Noname manuscript No.
(will be inserted by the editor)

Time-Slice Epistemology for Bayesians

(forthcoming in *Inquiry*, please cite published version)

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, over time, 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 sketch a new synchronic surrogate 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 argue 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 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, over time, mimic the behavior of diachronic norms.

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

This paper diverges from the way discussions about this topic usually go by taking 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 the

[1] For ease of exposition, in cases where this discussion refers to passages that appear in both Hedden (2015a) and Hedden (2015b), I will cite Hedden (2015b), since I assume this to be the later work.

reasons we have for thinking that there are only synchronic norms, I 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. Instead, the question this paper addresses 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 claiming this to be 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". Very roughly, 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 attitudes that feature in these epistemologies—those of credence and belief—we can also ask about the compatibility of other features of these frameworks. We can ask whether the norms and values that traditional epistemology cares about can be made compatible with the norms and values of Bayesian updating.

I want to suggest that the question of whether there is an adequate synchronic surrogate for Conditionalization is really a version of the previous question. It's the question of whether there is a formal updating rule that is able to accommodate the traditional epistemological concerns that motivate synchronism, with its focus on the agential perspective. Insofar as 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 about the extent to which formal epistemology succeeds at being epistemology at all.[2] To be clear at the outset, then, this is a paper in metaepistemology. The concerns raised and addressed in this paper ought to be interpreted as primarily metaepistemological concerns.

[2] For two classic discussions of this topic, see Harman (1986) and Christensen (2004).

Here's how the discussion will go. In §2, I describe some of the motivations for the synchronic view, as well as for the particular 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 sense. 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. This problem reveals something interesting about what we should expect a synchronic surrogate to do—there's a "goldilocks principle" that guides the question of whether some norm is an adequate surrogate for another. The ideal synchronic surrogate is one that mimics its target norm *just* enough, in just the right way. In §4 and §5, I defend a norm that does exactly this. My norm strikes the right balance of avoiding the problems faced by Diachronic Conditionalization, while retaining enough of its structure to be its surrogate.

2 The Synchronic View

Standard Bayesianism assumes both a synchronic and a diachronic norm. Its synchronic norm tells us that, at each time, our degrees of belief should obey the probability axioms. Our credence function should be a probability function. 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 your probability function at t. If E represents everything you learn in between t and t', then your credence at t' in each proposition H should be $p'(H)=p(H|E)$, if defined.

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 says that you should 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 upon your previous attitudes in this sort of way:

> **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: 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.[3] Again, this idea conflicts with the diachronic norm of Conditionalization, which crucially depends upon the personal identity relation. Hedden (2015a, p.456) describes this conflict 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

[3] For an earlier discussion of this idea in the Bayesian context, see Christensen (1991, p.246).

> 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 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.[4]

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.[5]

It's in light of these considerations that 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

[4] Hedden (2015b, p.33-35).
[5] Hedden (2015b, p.23).

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.[6]

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

Hedden (2015b, 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.[7] 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 probability 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.[8]

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, she use this evidence to update the credence function that she would hold if she were perfectly rational. 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 $p(H|E)$.[9]

[6] Meacham (2010, p.95).

[7] Hedden (2015b, p. 142).

[8] Hedden (2015b, 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 Kelly (2010), Titelbaum (2010), Kelly (2013), Schoenfield (2014) and Meacham (2014) and Titelbaum and Kopec (2016).

[9] Hedden (2015b, p.138). Following Hedden and convention, I will continue to refer to the function that we condition on our evidence as a prior function, even in cases where it is not temporally prior to our evidence.

Call the view implied by both Meacham and Hedden's accounts—that there are only synchronic norms for Bayesian updating—*the synchronic view*.[10]

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, *a priori* probability function. Meacham's account gives up diachronic coherence by handing the work that it does over to the agent's higher-order beliefs about her prior probability function. Both Hedden and Meacham trade off the prior function needed for diachronic coherence for a function that encodes stronger evidential constraints.

This tradeoff structure partly helps to explain what Hedden and Meacham take to be an important feature of these synchronic norms, which is 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.[11] 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.[12]

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

[10] I say 'implied' here because only Hedden endorses the synchronic view. Meacham does not endorse this view, but introduces it as a way that Bayesians that are so inclined can accommodate internalism.

[11] Meacham (2010, p. 95).

[12] Hedden (2015b, p. 139).

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.

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 sense. These problems arise out of a more general dilemma, one that the synchronic surrogate I propose in §4 is able to avoid.

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. (Note that the appeal to a 'seeming' here can be taken as a placeholder for any arbitrary mental state, as per Hedden's own weak internalist account of evidence.) 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 Some 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 also remember that it seemed to me earlier that E.

Reply. Hedden (2015b, 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 Hed-

den'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 argues 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 (2015b, 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 he defends is compatible with many accounts of evidence that are regarded as canonically externalist.[13] 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".[14] 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.[15] 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".[16]

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 endorsing the fairly strong requirement 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 fairly strong substantive account of evidence. Same

[13] 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.

[14] Hedden (2015b, p.22).

[15] Hedden (2015b, p.22, 28).

[16] See Hedden (2015b, p.142) and Hedden (2015b, p.147), respectively.

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 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—to the extent that we need to appeal to more than merely certain weak structural considerations like that evidence be a proposition—this account should not be ad hoc, but should be both continuous with the Bayesian framework and consistent with the very concerns that led us to synchronism in the first place.

Objection. Finally, one might concede that the conflicts problem is indeed a genuine problem, while nevertheless 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 be tempted to think 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 has previously gotten. Of course, unlike Hedden's norm, the diachronic framework 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.

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 of it. 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, I think we should expect a synchronic surrogate to do away with updates that are undermined by inconsistent evidence. To see this more clearly, consider 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 the wrong result that our previous evidence swamps our current seemings in this case where there is a conflict, Hedden's synchronic surrogate yields the desirable result that our current seemings swamp our previous evidence in cases of memory loss. What we are now in a position to see, however, is that this was only ever 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.

Earlier we said that the appeal to internalism and the appeal to problems about personal identity are two ways of appealing to the same general consideration: to the idea that being rational is a matter of believing in ways that are sensible from your perspective. Insofar as the appeal to a synchronic surrogate was appealed to *precisely* in order to capture this agential perspective, and insofar as the conflicts problem arises from the lack of a consistent perspective, this problem is an instance of *exactly* the worry 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 norm? 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 beliefs 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 credences and her higher-order credences, 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 of 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 of priors that one knows to be one's priors are a special case of updates of priors that one believes to be one's priors, since knowledge entails belief, but not vice versa. Moreover, since knowledge also entails factivity, updates of 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 to be its surrogate. As we've just noted, Meacham's norm tells us to update functions that we believe to be our priors, while the special case that Meacham

has in mind tells us to update functions that we know to be our priors, where "know" is taken to be shorthand for "certain of and right about", so that "A knows x iff a is certain of x and x is true."[17] Given that these are two qualitatively different updating scenarios, one might wonder whether Meacham's norm really is continuous with Diachronic Conditionalization.

I think the reason Meacham's assumption that being a special case of a norm is sufficient for surrogacy 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. One thing, A, is a limiting, special case of another thing, B, iff A is a special case of B, in virtue of taking the most extreme value that B allows. Correspondingly, for some update by a norm to be a limiting, special case of another update by a different 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, there is no qualitative difference between them; the only difference is one that can be expressed as a change in value. Therefore, where one norm is a limiting, special case of another, it seems reasonable to think they are continuous in a way that makes the latter a surrogate for the former.[18]

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.[19] 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 value that an update by Meacham's norm allows. It would need to be an update of a function that we believe with an extreme degree of certainty—a credence of one—is our prior function. But an update of 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. Since the class of updates that are limiting cases of updates by Meacham's norm are not updates by Diachronic Conditionalization, Diachronic Conditionalization is not a limiting case of Meacham's norm.

[17] Meacham (2010, fn. 12).

[18] 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.

[19] For the canonical description of Jeffrey Conditionalization, see Jeffrey (1965).

The higher-order beliefs Meacham substitutes for an agent's priors yields a norm that fails to reduce to Diachronic Conditionalization given any quantitative input.

We can take a step back and consider the general feature of Meacham's account that gives rise to the gap between his norm and the Bayesian framework. Meacham's norm says that the way we revise our beliefs ought to be entirely determined by our current evidence, including our higher-order evidence. But what distinguishes Bayesian updates—what differentiates such updates from classical statistical inference—is Bayesianism's commitment to a function distinct from our current evidence, whether this be a function that encodes some previous evidence, or the recommendations of some magical probability distribution in the sky, or whatever else. In taking our current higher-order evidence to be a substitute for our priors, Meacham's norm abandons the distinguishing feature of a Bayesian update. It's true that in the special case where our evidence is factive, we happen to recover this feature. But it still remains the case that a regular update by Meacham's norm lacks it entirely.

While this problem for Meacham's account is a subtle problem, it is not a small one. As I've noted already, I think the question of whether there can be a time-slice epistemology for Bayesians is really the question of whether formal and traditional epistemology can be made compatible. It's the question of whether Bayesianism can be preserved while accommodating the traditional epistemological values that motivate synchronism. While in the last section we saw that Hedden is able to preserve a connection with Diachronic Conditionalization at the cost of the agential perspective, Meacham's proposal makes the opposite tradeoff. Meacham is able to preserve the agential perspective, through his appeal to the agent's higher-order standpoint, at the cost of a connection with Diachronic Conditionalization.

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 happy with this result. For I think that even this weaker conclusion should leave us open to a better alternative.

4 Higher-Order Beliefs about Conflicting Evidence

I've argued that the question of whether some norm is an adequate surrogate for Diachronic Conditionalization is guided by a goldilocks principle. I've claimed 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, while these accounts are problematic in these ways when taken on their own, this section will argue that a combination of their features gets us what we are after. A combination of these features gets us a synchronic norm that avoids the problems of Diachronic Conditionalization while retaining enough of its structure to be its surrogate. 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. 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 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 as the higher-order beliefs an agent has about how *weighty* she now takes her remembered evidence to be. The idea I want to defend in this section is that an updating rule that appeals to higher-order beliefs to regulate conflicts in this way provides a better synchronic surrogate than the ones we've considered already.

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 these reasons are since it is the structure of the proposal that will allow it to succeed where both 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 judges that this evidence is *more reliable*. 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, and my proposal again does not turn on which of these understandings we adopt. The simplest way to go is perhaps 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 something like the following:

HO-Constraint: Let E supervene 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 are now, at t_n, producing her current seeming that the bird isn't blue ($\overline{B}$). 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 this conflicting evidence that supervenes upon her current mental states:

$$p_R(B)\text{=}.75$$
$$p_R(\overline{B}) = .25$$

Together, these values contribute to the weighted evidence partition that represents the agent's total evidence at t_n. Here is what an updating rule that draws upon this constraint 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 $\sum_i p(\mathrm{A}|\mathrm{E}_i)p_R(\mathrm{E}_i)$.[20]

[20] One might worry that there will not always be a coherent partition that is able to encode all the values generated by the HO-Constraint. For a proof that there always will be such a partition, in cases where the HO-Constraint is applicable, see Diaconis and Zabell (1982, §4.1-4.2). Say that the HO-Constraint yields the following values for the partitions $\{E_i\}$ and $\{F_i\}$: $p(\mathrm{E})$=.75, $p(\overline{E})$=.25, $p(\mathrm{F})$=.5, $p(\overline{F})$=.5. Then the weighted partition entailed by the HO-Constraint will be $\{EF_i\}$: i.e., $p(\mathrm{EF})$=.3, $p(\overline{E}\mathrm{F})$=.2, $p(\mathrm{E}\overline{F})$=.45, $p(\overline{EF})$=.05. This strategy treats E and F as independent and generates a partition that reflects their respective commitments. The assumption of independence is justified by the fact that the values of these partitions are determined independently of one another, i.e., by different beliefs about the reliability of the processes involved. (One might worry that there is one particular case that poses a problem for this method, which is where the agent gets as evidence some proposition like $E\wedge F$ that entails some further evidence proposition F. If the HO-Constraint assigns a value to the conjunction that is larger than the value it assigns to one of its conjuncts, it will be impossible to get a coherent evidence partition. However, it's difficult to imagine an example where my seeming delivers is a proposition whose negation is the proposition $\neg E\vee\neg F$. Instead, cases where it might seem that I

In tracking how relatively reliable we currently take our past seemings to be, our higher-order beliefs tell us how much weight we should assign to the evidence we now remember. This is exactly the normative role we should 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 number of ways we might go in such cases, plausibly, instances 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 about the relative reliability of this evidence. 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 first-order evidence that conflicts with the way that things seem to her now, either because she never received 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 cases that **A Higher-Order Synchronic Surrogate** overcomes the problems with Hedden's and Meacham's accounts:

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 already said that the synchronic view captures the idea that being rational is a matter of believing in ways that are sensible from your perspective. Like Meacham's norm, my synchronic surrogate doubles down on this idea by making rationality a matter of believing in ways that are sensible from your higher-order perspective.

get $E \wedge F$ are better understood, on the time-slice centric framework, as cases where it I get E and and I get F *at the same time*. There are no plausible cases, then, where the HO-Constraint yields values that won't generate a coherent partition.)

A Limiting, Special Case. Unlike Meacham's account, my synchronic surrogate is continuous with the Bayesian framework. 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. It is the limiting, special case of an update where 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 this solution, such a constraint would need to appeal to a uniform distribution over the *objects* of these beliefs. In the case of Meacham's updating rule, this would require appealing, 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.

One might object that I haven't done much to defend the particular measure I have used to define the HO-Constraint. Maybe some different constraint could do better. One thing my constraint does not take into account is an agent's quantitative beliefs about reliability: it does not take into account *how much more* reliable the agent judges herself to have been at the different times that she's gathered evidence. However, it seems reasonable to think that these more fine-grained beliefs about exactly how reliable an agent takes herself to have been won't be beliefs the agent has very often. An updating rule designed to have these elusive beliefs play a central role in our updates seems like a poor choice then.

The more important point, however, is that though I take the particular formulation of the HO-Constraint I have described in this section to be plausible, I have appealed to it mainly for illustrative purposes. The aim of this section has been to defend the idea that an updating rule that appeals to an agent's higher-order beliefs to regulate conflicts in her evidence provides a more adequate time-slice epistemology for Bayesians than either of the alternatives we've considered. It's the structure of my proposal that has allowed it to succeed where both Hedden and Meacham have failed. One might choose to define 'relative reliability' in some different way. Or one might forgo the appeal to this concept altogether while still holding onto what is

essential to the proposal—to the idea that a synchronic surrogate that includes an appeal to our higher-order beliefs about our evidence provides us with the synchronic surrogate we are after. It's a virtue of my proposal that it is compatible with many different ways of understanding what it might mean to be relatively reliable and, more generally, with many different ways of understanding what it might mean to judge some piece of evidence to be weightier than another. It's a virtue of my proposal that the strategy it advances is formal enough to be filled out in a number of different ways.

5 A Solution to the Goldilocks Problem

I've suggested that the question of whether some synchronic rule is an adequate surrogate for Diachronic Conditionalization is really the question of whether this rule can accommodate the norms of both formal and traditional epistemology. I've argued that neither of the synchronic rules in the literature are able to do this. The surrogates we've considered either cleave too closely to the norms they are intended to replicate, thereby reproducing the very problems they are intended to resolve; or else they diverge from these norms in subtle but significant ways. Just as attempts to assimilate the attitudes of formal and traditional epistemology yield unacceptable results in the form of the lottery paradox (*cf.* Kyburg (1961)), attempts to assimilate the updating norms of formal and traditional epistemology yield unacceptable results in the form of the goldilocks problem.

In the last section, I argued that the HO-Constraint provides a solution to this goldilocks problem. In avoiding the objections raised for the other approaches we've considered, it perfectly integrates the values of formal and traditional epistemology. It's natural to wonder, however, whether the HO-Constraint really does succeed in this way. Here's one last reason to think it might not. Earlier we saw that one way of succumbing to the goldilocks problem is to appeal to a substantive account of evidence—for example, to a factive account of evidence. Recall it was argued that such an account compromises the formality of our synchronic surrogate. One might now worry that the HO-Constraint *just is* this kind of substantive account of evidence we've been trying to avoid all along.

However, the HO-Constraint is no substantive account of evidence. It's a coherence constraint on our evidence since it tells us that our evidence is determined by how the way things currently seem coheres with our current higher-order beliefs. While I've offered a way of understanding this coherence constraint in terms of relative reliability, I've emphasized throughout

that the general strategy I've advanced is formal enough to be filled out in a number of different ways. As a formal, coherence constraint on our evidence, the HO-Constraint more closely approximates the ideal synchronic surrogate than any substantive account of evidence ever could by borrowing from a number of other Bayesian constraints. Like Lewis (1980)'s Principal Principle and van Fraassen (1984)'s Reflection Principle, it makes our rational credences a matter of deferring to an expert. Like Jeffrey Conditionalization, it assumes that our total evidence is a partition of propositions. The HO-Constraint is a hybrid of these commitments. It yields a more complicated picture of our evidence by appealing to the agent's present judgment about her past expertise.

Like these other paradigmatically Bayesian constraints, my synchronic surrogate posits a formal, coherence relation between our first-order beliefs and our higher-order beliefs. It's therefore clearly continuous with the Bayesian framework. It is also consistent with the values of traditional epistemology and, in particular, with those values that led us to synchronism in the first place. 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 an agent's perspective. It does this by invoking an agent's higher-order perspective. My synchronic surrogate's formal account of evidence encodes the very value that eludes the orthodox Bayesian—which is exactly the behavior we would expect of the ideal synchronic surrogate, one that mimics Diachronic Conditionalization in just the right way.

6 Acknowledgments

Thanks to Eric Campbell, Mike Nance, Jessica Pfeifer, Whitney Schwab, Steve Yalowitz, an audience at the University of Maryland, Baltimore County, and several anonymous referees, especially one from *Inquiry*, for very helpful comments and suggestions.

References

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

Christensen D (1991) Clever bookies and coherent beliefs. Philosophical Review 100(2):229–247

Christensen D (2004) Putting logic in its place: Formal constraints on rational belief. Oxford University Press on Demand

Diaconis P, Zabell SL (1982) Updating subjective probability. Journal of the American Statistical Association 77(380):822–830

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

Harman G (1986) Change in view: Principles of reasoning. The MIT Press

Hedden B (2015a) Time-slice rationality. Mind 124(494):449–491
Hedden B (2015b) Reasons without persons: rationality, identity, and time. OUP Oxford
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 (2013) 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: Ifs, 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
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
Titelbaum MG (2010) "not enough there there". evidence, reasons and language independence. Philosophical Perspectives 24:477–528
Titelbaum MG, Kopec M (2016) The uniqueness thesis. Philosophy Compass 11(4):189–200
White R (2005) Epistemic permissiveness. Philosophical perspectives 19(1):445–459
Williamson T (2000) Knowledge and its Limits. Oxford University Press