

# Time-Slice Epistemology for Bayesians\*

**Abstract** Recently some have challenged the idea that there are genuine norms of diachronic rationality. Part of this challenge has involved offering replacements for diachronic principles. Skeptics about diachronic rationality believe that we can provide an error theory for it by appealing to synchronic updating rules that mimic the behavior of diachronic norms. In this paper, I argue that the most promising attempts to develop this position within the Bayesian framework are unsuccessful. I defend a new synchronic surrogate of Conditionalization that draws upon some of the features of each of these earlier attempts. At the heart of this discussion is the question of what exactly it means to say that one norm is a surrogate for another. I suggest that surrogacy, in the given context, can be taken as a proxy for the degree to which formal and traditional epistemology can be made compatible.

## 1 Introduction

For many, being epistemically rational requires having beliefs that display a certain sort of coherence over time. However, some have challenged the idea that there are genuine norms of diachronic rationality. Part of this challenge has involved offering replacements for diachronic principles. Skeptics about diachronic rationality believe that we can provide an error theory for it by appealing to synchronic updating rules that mimic the behavior of diachronic norms.

In this paper, I argue that the most promising attempts to develop this position within the Bayesian framework are unsuccessful. I argue that the synchronic updating rules that Meacham (2010) and Hedden (2015b,a)<sup>1</sup> propose as surrogates for the norm of Conditionalization are inadequate in different ways. I conclude by proposing a new synchronic surrogate of Conditionalization that draws upon some of the features of each of these earlier attempts.

This discussion will take for granted that the reasons offered by those advocating for synchronism—or what others have called “time-slice epistemology”—are decisive reasons. In other words, while I’ll consider some of the reasons we have for thinking that there are only synchronic norms, I

---

\* All acknowledgments of help have been temporarily removed in order to preserve blind review.

<sup>1</sup> For ease of exposition, in cases where this discussion refers to passages that appear in both Hedden (2015b) and Hedden (2015a), I will cite Hedden (2015a) only, since I assume this to be the later work. All passages cited in Hedden (2015b) refer to ideas that are also expressed, in different form, in Hedden (2015a).

won't consider the reasons on the other side. I won't consider the reasons we might have for wanting to hold onto our diachronic norms.<sup>2</sup>

This paper, then, diverges from the way that discussions about time-slice epistemology usually go. The question this paper will address is what a synchronic surrogate should look like assuming that we have reason to want one. More specifically, the question this paper will address is what a Bayesian synchronic surrogate should look like assuming that we have reason to want one. This discussion assumes there's an interesting question to be asked, not only about whether there exists a synchronic norm that makes reasonable recommendations, but about whether this synchronic norm can lay claim to being a genuinely Bayesian norm: a genuine surrogate for the diachronic norm of Conditionalization. I take seriously the title of this paper, then, in maintaining that this is a discussion about whether there can be a time-slice epistemology *for Bayesians*. Exactly what this means will become clearer as we move forward. But it's worth saying something now about why this is a question worth asking.

As we'll see in just a bit, the kinds of considerations that motivate the time-slice-centric movement belong to what many refer to as "traditional epistemology". Roughly speaking, traditional epistemology has to do with the sorts of questions about knowledge, justified belief and agency that epistemology has long been in the business of asking. There is a live dispute about the extent to which the answers to these question can be made compatible with formal epistemologies like Bayesianism. While many of these discussions center on the compatibility of the doxastic attitudes that feature in these epistemologies—those of credence and belief—we can also ask about the compatibility of other aspects of these systems. We can ask whether certain *norms* that traditional epistemology cares about can be made compatible with the norms that define Bayesian updating.

I want to suggest that the question of whether there is an adequate synchronic surrogate for Bayesianism's main updating rule is really the question of whether the Bayesian framework can accommodate the traditional epistemological concerns that motivate synchronism. It's the question of whether we can get a synchronic surrogate that is both consistent with the agential perspective, and continuous with the Bayesian framework. To the extent that surrogacy serves as this proxy for the extent to which we can bridge the gap between formal and traditional epistemology, this paper belongs to a more general discussion that most would agree is worth having. My hope is that this paper will provide a model for how we might extend this discussion in a different direction.

---

<sup>2</sup> For some recent arguments in favor of this position, see Carr (2015), Podgorski (2016) and Schwarz (ms.).

Here's how the paper will go. In §2, I describe some of the motivations for the synchronic view, as well as for the synchronic surrogates that Meacham and Hedden propose. In §3, I raise problems for each of these surrogates. I argue that Hedden's norm gives rise to *conflict cases*: cases where an agent's mental attitudes justify conflicting bodies of evidence. I then go on to argue that Meacham's norm is not a surrogate for Conditionalization in any interesting way. The general problem this section uncovers is the following: while Hedden's norm mimics Diachronic Conditionalization *too* closely by mimicking even its defects, Meacham's norm does not mimic Diachronic Conditionalization closely enough. In §4 and §5, I propose and defend a norm that resolves this "goldilocks problem". My norm strikes just the right balance of avoiding the problems faced by Diachronic Conditionalization, while retaining enough of its structure to be a genuine surrogate for it.

## 2 The Synchronic View

Standard Bayesianism assumes both a synchronic and a diachronic norm. Its synchronic norm tells us that, at each time, our credence function should be a probability function. Our degrees of belief should obey the probability axioms. Arguably, however, the source of Bayesian epistemology's power and its appeal is its diachronic constraint—its updating rule. Bayesian epistemology tells us that our beliefs should evolve over time in the following way:

**Diachronic Conditionalization:** Let  $p$  be the agent's prior probability function. If at time  $t'$  you get evidence  $E$ , your credence at  $t'$  in each proposition  $H$  should be  $p'(H)=p(H|E)$ .

In requiring an agent's posterior credence distribution to be a function of her earlier credence distribution—her priors—Diachronic Conditionalization entails that what attitudes you ought to have at a time directly depends upon what attitudes you have at other times. Diachronic Conditionalization entails that you ought to update sequentially from one time to another by conditioning the priors that resulted from your last update on your current evidence, where these priors were the result of having conditioned the priors you had before those on your previous evidence, and so on. What all skeptics about diachronic rationality have in common is that they deny that this is a requirement of rationality. More generally, what all skeptics about diachronic rationality have in common is that they deny that what attitudes you ought to have directly depend on your previous attitudes:

**Skepticism about Diachronic Coherence:** The question of what attitudes you ought to have at a time does not directly depend upon what attitudes you have at other times.

There are a couple of reasons for thinking that skepticism about diachronic coherence might be warranted. Some have argued that diachronic norms are incompatible with internalism: roughly, the view that the justification for our epistemic state is determined by factors that are internal to the agent. Here's Meacham (2010, p. 94-95)'s description of the problem:

...there's a deep tension between internalism and diachronic credence constraints, like conditionalization. Diachronic credence constraints place restrictions on what our current credences can be, relative to our credences at other times. But our credences at other times are external to our current state, in any of the senses relevant to internalism: they needn't supervene on what we currently have access to, our current mental or intrinsic states, and so on. So internalism and diachronic credence constraints are incompatible.

The attractiveness of internalism might, then, be one reason to give up on diachronic norms. Another reason derives from puzzle cases about the nature of personal identity, like those made famous by Derek Parfit (1971, 1984). Some have claimed that these cases motivate the view that the relation one bears to one's past or future self is similar to the relation one bears to other persons. Therefore, just as one should not feel bound by the commitments of some other person, one should not feel bound by the commitments of prior instances of oneself.<sup>3</sup> Again, this idea conflicts with the diachronic norm of Conditionalization, which crucially depends upon the personal identity relation. Hedden (2015b, p.456) illustrates this with the following example:

One person (call her 'Pre') enters the teletransporter in New York. Her body is scanned, and at the moment her body is vaporized, two different molecule-for-molecule duplicates of her are created, one in Los Angeles and the other in San Francisco. Call them 'Lefty' and 'Righty', respectively. Lefty and Righty are qualitatively just like Pre in all physical and mental respects. Now, there is a debate about whether Lefty, or Righty, or both, or neither is the same person as Pre. But what I want to emphasize is that in order to determine what Lefty and Righty ought to believe, following the double teletransportation, we do not have to first settle this debate about personal identity over time. If Lefty appears and immediately gains some new evidence, we do not first have to figure out the correct theory of personal identity in order to determine what Lefty ought to believe. All that matters is what Lefty's present evidence is. But Conditionalization conflicts with this datum. It only

<sup>3</sup> For an earlier discussion of this idea in the Bayesian context, see Christensen (1991, p.246).

says that Lefty's credences ought to be constrained by Pre's credences if Lefty is the same person as Pre; it is silent if Lefty and Pre are not the same person.

Parfit himself took the lesson of these sorts of cases to be that it is continuity, rather than personal identity, that is important. However, Hedden gives us several reasons to reject a continuity relation. First, he claims that, like personal identity, such a relation could not help but be arbitrary from a normative point of view. Second, he claims that such a relation would have to be one that comes in degrees. And it's difficult to see how a Bayesian agent could come in degrees.<sup>4</sup>

It should be clear that the appeal to internalism and the appeal to problems about personal identity, or continuity, are two ways of appealing to the same general consideration. The reason we hesitate to take seriously other temporal instances of ourselves is that they are bound by different commitments than our present self, if we take seriously a weak version of internalism and take our epistemic commitments to supervene on our current mental states. Internalism gives us positive reason for thinking that the present is important, while problems about continuity and the nature of personal identity give us reason for thinking that the present is important, in virtue of establishing the past to be unimportant. Together they deliver the idea that being rational is a matter of believing in ways that are sensible from your perspective.<sup>5</sup>

In light of this, both Meacham and Hedden have proposed replacing Diachronic Conditionalization with synchronic surrogates of this norm. These surrogates are norms that they take to be similar in spirit to regular Conditionalization, but which entail no commitment to diachronic rationality.

Take Meacham (2010)'s account first. His synchronic surrogate says that our current credences should be a function, not of our prior credences, but of our current beliefs *about* our prior credences. More specifically, Meacham's account requires our current credences to be the weighted average of what we believe our previous credence function recommends about how we ought to revise our beliefs, in light of our current evidence:

**Meacham's Synchronic Surrogate:** If the strongest evidence you get is the proposition  $E$ , then your new degree of belief in  $A$ , for any  $A$ , should be  $cr_E(A) = \sum_i cr_E(\langle cr = p_i \rangle) \cdot p_i(A|E)$ , if defined, where  $i$  ranges over the space of probability functions, and  $\langle cr = p_i \rangle$  is the proposition that our previous credence function was  $p_i$ .<sup>6</sup>

If we assume an internalist account of evidence, Meacham's updating rule entails that an agent's current credences should be a function only of her current mental states. Meacham's

<sup>4</sup> Hedden (2015a, p.33-35).

<sup>5</sup> Hedden (2015a, p.23).

<sup>6</sup> Meacham (2010, p.95).

amendment to the traditional Bayesian formalism leaves us with a norm that governs the agent at each time, rather than over time.

Hedden (2015a, p. 23) also assumes a weak internalist constraint on evidence. The only constraints on evidence that he imposes is that it supervene on our current mental states and that it be a proposition.<sup>7</sup> Both Hedden and Meacham assume roughly the same picture of evidence then. But their accounts differ when it comes to priors. Rather than appealing to our beliefs about our credence functions at earlier times, Hedden’s synchronic surrogate appeals to a uniquely rational prior function. His account assumes the following principle:

**Uniqueness:** Given a body of total evidence, there is a unique doxastic state that it is rational to be in.<sup>8</sup>

Instead of requiring that the agent condition whatever credence function she had after her last update on her current evidence, Hedden’s account requires that, at every moment an agent has some evidence, we use this evidence to update the credence function that we would hold if we were perfectly rational. On this picture of things, an agent ought to update her credences according to what a “magic probability function in the sky” would recommend.<sup>9</sup> Here, too, then, we get a norm that can be satisfied at every moment, given a weak internalist constraint on evidence:

**Hedden’s Synchronic Surrogate:** Let  $p$  be the uniquely rational prior probability function. If at time  $t$  you have total evidence  $E$ , your credence at  $t$  in each proposition  $H$  should equal  $p(H|E)$ .<sup>10</sup>

Call the view implied by both Meacham and Hedden’s accounts—that there are only synchronic norms for Bayesian updating—the *synchronic view*.<sup>11</sup>

One way of understanding how the synchronic view can get away with jettisoning a coherence constraint is to see that it trades it in for a stronger constraint elsewhere. Hedden’s account gives up diachronic coherence, at the cost of a uniquely rational prior probability function. Since uniqueness entails that we condition an *a priori* credence function on our evidence, conditioning need not make an agent’s credences depend upon her prior function at some earlier time. By contrast, Meacham’s account gives up diachronic coherence by handing the work that it does

<sup>7</sup> Hedden (2015a, p. 142).

<sup>8</sup> Hedden (2015a, p.130). For some recent arguments in favor of uniqueness, see White (2005), Feldman (2007) and Dogramaci and Horowitz (2016). For some recent arguments against uniqueness, see for instance, Kelly (2010), Kelly (2013b), Schoenfield (2014) and Meacham (2014).

<sup>9</sup> See Kelly (2013a) and Schoenfield (2014) for this description of a uniquely rational prior.

<sup>10</sup> Hedden (2015a, p.138). Following Hedden and others, I will continue to refer to the function that we condition on our evidence as a prior function, even in cases like Hedden’s where we are assuming that the function we condition our evidence on is not one that is temporally prior to our evidence.

<sup>11</sup> I say ‘implied’ here because only Hedden endorses the synchronic view. Meacham does not endorse this view, but introduces it as a way that Bayesianism can accommodate internalism by those who are so inclined.

over to the agent's higher-order beliefs about her prior probability function. Appealing to an agent's higher-order beliefs about her priors, again, entails that conditioning need not make an agent's credences depend upon her prior function at some earlier time. Both Hedden and Meacham trade off the prior function needed for diachronic coherence for a function that encodes stronger evidential constraints.<sup>12</sup>

This tradeoff structure helps to explain what Hedden and Meacham take to be an important feature of these synchronic norms: that they reduce to Diachronic Conditionalization in certain special cases. We've just said that what does all the work for Meacham's synchronic surrogate is the constraint that our current credences be a function of our current credences about our prior credence function. Meacham argues that this assumption entails that where we *know* what our prior credences were, his updating rule reduces to Diachronic Conditionalization.<sup>13</sup> We'll consider this reasoning more carefully in §3.2.

What about Hedden's synchronic surrogate? As Hedden notes, his updating rule entails Diachronic Conditionalization in the special case where we haven't lost any evidence because we remember all of our previous evidence:

To see this, suppose that at  $t_1$  you have total evidence  $E_1$  and at  $t_2$  you gain evidence  $E_2$ , so that your total evidence is now  $E_1 \wedge E_2$ . According to Synchronic Conditionalization, your  $t_1$  credences ought to be  $p_1(\cdot) = p(\cdot | E_1)$  while your  $t_2$  credences ought to be  $p_2(\cdot) = p(\cdot | E_1 \wedge E_2)$ . But  $p_2$  is the probability function that results from taking  $p_1$  and conditionalizing on  $E_2$ . So when your evidence grows monotonically from  $E_1$  to  $E_1 \wedge E_2$ , Synchronic Conditionalization yields the same recommendations as (diachronic) Conditionalization.<sup>14</sup>

Recall that Hedden takes our evidence to supervene on our current mental states. This means that when our current mental states include our current memories of past events, these past events will be part of our total evidence. In the special case where you haven't lost any evidence because you remember all of your previous evidence, then, your total evidence will include your previous evidence. This means that, in the special case where you haven't lost any evidence, your previous evidence will be reflected in your current update—just as it would have been had you updated by Diachronic Conditionalization. Therefore, an update by Hedden's Synchronic Surrogate mimics an update by Diachronic Conditionalization in the special case where you haven't lost any evidence.

<sup>12</sup> Cf. Hedden (2015a, p. 138).

<sup>13</sup> Meacham (2010, p. 95).

<sup>14</sup> Hedden (2015a, p. 139).

### 3 Some Problems with the Synchronic View

Both Hedden and Meacham defend synchronic surrogates of Diachronic Conditionalization. In this section, I'll raise problems for each of their accounts. I'll argue that Hedden's norm gives rise to cases where the agent has conflicting evidence. I'll then go on to suggest that Meacham's norm may not actually have Diachronic Conditionalization as a special case in any interesting way. In §4, I'll use these observations to motivate a more promising synchronic surrogate of Diachronic Conditionalization, one that is able to accommodate the traditional epistemological concerns of the synchronist, while preserving the formality of Bayesian updating.

#### 3.1 Hedden's Account

##### 3.1.1 *The Conflicts Problem*

Hedden's synchronic norm takes as input the agent's *total evidence*, which is all of the evidence she has at the present moment. As we've just seen, Hedden claims that where an agent gets evidence  $E_1$  at  $t_1$  and  $E_2$  at  $t_2$ , her credences at  $t_2$  ought to be  $p_2(\cdot) = p(\cdot | E_1 \wedge E_2)$ . Her evidence at  $t_2$  is  $E_1 \wedge E_2$ , provided that she has not lost any evidence. The idea that an agent is able to lose evidence is an important component of Hedden's view, insofar as it is part and parcel of the idea that an agent is not beholden to commitments she does not currently hold. But this is, of course, consistent with thinking that an agent will retain her earlier evidence much of the time. I want to argue that this phenomenon of retaining evidence, and so of conditioning on one's total evidence, poses a problem for the account. This way of understanding an agent's total evidence leads to cases where it looks as though the agent has *conflicting evidence*.

It's easy to see why this is. Consider an agent who learns  $E$  at  $t_1$  because this is the way that things seem to her then. But a moment later, at  $t_2$ , it seems to her that  $\neg E$ . If, at  $t_2$ , she remembers that she learnt  $E$  at  $t_1$ —she remembers *now* that it seemed to her that  $E$  at  $t_1$ —it now follows that she must have *both*  $E$  and  $\neg E$  as evidence at  $t_2$ . But, of course, an agent cannot have as evidence both some proposition and its negation, since an agent cannot update on both some proposition and its negation. Call this *the conflicts problem*.

##### 3.1.2 *Objections*

The conflicts problem might seem trivial or like it can be easily resolved. I think the following responses to some initial objections to it suggest that it is neither of these things.

*Objection.* One might object that the potential for the sorts of conflict cases I've just described will not be realized very often. If memory is required to retain evidence, as Hedden's account assumes, conflict cases will be few, if any. For there won't be many instances where I remember most of what I have learnt in the past. And if there aren't many instances where I remember most of what I learnt in the past, then there won't be many instances where it now seems to me that  $\neg E$ , but I now remember that it seemed to me that  $E$  earlier.

*Reply.* Hedden (2015a, p.145-146) is quick to point out that 'memory' need not mean 'occurrent memory'. Memory can simply be a mental state, hanging out in the background. It's crucial that Hedden be able to conceive of memory in this way. As we've already seen, while Diachronic Conditionalization is a norm of coherence, Hedden's synchronic norm is, for all intents and purposes, an evidential norm: it is a norm that makes what it is rational for an agent to believe a matter only of the agent's evidence and what this evidence objectively supports.

Given this important difference between Hedden's norm and orthodox Diachronic Conditionalization, it's important that Hedden be able to establish that Diachronic Conditionalization is a special case of his updating norm, if his updating norm is to have any claim to being a Bayesian norm at all. The fact that his updating norm allows the agent to update on evidence that she remembers from some earlier time is precisely what enables this by ensuring that there will be at least some cases where an update by Hedden's synchronic surrogate mimics an update by Diachronic Conditionalization.

There is a tension, then, between the need to avoid conflicts and the need to mimic Diachronic Conditionalization. To the extent that the requirements on memory are weak enough to secure cases where an agent will retain evidence over time, and so mimic the verdicts of Diachronic Conditionalization, they will also make it more likely that an agent will retain evidence that contradicts the way that things seem to her now. To the extent that the requirements on memory are strong enough to make it less likely that this will happen, these more stringent requirements will also tend to preclude situations where the verdicts of Hedden's norm coincide with those of Diachronic Conditionalization because the agent retains the evidence she has gained at some earlier time.

*Objection.* One might object that we can avoid conflict cases, while allowing that Hedden's norm might mimic Diachronic Conditionalization some of the time, by requiring that we discount our current memory *only* in cases where a conflict arises. This wouldn't undermine our ability to retain evidence more generally.

*Reply.* But this seems like an *ad hoc* restriction. There isn't any principled ground upon which to argue that memory should be entirely discounted when, and only when, there is a conflict. Of course, the existence of a conflict suggests that a mistake has been made. And there should be some mechanism for correcting this mistake. In §4, I defend an account that includes such a mechanism. For now, let me just point out that it's arbitrary to think, without any further information, that the mistake is in our first piece of evidence rather than in the evidence that supervenes upon our current seeming. It's arbitrary to discount our current memory rather than our current seeming, without any reason.

*Objection.* We've assumed with Hedden that the only constraint on our evidence is that it supervene on our current mental states. Plausibly, however, we could strengthen our account of evidence in a way that blocks conflicts. One strategy along these lines would be to take our seemings at face value. Instead of assuming that we get as evidence the propositions  $E$  and  $\neg E$ , one might argue that what we really get as evidence, in cases like the one described above, are propositions of the form *it seems to me at  $t$  that  $E$*  and *it seems to me at  $t'$  that  $\neg E$* . No conflict there.

Even if the previous suggestion seems a little arbitrary, there is a more plausible account of evidence in the neighborhood. One might argue that all one needs to block conflicts is to adopt a factive account of evidence. A factive account of evidence would, of course, make it impossible to get as evidence both some proposition and its negation, since both could not possibly be true.

*Reply.* While this solution seems promising, it resolves conflicts in the wrong way. We said earlier that one of the motivations for the synchronic view is to make the Bayesian framework compatible with internalist accounts of evidence. There are two ways of understanding why we might want this. On the one hand, we might want to show that the Bayesian framework is compatible with evidential internalism because we think this is the right way to think about evidence. On the other hand, we might want to show that the Bayesian framework is compatible with evidential internalism because we want the Bayesian framework to be neutral between competing, substantive accounts of evidence. We want the Bayesian framework, which has previously only been able to accommodate evidential externalism, to be able to accommodate *both* evidential externalism *and* evidential internalism.

It's clear that it's the second motivation that drives both Hedden's and Meacham's accounts. Meacham does not endorse the synchronic view. He notes that the standard Bayesian framework assumes that we have externalist intuitions about evidence, and notes that problems arise for cases that draw upon internalist intuitions about evidence. Meacham (2010, p.95) offers his

synchronic surrogate as a way of accommodating the latter sorts of intuitions “without taking a stand on which of these options the Bayesian should adopt”. While Hedden (2015a, p.22) does endorse the synchronic view, he argues that it is “strictly speaking, compatible with both externalism and internalism”. Hedden argues that the very weak internalist constraint on evidence that he defends is compatible with many accounts of evidence that are regarded as canonically externalist.<sup>15</sup> Moreover he claims that even those externalists who reject his weak internalist constraint can still endorse the synchronic view “by holding that facts about your past are not among the external factors that affect how you ought to be now”.<sup>16</sup> While internalists have special reason to endorse the synchronic view, in light of the seriousness with which it takes the agent’s perspective, Hedden thinks that externalists still have reason to adopt the synchronic view, on the basis of puzzle cases for personal identity over time.<sup>17</sup> More generally, Hedden takes it to be an advantage of his account that his synchronic surrogate “is compatible with any account of evidence whatsoever” and that “whatever your favored account of evidence, it can most likely be plugged in to the framework provided by Time-Slice Rationality”.<sup>18</sup>

Since both Hedden’s and Meacham’s synchronic surrogates are motivated in part by the desire to make the Bayesian framework neutral between competing accounts of evidence, a solution to the conflicts problem that compromises this evidence neutrality by requiring that evidence be factive undermines a good part of the reason for appealing to these surrogates in the first place. While a factivity condition on evidence is compatible with both evidential internalism and evidential externalism, it still renders the Bayesian framework dependent upon a substantive account of evidence. Same goes, of course, for an account that precludes our getting as evidence anything but propositions about the way that things *seem* to us at particular times.

To be clear: I’m sure many philosophers would be happy with a factive account of evidence. But a formal epistemology should, by definition, be neutral between substantive accounts of evidence. Both Hedden and Meacham take this thought seriously, and so should we. To have evidence do all the heavy lifting in this case where we have already conceded a diminished role for our priors is, in effect, to get a solution to the conflicts problem by abandoning Bayesianism entirely. I want to suggest that, to the extent that we do need to appeal to an account of evidence to resolve the conflicts problem, this account should not be *ad hoc*, but should be both continuous with the

---

<sup>15</sup> For instance, Hedden notes that since Williamson (2000)’s E=K account of evidence takes knowledge to be a mental state, it is compatible with the idea that an agent’s evidence directly depends upon her current mental states, and only *indirectly* depends upon facts about the past that determine whether some mental state is knowledge.

<sup>16</sup> Hedden (2015a, p.22).

<sup>17</sup> Hedden (2015a, p.22, 28).

<sup>18</sup> See Hedden (2015a, p.142) and Hedden (2015a, p.147), respectively.

Bayesian framework and consistent with the very concerns that led us to synchronism in the first place. I'll return to consider these issues again in §5.

*Objection.* Finally, one might concede that the conflicts problem is indeed a genuine problem, while still insisting that it isn't a problem for Hedden's norm, qua synchronic surrogate. We noticed a moment ago that the conflicts problem arises in virtue of the very feature that allows Hedden's updating rule to have Diachronic Conditionalization as a special case. Given this, one might argue that conflict cases are the result of Hedden's norm mimicking Diachronic Conditionalization perfectly. It's well-known that Diachronic Conditionalization also cannot accommodate updates on inconsistent evidence. On the diachronic framework, one cannot condition on a proposition that conflicts with some evidence one already has. Of course, unlike Hedden's norm, the diachronic framework only lets us condition on the piece of evidence that we've gotten first, since it makes an update on the second piece of evidence undefined. Nevertheless, one might insist that the underlying problem for these frameworks is the same. Neither can accommodate two conflicting pieces of evidence. Given that the conflicts problem is a problem shared by both Hedden's norm and Diachronic Conditionalization, one might argue that Hedden's updating norm, qua synchronic surrogate, gets things exactly right. It perfectly mimics even the defects of Diachronic Conditionalization!

*Reply.* I think this objection reveals something interesting about what we should expect a synchronic surrogate to do. There seems to be a sort of "goldilocks principle" that guides the question of whether some norm is an adequate surrogate for another. We want our surrogate to mimic our target norm. But we don't want our surrogate to mimic our target norm *too* closely, lest it fail to accomplish the task for which it was created. We want our surrogate norm to mimic our target norm just enough, in just the right way.

I assume that Hedden's aim is to provide us with a norm that mimics Diachronic Conditionalization closely enough to be considered its surrogate, while doing away with certain undesirable features. Hedden's norm does away with updates that are undermined by puzzle cases about personal identity. It does away with updates that are in tension with internalism. In a similar way, and for similar reasons, we should expect a synchronic surrogate to do away with updates that are undermined by inconsistent evidence.

To see this more clearly, we need only recognize that the sorts of problem cases that underwrite the appeal to internalism and skepticism about personal identity, and which both Hedden and Meacham use to motivate the synchronic view, are cases that feature two inconsistent perspectives that the synchronic view is supposed to be able to reconcile. In Arntzenius (2003, p.356)'s Shangri-

La case, which both Hedden and Meacham use to motivate their accounts, our intuitions about what credences we should have that a certain coin has landed heads are different at different times, due to the way things seem to us at these different times. While Diachronic Conditionalization entails that our previous evidence, in effect, swamps our current seemings, in cases where there is a conflict, Hedden's synchronic surrogate yields the opposite result. Desirably, Hedden's synchronic surrogate yields the result that our current seemings swamp our previous evidence in cases of memory loss. However what we are now in a position to see is that this is only a partial solution to the problem of inconsistent perspectives. While Hedden's synchronic surrogate resolves the problem of inconsistent perspectives, in cases of memory loss, by identifying rationality with what is rational from one's current perspective, it breaks down in cases where our current perspective includes memories of past events. It breaks down in cases where our current perspective *includes* our past perspective.

We said earlier that the appeal to internalism and the appeal to problems about personal identity are two ways of appealing to the same general idea: to the idea that being rational is a matter of believing in ways that are sensible from your perspective. It's an appeal to the agential perspective. The appeal to a synchronic surrogate was supposed to capture this. Insofar as the conflicts problem arises from the lack of a consistent perspective, it is an instance of *exactly* the problem that we had expected the synchronic view to have resolved. The conflicts problem should, then, leave us open to a better synchronic surrogate than the one that Hedden defends.

Some might want to insist that, regardless of the issue of surrogacy, we should want our updating rule to crash if this crash is due to the agent trying to update on inconsistent information. However I think that an updating rule that can avoid crashing by incorporating *even more* of an agent's evidence into her update should be preferred to one that does not. Such a rule is capable of guiding the agent in a greater number of epistemic situations. In §4, I propose an updating rule that does just this.

### 3.2 Meacham's Account

Should we reject Hedden's norm in favor of Meacham's? Recall that Meacham's norm takes as input, not our total evidence, but our most recently gathered evidence. And it tells us that we should use this evidence to update, not a uniquely rational prior function, or the prior function that we had after our last update, but the prior function we *believe* ourselves to have had after our last update, weighted by how strongly we believe ourselves to have had it.

But Meacham's account is not without its own problems. Notice first that Meacham's updating norm is silent in any case where the agent doesn't have an opinion about her priors. And cases like these seem ubiquitous. We often don't have higher-order opinions about our past beliefs.

There's a much more serious worry for Meacham's account. Like Hedden's norm, Meacham's norm seems on its face to be quite different than Diachronic Conditionalization. We typically don't think of Bayesian updating as a relation between an agent's first-order beliefs and her higher-order beliefs, as Meacham's account would have us do. But, like Hedden, Meacham is able to claim a connection with Diachronic Conditionalization by arguing that Diachronic Conditionalization is a special case of his updating rule. Generally speaking, one thing,  $A$ , is a special case of another thing,  $B$ , iff every instance of  $A$  is also an instance of  $B$ , but not vice versa. For updates on priors that we *know* to be our priors to be a special case of Meacham's norm, then, these updates would need to be a proper subset of updates by Meacham's norm. And they are: updates on priors that one knows to be one's priors are a special case of updates on priors that one believes to be one's priors, since knowledge entails belief, but not vice versa. Moreover, since knowledge also entails factivity, updates on priors that we know to be our priors are updates by Diachronic Conditionalization.

Despite this, I want to suggest that Meacham's norm may not be an adequate surrogate for Diachronic Conditionalization. Meacham assumes that the fact that one norm is a special case of another is sufficient for the latter norm to be a surrogate for the former. But one might reasonably ask whether Meacham's norm shares enough in common with the special case of Diachronic Conditionalization he has in mind to be an adequate surrogate for it. Meacham's norm tells us to update on functions that we *believe* to be our priors. The version of Diachronic Conditionalization that Meacham has in mind tells us to update on functions iff we have (something like) *true justified beliefs* that they are our priors. These are two qualitatively different updating rules. It's trivial that we can turn any norm into a special case of another if we add to it, or delete from it, enough qualitative properties.

I think the reason that Meacham's assumption that being a special case of a norm is sufficient for being its surrogate seems so plausible is that, more often than not, when we think of one norm as a special case of another, we have in mind that it is a *limiting*, special case of that norm. Generally speaking, one thing,  $A$ , is a limiting, special case of another thing,  $B$ , if  $A$  is a special case of  $B$ , in virtue of taking the most extreme possible value that  $B$  allows. Correspondingly, for some update by a norm to be a limiting, special case of another update by a norm, the former update must be a special case of the latter, in virtue of taking an extreme quantitative input.

Where one norm is a limiting, special case of another, then, there is no qualitative difference between them; the only difference is one that can be expressed as a change in value.<sup>19</sup> Therefore, where one norm is a limiting, special case of another, it seems reasonable to think that they are continuous in a way that makes the latter a surrogate for the former.

The most well-known example of a limiting, special case, in the Bayesian literature, is the way that Diachronic Conditionalization is a limiting, special case of Jeffrey Conditionalization.<sup>20</sup> Diachronic Conditionalization is the case of Jeffrey Conditionalizing on evidence that we hold with a credence of one. By contrast, Diachronic Conditionalization is *not* a limiting, special case of Meacham's norm. For Diachronic Conditionalization to be a limiting special case of Meacham's norm, it would need to be an update by Meacham's norm, in virtue of taking the most extreme possible value of an update by Meacham's norm. It would need to be an update on a function that we believe with an extreme degree of certainty—plausibly, a credence of one—is our prior function. But an update on a function that we believe with an extreme degree of certainty is our prior function is *not* an update by Diachronic Conditionalization. The mere fact that an agent is certain that some prior function is her prior function does not entail that it *actually is* her prior function. Unlike knowledge, certainty does not entail factivity. The class of updates that are limiting cases of updates by Meacham's norm are not updates by Diachronic Conditionalization.

Here's the dilemma then. Either the updates by Meacham's norm that are supposed to be instances of the special case of Diachronic Conditionalization appeal to functions that we are *certain* are our priors, in which case they are not actually updates by Diachronic Conditionalization. Or else the updates by Meacham's norm that are supposed to be instances of the special case of Diachronic Conditionalization appeal to functions that we *know* are our priors, in which case they are not a *limiting* special case of Meacham's norm, since they are qualitatively different updates. In the latter case, it's not clear that these updates by Diachronic Conditionalization are continuous with Meacham's norm in a strong enough sense for Meacham's norm to be a surrogate for Diachronic Conditionalization.

We can take a step back and identify the general feature of Meacham's norm that gives rise to this problem. Meacham's norm has the consequence that the way we revise our beliefs need not be grounded in our actual priors at all, but in some subset of our current beliefs. But what unifies Bayesians—epistemologists, scientists and statisticians alike—is their commitment to either subjectively or objectively rational priors. As I suggested in the beginning of this discussion,

<sup>19</sup> One might point to degenerate, limiting special cases as examples of limiting cases that do issue in qualitative differences. A line segment, for instance, is a degenerate case of a triangle. But the point is that, even in these cases, the qualitative difference can be represented as a quantitative difference: as a change in the value of the parameter in question.

<sup>20</sup> For the canonical description of Jeffrey Conditionalization, see Jeffrey (1965).

I think the question of whether there can be a time-slice epistemology for Bayesians is really the question of whether Bayesianism can be preserved while accommodating the values of traditional epistemology. It's the question of whether formal epistemology and traditional epistemology can be made compatible. I'll have more to say about this in §5. For now, I'll simply note that the problem I've identified for Meacham's account is a subtle problem rather than a small one since it tracks the more fundamental problem this paper is concerned with resolving.

One might insist that being a non-limiting, special case of some norm actually is sufficient for surrogacy and that, at most, what I have shown is that a rule that has Diachronic Conditionalization as a limiting case will be *more* of a surrogate than Meacham's norm. At the end of the day, I would be very happy with such a result. For I think that even this weaker conclusion should leave us open to a better alternative.

#### 4 A Proposal: A Higher-Order Synchronic Surrogate

I've suggested that the question of whether some norm is an adequate surrogate for Diachronic Conditionalization is guided by a goldilocks principle. I've argued that both Hedden and Meacham's norms fail by the lights of this principle. While Hedden's norm mimics Diachronic Conditionalization too closely by mimicking even its defects, Meacham's norm does not mimic Diachronic Conditionalization closely enough. However, though these accounts are problematic in these ways, when taken on their own, I think a combination of their features gets us what we are after: a synchronic norm that avoids the problems of Diachronic Conditionalization while retaining enough of its structure to be a genuine surrogate for it. The key will be to reconceive the role played by higher-order beliefs.

Higher-order beliefs are sometimes appealed to in order to provide a constraint on an agent's priors.<sup>21</sup> We get an extreme version of this on Meacham's account where higher-order beliefs constrain an agent's priors by replacing them entirely. But, plausibly, higher-order beliefs can also be appealed to in order to constrain, not an agent's priors, but her *evidence*. Consider, again, the case where I have some mental states at  $t$  that yield as evidence,  $E$ , and some mental states at  $t'$  that yield as evidence,  $\neg E$ . As before, if I remember my evidence  $E$  at  $t'$ , then what we have is a conflict case. Since memory has a higher-order feel to it, a natural thing to do here, it would seem, is to represent the normative role that memory plays in the higher-order beliefs an agent has about how weighty she *now* takes her remembered evidence to be.

---

<sup>21</sup> See Lewis (1980)'s Principal Principle and van Fraassen (1984)'s Rational Reflection for two canonical examples of this.

There are a number of reasons we might judge some piece of evidence to be weightier than another. For our purposes, it does not matter what that reason is. As we will see in §5, it is the structure of the proposal that will allow it to succeed where Hedden and Meacham have failed. For the sake of concreteness, however, let us assume that an agent judges some piece of evidence to be weightier than some other piece of evidence to the extent that she is confident that this evidence is *more reliable*.

While there are a number of ways of spelling out what it might mean to judge that some piece of evidence is more reliable than some other piece of evidence, perhaps the simplest is to say that I judge that some piece of evidence is more reliable than another—more *relatively reliable*—if I am **now** more confident that the process producing the seeming upon which the former piece of evidence **initially supervened** is more reliable than the process producing the seeming upon which the latter piece of evidence initially supervened. Thus the normative constraint that governs each piece of evidence I have at the present moment will look like the following:

**HO-Constraint:** Let  $E$  be a proposition that supervenes upon the agent's current mental states, and let  $p_n(R_E)$  represent how confident the agent is, at the present moment,  $t_n$ , that  $E$  was initially formed by a process that was more reliable than the process that underwrites the initial formation of any conflicting evidence that also supervenes upon the agent's current mental states. Then  $p_R(E)$ , which represents the credence the agent ought to have in  $E$  at  $t_n$ , will be determined in the following way:

$$p_R(E) = p_n(R_E)$$

To see the HO-Constraint in action, suppose an agent now believes that her seeming at  $t_{n-1}$  that the bird is blue ( $B$ ) was produced by certain cognitive processes that were more reliable than the cognitive processes that produced her current seeming now, at  $t_n$ , that  $\bar{B}$ . Indeed, suppose that the agent is *three times* more confident that the former processes are more reliable than the current ones. Then the HO-Constraint will yield the following two values for these conflicting propositions that supervene upon her current mental states:

$$p_R(B) = .75$$

$$p_R(\bar{B}) = .25$$

Together, these values contribute to the weighted evidence partition that represents the agent's total evidence at  $t_n$ .<sup>22</sup> Here's what an updating rule that draws upon these observations ends up looking like:

**A Higher-Order Synchronic Surrogate:** Where  $p$  is the uniquely rational prior probability function, and where  $p_R$  is a probability distribution over the partition,  $\{E_i\}$ , which represents the agent's total evidence at the present moment, as determined by the HO-Constraint, the agent's degree of belief in A, for any A, should be  $\sum_i p(A|E_i)p_R(E_i)$ .

In tracking how relatively reliable we currently take our past seemings to be, our higher-order beliefs tell us how much normative weight we should assign to the evidence that we now remember. This is exactly the role we would want our memory to play.

One might worry that certain cases will cause trouble for my updating rule. For instance, one might wonder about cases where the agent doesn't have any opinion about the relative reliability of her mental states at other times, even though she remembers those other times. While there are a couple of ways one might go here, one option is to say that cases where the agent currently has no opinion about which of her mental states were most reliably formed are best represented as cases where the HO-Constraint delivers a uniform distribution over the partition,  $\{E_i\}$ , which represents an attitude of suspended judgment, or indifference, about the relative reliability of this evidence. Plausibly to maintain no information about the relative reliability of one's evidence is, in effect, to judge that all of one's evidence is equally likely to have been the most reliably formed.

One might also worry about cases where the agent does have higher-order evidence about how reliable she was at other times, but *doesn't* have any evidence that conflicts with the way things seem to her now, either because she never got any conflicting evidence, or because she forgot some of what she previously learned. However, cases where the agent has no conflicting evidence at the time that she updates, either because she hasn't gotten any or because she has forgotten some of it, are cases where her higher-order beliefs simply weight with a credence of one (i.e.,  $p_n(R_E) = 1$ ) a unique proposition representing how things seem to her at the present time that is, trivially, judged to be the most relatively reliable.

It is at least partly in virtue of being able to accommodate these situations that **A Higher-Order Synchronic Surrogate** overcomes the problems we've identified with Hedden's and Meacham's accounts:

---

<sup>22</sup> It's easy to see how this will generalize to cover more complicated cases. Say that the HO-Constraint yields the following values:  $p(E)=.75$ ,  $p(\bar{E})=.25$ ,  $p(F)=.5$ ,  $p(\bar{F})=.5$ . Then the weighted partition generated by the HO-Constraint will be  $\{EF_i\}$ : i.e.,  $p(EF)=.3$ ,  $p(\bar{E}F)=.2$ ,  $p(E\bar{F})=.45$ ,  $p(\bar{E}\bar{F})=.05$ .

*No Conflicts.* Unlike Hedden's account, **A Higher-Order Synchronic Surrogate** will not yield conflicts. By appealing to our higher-order evidence to resolve disagreement in our first-order evidence, we recover the consistency of our evidence without arbitrarily throwing any of it away. We've said that the synchronic view captures the idea that being rational is a matter of believing in ways that are sensible from your perspective. **A Higher-Order Synchronic Surrogate** doubles down on this idea by making rationality a matter of believing in ways that are sensible from your higher-order perspective.

*A Limiting, Special Case.* Unlike Meacham's account, **A Higher-Order Synchronic Surrogate** has Diachronic Conditionalization as a limiting, special case. An update by Diachronic Conditionalization is the limiting, special case of an update by **A Higher-Order Synchronic Surrogate** where there is no lost or conflicting evidence. Our higher-order beliefs assign a credence of one to a unique proposition that is, trivially, judged to be the most relatively reliable.

*Does Not Fall Silent.* Finally, unlike Meacham's account, **A Higher-Order Synchronic Surrogate** is not vulnerable to the worry that it will fall silent any time an agent lacks higher-order beliefs. If the agent has no higher-order beliefs, the account entails a fallback constraint: it entails that our first-order evidence be equally weighted. Meacham's account does not entail this sort of fallback constraint; nor is it compatible with such a constraint. In order to deal with cases where an agent lacks higher-order beliefs, in a way that is analogous to my solution, such a constraint would need to appeal to the *object* of these beliefs, in order to construct a uniform distribution over these objects. In the case of Meacham's updating rule, this would be an appeal, not to an agent's actual evidence, but to an agent's *actual priors*. Of course, a constraint that appealed to an agent's actual priors would not be permitted by the synchronic view.

## 5 Surrogacy as a Bridge Between Formal and Traditional Epistemology

I began this discussion by suggesting that the question of whether some synchronic norm is an adequate surrogate for Diachronic Conditionalization tracks the question of whether formal and traditional epistemology can be made compatible. We are now in a better position to appreciate this claim. The norms we've considered and rejected either reproduce problems similar to those they were intended to resolve, or else diverge in subtle but significant ways from the norm they were intended to replicate. In the same way that certain attempts to reduce rational belief to rational credence yield unacceptable results like the Lottery Paradox (*cf.* Kyburg (1961)), attempts to reduce the synchronic, agential perspective characteristic of traditional epistemology to the

updating constraint characteristic of Bayesian epistemology also yield unacceptable results in the form of the goldilocks problem.

With this in mind, I want to address a final concern. I've suggested throughout that we ought to reject any appeal to a substantive account of evidence. I've argued that the Bayesian framework can do this if it is supplemented by the HO-Constraint. But one might worry that this is precisely because the HO-Constraint *is* a substantive account of evidence. Even if it isn't, one might ask why the HO-Constraint ought to be preferred to such accounts.

Our observations so far entail a picture of what the ideal synchronic surrogate—that which avoids the goldilocks problem—will look like. The ideal surrogate will be *consistent* with the agential perspective upheld by traditional epistemology, and *continuous* with the structure of the Bayesian framework. Hedden's rule fails by the lights of this ideal in virtue of failing to be consistent, while Meacham's rule fails to be continuous. It should be clear why the HO-Constraint more closely approximates this ideal than any substantive account of evidence ever could. Unlike accounts that take evidence to be factive, etc., my higher-order constraint is a coherence constraint, one that has features in common with several other Bayesian constraints. Like Jeffrey Conditionalization, it assumes a more complicated picture of the structure of evidence than the orthodox Bayesian picture. It assumes that our total evidence is a partition of propositions. Like Lewis (1980)'s Principal Principle and van Fraassen (1984)'s Rational Reflection, it makes our rational credences a matter of deferring to an expert. The mandate of the HO-Constraint is a hybrid of these commitments. It aims to regulate conflicts in an agent's evidence by means of her judgment about her own expertise.

Like other paradigmatically Bayesian constraints, then, **A Higher-Order Synchronic Surrogate** posits a systematic, or formal, relation between the content of our evidence and the credences we assign our beliefs. Therefore, it is clearly continuous with the Bayesian framework. It is also consistent with the perspective and values of traditional epistemology. Unlike a factive account of evidence, it both captures and deepens the fundamental insight behind synchronism, which is that rationality is a matter of believing in ways that are sensible from the agent's perspective. It does this by invoking the agent's higher-order perspective. **A Higher-Order Synchronic Surrogate's** formal account of evidence encodes the very values that elude the orthodox Bayesian. This is exactly the sort of behavior we should expect of the ideal synchronic surrogate, one that mimics Diachronic Conditionalization in just the right way.

## References

Arntzenius F (2003) Some problems for conditionalization and reflection. *The Journal of Philosophy* 100(7):356–370

- Carr JR (2015) Don't stop believing. *Canadian Journal of Philosophy* 45(5-6):744–766
- Christensen D (1991) Clever bookies and coherent beliefs. *Philosophical Review* 100(2):229–247
- Dogramaci S, Horowitz S (2016) An argument for uniqueness about evidential support. *Philosophical Issues* 26(1):130–147
- Feldman R (2007) Reasonable religious disagreement. *Philosophers without gods: Meditations on atheism and the secular life*
- van Fraassen B (1984) Belief and the will. *The Journal of Philosophy* 81(5):235–256
- Hedden B (2015a) *Reasons without persons: rationality, identity, and time*. OUP Oxford
- Hedden B (2015b) Time-slice rationality. *Mind* 124(494):449–491
- Jeffrey R (1965) *The logic of decision*. Tech. rep., MCGRAW-HILL BOOK CO INC NEW YORK
- Kelly T (2010) Peer disagreement and higher order evidence. *Social Epistemology: Essential Readings*
- Kelly T (2013a) Evidence can be permissive p 298
- Kelly T (2013b) How to be an epistemic permissivist. *Contemporary debates in epistemology* pp 298–311
- Kyburg HE (1961) Probability and the logic of rational belief
- Lewis D (1980) A subjectivist's guide to objective chance. In: *Ips*, Springer, pp 267–297
- Meacham CJ (2010) Unravelling the tangled web: Continuity, internalism, uniqueness and self-locating belief. *Oxford studies in epistemology* 3:86
- Meacham CJ (2014) Impermissive bayesianism. *Erkenntnis* 79(6):1185–1217
- Parfit D (1971) Personal identity. *The Philosophical Review* 80(1):3–27
- Parfit D (1984) *Reasons and persons*. OUP Oxford
- Podgorski A (2016) A reply to the synchronist. *Mind* 125(499):859–871
- Schoenfield M (2014) Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs* 48(2):193–218
- Schwarz W (ms.) Evidentialism and conservatism in bayesian epistemology. Unpublished manuscript
- White R (2005) Epistemic permissiveness. *Philosophical perspectives* 19(1):445–459
- Williamson T (2000) *Knowledge and its Limits*. Oxford University Press