Bayesian Coherentism

L. Cassell

(draft—please don't cite without permission)

Abstract. This paper considers a problem for Bayesian epistemology and goes on to propose a solution to it. On the traditional Bayesian framework, an agent updates her beliefs by *Bayesian conditioning*, a rule that tells her how to revise her beliefs whenever she gets evidence that she holds with certainty. In order to extend the framework to a wider range of cases, Richard Jeffrey (1965) proposed a more liberal version of this rule that has Bayesian conditioning as a special case. *Jeffrey conditioning* is a rule that tells the agent how to revise her beliefs whenever she gets evidence that she holds with any degree of confidence. The problem? While Bayesian conditioning has a foundationalist structure, this foundationalism disappears once we move to Jeffrey conditioning. If Bayesian conditioning is a special case of Jeffrey conditioning then they should have the same normative structure. The solution? To reinterpret Bayesian updating as a form of coherentism.

Foundationalism and coherentism are competing views about the structure of epistemic justification. It's surprising, then, that they co-exist on the Bayesian framework. The explanation: Bayesianism is committed to norms that govern our degrees of belief—our credences—in propositions that stand in particular logical relations to each other at each time. It's also committed to norms that govern how these credences change over time in response to new evidence. Traditional Bayesian epistemology is coherentist with respect to the first set of norms. It's foundationalist with respect to the second. It has two strands of justification running through it.

This paper considers a problem for Bayesianism's second strand of justification, and goes on to propose a solution to it. On the traditional Bayesian framework, an agent updates her beliefs by *Bayesian conditioning*, a rule that tells her how to revise her beliefs, whenever she gets evidence that she holds with certainty. In order to extend the framework to a wider range of cases, Richard Jeffrey (1965) proposed a more liberal version of this rule. *Jeffrey conditioning* is a rule that tells the agent how to revise her beliefs, whenever she gets evidence that she holds with any degree of confidence. Jeffrey claimed that his rule has Bayesian conditioning as a special case. This claim is now a truism of Bayesian epistemology.

The problem? While Bayesian conditioning has a foundationalist structure, this foundationalism disappears once we move to Jeffrey conditioning. But if Bayesian conditioning is a special case of Jeffrey conditioning, then these two updating rules should have the same normative structure. We are then left with the following inconsistent triad: (1) If one norm is a special case of another, then they should have the same normative structure, (2) Bayesian conditioning is a special case of Jeffrey conditioning, (3) Bayesian conditioning and Jeffrey conditioning have different normative structures.

In this paper, I'll argue for an interpretation of the Bayesian framework that resolves this inconsistency by rejecting (3). I'll argue that both regular Bayesian updates and Jeffrey updates proceed from a common framework—one with a coherentist structure. My strategy will be to appeal to what has long been deemed to be a defect of Jeffrey conditioning: the fact that its updates aren't guaranteed to commute. To say that Jeffrey updates aren't guaranteed to commute is to say that an agent's credences after a sequence of updates will sometimes be determined by the order in which this evidence has been received. This feature of Jeffrey updates is a defect because the order in which some evidence has been received seems irrelevant to the impact it ought to have. While the fact that the Jeffrey framework can't guarantee that its updates will commute is standardly taken to show that the framework fails to satisfy an important desideratum for an updating rule, in this paper, I propose that we take the commutative property to play a more fundamental role. I propose that we take the commutative norm that Bayesianism is committed to to be one that grounds *particular updates*. In other words: some set of updates will be justified to the extent that they commute. Since the sort of consistency this norm encodes is to updates what the norm of evidential consistency from traditional formulations of coherentism is to beliefs, it looks as though the best way of understanding the structure of Bayesian updating is not as a form of foundationalism, but as a form of coherentism.

Before we get started, let me say a bit more about our inconsistent triad. The claim that Bayesian conditioning and Jeffrey conditioning have different normative structures will be argued for below. And the claim that Bayesian conditioning is a special case of Jeffrey conditioning is a mathematical fact, as we will also see in just a moment. Before we move on, however, I want to briefly defend (1): the claim that if one norm is a special case of another, then they should have the same structure. They should function in the same way. It might seem as though counterexamples to this idea aren't difficult to find. Consider a rule that tells you that you are permitted to drive in the carpool lane if you have any number of people in the car—except for just one. Here the special case of having just one person in the car calls for something different than the general case of having any number of people in the car. But this doesn't seem unreasonable. Or consider a law that requires a unanimous guilty verdict from a jury in order for the death penalty to be imposed. Though this law might be contentious, again, it's not clear that this is due to the fact that it treats the special case of unanimity as different from the more general case of a jury reaching its decision.[1]

It's true that, taken on its own, the special case does not seem unreasonable in either of these examples. But I think there's at least a sense in which, considered side by side with an instance of the more general case, the special case calls for an explanation. In both of these examples, it seems reasonable to question the sharp cut-off between the special case and the one just before it. It seems reasonable to ask why a car with only one person in it is so different from a car with two people in it, and why a unanimous jury is different in kind from a jury with only one dissenter.

In the same way, I want to suggest that the fact that the special case of Bayesian conditioning has a certain normative structure, while not problematic when considered on its own, is odd when considered together with the fact that Jeffrey conditioning, in general,

[1] Thanks to [redacted] for the last example. Thanks also to [redacted] and [redacted] for independently pressing this point.

does not have this same structure. The oddness in all of these cases may be defeasible. If one were to explain that unanimous agreement is necessary to establish that there is no reasonable doubt about a defendant's guilt, one might then see what makes a unanimous jury qualitatively different from a jury with one dissenter. In the same way, there are assumptions we might make about the Bayesian formalism that would justify treating updates on certain evidence as qualitatively different than updates in general. In the last section of this paper, I'll briefly consider some of these assumptions. But, for now, I'll assume that there's at least something to be said for an interpretation of Bayesian epistemology that yields the result that certain and uncertain updates proceed from a framework with the same normative structure.

One more thing before moving forward. Since the Jeffrey framework includes both updates on certain and uncertain evidence, one might object that there is no difference in normative structure between the Jeffrey framework and the regular Bayesian framework. Jeffrey conditioning is constrained in the very same way as Bayesian conditioning: like Bayesian conditioning, Jeffrey conditioning imposes a foundationalist constraint on evidence that receives a credence of one. However, the interesting question is not whether the Bayesian framework and the Jeffrey framework include the same constraint. The interesting question is whether this constraint governs all of the updates that proceed from these frameworks. That is, the interesting question is whether the constraint that governs updates made on evidence that we hold with certainty also governs updates on evidence that we hold with *any* degree of confidence. On the orthodox way of understanding the Bayesian updating framework, it does not. And this seems objectionable. If the change in the values that our evidence receives is meant to represent a change in the certainty of our evidence, this change in value should not entail a difference in anything else.

Here's how the discussion will go then. In §1, I give some background. In §2, I describe the sense in which regular Bayesian updating has a foundationalist structure. In §3, I explain why adopting Jeffrey conditioning entails abandoning this foundationalism. In §4, I propose a constraint that looks like a version of coherentism about updating and argue that it can make sense of the truism that Bayesian conditioning is a special case of Jeffrey conditioning. In §5, I give the formal details of this constraint. Finally, in §6, I revisit the motivation for this discussion.

1 Some Background

1.1 Diachronic Coherence for Bayesians

I've suggested that it's possible to ask whether Bayesian updating is a form of foundationalism or a form of coherentism. Let's begin by getting clear on exactly what this question means.

Standard Bayesianism assumes that an agent's credences in the propositions she entertains can be represented as an assignment of real numbers to those propositions. It further assumes that two norms of coherence govern this assignment. First, Bayesianism is committed to the constraint that, at each time, the agent's credences be a probability function. To say that a Bayesian agent is *synchronically coherent*, then, is to say that she conforms to Probabilism: 1) she assigns every proposition her credence function is defined over a non-negative value, 2) she assigns a credence of one to any tautology and, 3) for any mutually exclusive propositions, A and B, that her credence function is defined over, $cr(\text{A}) + cr(\text{B}) = cr(\text{A} \vee \text{B})$.

Second, Bayesianism is committed to the constraint that the agent's beliefs evolve over time in accordance with her conditional probabilities. If my credence in the proposition that I will play baseball tomorrow is .3, and my credence in the proposition that I will play baseball tomorrow *conditional* on it not raining is .7, then when I learn that it will not rain—when I get this as evidence—my credence that I will play baseball tomorrow should shoot up from .3 to .7. To say that a Bayesian agent is *diachronically coherent*, then, is to say that her conditional probabilities guide her belief revisions. We can describe this as the requirement that an agent's current probabilities be determined by her prior conditional probabilities, in the way that we've just done. A different, though equivalent way of describing this coherence constraint is as the requirement that the values of the conditional probabilities that yield the agent's current probabilities be the same before and after the update. We can think of the agent's conditional probabilities as arrows that proceed from her evidence and that guide the propagation of the rest of her probabilities. To serve this guiding function, they must remain fixed.[2]

As it happens, every probabilistic belief transition has a set of arrows. For any probabilistic belief transition, there will be *some* information—like learning that it will not rain—that each of my beliefs are conditional on in the same way before and after this transition. More formally, there will always be a partition (a set of mutually exclusive and exhaustive propositions, like {RAIN, ¬RAIN}) that is sufficiently fine-grained to represent this transition as an update that is conditional on that partition:

> **Descriptive Diachronic Coherence for Bayesians:** There is a sufficient partition, $\{\text{B}_\text{i}\}$, for every probabilistic belief transition.[3]

Or, equivalently, where $S = \{\text{B}_1, ..., \text{B}_\text{n}\}$ is a set of beliefs that form a partition, and where an agent has an experience that causes her to revise her beliefs, the transition between the agent's prior probability distribution, p, and posterior probability distribution, p', at t and t', respectively, can be formulated in a way that underlines that there will always be some conditional probability that can be understood to guide her belief revision by remaining the same before and after the update:

> **Descriptive Diachronic Coherence for Bayesians:**
> $\forall p \forall p' \exists S (\forall \text{B}_\text{i} \in S)\ \forall A (p(\text{A} \mid \text{B}_\text{i}) = p'(\text{A} \mid \text{B}_\text{i})\ )$, if defined.

[2] The arrow analogy is borrowed from Weisberg (2015). This guiding feature of our conditional probabilities is often referred to as 'rigidity' (see Jeffrey (1965)).

[3] For the proof of this, see Blackwell and Girshick(1979) and Diaconis and Zabell (1982, p. 824). 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.

Since **Descriptive Diachronic Coherence for Bayesians** holds for any two probability functions, it is no stronger than Probabilism. Since it is no stronger than Probabilism, one might rightly suspect that it will be too weak to capture any interesting notion of diachronic coherence. To see this more clearly, consider the very simple agent who only has beliefs about whether or not she will play baseball tomorrow: her credences are only defined over the partition $\{P, \neg P\}$. Suppose that these credences are $p(P)=.4$ and $p(\neg P)=.6$. Suppose further that the agent revises these credences to $p(P)=.7$ and $p(\neg P)=.3$. In this case, $\{P, \neg P\}$ is a sufficient partition for the update. Therefore, this belief transition satisfies **Descriptive Diachronic Coherence for Bayesians**. Intuitively, however, this case doesn't look much like an agent responding to her evidence. Among other things, we tend to think that updating in accordance with one's evidence happens when we come to change our belief in some proposition, on the basis of some *different* information. It happens when our views about whether we will play baseball tomorrow change in response to listening to the weather forecast and learning about the chance of rain. The lesson is that some belief transitions that satisfy **Descriptive Diachronic Coherence for Bayesians** don't look much like an agent being diachronically coherent at all, if we take such coherence to involve the agent getting evidence. Instead, what they look like is an agent swapping one set of probabilities for another.

We can remedy this by strengthening our account of diachronic coherence. We can do this by stipulating that it is only when the agent conditions her beliefs on partitions that meet some additional constraint for being evidence that she is diachronically coherent:

> **Normative Diachronic Coherence for Bayesians:** A probabilisitic belief transition ought to be such that,
>
> (a) there is a sufficient partition, $\{B_i\}$, for the transition, and
>
> (b) $\{B_i\}$ satisfies the conditions for being evidence.

Or, equivalently, where $E=\{B_1, ..., B_n\}$ meets the criteria for being evidence, 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:

> **Normative Diachronic Coherence for Bayesians:**
> $\forall p \forall p' \forall (B_i \in E)\, \forall A (p(A|B_i) = p'(A\,|B_i)\,)$, if defined.

What makes this formulation normative is that it is stronger than Probabilism: an agent might transition from one probability function to another in a way that violates it. What makes this formulation a norm of coherence is that it is defined over a set of credence functions. Finally, what makes this norm of coherence diachronic is that these credence functions are indexed to different times.

Before moving on, I want to mention one last way of understanding diachronic coherence for Bayesians: a middle path between descriptive and normative diachronic coherence. Instead of overcoming the weaknesses of the former by restricting the conditions under which some partition is evidence, we might simply take for granted the existence of an evidence partition, and ask about what follows from it. In other words, we might take the Bayesian agent's diachronic obligations to consist in how she ought to proceed, *assuming* that she has a certain piece of evidence. On this picture of things, the evidence partition is not normatively determined, but *causally* determined. It is "an internal or psychological condition" for which "mathematics has nothing to offer".[4]

This understanding of diachronic coherence looks like a more modest way of getting us what we are after. By stipulating that some partition of propositions constitutes the agent's evidence, it avoids the worry that it is too weak to capture any interesting notion of diachronic coherence. But since this understanding of diachronic coherence does not require an agent's evidence to satisfy any additional normative constraints, it is also no stronger than **Descriptive Diachronic Coherence for Bayesians**. Moreover, it makes sense of the way that people tend to talk about Bayesian updating: as a norm that governs the agent, *conditional* on our assuming her to have some piece of evidence. Some might even call this the default view of diachronic coherence for Bayesians.

I want to defer saying anything more about the default view for the moment. It will become clear a bit later on why this account of diachronic coherence cannot be used to unify Bayesian updates, in the way that we are looking to do.

1.2 A Formal, Deflationary Account of Evidence

For now, then, let us assume that the aim of this paper will require that we adopt a normative account of diachronic coherence. And this will require that we adopt an account of evidence. There are a couple of ways that we might go about this. The most familiar of these ways is to appeal to a *substantive* account of evidence: for instance, to the requirement that evidence be what one knows, or be related to what one has internal access to, or be formed by a reliable process, etc. What makes such accounts substantive ones is that Bayesians, qua Bayesians, aren't committed to the normativity of knowledge, or of access, or of reliability, etc. A substantive account of evidence, then, is a constraint on evidence formulated in terms of a property that is not already part of the Bayesian formalism.

This paper will take a different approach by defending a *formal* account of evidence. An account of evidence is formal just in case it's not substantive: just in case it *is* formulated in terms of some feature of the Bayesian formalism. A Bayesian formalist about evidence will hold that it is in virtue of the sufficient partition of an update being assigned certain values, or weights, that the agent can be said to have evidence, where these values are not further justified by the sorts of substantive considerations just mentioned.[5] Ex-

[4] Diaconis and Zabell (1982, p. 825).

[5] Perhaps an easier way of understanding what makes knowledge and access and reliability substantive, rather than formal, requirements is that they can be made sense of out of context: they can be defined independently of the other epistemic commitments one happens to hold. Formal norms are different in

actly what it will look like for a formal constraint on evidence to be satisfied will become clearer in the next section. But, for now, notice that the appeal to a formal account of evidence leaves us able to understand how it is possible to ask whether Bayesian updating is a form of foundationalism or a form of coherentism. Since what we will be after is a formal, or structural, account of evidence, and since foundationalism and coherentism are both structural norms, they will both be candidates for such an account.

Why defend a formal account of evidence in the first place? I think it's of interest to consider how much normativity can be defined of in terms of the commitments that Bayesians already hold. It's worth emphasizing again that such an account of evidence will be a deflationary one. Insofar as we are tempted to talk of partitions as "being evidence", it is because the weights these partitions get assigned are what determine the extent to which our constraint on evidence gets satisfied. This deflationary picture of evidence will yield an account of normative diachronic coherence with two distinguishing features:

- Agents aren't diachronically coherent, full-stop. Instead, diachronic coherence comes in *degrees*.
- Degrees of diachronic coherence aren't defined over updates. Instead, they are defined over *sets* of updates.[6]

In the previous section, we said that an agent's standing as synchronically coherent will depend upon how her credence functions are related to each other. In defining diachronic coherence over sets of updates, my account entails that an agent's standing as diachronically coherent will depend upon how her *updates* are related to each other. More specifically, it will depend upon how the pieces of evidence that underwrite these updates are related to each other. My account, then, leaves us with an interpretation of Bayesian epistemology that is coherentist, with respect to both strands of justification that run through it.

1.3 An Assumption

Finally, an assumption. Epistemic theories can be given one of two interpretations. On the one hand, we might think that what any such theory provides is *guidance* for how a rational agent ought to act. On the other hand, we might think that what any such theory provides is a way of *evaluating* an agent's actions, whether or not we would want to

this regard. Take, for instance, the formal norm of consistency. The way that a Bayesian defines consistency will differ from the way that a defender of Dempster-Shafer theory does this. While the Bayesian will spell out her notion of consistency by means of probability functions (e.g., "X is consistent if it treats these probability functions the same"), a Dempster-Shafer theorist, who trades in belief functions and mass functions, will define her notion of consistency in terms of these. Unlike knowledge or access or reliability, the norm of consistency is so thin that it isn't complete without a framework already in place to give it content.

[6] A consequence of this is that **Normative Diachronic Coherence for Bayesians** is more perspicuously formulated in terms of a pair of evidence partitions, instead of just one. How this can be done will become clearer in §4.

say that an agent *ought* to have done what she did. Bayesian epistemology, when understood in the first way, has received its fair share of criticism. This is because the sorts of idealizing assumptions that we need to get the framework off the ground require of ordinary agents that they perform operations that are computationally intractable. Here's Harman (1988, p. 25-26) on this:

> One can use conditionalization to get a new probability for P only if one has already assigned a prior probability not only to E but to P∧E. If one is to be prepared for various possible conditionalizations, then for every proposition P one wants to update, one must already have assigned probabilities to various conjunctions of P together with their denials. Unhappily, this leads to combinatorial explosion, since the number of such conjunctions is an exponential function of the number of possibly relevant evidence propositions.

And Earman (1992, p. 56):

> 'Ought' is commonly taken to imply 'can', but actual inductive agents can't, since they lack the logical and computational powers required to meet the Bayesian norms. The response that Bayesian norms should be regarded as goals toward which we should strive even if we always fall short is idle puffery unless it is specified how we can take steps to bring us closer to the goals.

In light of these sorts of criticisms, I will assume that Bayesianism is best understood as a set of evaluative norms, rather than as a set of action-guiding norms. This means that although Bayesian epistemology sets certain standards, there are no obligations issued by the theory. Just as we can say that cars are good, insofar as the brakes work, and bad insofar as they don't, without imposing any obligations on anyone to do anything, we can say that updates are good or bad, in virtue of certain features of them, without imposing any obligations on anyone to do anything.

The way that we've set things up in this section already encourages us to conceive of Bayesian epistemology as an evaluative theory. The natural question to ask on the action-guiding approach is: given what I take my evidence to be, how should I update? By contrast, the natural question to ask on the evaluative approach is: does my update have the right features? By defining evidence *in terms* of the update that it triggers, rather than vice versa, we set ourselves up to pursue the second of these questions. An agent's update will be good insofar as the formal constraint on evidence defended in this paper is satisfied, and bad insofar as it isn't. However, this does not guide the agent, or obligate her to update, in any particular way.

2 Bayesian Foundationalism

I claimed in the introduction that Bayesian conditioning can be understood as a form of foundationalism about updating.[7] To better understand the sense in which Bayesian conditioning has sometimes been taken to be a form of foundationalism, it will help to first get clear on what foundationalism amounts to in its more traditional guise—as a structure that applies to beliefs.

Traditional foundationalism about epistemic justification says that the ultimate source of the justification of all our beliefs is some privileged set of cognitive states that is the locus of this justification, but that can't be the target of it. It's the conjunction of the claims: (1) that some cognitive states are basic, in the sense of their being justified not in virtue of their relations to other cognitive states and, (2) that all non-basic states are justified in virtue of some relation that they bear to basic states. In addition, classical foundationalism assumes (3) that the distinguishing mark of basic states is their infallibility.[8] As should be clear from this description of it, foundationalism is *just* a structure. It takes no stand on what features some cognitive state must have in order to be a basic state, other than the role that it occupies on the foundationalist's picture of things. Since foundationalism is just a structure, we should be able to ascribe it to a theory like Bayesianism, which also does not take a stand on what makes something a basic state.

It's perhaps with this in mind that some have assumed that Bayesian conditioning's account of evidence makes Bayesian updating a form of foundationalism. Consider the following typical formulation of Bayesian conditioning:

> **Bayesian conditioning:** If the strongest evidence you get raises your credence in B to one, then your new degree of belief in A, for any A, should be $p'(A)=p(A|B)$, where A and B are propositions.

Where we assume the sort of Cartesian foundationalism that identifies a basic state with an infallible one, Bayesian conditioning satisfies (3) by requiring that the agent's evidence be a proposition of which she is certain.[9] The previous formulation clearly sat-

[7] References to Bayesian conditioning as a form of foundationalism are ubiquitous. Those who have explicitly taken standard Bayesian conditioning to instantiate a foundationalist structure include Christensen (1992), Skyrms (1997), Bradley (2005), Weisberg (2009) and Titelbaum (forthcoming), among others. The extent to which each of these discussions would endorse the picture of Bayesian foundationalism that I argue for in this section is unclear. However, the picture of foundationalism set out in this section is certainly inspired by these discussions.

[8] By contrast, many recent, non-classical foundationalist accounts, like Goldman (1988)'s reliabilism, Plantinga (1991)'s proper basicality, Pryor (2000)'s dogmatism and Huemer (2007)'s phenomenal conservativism defend some form of fallible foundationalism. That is, they maintain that the property that makes beliefs basic is something other than their infallibility.

[9] One might object that to identify a basic state with an infallible state is to make precisely the sort of substantive assumption that a formal account of evidence is supposed to avoid. While this is true, this assumption seems less substantive (i.e., controversial) than the ones that were appealed to earlier as paradigmatically substantive assumptions—e.g., that internalism or externalism is the correct theory of justification, etc. Nevertheless, if one is uncomfortable with this assumption, one can do without it since, as we've already said, (3) is not a necessary component of a foundationalist account. Thanks to [redacted] for pressing me to address this point.

isfies (2) as well: the values we assign the rest of our beliefs depend upon our evidence, qua basic state. What about (1)? While the agent uses her evidence proposition to infer the credences she holds in other propositions, the evidence proposition itself cannot receive this sort of support. This is because propositions that receive a credence of one cannot have their values changed by Bayesian conditioning at some later time. That certainties stay certainties is a mathematical feature of the Bayesian formalism. Crucially, then, *once* a belief becomes a basic state—once it becomes evidence—it is no longer able to receive the same sort of inferential support that it offers.[10]

Unlike traditional foundationalism, Bayesian conditioning is a diachronic constraint: it tells us what our beliefs ought to look like in the future, rather than whether they are justified at any given moment. Moreover, unlike traditional foundationalism, Bayesian conditioning justifies the *values* we assign our beliefs, rather than their contents. Nevertheless, foundationalism looks like it's a structure that can be applied to updates just as easily as it can be applied to beliefs. Traditional foundationalism makes justification a function of whether some belief (P) is in the set of beliefs justified by the agent's basic state (S):

Traditional foundationalism: f_S: P $\rightarrow \{0, 1\}$,

(1 if p $\in$ S, and 0 otherwise)

By contrast, where we take an update (UP) to be a probabilistic belief transition, Bayesian conditioning—or Bayesian foundationalism—makes justification a matter of whether this update is in the set of updates justified by the evidence that represents the agent's basic state (E):

Bayesian foundationalism: f_E: UP $\rightarrow \{0, 1\}$,

(1 if up $\in$ E, and 0 otherwise)

Though Bayesian foundationalism is an account of evidence, in the sense of providing an account of the properties that our evidence must have, Bayesian foundationalism is clearly not the sort of epistemic constraint that is going to get us knowledge. This is a common complaint about Bayesian epistemology, and the discussion of this paper isn't intended to resolve it. Instead, Bayesian foundationalism begins from the assumption that, though Bayesianism does not entail a substantive account of epistemic reasons, we can nevertheless still ask about the structure that would guide these reasons, if they were there. And we can still say that this structure has a good-making feature. A Bayesian update on certain evidence has the formal, good-making feature of forever eliminating any beliefs that are inconsistent with those beliefs that we hold with certainty. This structure is Bayesianism's formal constraint on evidence, and it defines the sense in which Bayesian conditioning is commonly assumed to be a form of foundationalism.

[10] I take this to be consistent with the idea that *before* a belief becomes evidence—when it is merely a proposition that the agent has some credence in—it is able to receive some inferential support.

3 Foundationalism Undermined

Most take the fundamental idea behind Jeffrey conditioning to be the thought that, as Jeffrey (1983, p. 171) himself put it: "it is rarely or never that there is a proposition for which the direct effect of an observation is to change the observer's degree of belief in that proposition to one." Most of the time we have an experience that changes our credence in some proposition without making us sure of it. We get a quick glimpse of color on the floor that makes us think that the sock might be red. But maybe it's really brown. Or maybe it's purple.

To capture this more realistic class of cases, we need a rule that tells us how we ought to revise our beliefs whenever we get this sort of uncertain evidence. Jeffrey (1965) introduces a rule that does just this by assuming that our evidence is a partition of propositions that can be assigned values other than zero and one:

> **Jeffrey conditioning:** If experience directly changes your credences over a partition $\{B_i\}$ from $p(B_i)$ to $p'(B_i)$, then your new degree of belief in A, for any A, should be $p'(A) = \sum_i p(A | B_i) p'(B_i)$.

This formulation makes it clear that, formally, Jeffrey conditioning has Bayesian conditioning as a special case. Both updating rules say that we should revise our beliefs in accordance with the conditional probability that our evidence determines. Assuming our evidence to be a partition allows us to accommodate the uncertainty of some pieces of evidence by allowing us to assign probabilities other than one and zero to the possibility that the sock is red, and to the possibility that it is brown, and to the possibility that it is purple—which, together, will sum to one. Assuming our evidence to be a partition also allows us to accommodate the certainty of some pieces of evidence by allowing us to assign probability one to the possibility that the sock is red and probability zero to the possibility that it isn't.

But although Bayesian conditioning is, in this sense, a special case of Jeffrey conditioning, it is a degenerate special case of it. Jeffrey conditioning lacks the foundationalist constraint that we've just described Bayesian conditioning as having. There are a couple of ways of understanding how this constraint is lacking. First, assume that we take the agent's basic state to be the evidence partition that she updates on. Since the propositions in this partition can receive a value of less than one—less than complete certainty—this partition does not include an infallible belief. More importantly, since the propositions in this evidence partition can receive a value of less than one, they are able to have their values changed by means of the same sort of inferential support that they offer by a future update. Therefore, the beliefs that comprise these evidence partitions violate the first and third conditions of foundationalism identified above.

Of course, many non-classical foundationalists also reject the first and third conditions identified above. Goldman (1988)'s reliabilism is an example of an externalist account that does this. Pryor (2000)'s phenomenal conservatism is an example of an internalist account that does this. Given that there is a way of understanding foundationalism that is quite common, and that violates the same constraints as Jeffrey conditioning, one might object that we cannot rule out Jeffrey conditioning as a form of foundationalism.

But notice that though the accounts just mentioned allow that an agent's current foundational belief can be changed later on, they are each able to provide *some* feature that their foundations must have now. Like classical foundationalists, non-classical foundationalists assume that there is some property that picks out a belief as being a foundational belief—whether this is being a belief that was formed by an unconditionally reliable process, or being a belief that is an undefeated seeming, etc. By contrast, Jeffrey conditioning makes no such assumption; on the Jeffrey framework, evidence has no such distinguishing property.[11]

So far, we've assumed that Jeffrey conditioning fails to be a form of foundationalism, in virtue of failing to satisfy the first and third conditions identified above. However, there's a different way of understanding how Jeffrey conditioning fails to be a form of foundationalism, one that understands this failure as a violation of the second condition.[12] Earlier we said that if there is no constraint on the evidence partition we update on—if there is no norm that tells us what this partition, or the values it gets assigned, ought to look like—then our account of updating is no stronger than Probabilism or, equivalently, **Descriptive Diachronic Coherence for Bayesians**. And this is precisely what Jeffrey conditioning does: it places no constraints on the weights that our evidence partition gets assigned. However, a different way of satisfying **Normative Diachronic Coherence for Bayesians** is to take the *experience* that gives rise to our evidence to be our basic state, and take Bayesianism's fundamental updating constraint to be that we revise our beliefs in the way that this experience prescribes.

But reconceiving the structure of the Bayesian formalism in this way isn't going to help us. This is because, while the Bayesian formalism includes a rule that regulates how a weighted evidence partition gives rise to an update, it does not include a norm that regulates how an experience gives rise to an update. Since experiences lack the inferential relation to updates that a basic state bears to non-basic states, they aren't better candidates for the role we are looking to fill.

Given all this, it looks like Jeffrey conditioning is strictly weaker than Bayesian conditioning. The latter includes a formal constraint on evidence that the former lacks.[13] These considerations also show us why the default view of Bayesian diachronic coherence that we briefly considered near the beginning of our discussion cannot help us unify updates on certain and uncertain evidence. Recall this view interprets Bayesianism's fundamental updating rule as one that does not take a stand on what our evidence is, but merely tells us what we ought to do once we have it in hand. The problem is that, even if we assume this interpretation of the Bayesian formalism, an update on certain evidence will *entail* certain other restrictions on our probability distribution in the future that Jeffrey conditioning does not entail. In particular, it will entail that we can no longer change

[11]Thanks to [redacted] for prompting me to be clearer about this.

[12]The following line of argument has a steady, if diffused, presence in the literature on Jeffrey conditioning. There are references to it as early as Carnap (1957) (reprinted in Jeffrey (1975)) and as late as Weisberg (2009).

[13]Again, one might object that Jeffrey conditioning is not entirely unconstrained, since it includes the foundationalist constraint for updates on certain evidence. Again, the appropriate response is to point out that this constraint fails to apply to Jeffrey updates in general.

our credence in any proposition that is inconsistent with our evidence. Therefore, adopting the default view will not allow us to avoid the strange consequence that updates on certain and uncertain evidence pattern in very different ways.

Finally, one might object that if Jeffrey conditioning fits uneasily within the Bayesian updating paradigm because it lacks a constraint on evidence, then so much the worse for Jeffrey conditioning. Most discussions about Bayesian updating take Jeffrey conditioning to be a literal footnote to the Bayesian program. But this objection misunderstands why discussions about Bayesian updating tend not to focus on Jeffrey conditioning. We idealize away from the sorts of cases that Jeffrey conditioning covers, not because they aren't important, but because we assume that they will be covered in the same way as regular Bayesian conditioning. To show that they are not covered in this way, then, is to undermine the Bayesian program in its entirety. Not only is it true that the cases we find ourselves in most often are those where we *do* get uncertain evidence, but if we take on the assumption that beliefs come in degrees, it seems arbitrary to withhold that assumption from those beliefs that constitute 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".

4 Bayesian Coherentism

4.1 Commutativity as Coherence

How *do* we unify Bayesian updates then? Since Bayesian conditioning entails a constraint that makes it stronger than Jeffrey conditioning, putting these two updating rules on a par will require making Jeffrey conditioning stronger. But we've just seen that we can't make Jeffrey conditioning stronger by making it a form of foundationalism. Putting these two updating rules on a par, then, will require reinterpreting the norm that supervenes on the formal property that makes Bayesian conditioning so strong. In this section, I'll argue that we can do this by reconceiving Bayesian updating as a form of coherentism.

The fundamental difference between regular Bayesian conditioning and Jeffrey conditioning has always been assumed to be that the latter generalizes the certainty of evidence. A second notable difference in these frameworks is that only the regular framework is commutative over weighted evidence partitions: only the regular framework makes the order in which we get evidence irrelevant to the credence distribution we end up with at each and every time that we update. The non-commutativity of Jeffrey conditioning is a significant mark against it. Say I come to learn that my nephew's baseball team won their game. And then I come to learn that my boss has given me a raise. It does not seem as though receiving these pieces of evidence in reverse order should change the credences I end up with. Consistency seems to require that identical pieces of information be treated the same, no matter the order in which they are received. If we have the same information in two sequences of updates, just arranged a little differently, this mere difference in ordering should make no difference.[14] (One might object that if we con-

[14] See Domotor (1980) and Doring (1999) for discussions that explicitly criticize Jeffrey conditioning for being non-commutative.

ceive of the inputs to the Jeffrey updating process as something other than weighted partitions, we don't get the result that the framework is non-commutative. We'll consider the relevance of this in the following section.)

While much discussed in the literature, the non-commutativity of Jeffrey conditioning has never been assumed to be a defining feature of it, in the way that the uncertainty of evidence has been so understood. Instead, it has been assumed to be an unfortunate, but non-essential property of the Jeffrey framework: the Jeffrey framework has the property of sometimes yielding updates that aren't commutative over weighted evidence partitions. This suggests an intriguing possibility: why *not* take the fundamental norm that grounds all Bayesian updates, including Jeffrey updates, to be that they minimize the defect of failing to commute. Why not take the norm for evidence that governs all updates to be, not that these updates be grounded in a certainty, as Bayesian foundationalism would have it, but that they be minimally non-commutative. This would mean understanding the formal norm for evidence that governs updates to be the requirement that the values these updates yield be as insensitive as possible to the order in which these updates were made. It would mean evaluating updates based upon the extent to which they are consistent in this way.

Whether or not this proposal for grounding Bayesian updates is reasonable depends upon whether we think that minimizing the extent to which updates fail to commute is a norm that Bayesians ought to be interested in. Given that so much has been made of the commutative property in the Bayesian literature, it's clear that it *is* a norm that Bayesians ought to be interested in. And I think we can say something even stronger than this; I think we can give the commutative norm I've just described an interesting gloss. I think that a norm that requires that we minimize the extent to which Bayesian updates fail to commute makes the Bayesian framework look like a form of coherentism about updating. In order to see this, it will again be useful to consider what coherentism looks like as a structure of justification that applies to beliefs.

Like traditional foundationalism, traditional coherentism assumes that the target of justification is a set of beliefs. Unlike traditional foundationalism, traditional coherentism assumes that the locus of justification is not some particular cognitive state, but is instead the relation that our beliefs stand in with one another. Traditional coherentism maintains that some set of beliefs is justified exactly when its component beliefs fit correctly, or cohere, with one another. On a standard coherentist account, probabilistic coherence, logical coherence, and evidential coherence are each measures that contribute to a belief set's coherence.[15] Do these measures translate to the Bayesian framework? Logical coherence and probabilistic coherence are both preserved by Bayesianism's synchronic constraints: they are preserved no matter how we understand the structure of diachronic coherence for Bayesians. The interesting question, then, is what an account of

[15] There are, of course, many different kinds of coherentist accounts of justification, both historically and contemporarily. And many earlier coherentists did not endorse all three of these constraints. Ewing (1934), for instance, takes coherence to be a matter of logical coherence alone, while Lewis (1946) takes coherence to be a matter of probabilistic coherence alone. Notably, Bonjour (1985) takes coherence to be a matter of logical and probabilistic coherence, as well as a number of other requirements that might be held to fall under the heading of evidential coherence (see pp. 97-99 for the details).

evidential coherence will amount to in a Bayesian updating setting. What could it mean for a set of weighted evidence partitions to be consistent? And what could it mean for them to be consistent over updates?

It's well-understood what evidential coherence amounts to in the traditional setting: it is a measure of the degree to which some proposition confirms each other belief in the set to which it belongs. It is a measure of the degree to which every proposition in a set is *evidence for* every other proposition in that set. If I hold the belief that it will rain in a few hours (P_1), and also the belief that the owner of the shop down the street just put out her umbrella stand (P_2), then the belief that the baseball game will be rained out this afternoon (P_3), if it increases the proportion and strength of the inferential connections between the beliefs in this set, increases the evidential coherence of this set of beliefs.[16] Traditional coherentism makes justification a function of the degree to which some set of propositions is consistent.

I think we are now in a position to understand what Bayesian evidential coherence might amount to. We can triangulate on an account of Bayesian coherentism from the descriptions of Bayesian foundationalism and traditional coherentism we already have. From our earlier description of *Bayesian foundationalism*, we borrow the idea that what we are interested in are updates:

$$f_E\colon \text{UP} \rightarrow \{0, 1\}$$

From the above description of *traditional coherentism*, we borrow the idea that the relation of justification we are interested in is consistency:

$$f\colon \{P_1, P_2, ...P_n\} \rightarrow \mathbb{R}^+.$$

Together, these commitments entail that the ideal of justification for the Bayesian coherentist is a consistent set of updates, UP, where, again, consistency comes in degrees:

$$f\colon \{UP_1, UP_2, ...UP_n\} \rightarrow \mathbb{R}^+$$

The ideal of *Bayesian coherentism*, then, is the requirement that one's updates be consistent. Since the only way for an updating framework to treat consistently the weighted evidence partitions for a set of updates is to require that they yield the same updates whenever we get them, Bayesian coherentism is just the requirement that our updates commute.

If all this is right, then an alternative to understanding Bayesian updating as a form of foundationalism is to understand it as a form of coherentism. I'll go on to say more in the next section about what Bayesian coherentism might amount to by offering a proposal

[16] Many contemporary coherentist accounts spell out the notion of an inferential connection probabilistically. (There are many such accounts. For one notable account, see Fitelson (2003).) For instance, a very simple view might hold that, in the previous case, what accounts for the increased coherence provided by P_3 is that $p(P_1 \mid P_2) < p(P_1 \mid P_2 \wedge P_3)$ and $p(P_2 \mid P_1) < p(P_2 \mid P_1 \wedge P_3)$.

for how commutativity can be represented as a relation that comes in degrees. Before we do that, however, it will be helpful to get a feel for where we are right now. We can state Bayesian foundationalism and Bayesian coherentism side by side in a way that illustrates that each provides us with a different version of **Normative Diachronic Coherence for Bayesians**:

> **Bayesian foundationalism:** A probabilisitic belief transition will be such that:
>
> (a) There is a sufficient partition, $\{E_i\}$, for the transition.
>
> (b) It is diachronically coherent iff some E_i is held with certainty.
>
> **Bayesian coherentism:** A pair of probabilisitic belief transitions will be such that:
>
> (a) There are sufficient partitions, $\{E_i\}$, $\{F_j\}$, for each of these transitions, and
>
> (b) They are diachronically coherent to the extent that they minimize evidential incoherence (in a sense that will be made more precise in the following section).

As I noted earlier, an interesting feature of Bayesian coherentism is that, unlike either Bayesian conditioning (i.e., Bayesian foundationalism) or Jeffrey conditioning, it is undefined for a single update. Therefore, it is not entailed by either Bayesian conditioning or Jeffrey conditioning. While this makes Bayesian coherentism an amendment to the traditional Bayesian framework, it is not an amendment that requires the Bayesian to take on any additional substantive commitments. No matter what other commitments one happens to hold, inconsistency will always be a *prima facie* defect. This explains the importance that Bayesians, and formal epistemologists in general, have placed on the commutative property. In effect, Bayesian coherentism represents a way of articulating a commitment that Bayesianism, as well as every other normative theory, already holds.

It's important to be clear once again that, like Bayesian foundationalism, Bayesian coherentism isn't going to get us knowledge. The sorts of substantive reasons that we would need for this are precisely those that a formal account of evidence is *not* going to yield. But though Bayesian coherentism (like Bayesian foundationalism) isn't an account of epistemic reasons, it earns its name by drawing upon one of coherentism's structural features to define a formal account of evidence. Bayesian coherentism has the formal, good-making, coherentist feature of making justification depend upon the consistency of evidence. This structure defines the sense in which Bayesian updating might be assumed to be a form of coherentism.

4.2 *A Unified Account of Bayesian Updating*

The final piece of the puzzle is to see how adopting Bayesian coherentism helps us with the problem of being able to say that both updates on certain and uncertain evidence proceed from the same normative structure. For, at first glance, it looks as though our initial problem persists: it looks as though Jeffrey conditioning bears the same relation to Bayesian coherentism that it bears to Bayesian foundationalism. Updates on uncertain evidence fail, in general, to be updates on basic states. But they also fail, in general, to be updates that commute. So, is appealing to a commutative norm really any different than appealing to a foundationalist norm?

There *is* a relevant difference between these two sorts of appeals. What makes Bayesian foundationalism problematic is that adopting it would mean having to say that every update on uncertain evidence, qua update on uncertain evidence, is incapable of making the agent diachronically coherent. Trivially, updates on uncertain evidence aren't capable of being updates on certain evidence. And updates on certain evidence are the only updates that have foundationalist properties.

But Bayesian coherentism does not have this same feature. This is because both updates on certain *and* uncertain evidence are capable of commuting. If we are looking for a norm to unify these two types of updates, then, a norm that makes diachronic coherence a matter of updates commuting is capable of fulfilling this function. The fact that updates on uncertain evidence, qua updates on uncertain evidence, are capable of satisfying the norm to commute, suggests that the best interpretation of why some set of updates on uncertain evidence have failed to commute is that they have failed to conform to Bayesian coherentism. By contrast, the fact that updates on uncertain evidence, qua updates on uncertain evidence, *aren't* capable of satisfying the norm to be an update on a basic state, means that the *only* explanation for why updates on uncertain evidence fail to proceed from a basic state is that the foundationalist norm that we would need to render this verdict just isn't there.[17]

In short, then, the fact that updates on uncertain evidence can't conform to a norm formulated in terms of a basic state entails that such updates aren't governed by Bayesian foundationalism. It entails that there is no such norm. By contrast, the fact that updates on uncertain evidence *are* capable of conforming to Bayesian coherentism suggests that they *can* be governed by Bayesian coherentism. It suggests that Bayesian coherentism might be a norm for such updates. Given what we've seen in this section, I think we can say that Bayesian coherentism is plausibly a norm for such updates.

Indeed, I think we can say something even stronger than this. I think we can say that, not only are all updates on uncertain evidence capable of satisfying Bayesian coherentism, but that all updates on uncertain evidence *do*, to some extent at least, satisfy Bayesian coherentism. As I've alluded to already, identifying commutativity with coherence makes it natural to want to give commutativity a degree-theoretic interpretation. I do this in §5. This move will let us say that all Bayesian updates are diachronically coherent, to a degree. Before we get to that, however, we need to consider an important worry

[17] This follows from standard deontic logic, which says that a norm can't require X if X is logically impossible. Thanks to [redacted] for helping me to clarify this point.

for the proposal.

4.3 An Objection

We've assumed that some Jeffrey updates are non-commutative in a way that ought to be minimized or avoided. But is this right?

Lange (2000) famously argues that although the Jeffrey framework is non-commutative over weighted evidence partitions, this feature of it does not make it defective. This is because the framework does commute the *experiences* that underwrite belief revisions. More precisely, it commutes the Bayes factors we might plausibly identify with these experiences. If we take experiences, rather than weighted evidence partitions, to be the inputs to the updating process, then we get the result that the Jeffrey framework does indeed commute its inputs. Moreover, since this picture of the Jeffrey framework has Bayesian conditioning as a limiting case of Jeffrey conditioning, but not as a special case of it, this picture of the Jeffrey framework undermines Bayesian coherentism, not only by eliminating the need to regulate weighted evidence partitions, but by eliminating our inconsistent triad.[18]

It's true that there's an understanding of the Jeffrey framework that makes my solution unnecessary and, so, unmotivated. But it does this at the cost of denying that the Bayesian updating framework should be taken to be an account of epistemic justification in the first place. We saw in the last section that a problem with taking experiences to be our foundations is that the Bayesian framework does not entail any rational way of mapping these experiences to our updates. It's for this reason that Field argued that an account of how experience figures into the updating process should not be conceived of as a rational account, but as "the problem of giving a complete psychological theory for a Bayesian agent" (p. 364). This picture of the structure of the Bayesian framework stands in tension with our inconsistent triad, which assumes that we have reason to want to square the formal features of Bayesian conditioning and Jeffrey conditioning with their normative features. While it's true that we can avoid an inconsistency between these two sets of features by insisting that the latter do not exist, I assume that, for many, this would be to give up the game.[19]

There may be other ways of undermining Bayesian coherentism. For instance, whether or not we would want our weighted evidence partitions to commute may depend, in part, upon the sort of content we take our evidence propositions to have. It would take a much longer discussion to sort out that issue. While we'll consider a more general objection to Bayesian coherentism in the next section, it's worth noting now a point that will be returned to there, which is that the aim of this paper isn't to provide a thoroughgoing

[18]For the point that Bayesian conditioning is not a special case of Jeffrey conditioning, on this picture of things, see Field (1978, p. 365). Lange (2000, p. 397) also gestures towards this idea.

[19]It's worth noting that even if there were some way of revising our normative framework, so as to make it an account of epistemic justification, it's not clear that this would be enough to undermine Bayesian coherentism. The inputs to the updating process cannot be *both* experiences and weighted evidence partitions. Therefore, it's clear that we don't need both experiences and weighted evidence partitions to commute. But this does not entail that we should choose the former over the latter, in cases where some revision to the framework needs to be made.

defense of Bayesian coherentism. The claim this paper advances is that, *if* we want to unify all Bayesian updates under a single norm, Bayesian coherentism is, plausibly, the form that this norm should take. While the aim of unifying Bayesian updates gives us some reason to adopt Bayesian coherentism, it is consistent with there being reasons, even very weighty reasons, that run in the other direction.

5 Evidential Incoherence

In the last section, I proposed an interpretation of the Bayesian framework that allows us to say that updates on certain and uncertain evidence proceed from the same normative structure. This proposal assumes that we can assess the incoherence of sets of updates based upon the extent to which they instantiate what has long been deemed to be a bad-making feature of the Bayesian formalism. If one wants to reject the proposal then one must either deny that (1) commutativity is an important feature for an updating rule to guarantee, or that (2) the fact that commutativity is an important feature for an updating rule to guarantee does not mean that it is an important feature for individual sequences of updates to have. Absent an argument for one of these claims, I assume that we have good reason to proceed with the question of how a norm that draws on this intuitive idea might be developed.[20]

An obvious place to begin is with a formal property that we can identify with commutativity. Diaconis and Zabell (1982, p. 825) have shown that the property of *Jeffrey independence* is both necessary and sufficient for commutativity.[21] Here's what this property amounts to:

> **Jeffrey Independence:** Let P be a probability function. And let $P_{\mathcal{E}}$ and $P_{\mathcal{F}}$ be the probability functions that result from updating P on the partitions $\mathcal{E}=\{E_i\}$ and $\mathcal{F}=\{F_j\}$, respectively. The partitions $\mathcal{E}$ and $\mathcal{F}$ are Jeffrey independent with respect to $\{p_i\}$ and $\{q_j\}$ if $P_{\mathcal{E}}(F_j)=P(F_j)$ and $P_{\mathcal{F}}(E_i)=P(E_i)$ holds for all *i* and *j*.
>
> Thus, Jeffrey independence says that Jeffrey updating on $\mathcal{E}$ with probabilities p_i does not change the probabilities on $\mathcal{F}$ and similarly with $\mathcal{E}$ and $\mathcal{F}$ interchanged.

The most straightforward way of developing the idea that commutativity is a form of evidential coherence is to require that a sequence of updates be Jeffrey independent. However, if we are looking to mimic the concept of evidential incoherence that we are borrowing from traditional epistemology, a degree-theoretic account of Bayesian evidential incoherence seems more appropriate.

[20] *Can* we reject one of these two claims? I've suggested that the first claim seems unimpeachable. The formal epistemology literature seems to care a lot about the commutativity of formal updating rules.

What about (2)? Perhaps one might want to argue that the kind of defect the non-commutativity of updates represents is not a defect of particular updates, but is a sort of inconsistency that inheres in the framework in general. However, it's difficult to imagine what it would mean for the framework in general to be defective, in a way that doesn't accrue to particular updates. I think, then, that we are safe in proceeding.

[21] For the theorem, see Diaconis and Zabell (1982), §3.3.

How do we get a degree-theoretic account of Bayesian evidential incoherence? Having identified complete evidential coherence with commutativity, and so with Jeffrey independence, an obvious approach would be to define evidential incoherence as the degree of a violation of Jeffrey independence for a sequence of updates. In other words, we might take the agent's degree of evidential incoherence to be a measure of how much the weighted evidence partitions in question fail to be independent of each other. There are undoubtedly many plausible ways of devising a measure of this violation. I won't consider in detail the merits of each of them. However, in the appendix, I sketch one plausible measure of a violation of Jeffrey independence, which I therefore also take to be a plausible measure of Bayesian evidential incoherence.

Assuming that my measure represents a plausible way of giving content to the notion of evidential incoherence, the next question to ask is how this measure should be put to use in a norm. Since the aim of this paper is merely to establish Bayesian coherentism as an alternative to Bayesian foundationalism, rather than to work out what the best version of Bayesian coherentism might turn out to be, I won't consider that question here. Perhaps we would want our norm to govern only pairs of sequential updates. Or maybe we would want our norm to govern larger sets of updates taken pairwise. But however we choose to go, it's clear that the norm we would want must be committed to the following:

> **Bayesian coherentism (Revised):** A pair of probabilistic belief transitions will be such that:
>
> (a) There are sufficient partitions, $\{E_i\}$, $\{F_j\}$, for each of these updates, and
>
> (b) They are diachronically coherent to the extent that they minimize evidential incoherence, where the latter will be a function of the degree to which $\{E_i\}$, $\{F_j\}$, relative to their weights, $\{p_i\}$ and $\{q_j\}$, are Jeffrey independent.

It's a common idea that there are degrees of probabilistic incoherence.[22] This discussion introduces the idea that there are also degrees of diachronic incoherence that aren't reducible to the latter by defending a degree-theoretic account of **Normative Diachronic Coherence for Bayesians** that isn't reducible to **Descriptive Diachronic Coherence for Bayesians**. Or, equivalently, by defending a degree-theoretic account of **Normative Diachronic Coherence for Bayesians** that isn't reducible to Probabilism. On the account that I've called Bayesian coherentism, perfect normative diachronic coherence is the special case where the agent's updates commute.

We've noted throughout that Bayesian coherentism is not going to get us knowledge. Given this, and given the formal nature of its constraint, one might assume that Bayesian coherentism will stand in need of more justification than we have given it here. Incoherence on the Bayesian framework has traditionally been taken to be more or less

[22] See Schervish, Seidenfeld, and Kadane (2000, 2002a, 2002b). For a more recent account, see Staffel (2015).

synonymous with the propensity for an agent's credal state to fail to maximize some sort of utility. Dutch book arguments famously show that where an agent fails to be coherent, her credal state fails to maximize practical utility by sanctioning a series of bets that would lead to a sure loss of money.[23] More recently, accuracy arguments have been taken to establish that, where an agent fails to be coherent, she fails to have a credal state that minimizes inaccuracy.[24] The Bayesian norms that prescribe against the sort of incoherence that these arguments target are taken to be justified by these very facts.

Whether Bayesian coherentism can be given this sort of defense is a question for another day. This question belongs to a more general discussion about whether formal evidential norms can be vindicated by dutch book arguments and accuracy arguments.[25] Again, the aim of this paper hasn't been to offer a comprehensive defense of Bayesian coherentism, but to argue that this is the structure the Bayesian formalism would need to have in order to unify updates on certain and uncertain evidence. This discussion has assumed, as I think we always have, that we can identify some intuitive features that we would want to ascribe to the Bayesian framework, and *then* go on to ask whether the norms these intuitive features give rise to are utility-promoting. This paper has carried out the first part of this project. If, in the end, it turns out that there is nothing to be said for Bayesian coherentism, other than that it does unify Bayesian updates—if it turns out that this norm does not in fact minimize inaccuracy or make the agent dutch book invulnerable—it may turn out that there is not an enormous amount of pressure to minimize incoherence in the way that Bayesian coherentism advises.

6 Final Thoughts

I want to conclude by reconsidering the motivation for this discussion. We have been assuming throughout that Jeffrey conditioning and regular Bayesian conditioning ought to be brought together; that updates on certain and uncertain evidence should proceed from frameworks with the same normative structure. But maybe they shouldn't. Maybe one can provide a principled explanation for why they don't. Maybe like Field (1978, p. 365) claims, when discussing his own variant of Jeffrey conditioning, we would want to say that Bayesian conditioning is too much of an idealization to ever be of any use. Or maybe like Lange (2000, p. 397), we would want to hold that the conditions under which we get uncertain evidence differ in relevant ways from those under which we get certain evidence. Lange argues that our background beliefs only make a difference to the probative value of an experience in cases where this experience gives rise to uncertain

[23] See de Finetti (1937, 1972) for a justification of Probabilism. See Lewis (1999) and Teller (1973), and Skyrms (1987) for justifications of Bayesian conditioning, and Jeffrey conditioning, respectively.

[24] See Joyce (1998, 2009) for justifications for Probabilism, and Greaves and Wallace (2006) for a justification for updating, and Leitgeb and Pettigrew (2010a, 2010b) and Pettigrew (2016) for justifications for both of Probabilism and updating.

[25] This question has been thrown into focus by much of the recent literature on accuracy. See, for instance, Meacham (forthcoming) for some worries that the accuracy framework might not be able to vindicate evidential norms. See Pettigrew (2016) for an argument that it can. Other norms that might fall into this category, insofar as they impose formal constraints on the agent's credence function before the agent updates, include Lewis's (1980) Principal Principle and van Fraassen's (1984) Rational Reflection.

evidence. Though he does not elaborate on why he thinks this, he does note at one point that cases where we update on evidence to which we have assigned a credence of one are cases where our background beliefs fail to function as "extended sense organs", in the way that they do when we assign our evidence some other value.[26]

Despite the weird imagery, this does not seem like a crazy suggestion. For starters, it does seem as though some propositions, though they might be triggered by experience, are not *justified* by experience. When I change my credence in the proposition that a difficult math proof is correct, though this change may be accompanied by certain sensory experiences that are brought about by introspection, these experiences do not seem to be what *justify* these revisions, in the way that my belief that the sky is blue is justified by an experience with a certain phenomenal character.[27] If this is the case—and if it is also the case that updates on certain evidence are exactly those updates that aren't justified by experience—then the agent's prior expectations won't have a hand in determining what her experience justifies since, in these cases, the agent's experience does not justify anything at all.

More generally, the previous passage raises a possibility that we have not yet considered, which is that the type of content to which we happen to be justified in assigning a credence of one might differ in some relevant way from the type of content to which we happen to be justified in assigning a lesser value. If there is indeed this difference in content between our certain and uncertain evidence, the unified account that we have been after in this paper may be inappropriate. For while it may be implausible that there is a sharp cut-off between certain and uncertain evidence, there may very well be a sharp cut-off between the different types of propositions this evidence corresponds to. Updates on certain evidence might be special, then, in the way that a verdict that precludes all reasonable doubt is special. If they are, there may be reason, after all, to think that updates on certain and uncertain evidence proceed from frameworks with different normative structures.[28]

Of course, those who tell this sort of story owe us an account of why we might be justified in assigning only some particular class of propositions a credence of one. This might be a considerable task. Or it might not be. A modest proposal along these lines would be to appeal to the principle of Continuing Regularity. This principle says that we should assign probability one only to logical truths and zero only to contradictions (or, to necessary and impossible propositions, respectively). While not universally endorsed, this principle is believed by many Bayesians to be quite plausible. And since there is more or less agreement about which propositions are necessary and impossible, we would easily be able to identify the sorts of propositions to which we are justified in assigning a credence of one.

Maybe, then, there's some argument from Continuing Regularity to the conclusion that certain and uncertain evidence ought to proceed from frameworks with different normative structures. It would be interesting if there were. I think it's of interest to con-

[26] Lange (2000, p. 400).

[27] See [redacted] for further discussion.

[28] Of course, this will depend upon how they are different. In the end, such a difference might very well turn out to be irrelevant as well.

sider all the possible ways that the core commitments of the Bayesian framework can be articulated. One such way, which we have been considering here, is suggested by what appears as a footnote in nearly every paper on Bayesian epistemology. This is the assumption that Bayesian conditioning and Jeffrey conditioning are perfect parallels with respect to their formal structures. Since Bayesian updating is a normative theory, I have argued that it makes some sense to ask what it would mean for these updating rules to also be perfect parallels with respect to their normative structures. This paper has tried to answer this question. Maybe there's not much going for the answer that's been provided besides its connection to the apparent truism that appears in all of these footnotes. But I do think it's interesting—and, also, surprising—to discover what this apparent truism ends up committing us to.

Appendix: A Measure of Evidential Incoherence

There are different formal measures we might use to give content to the notion of evidential incoherence. In this appendix, I develop one such measure, which draws upon the seminal work of Diaconis and Zabell (1982) on the formal properties of Jeffrey conditioning. Since this measure is intended to track the degree to which some set of evidence fails to commute, it draws upon a property that is both necessary and sufficient for commutativity (p. 825):

> **Jeffrey Independence:** Let P be a probability function. And let $P_{\mathcal{E}}$ and $P_{\mathcal{F}}$ be the probability functions that result from updating P on the partitions $\mathcal{E}=\{E_i\}$ and $\mathcal{F}=\{F_j\}$, respectively. The partitions $\mathcal{E}$ and $\mathcal{F}$ are Jeffrey independent with respect to $\{p_i\}$ and $\{q_j\}$ if $P_{\mathcal{E}}(F_j)=P(F_j)$ and $P_{\mathcal{F}}(E_i)=P(E_i)$ holds for all *i* and *j*.
>
> Thus, Jeffrey independence says that Jeffrey updating on $\mathcal{E}$ with probabilities p_i does not change the probabilities on $\mathcal{F}$ and similarly with $\mathcal{E}$ and $\mathcal{F}$ interchanged.

It should be clear how the lack of Jeffrey independence undermines the commutativity of updates. Since Jeffrey updating fixes the probabilities on the evidence partition, the values for the members of a partition that we get when we update on it first will be the same values that we get when we update on this same partition second. This means that the order in which we update on a partition will make a difference to the agent's final credence distribution. More carefully, if $P_{\mathcal{E}}(F_j)\neq P(F_j)$, then where we get $\mathcal{F}$ first, our second update on $\mathcal{E}$ will change the values along $\mathcal{F}$. Therefore, these values will differ from those that we would have gotten by updating on $\mathcal{F}$ second. Similarly, if $P_{\mathcal{F}}(E_i)\neq P(E_i)$, then where we get $\mathcal{E}$ first, our second update on $\mathcal{F}$ will change the values along $\mathcal{E}$. Therefore, these values will differ from those that we would have gotten by updating on $\mathcal{E}$ second.

Given the connection between Jeffrey independence and commutativity, if we are looking to assess the degree to which some set of updates fails to commute, understanding what it would mean for some set of evidence to fail to be Jeffrey independent *to a degree* seems like a good place to start. Given what we've just said, one plausible way of interpreting the degree of a violation of Jeffrey independence is as the sum of the amount by which updating on $\mathcal{E}$ with probabilities $\{p_i\}$ changes the probability of each F_j, taken together with the sum of the amount by which updating on $\mathcal{F}$ with probabilities $\{q_j\}$ changes the probability of each E_i.

To begin to formalize this idea, notice that since Jeffrey independence says that $P_{\mathcal{E}}(F_j) = P(F_j)$, for all *j*, and $P_{\mathcal{F}}(E_i) = P(E_i)$, for all *i*, it will hold whenever $\frac{P_{\mathcal{E}}(F_j)}{P(F_j)}=1$, for all *j*, and $\frac{P_{\mathcal{F}}(E_i)}{P(E_i)}=1$, for all *i*. Plausibly, then, the degree of a violation of Jeffrey independence will correspond to the amount by which each of these diverges from 1, for all *j* and for all *i*, respectively. To formulate a measure that can account for this in a perspicuous way, we can note, with Diaconis and Zabell, that $\frac{P_{\mathcal{E}}(F_j)}{P(F_j)}=\sum_i p_i r_{ij}$ where,

$$r_{ij} = \frac{P(E_iF_j)}{P(E_i)P(F_j)}.$$

Given this, the degree of a violation of Jeffrey independence will correspond to the sum of the amount by which $\sum_i p_i r_{ij}$ diverges from 1, for each *j* (i.e., $\sum_j(|1\text{-}\sum_i p_i r_{*,j}|)$), taken together with the sum of the amount by which $\sum_j q_j r_{ij}$ diverges from 1, for each *i* (i.e., $\sum_i(|1\text{-}\sum_j q_j r_{i,*}|)$).[29] The following is a measure of evidential incoherence that follows from this idea:

> **Evidential Incoherence**: Let P be a probability function and let $P_{\mathcal{E}}$ be the probability function that results from updating P on $\mathcal{E}$ with probabilities p_i. Further, let $r_{ij}=\frac{P(E_iF_j)}{P(E_i)P(F_j)}$, $r_{ij'}=\frac{P_{\mathcal{E}}(E_iF_j)}{P_{\mathcal{E}}(E_i)P_{\mathcal{E}}(F_j)}$, if defined, and 1 otherwise.
>
> A sequence of updates on the partitions, $\mathcal{E}$ and $\mathcal{F}$, is coherent to the extent that it minimizes $\sum_j(|1\text{-}\sum_i p_i r_{*,j}|)+\sum_i(|1\text{-}\sum_j q_j r_{i,*'}|)$.

We can run through a couple of examples to see how this measure will work.

(1) Consider the following initial credence distribution $P(E_iF_j)$:

	F_1	F_2	F_3	
E_1	.25	.125	.125	.5
E_2	.125	0	.125	.25
E_3	.125	.125	0	.25
	.5	.25	.25	

Now assume that the agent gets as evidence, $p(E_1)$=**.5**, $p(E_2)$=**.2**, $p(E_3)$=**.3**, and $q(F_1)$=**.2**, $q(F_2)$=**.4**, $q(F_3)$=**.4**, in turn. After the first update, we are left with the values on the left. After the second update, we are left with the values on the right:[30]

	F_1	F_2	F_3	
E_1	.25	.125	.125	**.5**
E_2	.1	0	.1	**.2**
E_3	.15	.15	0	**.3**
	.5	.275	.225	

	F_1	F_2	F_3	
E_1	.1	.182	.22	.502
E_2	.04	0	.18	.22
E_3	.06	.218	0	.278
	.2	**.4**	**.4**	

[29] Where we take (r_{ij}) to be a matrix of the relations of dependence that hold between $\{E_i\}$ and $\{F_j\}$, $r_{i,*}$ refers to the i^{th} row of (r_{ij}), and $r_{*,j}$ refers to the j^{th} column of (r_{ij}).

[30] Some of the values in the second matrix are approximations. However, since they are not inputs into our measure, this makes no difference.

Before the first update, (r_{ij}) represents the relations of dependence that hold between $\{E_i\}$ and $\{F_j\}$. Before the second update, $(r_{ij'})$ represents the relations of dependence that hold between $\{E_i\}$ and $\{F_j\}$:

$$(r_{ij}) = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 2 \\ 1 & 2 & 0 \end{pmatrix} \quad (r_{ij'}) = \begin{pmatrix} 1 & .91 & 1.11 \\ 1 & 0 & 2.22 \\ 1 & 1.82 & 0 \end{pmatrix}$$

Multiplying each member of each column of (r_{ij}) by p_i, and taking the amount by which the sum of the values of each column diverges from one, and multiplying each member of each row of $(r_{ij'})$ by q_j, and taking the amount by which the sum of the values of each row diverges from one yields:

$$\sum_j(|1\text{-}\sum_i r_{*,j}p_i|)+\sum_i(|1\text{-}\sum_j r_{i,*'}q_j|)\approx .37$$

This is a measure of the evidential incoherence of the updates described.

(2) Now assume that, given the same initial credence distribution, the agent gets as evidence, $p(E_1)=\mathbf{.6}$, $p(E_2)=\mathbf{.2}$, $p(E_3)=\mathbf{.2}$ and, $q(F_1)=\mathbf{.4}$, $q(F_2)=\mathbf{.3}$, $q(F_3)=\mathbf{.3}$, in turn. After the first update, we are left with the values on the left, and after the second update, we are left with the values on the right:

	F_1	F_2	F_3	
E_1	.3	.15	.15	**.6**
E_2	.1	0	.1	**.2**
E_3	.1	.1	0	**.2**
	.5	.25	.25	

	F_1	F_2	F_3	
E_1	.24	.18	.18	.6
E_2	.08	0	.12	.2
E_3	.08	.12	0	.2
	.4	**.3**	**.3**	

Before the first update, (r_{ij}) represents the relations of dependence that hold between $\{E_i\}$ and $\{F_j\}$. Before the second update, $(r_{ij'})$ represents the relations of dependence that hold between $\{E_i\}$ and $\{F_j\}$:

$$(r_{ij}) = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 2 \\ 1 & 2 & 0 \end{pmatrix} \quad (r_{ij'}) = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 2 \\ 1 & 2 & 0 \end{pmatrix}$$

Multiplying each member of each column of (r_{ij}) by p_i, and taking the amount by which the sum of the values of each column diverges from one, and multiplying each member of each row of $(r_{ij'})$ by q_j, and taking the amount by which the sum of the values of each row diverges from one yields:

$$\sum_j(|1\text{-}\sum_i r_{*,j}p_i|)+\sum_i(|1\text{-}\sum_j r_{i,*'}q_j|)=0$$

This is a measure of the evidential incoherence of the updates described. This set of updates exhibits perfect evidential coherence or, equivalently, perfect diachronic coherence.

References

[1] BonJour, Laurence. 1985. *The Structure of Empirical Knowledge.* Harvard University Press.

[2] Bradley, Richard. 2005. "Radical Probabilism and Bayesian conditioning." *Philosophy of Science.*

[3] Carnap, Rudolf. 1950. "Empiricism, Semantics, and Ontology." *Revue Internationale de Philosophie* 4 (11):20-40.

[4] Chrisman, Matthew. 2008. "Ought to Believe". *Journal of Philosophy* 105 (7):346-370.

[5] Christensen, David. 1992. "Confirmation Holism and Bayesian Epistemology." *Philosophy of Science*, 59(4): 540-57.

[6] de Finetti, Bruno. 1937. "Foresight: Its Logical Laws, Its Subjective Sources", in Henry E. Kyburg and Howard E.K Smokler (eds.), Studies in Subjective Probability, Huntington, NY: Robert E. Kreiger Publishing Co.

[7] de Finetti, Bruno. 1972. *Probability, Induction and Statistics*. New York: Wiley.

[8] Diaconis, Persi and Sandy Zabell. 1982. "Updating Subjective Probability". *Journal of the American Statistical Association* 77.380: 822-830.

[9] Domotor, Zoltan. 1980. "Probability Kinematics and Representation of Belief Change". *Philosophy of Science*, 47: 384-403.

[10] Doring, Frank. 1999. "Why Bayesian Psychology is Incomplete". *Philosophy of Science* (Proceedings), 66: 379-S389.

[11] Ewing, A. C. 1934. *Idealism: A Critical Survey.* London: Methuen.

[12] Feldman, Richard. 2001. "Voluntary Belief and Epistemic Evaluation." In Matthias Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press 77–92.

[13] Fitelson, Branden. 2003. "A Probabilistic Measure of Coherence." *Analysis*, 63: 194-199.

[14] Goldman, Alvin. 1988. "Strong and Weak Justification", in *Philosophical Perspectives 2: Epistemology*, James Tomberlin (ed.), Atascadero, Calif.: Ridgeview Publishing Co., pp. 51-69.

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

[16] Huemer, Michael, 2007. "Compassionate phenomenal conservativism." *Philosophy and Phenomenological Research* 74 (1): 30-55.

[17] Jeffrey, Richard. 1965. *The Logic of Decision*. Chicago: Chicago University Press.

[18] Jeffrey, Richard. 1983. "Bayesianism with a human face", in John Earman (ed.), *Testing Scientific Theories*, 133-56. University of Minnesota Press.

[19] Jeffrey, Richard. 1975. "Carnap's Empiricism", in G. Maxwell and R.M. Anderson, Jr., (eds.), *Minnesota Studies in the Philosophy of Science*. Vol. 6, *Induction, Probability, and Confirmation*. Minneapolis: University of Minnesota Press, pp. 37-49.

[20] Joyce, James M. 1998. "A nonpragmatic vindication of probabilism." *Philosophy of Science* 65 (4):575-603.

[21] Joyce, James. 2009. "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In Franz Huber Christoph Schmidt-Petri (eds.)" *Degrees of Belief*. Synthese. pp. 263-297.

[22] Keynes, John Maynard. 1921. *A Treatise on Probability.* Dover Publications.

[23] Lange, Marc. 2000. "Is Jeffrey Conditioning Defective in Virtue of Being Non-Commutative." *Synthese*

[24] Laplace, P. S. "Essai Philosophique sure les Probabilites." Courcier, New York, 1814. English translation by FW Truscott and FL Emory.

[25] Leitgeb, Hannes and Richard Pettrigrew. 2010a. "An Objective Justification of Bayesianism I: Measuring Inaccuracy." *Philosophy of Science* 77 (2):201-235.

[26] Leitgeb, Hannes and Richard Pettigrew. 2010b. "An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy." *Philosophy of Science* 77 (2): 236-272.

[27] Lewis, C.I. 1946. *An Analysis of Knowledge and Valuation.* Open Court.

[28] Lewis, David. 1999. *Papers in Metaphysics and Epistemology: Volume 2*. Cambridge University Press.

[29] Meacham, Christopher J. G. forthcoming. "Can All-Accuracy Accounts Justify Evidential Norms?" In Kristoffer Ahlstrom-Vij Jeff Dunn (eds.), *Epistemic Consequentialism.* Oxford University Press.

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

[31] Plantinga, Alvin. 1991. "Warrant and Designing Agents: A Reply to James Taylor." *Philosophical Studies* 64 (2):203 - 215.

[32] Pryor, James. 2000. "The Skeptic and the Dogmatist," *Nous*, 34 (4): 517-549.

[33] Teller, Paul. 1976. "Conditionalization, Observation, and Change of Preference", in W. Harper and C.A. Hooker (eds.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, Dordrecht: D. Reidel.

[34] Schervish, Mark, Seidenfeld, Teddy and Kadane, Joseph. 2000. "How Sets of Coherent Probabilities May Serve as Models for Degrees of Incoherence." *International Journal of Uncertainty, Fuzziness and Knowlegde-based Systems* 8: 347-355.

[35] Schervish, Mark, Seidenfeld, Teddy Kadane, Joseph. 2002a. "Measuring Incoherence." Sankhya: The Indian Journal of Statistics 64: 561-587.

[36] Schervish, Mark, Seidenfeld, Teddy Kadane, Joseph. 2002b. "Measures of Incoherence: How not to Gamble if You Must," *Bayesian Statistics* 7, ed. Jose Bernardo, et. al., 385-401. Oxford: Oxford University Press.

[37] Skyrms, Brian. 1987. "Dynamic coherence and probability kinematics." *Philosophy of Science* 54 (1):1-20.

[38] Skyrms, Brian. 1997. "The Structure of Radical Probabilism." *Erkenntnis*, 45: 285-97.

[39] Staffel, Julia. 2015. "Measuring the Overall Incoherence of Credence Functions." *Synthese*.

[40] Teller, Paul. 1973. "Conditionalization and Observation." *Synthese*, 26: 218-258.

[41] Titelbaum, Michael. forthcoming. "Normative Modeling". *Methods in Analytic Philosophy: A Contemporary Reader.* J. Horvath ed. Bloomsbury Academic Press.

[42] Weisberg, Jonathan. 2009. "Commutativity or Holism? A Dilemma for Conditionalizers." *British Journal for the Philosophy of Science*, 60(4): 793-812.

[43] Weisberg, Jonathan. 2015. "Updating, Undermining and Independence." *British Journal for the Philosophy of Science*.

[44] Wolterstorff, Nicholas. 2010. *Practices of Belief.* Cambridge University Press.