The Retrospective Account of Bayesian Updating

Abstract. This paper argues for a new account of Bayesian updating by taking a retrospective approach to diachronic coherence. This approach says that an agent is diachronically coherent whenever the information she has revised her beliefs on satisfies whatever constraint we would want our evidence to satisfy. This approach contrasts with a common way of thinking about the Bayesian framework, according to which it treats evidence as a black box. The aim of this paper is to provide a different interpretation of Bayesianism's main updating constraint by filling in this black box with a Bayesian account of evidence.

1 Introduction

Say I come to you and I ask: How should I revise my beliefs when I learn something new? It's likely you would recommend one of two things. On the one hand, you might tell me that when I get evidence, I should revise my beliefs using *the right updating procedure*. I should update my beliefs using a rule that is likely to result in beliefs that are more accurate or that have some other valuable property. On the other hand, you might tell me that when I update my beliefs, I should revise my beliefs using *the right information*: I should update on information that is likely to result in beliefs that are more accurate or more valuable. This is similar advice given from two different perspectives. In the first case, you assume that I already have some evidence in hand, and you tell me what I should do with it. In the second case, you assume that I already have some updating rule in mind, and you tell me what I should apply it to. In both cases, you assume that what I am after is a more accurate, or valuable, set of beliefs.

Bayesians have traditionally adopted the first perspective. According to Bayesian orthodoxy, evidence is a black box that gives rise to the norm of conditionalization.[1] *Whatever* it means for an agent to get evidence, when this happens, it triggers the requirement to condition on it. This paper argues for a new account of Bayesian updating by defending what I will call *the Retrospective Account*. The Retrospective Approach takes the second perspective described above. Instead of claiming that an agent is diachronically coherent when, assuming she has some piece of evidence, we can infer that she has revised her beliefs by conditioning on it, the Retrospective Approach claims that an agent is diachronically coherent when, assuming she has conditioned on some information, we can infer that this information satisfies whatever constraint we would want our evidence to satisfy.

On the Retrospective Approach, then, whether or not an agent has updated at all, properly speaking, is determined by whether she has conditioned her beliefs on the right information. I argue that one natural interpretation of "the right information" is information that is not inconsistent. In conditioning on inconsistent information, we violate a norm that generalizes Probabilism. On the view I propose, then, synchronic coherence and diachronic coherence turn out to be norms of the very same kind.

This paper makes three assumptions that I should flag up front. The first assumption this discussion will make is that the best interpretation of Bayesian updating is something that is, in some sense, up for grabs. It's often noted that there are many different ways of articulating Bayesianism's main commitments.[2] Given this, a reasonable question to ask, though one that is not asked very often, is what formal properties a Bayesian would want her evidence to have. The aim of this paper is to offer an answer to this question. In answering this question, my paper resolves a certain theoretical puzzle by making sense of the second perspective described above.

There are two more concrete problems that the second perspective described above also allows us to resolve. First, as we'll see in the next section, the literature has struggled to

[1] For this description, see, for instance, Skyrms [1987].

[2] Famously, I.J. Good [1971] claimed that there are at least 46,656 varieties of Bayesians.

provide a satisfying accuracy-based defense of orthodox accounts of Bayesian updating. By contrast, the Retrospective Account I argue for can be given a very simple accuracy argument. This is one of the reasons in its favor.

There's a second reason for favoring my Retrospective Account. Plausibly, a desiderata for any updating rule is that it satisfy the following constraint:

> **Non-Arbitrariness:** The question of *when* an agent revises her beliefs should be irrelevant to whether or not she is diachronically coherent.

This principle captures the plausible thought that an agent's standing as diachronically coherent is a function only of her evidence and her priors. As we will see in §3, while the orthodox account of Bayesian updating fails by the lights of this principle, the Retrospective Approach is in a position to satisfy it. Thus, the first assumption this discussion will make is that there is a better account of Bayesian updating to be had.

The second assumption this discussion will make concerns the status of Bayesianism's norms. It's well-known that Bayesian epistemology is accused of placing unrealistic demands on the epistemic agent. Bayesian epistemology requires that agents be logically omniscient: that they know all logical truths. Bayesian epistemology also requires agents to have computational capacities that far outstrip those of any actual agent.[3] These problems have generated an entire literature of responses, which are well worth a paper of their own. But even without evaluating each of them, the general lesson is clear. We must either revise Bayesian epistemology's norms or we must revise the way we interpret these norms.

I propose that we do the latter. Rather than understand Bayesian epistemology as a set of action-guiding norms, I propose that we understand Bayesian epistemology as a set of evaluative norms. Just as we can say that a knife is good insofar as it's sharp and bad insofar as it isn't, without imposing any obligations on anyone to do anything, we can say that certain credence distributions, or sets of credence distributions, have certain features that make them good or bad without imposing any obligations on anyone to do anything.[4]

[3] For canonical expressions of this complaint, see Harman [1986, p.25-26] and Earman [1992, p. 56].

[4] Note that I will use the terms "evaluative" and "normative" interchangeably in some contexts. Thus, I won't

The Retrospective Approach *entails* the evaluative interpretation of Bayesian epistemology. If whether I have evidence is determined by how I have revised my beliefs already, as the Retrospective Approach implies, there's no sense in which this approach can offer any guidance about how to update. There's no sense in which I can be told what credences I should update to if it is not known in advance what kind of evidence I have. But, not being able to be guided in this way, I can nevertheless still retrospect over my belief revision and ask whether I have revised my beliefs in a way that is good or that is best—in a way that implies that the information I have revised my beliefs on has those formal properties that make it evidence.

Finally, the third assumption this discussion will make is that when deciding on the best interpretation of Bayesian updating, we should consider Bayesian updating in its most general form. Although most discussions about diachronic coherence begin from the assumption that we get propositional evidence, or certain evidence, I'm going to assume that Bayesian epistemology includes norms that govern cases where we get *uncertain* evidence as well. While this commitments is normally taken to distinguish the view many refer to as "Jeffrey conditioning" (see Jeffrey [1965]), my proposal will distance itself from Jeffrey conditioning by distancing itself from many of its motivations. My account develops some unexplored possibility space, in the realm of Bayesian updating, by combining some of the formal features of Jeffrey conditioning with a normative commitment of regular, Bayesian conditioning: the commitment to update on evidence that is unrevisable.

An updating framework that can accommodate uncertain evidence shouldn't need too much defending. Arguably, the situations we find ourselves in most often are cases where we *do* seem to have uncertain evidence. I assume the reason we idealize away from the sorts of cases that Jeffrey conditioning covers isn't because they aren't important, but because we assume they will be covered in the same way as regular Bayesian conditioning. It's good practice to abstract away from the uncertainty of evidence when this detail isn't important for whatever purpose the framework is being put towards. However, when the problem at issue

assume the normative to be necessarily tied to action-guidingness.

is how we should interpret the Bayesian framework itself, the structure of uncertain evidence is surely a detail worth considering.

If we assume that beliefs come in degrees, then, it seems arbitrary to withhold this assumption from our evidence. As Jeffrey [1992, p.11] put it, opinion, no matter how it comes to us, should be probabilities "all the way down, to the roots." Jeffrey's rule represents a particularly strong way of reflecting the idea that what makes some view Bayesian is its commitment to Probabilism. The account of diachronic coherence defended in this paper does this as well; on my account of it, diachronic coherence is Probabilism all the way down. However, though my account is best articulated from a perspective that assumes that we sometimes get uncertain evidence, the results of this discussion don't depend upon this assumption. One can reject this assumption and still accept my proposal.

Here's how the discussion will go. In §2, I describe the Retrospective Approach to updating, which combines features of two accounts of diachronic coherence commonly found in the literature. On the Retrospective Approach, an agent is diachronically coherent whenever the information she has conditioned her beliefs on satisfies whatever constraint we would want our evidence to satisfy. In §3, I show that **Non-Arbitrariness** on the Retrospective Approach implies a very plausible constraint on evidence. It implies the constraint that the evidence an agent has accumulated over time be representable as a probability function. I call this constraint *Diachronic Probabilism*. In §4, I consider some objections to Diachronic Probabilism and argue that it accommodates certain externalist intuitions many have about evidence. In §5, I argue that Diachronic Probabilism can also be vindicated by an accuracy argument. I conclude in §6.

A quick word about terminology. In what follows, I will assume that "diachronic coherence" is synonymous with "updating" and also with "conditioning or revising on one's evidence". I will assume that these are all success terms, in contrast with "conditioning or revising on one's information", though I will sometimes violate these linguistic stipulations when context would make it too awkward to do otherwise.

2 The Retrospective Approach

Standard Bayesianism assumes that the degrees of confidence, or credences, that an agent assigns to the members of the set X of propositions she entertains can be represented as an assignment of real numbers to those propositions. It further assumes that, for all $A \in X$, the following synchronic norm of **Probabilism** governs this assignment, with its three basic axioms:

1. Non-negativity $(cr(A) \geq 0$, for all $A)$
2. Normality (if A is a tautology, then $cr(A) = 1)$
3. Finite Additivity (if A and B are incompatible, then $cr(A) + cr(B) = cr(A \vee B))$

Most Bayesians also take an agent's credences to satisfy a diachronic constraint. This constraint tells us how we should revise our beliefs when we get some new evidence. It tells us that we should be guided by our conditional probabilities. Normally, this is interpreted as the requirement that our current unconditional probabilities should equal our prior probabilities conditional on the evidence that we've gotten. However, as Blackwell and Girshick [1979, p.221] show, this rule falls out of the more general requirement that our conditional probabilities remain fixed, or "rigid", on the evidence that we've gotten in between this transition.[5] Some proposition, A, is rigid on E, for the transition from p to p', just in case the following holds:

Rigidity: $p(A \,|\, E) = p'(A \,|\, E)$, if defined.

If our conditional probabilities are to guide our updates, they must remain the same throughout these updates. That's what Rigidity says. While any Bayesian account of diachronic coherence must be committed to Rigidity, there are different accounts one can give of how this

[5] For discussion of this result, see Diaconis and Zabell [1982, p.824]. For other discussions that emphasize this guiding feature of our conditional probabilities, see Jeffrey [1965], Jeffrey [1970], Skyrms [2006] and, more recently, Weisberg [2014].

mathematical property is related to the normative property of diachronic coherence. This section considers three ways of spelling out this relationship.

On the one hand, we might think that being diachronically coherent over a belief transition simply requires that there be *some* way of describing this transition as rigid. This idea gives rise to what we might call the Deflationary Account of updating, according to which there is nothing more to being diachronically coherent than to remaining probabilistically coherent over time. This account is inspired by the fact that, as Blackwell and Girshick [1979] famously show, Rigidity is simply a consequence of the probability axioms. For any probabilistic belief transition, there is *some* partition that is sufficiently fine-grained (a "sufficient partition") to represent the transition in question as rigid on this partition.[6] Such an account of diachronic coherence can be formulated in the following way:

> **A Deflationary Account:**
>
> Where $\mathcal{S}=\{B_1, ..., B_n\}$ is a set of beliefs that form a partition, and where, p, is the agent's prior probability distribution, and, p', is the agent's posterior probability distribution, an agent is diachronically coherent just in case the following holds:
>
> $\forall p \forall p' \exists \mathcal{S} (\forall B_i \in \mathcal{S})\ \forall A (p(A \left| B_i) = p'(A \right| B_i)$), if defined.

The Deflationary Account entails that "getting evidence" and "updating on evidence" are one and the same process. Moreover, it entails that diachronic coherence on the Bayesian framework is no stronger than Probabilism. These last two consequences are at odds with the way we normally think about Bayesian updating. A more familiar way of understanding the relationship between getting evidence and revising in accordance with Rigidity takes these to be two different steps in the updating process. We get some piece of evidence *first* and then the only right way to revise one's beliefs is in a way that is rigid on this *particular* piece of

[6] See also Diaconis and Zabell [1982, p.824] and Jeffrey [1992, p.122-128] for discussion of this result and related issues. As Diaconis and Zabell note, there will be cases where our conditional probabilities are undefined for some partition—namely, where we assign a member of our partition a credence of zero. However, their result still holds for all updates, if we take a sufficient partition to be a partition that is sufficient to represent a probabilistic belief transition as an update that is conditional on every proposition in this partition for which a conditional probability is defined (*Cf.* Blackwell and Girshick [1979, §8.4.3]).

evidence. According to this more orthodox approach to updating, which was sketched in the opening paragraphs of this discussion, the process of getting evidence is the exogenous part of an update: it is the assumption we hold fixed in order to determine whether our updating norm has been satisfied. By contrast, the process of revising one's beliefs on this evidence in accordance with Rigidity is the endogenous part of an update. It is the norm-governed part of the update.[7] Whatever it means for an agent to get evidence, when this happens, it triggers the requirement to condition on it. This picture of updating can be captured by the following conditional norm:

The Orthodox Account:

Where $\mathcal{S}=\{B_1, ..., B_n\}$ is a partition of propositions whose values have changed non-inferentially, and where, p, is the agent's prior probability distribution, and, p', is the agent's posterior probability distribution, an agent is diachronically coherent just in case the following holds:

$\forall p \forall p' \forall B_i(\exists S(B_i \in S) \rightarrow \forall A(p(A|B_i) = p'(A|B_i)))$, if defined.[8]

The Orthodox Account says that an agent is diachronically coherent whenever she gets evidence and revises her beliefs in a way that is rigid on this evidence. Since a probabilistic agent might fail to be coherent in this way—since she might fail to revise her beliefs on the particular piece of evidence she has in hand—the Orthodox Account is stronger than Probabilism, which seems like something we would want.

There are worries for the Orthodox Account though. A constraint on the adequacy of any Bayesian rule is that it be able to be provided with an accuracy-based defense. The program of "accuracy-first epistemology" is charged with doing precisely this. It provides us with theorems that show that, relative to some reasonable class of inaccuracy measures, many Bayesian-type norms satisfy analogues of the decision-theoretic norms of dominance

[7] The endogenous/exogenous distinction was introduced by Howson and Urbach [2006]. Note that exogenous here need not mean causal. I am simply taking it to mean 'outside the domain of the normative'.

[8] For some canonical examples of the Orthodox Account, see Skyrms (1987), Howson and Urbach (1989) and Jeffrey (1992), among many others.

(namely, accuracy-dominance) and maximize expected utility (namely, minimize expected inaccuracy).[9]

Unfortunately, the updating rules that are vindicated by accuracy-first epistemology only very loosely track the Orthodox Account. While Leitgeb and Pettigrew [2010b] defend a diachronic Bayesian updating rule, their rule doesn't take updates to be rigid.[10] Conversely, while both Greaves and Wallace [2006] and Pettigrew [2016] provide an accuracy argument for a rigid Bayesian updating rule, this updating rule, which is justified by the norm to minimize expected inaccuracy, differs from the Orthodox Account, and thus imposes additional restrictions on the Orthodox Account, in 1) not permitting the agent to update on uncertain evidence and in 2) requiring that the agent know in advance the partition that she updates on.[11] This last feature is especially problematic, for it entails that Bayesianism cannot accommodate cases of "reasoning in the wild". While I sometimes update my beliefs in response to some particular question that I am trying to settle, at other times, I am simply struck by some information. I go outside, not intending to settle the question of whether or not it is raining, and yet I nevertheless come to discover that it is and, so, go on to adjust my beliefs accordingly.

There is a third approach to diachronic coherence that avoids both of these problems. It does this by reversing the relation that the Orthodox Account prescribes. Instead of requiring that Rigidity apply, *conditional* on us having some evidence, this approach requires that the information we have revised our beliefs on satisfy some evidential constraint, *conditional* on us having revised by Rigidity. More carefully, this approach claims that some probabilistic belief revision constitutes an update just in case, (1) there is a sufficient partition for that transition and, (2) the values over that sufficient partition satisfy whatever constraint we would want our evidence to satisfy:

[9] For some of the most prominent of these arguments, see Greaves and Wallace [2006], Joyce [2009], Leitgeb and Pettigrew [2010a], Leitgeb and Pettigrew [2010b] and Pettigrew [2016].

[10] *Cf.* Pettigrew [2016].

[11] For this reason, versions of this norm are sometimes referred to as "Plan Conditionalization" (*Cf.* Easwaran [2013]).

A Retrospective Approach to Updating:

Where $\mathcal{S}=\{\mathrm{B}_1, ..., \mathrm{B}_n\}$ is a set of beliefs that form a partition, and where, p, is the agent's prior probability distribution and, p', is the agent's posterior probability distribution, an agent is diachronically coherent just in case the following two conditions hold:

1) $\forall p \forall p' \exists S(\forall \mathrm{B_i} \in S)\ \forall A(p(\mathrm{A}\big|\mathrm{B_i}) = p'(\mathrm{A}\ \big|\mathrm{B_i})\ )$, if defined, and

2) $\mathcal{S}_1, \mathcal{S}_2, ..., \mathcal{S}_n$ have values that satisfy our constraint on evidence.

Like the Orthodox Account, the Retrospective Approach also has an endogenous part and an exogenous part—but they are reversed. Instead of assuming that an agent has some piece of evidence, and going on to ask whether she has arrived at the right posterior credences by conditioning, the Retrospective Approach assumes that an agent has arrived at her posterior credences by conditioning, and goes on to ask whether she has conditioned on the right information: on evidence.[12] Like the Orthodox Account, the Retrospective Approach is stronger than Probabilism. However, while the Retrospective Approach implies that there is some question that is answered by an agent's update, in virtue of this update implying a certain evidence partition, it takes no stand on whether this question was ever asked. By neither requiring, nor precluding, that we always reason in response to some question we have in mind in advance of our update, the Retrospective Approach is consistent with a variety of forms of reasoning. In addition, the Retrospective Approach both accommodates updates on uncertain evidence and, as we will see in §5, can be defended by a relatively simple accuracy argument.

[12] This approach is inspired by, though substantially more developed than the brief discussion of "retrospective conditioning" that appears in Diaconis and Zabell [1982, p.822]. This approach is also similar to the approach developed in Cassell [forthcoming], though as we will see, it yields an account of updating with a different structure that is motivated and defended on very different grounds.

Summing up, then, there are two perspectives from which to develop an account of diachronic coherence that is stronger than Probabilism. These perspectives correspond to two different questions we might ask. On the one hand, we might ask: *assuming that this is my evidence, have I revised my beliefs by conditioning on it?* If I have, then according to the Orthodox Account, I am diachronically coherent. Alternatively, we might ask: *assuming I have conditioned my beliefs on some information, does this information have features that make it evidence?* If it does, then according to the Retrospective Approach, I am diachronically coherent.

3 The Retrospective Account

We've considered three ways of understanding what it might mean for a Bayesian agent to be diachronically coherent. One of these ways is not like the others though. Unlike the Deflationary Account and the Orthodox Account, the Retrospective Approach is a schema. It leaves open what features some information must have in order to be evidence.

In this section, I propose and defend a way of filling out the Retrospective Approach. My aim in doing this is to provide a proof of concept for the Retrospective Approach: to show that there is a way of developing this approach into an account. I'll argue that the Retrospective Approach improves upon the Orthodox Account in its ability to satisfy the following constraint:

> **Non-Arbitrariness:** The question of *when* an agent revises her beliefs should be irrelevant to whether or not she is diachronically coherent.

The crucial insight will be that since the Retrospective Approach leaves open what it takes for some information to be evidence, we can simply use **Non-Arbitrariness** to fill this in. We can do this in a way that does not appeal to anything above and beyond what Bayesians are already committed to.

We can start to see how by drawing on a distinction from Meacham [2016a]. Meacham distinguishes what he calls *sequential updating* from *interval updating*. On the sequential model of updating, we revise our beliefs at each time that we get some new piece of evidence.

The sequential model of updating tells us to condition our current priors on our current information, where our current priors were generated by conditioning the priors that we had before those on the information we had a moment ago, and so on. Each of the approaches described in the previous section assumes that we update sequentially in this sort of way.

As Meacham notes, the sequential model of updating seems too strong. We should be able to relax the requirement that an agent must update at each and every time she gets some new piece of information.[13] If an agent has gotten information at several different times in between t_0 and t_6, she should be able to condition on the information she has gotten in between t_0 and t_3, and then on the information that she has gotten in between t_3 and t_6. Or she should be able to condition on the evidence that she's gotten in between t_0 and t_6 all at once. She should be able to condition on the evidence she's received over any interval of time.

Closely related to this last idea is the thought that sequential and interval models of updating should coincide. If an agent has gotten information at several different times in between t_0 and t_6, she should be able to condition on all of this information at once, or on each piece of this information in turn, without this making a difference to her posterior credence distribution. With this in mind, a natural way of articulating what **Non-Arbitrariness** amounts to on the Orthodox Account is to say that it amounts to the following:

> **Orthodox Non-Arbitrariness:** Our posterior credence distribution should be the same (uniquely correct one) regardless of whether we update using sequential or interval models of updating.

The Orthodox Account famously violates **Orthodox Non-Arbitrariness**. Say I get some evidence that leaves me uncertain about whether a difficult math proof is correct, and that this evidence leads me to change my credence in this proposition to .5. (Perhaps this evidence is my having run through this proof quickly in my head.) And then I get evidence that leads me to change my credence in this proposition to .6. (Perhaps this evidence is the testimony of a friend who is reliable about these matters sixty percent of the

[13] For this idea, see also Skyrms [1983], Titelbaum [2013], and Meacham [2016b].

time.) In this case, if I update sequentially, the probative value of my most recent evidence swamps that of my previous evidence. On the sequential approach, my credence that this proof is correct will be .6. By contrast, if I update using the interval model of updating, my update is undefined since we can't update on the same evidence partition with two different values. On the Orthodox Account, then, sequential and interval updating come apart. The Orthodox Account fails to satisfy **Non-Arbitrariness**.[14]

A natural way of articulating what **Non-Arbitrariness** amounts to on the Retrospective Approach is to say that it amounts to the requirement that the information an agent has revised her beliefs on satisfies whatever constraint we would want our evidence to satisfy, regardless of whether this constraint is imposed on all of this information at once, or on each piece of this information in turn. Suppose that H is the proposition that *it will hail today*. And suppose that, at t_1, an agent revises her beliefs on the weighted partition $\{(H, .5), (\overline{H}, .5)\}$ and then, at t_2, she revises her beliefs on the weighted partition $\{(H, .6), (\overline{H}, .4)\}$. One way of interpreting what **Non-Arbitrariness** requires is that if *each* of these two weighted partitions satisfies some normative constraint on evidence, then the total set $\{\{(H, .5), (\overline{H}, .5)\}, \{(H, .6), (\overline{H}, .4)\}\}$ should also satisfy this constraint:

> **Retrospective Non-Arbitrariness:** The information that we revise our beliefs on should satisfy our evidential constraint whether we use sequential or interval models of updating.

Since, unlike the Orthodox Account, the Retrospective Approach is just a schema, we can't say that it violates **Retrospective Non-Arbitrariness**. Indeed, we can *guarantee* that it doesn't simply by taking the latter to *be* our account of evidence:

[14] There are interesting questions about when the swamping feature of sequential, orthodox updating gets us the intuitively right result and when it doesn't. The most notable of these cases are ones where our updates seem to be justified by the phenomenal character of our experiences (see Lange [2000]). However, since the intuitions that underwrite these cases assume that the inputs to the framework are experiences rather than weighted evidence partitions, they assume the action-guiding interpretation of Bayesian updating that we have rejected. Thus, they aren't relevant to whether **Orthodox Non-Arbitrariness** is satisfied. (More on this in §4.)

A Retrospective Account of Updating:

Where $\mathcal{S}=\{B_1, ..., B_n\}$ is a set of beliefs that form a partition, and where p and p' are probability distributions that are held sequentially at times t and t', an agent is diachronically coherent just in case the following two conditions hold:

1) $\forall p \forall p' \exists S(\forall B_i \in S)\ \forall A(p(A|B_i) = p'(A\,|B_i)\,)$, if defined, and

2) $\{\mathcal{S}_1, \mathcal{S}_2, ..., \mathcal{S}_n\}$ have values that satisfy **Retrospective Non-Arbitrariness**

Taken together, the first and second constraint say that an agent who is looking back and evaluating information she has already revised her beliefs on sequentially should recognize that her constraint on evidence would have been satisfied, even if she had revised her beliefs at different times than she actually did. What it takes to comply with this requirement will, of course, depend upon what our first-order constraint on evidence happens to be. While there are many different constraints we could adopt, I want to offer a proposal that does not involve appealing to anything beyond what Bayesians are already committed to.

A norm that is trivially satisfied by all sequential accounts of evidence is that evidence be probabilistic. If we assume that our credences should, at every moment, satisfy the probability axioms, then an account of updating that tells us to update at every moment that we get some new evidence will entail that our evidence satisfies Probabilism. Indeed, the discussion of the previous section suggests that this is the *only* constraint that all sequential accounts of updating have in common. Even the Deflationary Account satisfies this constraint. If Probabilism is a constraint on sequential evidence, then by **Retrospective Non-Arbitrariness**, it must also be a constraint on the information accumulated over an interval of time.

Of course, Probabilism *can't* be a constraint on the information accumulated over an interval of time, since Probabilism is a constraint on a function, and evidence that has been accumulated over an interval of time is not even a function to begin with. However, there is a more general constraint in the neighborhood that has Probabilism as a special

case. Every probability function, and so all sequential evidence, is, trivially, *representable as a probability function*. Evidence is representable as a probability function just in case there is a probabilistically weighted partition that represents all the constraints that this evidence imposes. In addition to sequential evidence, it's easy to see that evidence that is accumulated over time is also sometimes representable as a probability function. An agent who gets $\{(E, .5), (\overline{E}, .5)\}$ at t_1, and then gets $\{(F, .5), (\overline{F}, .5)\}$ at t_2, has as evidence $\{(EF, .25), (\overline{E}F, .25), (E\overline{F}, .25), (\overline{EF}, .25)\}$. By contrast, consider again our agent from before who gets as evidence the proposition that a difficult math proof is correct with a credence of .5, and then gets the same proposition as evidence with a credence of .6. Such an agent who updates on $\{(H_1, .5), (\overline{H}_1, .5)\}$ at t_1 and then on $\{(H_2, .6), (\overline{H}_2, .4)\}$ at t_2 does not have evidence, properly speaking, since there is no way to represent this set of information as a probability function. Roughly, this is because there is no way of accommodating the inconsistent values assigned to the members of the partition in question. Less roughly, such cases are ones where Normality is violated across each of the propositions in this partition, when paired with their logical complement from a different time: i.e., $(H_1, .5) + (\overline{H}_2, .4) \neq 1$ and $(H_2, .6) + (\overline{H}_1, .5) \neq 1$. With this in mind, we might take the following principle to provide the weakest account of evidence that is compatible with **Retrospective Non-Arbitrariness**:

> **Diachronic Probabilism:** An agent's accumulated evidence over an interval of time should be representable as a probability function.

Cases where an agent violates Diachronic Probabilism are cases where she has inconsistent information.[15] It might seem less than obvious that being inconsistent in this way constitutes the violation of a norm. In the sections that follow, I will say more in defense of

[15] This proposal contains an important ambiguity in what it means for information to be consistent. On the one hand, one might take an agent's evidence to be consistent just in case the values over the evidence partition in question never change, either directly in virtue of updating on the same partition more than once *or* indirectly in virtue of making an update on a *different* evidence partition that changes the values along this earlier partition. On the other hand, one might maintain that evidence is consistent *even if* the second condition is violated: even if the values over an evidence partition received earlier change in virtue of an update over a different evidence partition received later. To remain diachronically coherent in the first, stronger sense, an agent must continually update on finer grained partitions—on $\{E_i\}$, and then on $\{EF_j\}$, etc. There are interesting consequences of both these approaches, which cannot be explored here. While I've been assuming the second account of evidential consistency, nothing in this paper turns on which of these accounts we adopt.

this idea. For now, let me simply gesture at this defense. My proposal draws upon what I take to be an intuitive idea, namely, that our evidence should have a standing that our credences don't have. While the Bayesian framework is driven by the idea that we should be able to change our minds about what we believe, it should *not* entail that we should be able to change our minds about our evidence. If we are to condition our beliefs on some piece of information, that information should meet the higher standard of being information that we don't ever have reason to revise.

I'll have more to say in the following section about why this is a plausible standard to adopt. The aim of this section has been to clear the way for such a defense by filling out the Retrospective Approach to updating with a constraint on evidence. I've suggested that Diachronic Probabilism is a good candidate for this constraint. Diachronic Probabilism takes what has traditionally been assumed to be a trivial assumption about the nature of evidence—that it is probabilistic—and generalizes it into a substantive diachronic constraint, without appealing to anything stronger than a principle about non-arbitrariness. We can now fill in our schema in the following way:

> **A Retrospective Account of Updating (Final)**:
> Where $\mathcal{S}$=$\{B_1, ..., B_n\}$ is a set of beliefs that form a partition, and where p and p' are probability distributions that are held sequentially at times t and t', an agent is diachronically coherent just in case the following two conditions hold:
>
> 1) $\forall p \forall p' \exists \mathcal{S} (\forall B_i \in \mathcal{S})\ \forall A (p(A \left| B_i \right.) = p'(A \left| B_i \right.)$), if defined, and
>
> 2) $\{\mathcal{S}_1, \mathcal{S}_2, ..., \mathcal{S}_n\}$ have values that satisfy Diachronic Probabilism.[16,17]

It's tempting to think that we violate Diachronic Probabilism anytime we violate or-

[16] I assume, here, that Diachronic Probabilism applies to the members of this set, taken pairwise.

[17] Adopting Diachronic Probabilism will require us to make a noteworthy adjustment to our schema from before. In taking diachronic coherence to be a function of whether sequential and interval frameworks coincide, Diachronic Probabilism assumes that diachronic coherence should be defined over *sets* of belief revisions, rather than individual belief revisions. I'll defend this assumption in the next section.

dinary Probabilism. However, I'm going to assume what all standard accounts of updating assume, which is that an updating rule is only a constraint on an agent who already conforms to ordinary Probabilism. This means that the only way for an agent to violate Diachronic Probabilism is for her to have updated on inconsistent information.

Despite the fact that we can't violate Diachronic Probabilism by violating ordinary Probabilism, there's clearly a very close connection between these two norms. Later I'll say more in support of this idea by showing that, just as we have extended the core commitments of Probabilism to Diachronic Probabilism, we can extend the justification that we have for conforming to Probabilism to provide a justification for Diachronic Probabilism. If we assume what seems reasonable—that our epistemic norms are defined by the reasons we have to conform to them—this means that Diachronic Probabilism and ordinary Probabilism turn out to be norms of the very same kind.

4 Objections and Replies

Here is a summary of the argument so far:

1. There is a plausible way of understanding diachronic coherence, according to which it is the requirement that the information we use to revise our beliefs satisfy some norm for evidence.

2. Our norm for evidence should be the same whether we update using the sequential or interval model. (By Non-Arbitrariness)

3. The information we revise our beliefs on should be representable as a probability function. Or, equivalently, Diachronic Probabilism is a norm for evidence. (By 2)

4. The Retrospective Account is a plausible way of understanding diachronic coherence. (By 1 and 3)

This ends my sketch of the Retrospective Account of updating. By way of defense and further elaboration, I turn now to consider a series objections and replies.

Objection. A natural worry about the sort of interval updating the Retrospective Account appeals to is that it seems to commit us to the view that the agent at time t_3 who has evidence, but has not updated on it, is fully rational. But this seems wrong. After all, her credences don't reflect her evidence. If you're a juror in a murder trial, it's not going to be much of a defense to say that, while you had the evidence, you were waiting until after the trial was over to update on it.[18]

Reply. While this is a worry for interval updating on the Orthodox Account, the Retrospective Account is explicitly designed to avoid this problem. The first condition of the account implies that a necessary condition for getting evidence is that it is information that an agent has *already* revised her beliefs on. What the second condition requires is that *had* the agent not updated on all of this information sequentially, it would *still* be the case that each of her updates would have satisfied Diachronic Probabilism, in the way that her sequential updates do trivially. The appeal to interval updating is meant to serve as a check on a kind of arbitrariness, then, rather than a way the agent might be permitted to avoid updating on her evidence.

Objection. One might object that the Retrospective Account encodes two ideas that are in tension: (1) that our evidence can be uncertain, and (2) that our evidence is unrevisable. Though the arbitrary swamping of uncertain evidence is surely a bad thing, being unable to *ever* revise uncertain evidence might seem just as bad. Isn't the point of being able to hold uncertain evidence that we can revise it later if we need to? More generally, one might worry about the notion of inconsistency at the heart of this account. It's clear why we should not hold propositional evidence that is inconsistent. It's clearly a bad thing to hold as evidence both some proposition and its negation. But one might ask whether it is really

[18] Thanks to [redacted] for raising this worry and for this example.

irrational to update on the same evidence partition more than once? And, if so, *why* is it irrational?

For a more concrete illustration of the worry, suppose my friend Kai is walking towards me. I start off with a low credence that it is Kai, because I think he's in another town. However, as he gets closer and closer, my visual evidence gets better and better, and increases my credence. The natural thing to say is that this is perfectly good evidence and I should update on it sequentially. Indeed, this is precisely the sort of example that typically motivates Jeffrey conditioning.[19]

Reply. As noted already, the Retrospective Account isn't an action-guiding account of rationality. Therefore, it rejects the sort of internalism about rationality that motivates the previous example. To the extent that our intuitions about how an agent ought to act track what it is rational for the agent to do—to the extent that they track the action-guidingness of our norm—these intuitions simply aren't relevant for the sort of account that is being defended. Belief states, or sets of belief states, can have different good-making properties. And not all of these good-making properties appeal to an agent's internal state. Just as an agent can conform to the norm of Probabilism without her internal state being a certain way, so too can she conform to the norm of Diachronic Probabilism without her internal state being a certain way.

Crucially, though we may not be able to give an argument for the irrationality of updates that violate the Retrospective Account, we can say something more about what makes such updates bad. Near the end of this section, I will say more in defense of the sort of externalist constraint that my account represents. In the next section, I'll argue that updates that fail to satisfy this constraint do worse, in an accuracy-related sense, than those that conform to this norm.

Finally, it should be noted that the commitments of the Retrospective Account, while

[19] Thanks to [redacted] for raising this worry and for this example.

unorthodox, aren't unprecedented. One finds something like an account of unrevisable, uncertain evidence in Sarah Moss [2018]'s recent, extremely influential account of probabilistic knowledge. According to Moss, the contents of attitudes like belief and assertion are sets of probability spaces over propositions—or, as she calls them, probabilistic contents. Moss uses this novel account of content to develop the idea that credences, just like full beliefs, are among the kinds of attitudes that can constitute knowledge.

Unlike Moss, I advance an account of evidence rather than of knowledge. And I jettison Moss's more complicated account of content in favor of a more complicated account of the agent's attitudes: while Moss takes the contents of our beliefs to be thoroughly probabilistic, I assume it is the weighted partitions that constitute our evidence that are thoroughly probabilistic. But my account is very much in the spirit of Moss's account. Like Moss, I appeal to the intriguing, and I think ultimately plausible thought, that the *certainty* with which we hold some informational content should be treated as orthogonal to its *normative status* as knowledge, evidence, etc. Consequently, beliefs like "It's probably five o'clock" can be unrevisable, either in virtue of being knowledge, as on Moss's account, or in virtue of being evidence, as on mine. Just as some beliefs with probabilistic contents constitute knowledge, in virtue of satisfying a version of those constraints that are familiar from traditional epistemology, we will see in the next section that some probabilistically weighted partitions constitute evidence, in virtue of satisfying a version of a familiar Bayesian accuracy norm.

Objection. This response is unsatisfying. What I *want* is an action-guiding account of rationality.

Reply. This objection is an objection to *any* Bayesian updating rule, not just mine. As noted earlier, while the Orthodox Account has the surface structure of an action-guiding account of rationality, it fails just as miserably to deliver such an account, in virtue of its

flagrant violations of the ought-implies-can principle. Applying the Orthodox Account requires that agents be logically omniscient and requires that they have computational capacities that far outstrip those of any actual agent. Both the Orthodox Account and the Retrospective Account fail as accounts of action-guiding rationality. The difference is that the Retrospective Account does not have the form of an action-guiding account. It is honest about what it is able to deliver. Unlike the Orthodox Account, the Retrospective Account is able to successfully provide *something* that is of value. It is able to tell us when a set of updates is suboptimal in virtue of having been made on inconsistent evidence.

Objection. A key move the Retrospective Account makes is to take diachronic coherence to be a relation that is defined not over belief revisions, but over *sets* of belief revisions. This seems at odds with the way we normally think about Bayesian updating.

Reply. While thinking about diachronic coherence in this way goes against the grain, it is not a crazy suggestion. There is precedence in the literature for accounts of Bayesian updating that don't take it to be a relation defined over individual credence functions that the agent has at two different times. Consider time-slice epistemology, which argues that we should adopt a synchronic surrogate of the Orthodox Account's updating rule.[20] Diachronic Probabilism is the inverse of these synchronic surrogates. Instead of taking the units over which our updating constraint is defined to be *less* temporally extended than the units assumed by the Orthodox Account, Diachronic Probabilism takes these units to be *more* temporally extended than the units assumed by the Orthodox Account. Instead of providing a synchronic surrogate for the diachronic constraint of Conditionalization, Diachronic Probabilism provides a diachronic surrogate for the synchronic constraint of Probabilism.

There's a more fundamental way that Diachronic Probabilism and what we might call "Synchronic Conditionalization" are inverses of one another. As I've alluded to already, one

[20] For two canonical examples of these synchronic surrogates, see Meacham [2010] and Hedden [2015a,b].

reason we have for thinking that getting inconsistent evidence is a bad thing appeals to an intuition that I think is widely shared, namely, that unlike our credences, our evidence should not be up for revision. Our evidence should have a standing that our credences don't have. Consider accounts of evidence that take it to be knowledge, or to be factive, or—a bit closer to home—that take evidence to be that which we get with a credence of one, so that we can never lose it (more on this in a moment). My approach does not require that evidence be any of these things. But in maintaining that an agent is diachronically coherent just in case her evidence is consistent over time, it preserves, in a formal way, an important feature of each of these accounts. My account preserves the idea that our evidence is *unrevisable* by entailing that we are diachronically incoherent—that we do not, properly speaking, have evidence at all—when the values over the partitions we have conditioned on have changed over time.

Some would take the unrevisability of evidence to be a bug, rather than a feature, of an updating framework built to accommodate it. But the fact that not everyone would endorse this picture of evidence is no mark against Diachronic Probabilism. Consider again its inverse, Synchronic Conditionalization. One of the problems that motivates Synchronic Conditionalization is that agents on the Orthodox, sequential framework are unable to lose certain evidence. As Meacham [2010, p.94] notices, the problem of not being able to lose certain evidence is, more fundamentally, the problem of there being a mismatch between our intuitions about evidence and our updating framework. Cases where it feels as though we should be able to lose certain evidence—cases of memory loss, for instance—are motivated by internalist intuitions about the nature of evidence. They are motivated by the thought that our evidence should supervene on the way that things seem to us at the present moment. But the Orthodox, sequential framework is committed to a form of externalism about evidence that we get with certainty, since it entails that anytime our current seemings conflict with the evidence that is encoded in our priors, the latter swamps the former. Synchronic surrogates of Conditionalization resolve this mismatch by positing priors that don't encode an agent's previous evidence, thereby bringing the framework in line with evidential internalism.

Diachronic Probabilism is a solution to a problem that is the mirror image of the problem of losing certainties: *the problem of retaining uncertainties*. Just as it's impossible to lose certain evidence on the Orthodox, sequential framework, when we get conflicting evidence, we saw in §3 that it's impossible to retain uncertain evidence in this same situation. Anytime our current uncertain evidence conflicts with our previous uncertain evidence, the former swamps the latter. This swamping feature commits the Orthodox, sequential framework to a form of internalism about uncertain evidence. It entails that our total evidence is determined by the way that things are for us at the present moment. To the extent that we are pulled by the externalist intuition that evidence is not the kind of thing that should be beholden to how things are for us at the present moment, Diachronic Probabilism resolves this mismatch between our intuitions and the framework by positing an account of evidence that entails that we are diachronically incoherent anytime our current evidence swamps our previous evidence—or, equivalently, anytime the values over the partitions we have conditioned on have changed over time.

From the point of view of the internalist, then, Orthodox, sequential updating encodes too much memory, while from the point of view of the externalist, Orthodox, sequential updating encodes not enough memory. Internalists resolve the former problem by updating a credence function that the agent has at each moment. I've argued that externalists can resolve the latter problem by making diachronic coherence a matter of updating on information that remains unchanged over time. Like Synchronic Conditionalization, Diachronic Probabilism is formally compatible with many different accounts of evidence. But it is most friendly to those who lean externalist, in the way that Synchronic Conditionalization is most friendly to those who lean internalist. It's no objection to Diachronic Probabilism, then, that it is not for everyone, any more than this is an objection to time-slice epistemology. The cost of a framework, like the Orthodox, sequential framework, that is not biased towards some substantive account of evidence is that it succumbs to counterexample—to agents who can't forget or to agents who can't remember—whenever it is paired with the intuitions that underwrite one of

these accounts. Diachronic Probabilism is a remedy for the second kind of counterexample, just as Synchronic Conditionalization is a remedy for the first.

5 An Accuracy Argument for the Retrospective Account

We've seen that the Retrospective Account, via Diachronic Probabilism, generalizes the norm of Probabilism. It should come as no surprise, then, that Diachronic Probabilism can be justified by an accuracy argument that generalizes the accuracy argument for Probabilism.

As we've noted already, the aim of accuracy-first epistemology is to vindicate various epistemic norms by showing that they are the means to the end of gradational accuracy. To get an idea of what the accuracy-dominance argument for Probabilism looks like, consider a weather forecaster who assigns a credence of .7 to the proposition that it will hail tomorrow and a credence of .5 to the proposition that it won't hail tomorrow. This agent has a probabilistically incoherent set of credences, since $.7+.5\neq 1$. It can be shown that, because this agent has a probabilistically incoherent set of credences, her credences are accuracy-dominated. According to a reasonable measure of inaccuracy, there exists a probabilistically coherent set of credences that will be less inaccurate than it is, no matter which state of the world comes about. In the case just described, these probabilistically coherent credences are $\{(H, .6), (\overline{H}, .4)\}$.[21]

By contrast, a probabilistically coherent set of credences would *not* be accuracy-dominated. It is not the case that there is some set of credences that would do better

[21] If we assume the Brier score, for instance, we get that if it does hail, $\{(H, .7), (\overline{H}, .5)\}$ will be more inaccurate than $\{(H, .6), (\overline{H}, .4)\}$:

$$(1 - .6)^2 + (0 - .4)^2 < (1 - .7)^2 + (0 - .5)^2$$
$$.32 < .34$$

And if it doesn't hail, $\{(H, .7), (\overline{H}, .5)\}$ will still be more inaccurate than $\{(H, .6), (\overline{H}, .4)\}$:

$$(0 - .6)^2 + (1 - .4)^2 < (0 - .7)^2 + (1 - .5)^2$$
$$.72 < .74$$

than it does, no matter which state of the world comes about. Since non-probabilistic credence assignments are accuracy-dominated, but probabilistic credence assignments are not accuracy-dominated, a decision-theoretic agent with an epistemic utility function does better by conforming to Probabilism.

Earlier we noted that an obstacle to thinking of an agent who violates Diachronic Probabilism as an agent who violates a probabilistic norm is that the information that Diachronic Probabilism targets is not even a function to begin with. We encounter an analogous problem when it comes to justifying Diachronic Probabilism. Since inconsistent evidence is not a function, and since accuracy-dominance can only be defined over functions, inconsistent evidence cannot be claimed to be accuracy-dominated.

But our analogous problem comes with an analogous solution. Just as Diachronic Probabilism is a norm that generalizes ordinary Probabilism, we can get a justification for Diachronic Probabilism that generalizes the justification for ordinary Probabilism.

We've just seen that the accuracy argument for Probabilism says that, no matter how the world turns out, an agent who violates Probabilism does better by not holding the credences that she holds. In a similar way, we might say that, no matter how the world turns out, an agent who violates Diachronic Probabilism does better by not holding *all of the evidence* that she holds. This is because, no matter how the world turns out, *some piece of evidence that she holds will be more inaccurate than some other piece of evidence that she holds*.

To see this, recall our example from §3—$\{\{(H_1, .5), (\overline{H}_1, .5)\}, \{(H_2, .6), (\overline{H}_2, .4)\}\}$. This is an example where we are probabilistically incoherent over each of these propositions, when paired with their logical complement from a different time: $(H_1, .5) + (\overline{H}_2, .4) \neq 1$ and $(H_2, .6) + (\overline{H}_1, .5) \neq 1$. This gives rise to the result that we've just described, namely, that no matter how the world turns out, one of these sets of values will inevitably be more inaccurate than the other. For if it turns out that H, then $\{(H_1, .5), (\overline{H}_1, .5)\}$ will be more inaccurate than $\{(H_2, .6), (\overline{H}_2, .4)\}$. And if it turns out that $\overline{H}$, then $\{(H_2, .6), (\overline{H}_2, .4)\}$ will

be more inaccurate than $\{(H_1, .5), (\overline{H}_1, .5)\}$.[22] No matter what happens, we will have reason to discard *some* of our evidence, contra the intuition, held by many, that evidence is not the sort of thing that we should have reason to discard.

It should be clear what's going on here. The only way to update over a partition more than once, in a way that assigns this partition inconsistent values, is for these values to change. If the same partition has been assigned different values, then one set of values will inevitably be more inaccurate than the other. That's what the argument for Diachronic Probabilism says.

By contrast, it's *not* the case that an agent who conforms to Diachronic Probabilism necessarily has evidence, such that one piece of evidence will inevitably be more inaccurate than the other, no matter what happens. Whether or not an agent has reason to discard some of her evidence will depend upon the details of the case, rather than upon the necessary, formal features of the evidence that she updates on. Say that we again have an agent who gets $\{(E, .5), (\overline{E}, .5)\}$ at t_1, and then gets $\{(F, .5), (\overline{F}, .5)\}$ at t_2. This agent has as evidence $\{(EF, .25), (\overline{E}F, .25), (E\overline{F}, .25), (\overline{EF}, .25)\}$. Here it is *not* the case that this evidence is such that the agent has reason to discard some of it, no matter which state of the world comes about. Say EF comes about. Then whether or not the agent has reason to discard $\{(E, .5), (\overline{E}, .5)\}$ or $\{(F, .5), (\overline{F}, .5)\}$, or both $\{(EF, .25), (\overline{E}F, .25), (E\overline{F}, .25), (\overline{EF}, .25)\}$, or neither, will depend upon whether this evidence brings her closer to EF than she was initially. If, say, her initial credences were $\{(EF, .1), (\overline{E}F, .1), (E\overline{F}, .1), (\overline{EF}, .7)\}$, then it is *not* the case that this agent has reason to discard *any* of her evidence.

[22] Again, assuming the Brier score, if it does hail, then $\{(H, .5), (\overline{H}, .5)\}$ will be more inaccurate than $\{(H, .6), (\overline{H}, .4)\}$:

$$(1 - .6)^2 + (0 - .4)^2 < (1 - .5)^2 + (0 - .5)^2$$
$$.32 < .5$$

And if it doesn't hail, then $\{(H, .6), (\overline{H}, .4)\}$ will be more inaccurate than $\{(H, .5), (\overline{H}, .5)\}$:

$$(0 - .5)^2 + (1 - .5)^2 < (0 - .6)^2 + (1 - .4)^2$$
$$.5 < .72$$

This, again, contrasts with evidence that violates Diachronic Probabilism. While it's possible for each *individual* piece of an agent's inconsistent evidence set to bring her closer to some world than she was initially, it's impossible for *all* of the pieces of an agent's inconsistent evidence set to do this. Whatever happens, she will always have reason not to hold some of her evidence.

While this argument for Diachronic Probabilism is not an accuracy-dominance argument, it generalizes the reasoning behind such arguments. Both the argument for Diachronic Probabilism and the argument for Probabilism say that no matter which state of the world comes about, an agent does better by discarding some of her credences (or evidence), in favor of some probabilistic credences (or evidence). The argument for Probabilism is the special case where, in addition, the probabilistic credences the agent has reason to favor is the same, no matter which state of the world comes about. Thus, we have the same general reason for conforming to Diachronic Probabilism that we have for conforming to ordinary Probabilism: in order not to have evidence that we will inevitably have reason to discard. An agent who has an evidence set that contains evidence that she will inevitably have reason to discard has evidence that lacks a valuable property it might otherwise have had.

This argument avoids the problems faced by the accuracy argument for the Orthodox Account. It maps perfectly onto the constraint defended in this paper; it does not require that the agent know in advance the partition that she updates on; and it does not preclude updates on uncertain evidence. The Retrospective Account may, then, turn out to be the norm that strikes just the right balance of being strong enough to impose constraints on a wide range of epistemic situations, while also recommending credences that promote some form of epistemic utility. While these remarks are too brief to establish this once and for all, they at least suggest that we cannot rule out the Retrospective Account as a serious contender for the updating norm that we should adopt.

6 Conclusion

I've argued that the Retrospective Account reverses a number of the features of orthodox updating, to great effect. Before closing, I want to offer one last remark on the account.

I've noted that the Retrospective Account of updating makes diachronic coherence continuous with synchronic coherence. Both types of coherence relations can be described as the norm that our evidence be representable as a probability function. An interesting consequence of this is that, despite the fact that synchronic coherence is a precondition for diachronic coherence, our diachronic norm is conceptually prior to our synchronic norm. Being a probability function is the special case of being representable as a probability function. It is the special case of having evidence that we don't have reason to discard no matter what. It is the special case of being unrevisable. It is the special case of avoiding arbitrariness, trivially, in virtue of representing our evidence at a single time. What it means to be synchronically coherent is to satisfy the norm for diachronic coherence in a particular way.

What should we think of this? While this account of the relationship between synchronic and diachronic coherence isn't a clear advantage of my approach, it may be a happy consequence for some. For many, being epistemically rational is most centrally a matter of having beliefs that display a certain sort of coherence over time. It's this kind of coherence that is most congenial to the project of helping us to achieve our long term goals. For some, then, it may be a welcome result that in reversing the relation that the Orthodox Account prescribes, the Retrospective Account ends up reversing the relationship between synchronic and diachronic coherence as well.

References

David A Blackwell and Meyer A Girshick. *Theory of games and statistical decisions*. Courier Corporation, 1979.

Lisa Cassell. Bayesian coherentism. *Synthese*, pages 1–28, forthcoming. [DOI].

Persi Diaconis and Sandy L Zabell. Updating subjective probability. *Journal of the American Statistical Association*, 77(380):822–830, 1982.

John Earman. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, ma: MIT Press, 1992.

Kenny Easwaran. Expected accuracy supports conditionalization—and conglomerability and reflection. *Philosophy of Science*, 80(1):119–142, 2013.

I.J. Good. Letter to the editor. *The American Statistician*, 1971.

Hilary Greaves and David Wallace. Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(459):607–632, 2006.

Gilbert Harman. *Change in view: Principles of reasoning.* The MIT Press, 1986.

Brian Hedden. Time-slice rationality. *Mind*, 124(494):449–491, 2015a.

Brian Hedden. *Reasons without persons: rationality, identity, and time*. OUP Oxford, 2015b.

Colin Howson and Peter Urbach. *Scientific reasoning: the Bayesian approach*. Open Court Publishing, 2006.

RC Jeffrey. The logic of decision. Technical report, MCGRAW-HILL BOOK CO INC NEW YORK, 1965.

Richard Jeffrey. *Probability and the Art of Judgment*. Cambridge University Press, 1992.

Richard C Jeffrey. Dracula meets wolfman: Acceptance vs. partial belief. In *Induction, acceptance and rational belief*, pages 157–185. Springer, 1970.

James Joyce. Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In Franz Huber and Christoph Schmidt-Petri, editors, *Degrees of Belief*, pages 263–297. Synthese, 2009.

Marc Lange. Is jeffrey conditionalization defective by virtue of being non-commutative? remarks on the sameness of sensory experiences. *Synthese*, 123(3):393–403, 2000.

Hannes Leitgeb and Richard Pettigrew. An objective justification of bayesianism i: Measuring inaccuracy. *Philosophy of Science*, 77(2):201–235, 2010a.

Hannes Leitgeb and Richard Pettigrew. An objective justification of bayesianism ii: The consequences of minimizing inaccuracy. *Philosophy of Science*, 77(2):236–272, 2010b.

Christopher JG Meacham. Unravelling the tangled web: Continuity, internalism, uniqueness and self-locating belief. *Oxford studies in epistemology*, 3:86, 2010.

Christopher JG Meacham. Understanding conditionalization. *Canadian Journal of Philosophy*, 45(5-6):767–797, 2016a.

Christopher JG Meacham. Ur-priors, conditionalization, and ur-prior conditionalization. *Ergo, an Open Access Journal of Philosophy*, 3, 2016b.

Sarah Moss. *Probabilistic knowledge*. Oxford University Press, 2018.

Richard Pettigrew. *Accuracy and the Laws of Credence*. Oxford University Press, 2016.

Brian Skyrms. Three ways to give a probability assignment a memory. 1983.

Brian Skyrms. Dynamic coherence and probability kinematics. *Philosophy of Science*, 54(1): 1–20, 1987.

Brian Skyrms. Diachronic coherence and radical probabilism. *Philosophy of science*, 73(5): 959–968, 2006.

Michael G Titelbaum. *Quitting certainties: A Bayesian framework modeling degrees of belief*. Oxford University Press, 2013.

Jonathan Weisberg. Updating, undermining, and independence. *The British Journal for the Philosophy of Science*, 66(1):121–159, 2014.