assumption divides propositions into three types: those that reason might settle (the *a priori* propositions), those that observation might settle (the propositions about the observed), and those that neither reason nor (heretofore) observation suffices to settle—propositions about the unobserved. This latter type is the proper target of poric Bayesianism. What credences in propositions about the unobserved are warranted by our means of inquiry? The argument for subjectivism presented in the previous section suggests that the answer to this question is: none. Our means of inquiry do not warrant any particular credences in propositions about the unobserved. For the poric Bayesian, and in accordance with widespread intuitions, this is a terrible conclusion: there is a domain of inquiry for which our means do not suffice to support any particular epistemic attitude. We are now in a position to appreciate the dialectical role of the idealisation assumption. By assuming that agents have their capacities of inquiry perfectly, the poric theorist is *conceding* more to agents than they in fact have. For if *even* agents who have these means of inquiry perfectly cannot have justified beliefs about the unobserved, agents who have them imperfectly certainly cannot.

So the poric Bayesian makes the idealisation assumption and thus rejects premise **4**, not because she is naive or confused about the capacities of real human agents, but because she wishes to concede to such agents more than they have, in order to determine what they might be able to achieve, epistemically. This is in sharp contrast with the situation of the telic Bayesian, for whom the idealisation assumption is not a concession but a injunction. Once we get the broader aims of telic and poric Bayesians plainly into view, it becomes clear that they should respond differently to the over-demandingness objection. The telic Bayesian indeed has no choice but to reject the necessity claim, for it makes an impossible demand to his agents. But the poric Bayesian should embrace the necessity claim, because for her, it plays the role not of a demand but of a concession: it gives to the agent more than she has, so as to see how far she could in principle go.

5.

This brings us to the other widespread objection to subjectivism: the under-demandingness objection. Just like with the over-demandingness objection in the previous section, my aim in this section is to show that, although there are good reasons to take the under-demandingness objection to be decisive against subjectivism on the telic approach, this is not so on the poric approach to rationality. Thus I will have accounted for the widespread endorsement of the under-demandingness objection to the sufficiency claim, shown that that the objection does not threaten the poric case for subjectivism, and confirmed that there is a genuine tension between the telic and poric approaches. Let me start by reproducing the argument.

Under-demandingness objection.

1. All that's rationally required of an agent is for her to be certain in the truth of trivial and

evidential propositions. (Sufficiency Claim)

2. A credence of .99 in the proposition that a coin will land heads is rational. (1)

3. It is not. (Premise)

4. So, the sufficiency claim is false.

We should begin by noting that this argument is less sophisticated than the over-demandingness objection. Its structure is very basic: we consider a direct consequence of the sufficiency claim, and rule on the basis of intuition that it is false. Now, a mere hunch can clearly not be reason enough to reject the sufficiency claim. We need (at least!) to explain why we have this intuition, in a way that exposes some philosophical consideration. Bayesians have attempted this in three ways—in what follows, I review each of them in turn.

5.1.

The thought at the heart of the first is that it is irrational for an agent to have a very high credence in the proposition that a coin will land heads because such a credence fails to match up with the objective chance that it will. We can call this version of the argument the ***chance-based version***. So far in this paper, we have used probability functions to represent epistemic phenomena: agents' degrees of belief. But these functions are also used to represent ontic, or mind-independent phenomena; so interpreted, they are known as ***chances***.[27] And a popular thought is that credences and chances are not unrelated—that the credences we ought to have are constrained in some way by ontic probabilities. To capture this thought, philosophers have formulated so-called ***bridge principles***, the most famous of which is Lewis' Principal Principle (1980). Which formulation of a bridge principle is the right one is a surprisingly complex and subtle matter, but the details are not important here.[28] What is important—what is at the core of all contestant bridge principles—is that an agent's credence p in a chancy proposition X, given that the correct chance function is ch, should be equal to the value that ch ascribes to X. So, where C_{ch} denotes the proposition that the correct chance function is ch:

$$p(X|C_{ch}) = ch(X)$$

Pettigrew has proposed a Dutch Book argument (2020) and several accuracy arguments (2012, 2013a, 2016a) for bridge principles. (As we saw in §2, these argument-types are telic.) But assuming that at least one of these arguments succeeds, and thus that a bridge principle indeed holds, it still does not follow that premise **3** is true—that an agent with a very high credence that the coin will land heads is irrational. For what bridge principles express is that there ought to be a particular relation between the agent's credence in the chancy proposition X and the proposition that the correct chances are given by ch. But if the agent does not know which chance function is the correct one, a bridge principle

[27] For an overview of different ontic interpretations of probability, see Hájek (2019).

[28] Pettigrew gives an illuminating presentation and critical discussion of numerous versions (2016a, pt. II).

alone cannot constrain her credences. This is often expressed by saying that bridge principles are mere coherence requirements—or, in my termonilogy from §1, they are formal norms. It follows that, for it to be possible to run the chance-based version of the under-demandingness objection on the poric approach, it must be the case that a Bayesian agent is in a position to become certain of which is the correct chance function, on the basis of her means of inquiry—and it is not.

Accounts of chance can be divided in two kinds.[29] On the *non-reductionist* accounts, such as the hypothetical frequency or the propensity accounts, chances are not reducible to facts that can be observed. On the hypothetical frequency account for instance, the chance of a coin landing heads is the limiting frequency with which it would land heads, were it flipped an infinite number of times. It is plain that chances on these accounts cannot be determined by observation.[30] By contrast, on the *reductionist* accounts, such as the finite frequency account, chances are a kind of summary of observable facts. For instance, on the finite frequency account, the chance of a coin which has been tossed ten times landing heads is the frequency with which it has actually landed heads. Chances on these accounts can be determined by observation.

The non-reductionist accounts of chance cannot help construct a poric version of the under-demandingness objection, since, by definition, these chances cannot be determined by observation. But, can we not conjoin a bridge principle with a reductionist account of chance to support premise **3**? We cannot. This is because, in order for the agent to be in a position to determine the (reductionistically interpreted) chance of a proposition, it must be for that proposition to be about the observed. But in order to undermine the sufficiency claim, it has to be the case that the proposition considered is one about the unobserved—one not determinable on the basis of observation. Indeed, the anti-subjectivist insists that there are propositions beyond the observed in which agents can have justified credences. It follows that, whether on a non-reductionist or a reductionist account of chance, there cannot be a poric chance-based version of the under-demandingness objection.

5.2.

This brings us to the second widespread reason for rejecting the sufficiency claim, which I take to be the most widely endorsed. According to this second group of Bayesians, premise **3** is true—it is irrational for an agent to have such a high credence that the coin will land heads—because, in the absence of evidence, having such a wild credence is taking an undue risk. We can call this version of the under-demandingness objection the *risk-based version.*

This version of the objection consists in a risk-based argument to the effect that an agent should

[29] I have defined the reductionist/non-reductionist distinction, not in the typical way to mean ir·reducible to *non-modal* facts, but to mean ir·reducible to *observable* facts. So, on my way of drawing the distinction, it is plausible to categorise the best-systems account as non-reductionist, for although it is reducible to non-modal facts, it is not reducible to observable facts—one cannot observe all facts past and future.

[30] Although, some authors have attempted to ground chanciness in empirical features of the world, by appealing to what is known as the *method of arbitrary functions.* See for example Strevens (2013). For a counter-argument, see "Objectivity and the Method of Arbitrary Functions", published as Canson (forthcoming).

have minimally opinionated credences—she should satisfy a *principle of indifference*. Such principles state that, in the absence of evidence, an agent's credence function should be flat; that is, agents should assign equal credences to (inconsistent) propositions. An indifference principle was first formulated by Bernoulli (1837/1954) and comes to Bayesians through Keynes (1921) and Carnap (1950); adherence to such a principle is the hallmark of what is known as *objective Bayesianism*.[31]

There are two risk-based arguments for the principle of indifference. The first, by Williamson (2010), goes as follows. An agent with non-indifferent credences will, because of the constitutive relationship between credences and action, be disposed to act in very risky ways. For instance, the person with a very high credence that the (unknown) coin will land heads will be disposed to bet a significant portion of his wealth on that outcome. But doing so, the argument goes, goes against what it would be best for the agent to do, and so the agent ought to have a milder opinion. The second risk-based argument, by Pettigrew (2016a,b), has a similar high-level structure (though it differs significantly in the details). An agent with a very high credence that the (unknown) coin will land heads takes a big epistemic risk that the agent with a middling credence in that proposition does not—she risks being very inaccurate. And, the argument goes, agents ought not take undue alethic risks: they ought to satisfy the principle of indifference.

Let us see how these risk-based arguments for the principle of indifference feature in a response to the sufficiency claim. Firstly, they vindicate premise **3** of our argument above: the agent is faced with two propositions, heads and tails, and the principle of indifference entails that she ought to have equal credence in both. Moreover, they are typically thought to vindicate intuitions of rationality more broadly. For consider an agent who has a low credence that the sun will rise tomorrow. Intuitively, this agent is irrational. Risk-based proponents of the principle of indifference explain this intuition in the following way. Before he had any evidence about the sun's movements, the rational credence function for him to have had would have been the indifferent one: this follows from the indifference principle. Since then, he has received plenty of evidence to the effect that the sun has risen every day. Given the principle of conditionalisation (and given some constraints on conditional credences), his credence that the sun will rise tomorrow should be greater than half.[32] It follows, let us grant, that risk-based arguments refute the sufficiency claim. But note that both of these arguments are telic: it is because not having indifferent credences risks being (pragmatically or alethically) bad for the agent that she ought to satisfy the indifference principle. Furthermore, it is hard to see how this telic argument could be reformulated in a poric way: it relies unavoidably on the expected goodness of adopting particular credences. As such, to the extent that a risk-based version of the objection is successful against the

[31] Objective Bayesianism is meant to be the opposite of what is sometimes known as Subjective Bayesianism, and which I have called subjectivism. But as I have defined it, subjectivism is the conjunction of two claims, the necessity claim and the sufficiency claim, and it is only the latter that objective Bayesians reject. Furthermore, they reject it by appealing to an indifference principle, but as we have just seen, there are more than just one way to reject it. We can say though that the so-called objective Bayesians are the most objective of Bayesians in one sense: they impose the most stringent constraints on all of an agent's credences.

[32] I argue against this line of reasoning in "Beliefs About the Unobserved".

sufficiency claim, it is so in virtue of being telic. While it may constitute a good reason for telic Bayesians to reject subjectivism, it does not threaten poric Bayesians' adherence to the position.

5.3.

This brings us to the third version of the under-demandingness objection. According to this third group of Bayesians, premise **3** is true—it is irrational for an agent to have such a high credence that the coin will land heads—because such an epistemic attitude fails to reflect the agent's evidence. Proponents of this view often cite Hume as asserting that "the wise man [...] proportions his beliefs to the evidence" (1748, 10.1–4). We can call this version of the argument the *evidentialist version.*

Evidentialists in the Bayesian tradition are particularly noted for their defence of the principle of indifference, which takes place in two steps. They start from the consideration that, since our agent does not have any evidence about the outcome of the coin flip, her evidence supports both outcomes—heads and tails—equally.[33] They continue with the contention that, as Keynes (1921) puts it, "if the evidence affords no ground for attributing unequal probabilities to the alternative predictions, it seems to follow that they must be equal" (p. 45). Or as White (2009) more recently puts it, agents' credences should "reflect their evidence (or lack of it)" (p. 171). Thus in the absence of evidence, agents should not assign a higher credence to one proposition over another. But evidentialist arguments can also be given beyond situations where the agent lacks evidence. For instance, an evidentialist might claim that our evidence that the sun has risen every day so far supports a high degree of confidence that it will rise again tomorrow, and conclude that an agent ought to have a high credence that it will.[34] This means that, if successful, these arguments have a potent anti-subjectivist potential: they can be used to vindicate credal constraints in a wide variety of scenarios.

In general, evidentialism can be charcaterised as the conjunction of two claims. **(1)** According to the descriptive claim, for any body of evidence, there is a unique credence function that is supported by that evidence.[35] **(2)** According to the normative claim, an agent ought to have the credences supported by her evidence. (Alternatively, her credences ought to "reflect" her evidence, in White's terminology (2009); or they ought to "respect" it, in Sliwa and Horowitz' words (2015).) In the rest of this subsection, I will examine whether evidentialism threatens my contention that the sufficiency claim is true on the poric approach. I will consider two interpretations of evidentialism. The first one, already mentioned

[33] I should note that this claim does not seem right to me; rather, I would be tempted to say that there is no particular credence that lack of evidence supports, which would be more in line with the poric spirit of this paper. But this will not be relevant for our purposes, as we will see.

[34] Thus evidentialists can avoid appealing to the credences the agent had before receiving evidence in order to determine what credences it is rational for them to have at a later time.

[35] Some think that it is not a unique credence function but a unique *representor* (a representor is a set of credence functions) that is supported by the evidence—they are the so-called *imprecise Bayesians.* The debate between precise and imprecise Bayesians concerns the level of grain of the epistemic state supported by the evidence; but what matters here is the assumption that there is a unique epistemic state which is indeed supported by the evidence. So I leave aside, here, and in fact elsewhere in this paper, concerns about the grain of epistemic attitudes. See Levi (1985) and Joyce (2005) for early texts, and Bradley (2019) for an overview.

in §2, is orthogonal to the means-ends conception of epistemology that I am working within, and as such, irrelevant to my purposes. The second one can be situated within means-ends epistemology, but I will argue does not threaten my contention.

Let us begin with the first interpretation of evidentialism. On this interpretation, it is *constitutive* of rationality that an agent's credences reflect the evidence. To deform Hume's disctum: to be wise just is to proportion one's beliefs to the evidence. On this interpretation, the normative claim **(2)** is analytic, and all the work is done by the descriptive claim **(1)**, which, it should be emphasised, has long been taken to be the Achilles heel of the position. Indeed, Keynes and his successors were widely criticised for failing to provide an account of evidential support, despite their views relying on the existence of one. Nonetheless, since many contemporary Bayesians—including Sliwa and Horowitz (2015), Schoenfield (2015, 2017b), or Dorst (forthcoming)—use the notion of evidential support, I will assume that whether an account of it can be formulated is still a live question. What is interesting, for my purposes, about this interpretation of evidentialism, is that it is neither telic nor poric; it in fact refuses what I called in §2 a means-ends conception of epistemology, whereby inquiring agents have a goal that they ought to achieve. Indeed in that section, I quoted Berker as claiming of Conee and Feldman, the most prominent evidentialists outside the Bayesian tradition, that "talk [of epistemic goals] can be excised from their program without major loss" (2013a, p. 380). Instead, for this kind of evidentialist, the epistemic agent is not attempting to determine what is the case, she is merely mirroring her evidence. Rationality on the constitutive interpretation is just a matter of calibrating one's credences to the evidence, and not a matter of seeking the truth. As Berker (2013b) puts it, it is "sideways-looking" (p. 377). As such, this interpretation is orthogonal to the framework within which my arguments in this paper operate.

Let us then turn to a second interpretation of evidentialism, consistent with the means-ends conception. How can evidentialism be formulated within such a framework?[36] A natural answer is that agent's beliefs ought to mirror their evidence because this is a good way of achieving the ultimate epistemic aim, namely, determining what is the case. So, on this second interpretation of evidentialism, believing in line with one's evidence is a means to achieving epistemic aims.[37] Note that this has a distinctly poric flavour. Might it be that evidentialism, so interpreted, yields a poric case against the sufficiency claim?

To answer this question, let us use a case proposed by Sliwa and Horowitz: "Anton is an anesthesiologist, trying to determine which dosage of pain medication is best for his patient: A or B. To figure

[36] For variants of this question, see Easwaran and Fitelson (2012), Schoenfield (2015) and Easwaran (2017).

[37] Pettigrew (2013b) writes, in relation to evidentialism, that "accuracy is not the only goal of credences: there is also the goal of matching one's credences to one's evidence" (p. 579). He goes on to consider how the goals might relate, and concludes that the only viable option if one is to remain within what I have called the means-ends framework is to hold that the "evidential goal is not an independent goal at all, but rather a byproduct of the goal of accuracy" (p. 579). There is a literal reading of Pettigrew's claims on which he disagrees with me: he takes reflecting the evidence to be a goal, albeit subsidiary to that of accuracy, whereas I take it to be a means of achieving accuracy. But I think there is less disagreement than it might appear: to be a subsidiary goal to X, and to be a means to X seem very similar things to me.

this out, Anton assesses some fairly complex medical evidence. When evaluated correctly, this kind of evidence determines which dose is right for the patient. After thinking hard about the evidence, Anton becomes highly confident that dose B is right. In fact, Anton has reasoned correctly; his evidence strongly supports that B is the correct dose" (p. 2836).[38] To fit this case in the two-step evidentialist recipe I provided above: **(1)** Anton's evidence supports a high credence that B is the correct dose; and **(2)** he should therefore adopt a high adopt a high credence in B.

Now, suppose that the descriptive claim **(1)** is indeed true; that is, that there is a (good) notion of evidential support such that Anton's evidence supports a high credence in B. Then, provided that **(2)** holds, it follows that the sufficiency claim is false: there are requirements of rationality beyond those maintained by subjectivists. But this interpretation of evidentialism is not poric. For if we interpret the notion of evidential support from the descriptive claim as one of poric warrant—that is, as one warranted by the agent's means of inquiry: observation and reason—it cannot be the case that Anton's evidence warrants a high credence in B. Indeed, whether B is an effective treatment is not observed, and it is not determinable *a priori*. Thus the claim that Anton's evidence support a high credence in B must operate with a non-poric conception of evidential support.[39] To conclude: if evidentialism underpins the under-demandingness objection, it is not poric; or contrapositionally, if evidentialism is poric, it does not underpin the under-demandingness objection.

Let me conclude on this section. The success of the under-demandingness objection relies on there being a reason to accept a claim like premise **3**: that it is irrational to have a very high credence that a coin about which nothing is known will land heads. In this section, I examined three versions of the objection, each putting forward a different reason for **3**. The chance-based version failed for being a mere coherence requirement. The risk-based version was more convincing, but was distinctly telic. And the evidentialist version was either non-means-ends (thus not relevant to us), non-poric, or non-successful. I started this section by referencing the strong and widespread intuition to the effect that the agent of premise **3** is indeed irrational, and by advocating for a philosophical account of this intuition. We are now in a position to do this. For the telic Bayesian, the intuition that the agent is irrational expresses a judgement, maybe to the effect that she would be badly off if she had this credence. But for the poric Bayesian, it expresses a kind of lament—how could the agent's means of inquiry warrant so little?

6.

My primary aim in this paper has been to provide an answer to the titular question: why subjectivism? The answer I have provided is that, on the poric account of rationality, according to which epistemic

[38] There is nothing special about this case, it is fairly standard within the formal evidentialist literature.

[39] On the poric interpretation of evidentialism, although **(1)** is false, we get **(2)** for free—to be justified for the poric just means to be warranted by the means of inquiry. But on the non-poric approach, the normative claim is non-analytic and must be defended: Why accept that believing in line with one's evidence is a good way of achieving accuracy?

states are justified to the extent that they are warranted by the agent's means of inquiry, subjectivism is an appealing view. Not only is there a convincing argument in favour of the position (§3), but the typical arguments against it fail to have the required grip (§4–5). But, alongside this, I have also tried, if not to vindicate, at least to countenance the widespread anti-subjectivist beliefs, by conceding that they might be appealing on the alternative (and widely assumed), telic account of rationality. It follows that the debate between subjectivists and anti-subjectivists is best seen, not as a debate about particular putative norms of rationality, but as a debate about which approach to epistemic justification to adopt.

A natural reaction to this state of affairs would be to rule the telic/poric dispute—and the subjectivist/anti-subjectivist dispute with it—as merely verbal. But I think we should resist this urge. In general, in order to ascertain whether two opponents are talking to or past one another, one must determine whether they are trying to answer the same question. In the case of the poric and the telic, it is plausible that they are. Firstly, note that for both, epistemology is a goal-directed endeavour. This is obvious in the case of the telic, but also holds of the poric. Indeed, the poric mandate to adopt the epistemic states warranted by one's means of inquiry cannot even be formulated without presupposing the existence of an epistemic aim: what is warranted is what can be built *up* from the means—built that is, in the direction of the aim. (There is an interesting question, of course, about what it is that is to be achieved—whether the goal is alethic, or pragmatic, or both—what shape it takes beyond this—but what matters here is a more abstract question about the structure of epistemic endeavours.) And secondly, note that achievement is a matter, not just of where one ends up, but also of what one started with. That the poric presupposes this is obvious, but the telic does too. Indeed the typical telic is not merely instructing agents to conform to the epistemic goal(s): he does not rule for instance that agents ought to have all and only true beliefs. Instead, he insists that agents ought to have the best beliefs that they could have, given their epistemic standing. In sum, the telic and the poric can plausibly be portrayed as trying to answer the same question: that of what agents can epistemically achieve on the basis of their means. And given that they are both trying to answer the same question, but that their responses differ, at most one of them can be right. This is why I think that a diagnosis of verbal dispute would be misguided.[40]

The telic could resist this, most obviously by denying that he sees epistemic endeavours through the lens of achievement. He would then, of course, have to explain what he thinks epistemic activity is about. But this insistence that the subjectivism/anti-subjectivism debate is merely verbal would not frustrate my main aim. Indeed, if I am to convince the reader to take subjectivism seriously, I need only to persuade her of two weaker claims: **(1)** that the means-ends conception of epistemic activity is a valuable one; and **(2)** that the poric account of epistemic justification plays a valuable role on this conception. And I think that these claims are clearly true. An important part of both scientific and quotidian endeavours is finding out about the world. And the way that we go about this kind of inquiry

[40] By contrast, the evidentialist who is involved in what Berker (2013b) calls a "sideways" endeavour, namely that of merely reflecting one's evidence, is not asking about what can be achieved. So were this evidentialist and the typical telic Bayesian engaged in a dispute, it would be a verbal one.

is using the means at our disposal: reason, the senses, etc. So whether one accepts my claim that the telic and the poric are actually disagreeing, one must accept my claim that poric epistemology is important—and thus, must take subjectivism seriously.

This is, in some sense, unfortunate. For subjectivism is a very pessimistic answer to the question of what one can epistemically achieve: if what I have said in this paper is right, then there are no epistemic attitudes that are warranted towards propositions in the relevant domain of inquiry, the empirical unobserved. But note that I have left one assumption unexamined: the dualism assumption, according to which agents have two means of inquiry—reason, and observation. In the rest of this paper, I examine various ways of rejecting the dualism assumption, and I conclude pessimistically.

There are three ways in which one could reject the dualism assumption. **(a)** One could argue that agents lack one of the two means posited by the assumption; agents lack either observation or reason as a means of inquiry. Now, not only will this strategy clearly not help us escape the grip of subjectivism, but it would also force us to abandon the building blocks of Bayesianism: agents must be able to reason if they are to satisfy probabilism, and to observe if they are to satisfy conditionalisation. This brings us to: **(b)** one could reject that observation and reason are two distinct capacities. Once again, this strategy will clearly not go very far in helping us refute subjectivism, and once again, it would force us to abandon some foundations of Bayesian epistemology; in particular, the contention that learning takes place through the acquisition of (pure) observational evidence.[41] This in turn brings us to: **(c)** one could reject that observation and reason are the only means of inquiry. This third option seems highly promising, for there is a widespread way by which people come to have justified beliefs about the world beyond reason and observation: testimony.

This has been extensively discussed in the Bayesian literature, under the banner of *deference principles*. The details are subtle, but we can gloss over them for our purposes. A deference principle says that an agent's credence p in a proposition X, given that an expert's credence function is p^+, should be equal to the value that p_E ascribes to X. So, where E_{p^+} denotes the proposition that the expert's credence function is p^+:

$$p(X|E_{p^+}) = p^+(X)$$

Where the expert is one's better informed future self, we call this a *reflection principle* (van Fraassen, 1984), and where the expert is a better informed other agent, we call it an *expert deference principle* (Elga, 2007). Various arguments have been proposed, including a Dutch Book argument by van Fraassen (1984), and an accuracy argument by Easwaran (2013). Now, like chance functions (as discussed in §5), deference principles can only help constrain an agent's credences if she in fact comes to know the shape of expert functions. This is where testimony might come in. But in fact, it might not need to: expert functions are possible to come to know by reason—reasoning about what one will believe in cases of reflection—and observation—hearing from experts about their credences in cases of expert

[41] The claim that observation is thought-laden is clearly plausible—and I think constitutes an under-explored challenge to Bayesian methodology. But this is for another time.

deference. One might interpret this as counting testimony as a kind of observation: the observation of others' credences.

But whether one sees testimony as a kind of observation or as an independent means of inquiry, one is unlikely to get very far. For if Bayesian epistemology applies to everyone—as it surely does if it applies to anyone—then that which we can come to know by testimony is limited to what others can come to know by observation. And many—most!—of the propositions of interest to an agent are not only unobserved by herself, but by all agents. (Will this coin land heads? Will the sun rise tomorrow? Will my loved ones become infected with Covid-19?) This is so in large part because credences guide action, and the propositions that matter to the choice of the deliberating agent are about the future. When the rain-averse agent wonders whether he ought to take his umbrella, what matters to him is whether it will rain after he leaves the house. When the policy-maker deliberates on climate action, what matters to her is what will happen at the time she plans for the policies to be implemented. And what additional means of inquiry that agents have could bridge that gap? So I conclude, pessimistically, that subjectivism is unfortunately true.

15th May 2020

References

Berker, S. (2013a). Epistemic Teleology and the Separateness of Propositions. *Philosophical Review 122*(3), 337–393.

Berker, S. (2013b). The Rejection of Epistemic Consequentialism. *Philosophical Issues 23*(1), 363–387.

Bernoulli, D. (1738/1954). Exposition of a New Theory on the Measurement of Risk. *Econometrica 22*(1), 23–36.

Bolker, E. (1966). Functions Resembling Quotients of Measures. *Transactions of the American Mathematical Society 124*(2), 292–312.

Bradley, R. (2017). *Decision Theory with a Human Face*. Cambridge University Press.

Bradley, S. (2019). Imprecise Probabilities. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.

Briggs, R. A. and R. Pettigrew (2020). An Accuracy-Dominance Argument for Conditionalization. *Noûs 54*(1), 162–181.

Bronfman, A. (2014). Conditionalization and Not Knowing That One Knows. *Erkenntnis 79*(4), 871–892.

Buchak, L. (2013). *Risk and Rationality*. Oxford University Press.

Canson, C. (forthcoming). Objectivity and the Method of Arbitrary Functions. *British Journal for the Philosophy of Science*.

Carnap, R. (1950). *Logical Foundations of Probability*. University of Chicago Press.

Carr, J. (2017). Epistemic Utility Theory and the Aim of Belief. *Philosophy and Phenomenological Research 95*(3), 511–534.

Chalmers, D. J. (2011). Frege's Puzzle and the Objects of Credence. *Mind 120*(479), 587–635.

Conee, E. and R. Feldman (2004). *Evidentialism: Essays in Epistemology*. Oxford University Press.

Davidson, D., J. C. C. McKinsey, and P. Suppes (1955). Outlines of a Formal Theory of Value. *Philosophy of Science 22*(2), 140–160.

de Bona, G. and J. Staffel (2018). Why Be Coherent? *Analysis 78*(3), 405–415.

de Finetti, B. (1937). La Prévision: Ses Lois Logiques, Ses Sources Subjectives. *Annales de l'Institut Henri Poincaré 7*, 1–68.

de Finetti, B. (1962). Does it Make Sense to Speak of Good Probability Appraisers? In I. J. Good (Ed.), *The Scientist Speculates. An Anthology of Partly-Baked Ideas*, pp. 357–64. Basic Books.

Descartes, R. (1641/1993). *Meditations on First Philosophy*. Hackett Publishing.

Dietrich, F., C. List, and R. Bradley (2016). Belief Revision Generalized: A Joint Characterization of Bayes's and Jeffrey's Rules. *Journal of Economic Theory 162*, 352–371.

Dorst, K. (forthcoming). Evidence: A Guide for the Uncertain. *Philosophy and Phenomenological Research*.

Easwaran, K. (2013). Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection. *Philosophy of Science 80*(1), 119–142.

Easwaran, K. (2017). The Tripartite Role of Belief: Evidence, Truth, and Action. *Res Philosophica 94*(2), 189–207.

Easwaran, K. and B. Fitelson (2012). An 'Evidentialist' Worry About Joyce's Argument for Probabilism. *Dialectica 66*(3), 425–433.

Elga, A. (2007). Reflection and Disagreement. *Noûs 41*(3), 478–502.

Feldman, R. and E. Conee (1985). Evidentialism. *Philosophical Studies 48*(1), 15–34.

Galavotti, M. C. (2011). The Modern Epistemic Interpretations of Probability: Logicism and Subjectivism. In S. H. Dov M. Gabbay and J. Woods (Eds.), *Handbook of the History of Logic*, Volume 10, pp. 153–203.

Galavotti, M. C. (2016). The Origins of Probabilistic Epistemology. In A. Hájek and C. Hitchcock (Eds.), *The Oxford Handbook of Probability and Philosophy*. Oxford University Press.

Galavotti, M. C. (2018). Who is Afraid of Subjective Probability? In A. Giordani and C. de Florio (Eds.), *From Arithmetic to Metaphysics: A Path Through Philosophical Logic*, pp. 151–158. De Gruyter.

Gallow, J. D. (2017). Diachronic Dutch Books and Evidential Import. *Philosophy and Phenomenological Research 99*(1), 49–80.

Gallow, J. D. (forthcoming). Updating for Externalists. *Noûs*.

Greaves, H. (2013). Epistemic Decision Theory. *Mind 122*(488), 915–952.

Greaves, H. and D. Wallace (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind 115*(459), 607–632.

Hacking, I. (1967). Slightly More Realistic Personal Probability. *Philosophy of Science 34*(4), 311–325.

Hájek, A. (2019). Interpretations of Probability. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.

Hume, D. (1748/1999). *An Enquiry Concerning Human Understanding*. Oxford University Press.

Jeffrey, R. C. (1965). *The Logic of Decision*. University of Chicago Press.

Jeffrey, R. C. (1968/1992a). Probable Knowledge. In *Probability and the Art of Judgment*. Cambridge University Press.

Jeffrey, R. C. (1991/1992b). Radical Probabilism. In *Probability and the Art of Judgment*. Cambridge University Press.

Joyce, J. (1998). A Non-Pragmatic Vindication of Probabilism. *Philosophy of Science 65*(4), 575–603.

Joyce, J. (1999). *The Foundations of Causal Decision Theory*. Cambridge University Press.

Joyce, J. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives 19*(1), 153–178.

Kelly, T. (2003). Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research 66*(3), 612–640.

Keynes, J. M. (1921). *A Treatise on Probability*. Wildside Press.

Konek, J. and B. A. Levinstein (2019). The Foundations of Epistemic Decision Theory. *Mind 128*(509), 69–107.

Lange, M. (1999). Calibration and the Epistemological Role of Bayesian Conditionalization. *Journal of Philosophy 69*(9), 294–324.

Leitgeb, H. and R. Pettigrew (2010a). An Objective Justification of Bayesianism I: Measuring Inaccuracy. *Philosophy of Science 77*(2), 201–235.

Leitgeb, H. and R. Pettigrew (2010b). An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of Science 77*(2), 236–272.

Levi, I. (1985). Imprecision and Indeterminacy in Probability Judgment. *Philosophy of Science 52*(3), 390–409.

Lewis, D. K. (1980). A Subjectivist's Guide to Objective Chance. In R. C. Jeffrey (Ed.), *Studies in Inductive Logic and Probability*, pp. 83–132. University of California Press.

Lewis, D. K. (1999). Why Conditionalize? In *Papers in Metaphysics and Epistemology*, pp. 403–407. Cambridge University Press.

Lyons, J. (2017). Epistemological Problems of Perception. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.

Pettigrew, R. (2012). Accuracy, Chance, and the Principal Principle. *Philosophical Review 121*(2), 241–275.

Pettigrew, R. (2013a). A New Epistemic Utility Argument for the Principal Principle. *Episteme 10*(1), 19–35.

Pettigrew, R. (2013b). Accuracy and Evidence. *Dialectica 67*(4), 579–596.

Pettigrew, R. (2016a). *Accuracy and the Laws of Credence*. Oxford University Press.

Pettigrew, R. (2016b). Accuracy, Risk, and the Principle of Indifference. *Philosophy and Phenomenological Research 92*(1), 35–59.

Pettigrew, R. (2018). Making Things Right. In H. K. Ahlstrom-Vij and J. Dunn (Eds.), *Epistemic Consequentialism*, pp. 220–239. Oxford University Press.

Pettigrew, R. (2019). Epistemic Utility Arguments for Probabilism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*.

Pettigrew, R. (2020). *Dutch Book Arguments*. Cambridge University Press.

Pettigrew, R. (forthcominga). Logical Ignorance and Logical Learning. *Synthese*.

Pettigrew, R. (forthcomingb). What is Conditionalization, and Why Should We Do It? *Philosophical Studies*.

Ramsey, F. P. (1926/1931). Truth and Probability. In *The Foundations of Mathematics*. Kegan Paul.

Roussos, J. (ms). Formal Philosophy as Modelling.

Russell, B. (2014). *A Priori* Justification and Knowledge. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.

Savage, L. J. (1954). *The Foundations of Statistics*. Wiley.

Schoenfield, M. (2015). Bridging Rationality and Accuracy. *Journal of Philosophy 112*(12), 633–675.

Schoenfield, M. (2017a). Conditionalization Does Not (in General) Maximize Expected Accuracy. *Mind 126*(504), 1155–1187.

Schoenfield, M. (2017b). The Accuracy and Rationality of Imprecise Credences. *Noûs 51*(4), 667–685.

Sliwa, P. and S. Horowitz (2015). Respecting *All* the Evidence. *Philosophical Studies 172*(11), 2835–2858.

Staffel, J. (2017). Should I Pretend I'm Perfect? *Res Philosophica 94*(2), 301–324.

Strevens, M. (2013). *Tychomancy*. Harvard University Press.

Sylvan, K. (2018). Veritism Unswamped. *Mind 127*(506), 381–435.

Sylvan, K. (2020). An Epistemic Non-Consequentialism. *Philosophical Review 129*(1), 1–51.

Teller, P. (1973). Conditionalization and Observation. *Synthese 26*(2), 218–258.

van Fraassen, B. (1983). Calibration: A Frequency Justification for Personal Probability. In R. Cohen and L. Laudan (Eds.), *Physics, Philosophy, and Psychoanalysis*. Reidel.

van Fraassen, B. (1984). Belief and the Will. *Journal of Philosophy 81*(5), 235–256.

van Fraassen, B. (1989). *Laws and Symmetry.* Oxford University Press.

White, R. (2009). Evidential Symmetry and Mushy Credence. *Oxford Studies in Epistemology 3*(1), 161–186.

Williamson, J. (2010). *In Defence of Objective Bayesianism.* Oxford University Press.

Zabell, S. (2016). Symmetry Arguments in Probability. In A. Hájek and C. Hitchcock (Eds.), *The Oxford Handbook of Probability and Philosophy.* Oxford University Press.