UR-PRIORS, CONDITIONALIZATION, AND UR-PRIOR CONDITIONALIZATION

CHRISTOPHER J. G. MEACHAM

University of Massachusetts Amherst

Conditionalization is a widely endorsed rule for updating one's beliefs. But a sea of complaints have been raised about it, including worries regarding how the rule handles error correction, changing desiderata of theory choice, evidence loss, self-locating beliefs, learning about new theories, and confirmation. In light of such worries, a number of authors have suggested replacing Conditionalization with a different rule—one that appeals to what I'll call *ur-priors*. But different authors have understood the rule in different ways, and these different understandings solve different problems. In this paper, I aim to map out the terrain regarding these issues. I survey the different problems that might motivate the adoption of such a rule, flesh out the different understandings of the rule that have been proposed, and assess their pros and cons. I conclude by suggesting that one particular batch of proposals, proposals that appeal to what I'll call *loaded evidential standards*, are especially promising.

1. Introduction

The standard Bayesian account of rationality provides us with a rule for how a subject should update her degrees of belief, or credences, in light of evidence. This rule—Conditionalization—tells us that upon receiving evidence, a subject's new credences should be equal to her previous credences conditional on that evidence. Conditionalization has been subject to a sea of complaints, including worries regarding how it handles error correction, changing desiderata of theory choice, evidence loss, self-locating beliefs, learning about new theories, and confirmation. Nevertheless, Conditionalization remains a widely endorsed updating rule.

One can find discussions of an alternative to this rule in the Bayesian litera-

Contact: Christopher J. G. Meacham <cmeacham@philos.umass.edu>

ture dating back to the 1950s.¹ Although there are differences of detail between these proposals, they all take the following form. First, they appeal to what I'll call an *ur-prior* function (also sometimes called *hypothetical priors*). This function might be understood in various ways, but common candidates include: the credences a subject should have if she had no evidence, a subject's initial credences, a subject's evidential standards, and any function that plays the right diachronic role. Second, they tell us that a subject's credences at a given time should be equal to her ur-priors conditional on her evidence at that time. I'll call any rule of this form *Ur-Prior Conditionalization*.

A number of people have motivated adopting Ur-Prior Conditionalization as a way to deal with the problems afflicting Conditionalization. But different authors have been motivated by different problems, and so have fleshed out the rule, and the notion of ur-priors it employs, in different ways. This makes it important to have a clear picture of the problems Ur-Prior Conditionalization is supposed to solve, and the extent to which different versions of the rule are able to address these problems. For it's only by doing so that we can get a clear picture of the pros and cons of different ways of developing Ur-Prior Conditionalization, and accurately assess whether such a rule is worth adopting at all.

My goal in this paper is to map out the terrain regarding these issues in three respects. First, I aim to survey the different problems that might motivate the adoption of Ur-Prior Conditionalization. Second, I aim to lay out the different understandings of ur-priors and Ur-Prior Conditionalization that have been proposed, and to clarify these sometimes vaguely specified proposals by giving them precise formulations. Third, I aim to provide an assessment of the relative pros and cons of these different understandings.

The rest of this paper will proceed as follows. In the second section I'll briefly sketch some background on Bayesian norms. In the third section I'll consider a number of complaints about Conditionalization that might motivate the move to Ur-Prior Conditionalization. In the fourth section I'll consider several ways of understanding ur-priors and Ur-Prior Conditionalization, and sketch some potential worries for these understandings. In the fifth section, I'll assess how well these different understandings do at addressing the complaints about Conditionalization raised earlier. In the sixth section I'll briefly sketch how the preceding discussion changes if one adopts Evidential Uniqueness, the thesis that every batch of evidence fixes a unique rational doxastic state. In the seventh section I'll conclude by arguing that one of these versions of Ur-Prior Conditionalization—a version which appeals to what I'll call *loaded evidential standards*—is especially promising.

1. For example, see Carnap (1950), Levi (1980), Skyrms (1983), Williamson (2000), Meacham (2008), Titelbaum (2013) and Hedden (2015).

2. Background: Bayesian Norms

Let a subject's *credence function* be a function which assigns real numbers to propositions in a way that represents her confidence in those propositions, where an assignment of 1 represents certainty that the proposition is true, and an assignment of 0 represents certainty that the proposition is false.² A credence function cr is *probabilistic* iff it satisfies the following two conditions. First, the set of propositions that cr assigns values to forms a Boolean Algebra—i.e., a set of propositions closed under conjunction, disjunction, and negation. Second, the values cr assigns to these propositions satisfies the following constraints (where $\top$ is the tautology and $\perp$ the anti-tautology):

1. $\forall A, cr(A) \geq 0$.
2. $cr(\top) = 1$.
3. $\forall A, B$, if $A \wedge B = \perp$, then $cr(A) + cr(B) = cr(A \vee B)$.

If cr is a probability function, cr assigns values to both $A \wedge B$ and B , and cr assigns a non-zero value to B , then we can define the *conditional probability* of A given B as follows:³

$$cr(A \mid B) := \frac{cr(A \wedge B)}{cr(B)}$$

Note that it follows from this definition that if cr is not a probability function, then $cr(A \mid B)$ is undefined.

With this in hand, we can characterize the two core Bayesian norms:

Probabilism: A subject's credences cr should be probabilistic.

Conditionalization: If a subject with credences cr gets evidence E , then her new credences should be $cr^{+E}(\cdot) = cr(\cdot \mid E)$, if defined.⁴

2. Or at least *virtual* certainty that the proposition is true/false, since one may hold that there are cases in which one can have a credence of 1/0 in a proposition without being certain that it's true/false (e.g., that at least one of/none of a countably infinite number of fair coin tosses will land heads). In what follows, I'll put a number of complications relating to this, and to various other infinity issues, to one side. For although these complications are important to several of the issues that will be discussed (e.g., if we're going to use ur-priors to handle self-locating beliefs in the manner discussed in Section 5.5, and there are an uncountably infinite number of times, then there's pressure to characterize ur-priors using primitive conditional probabilities or non-standard probabilities), I lack the space to adequately address these issues here. For discussions of some of these issues, see Weatherson (2005), Easwaran (2014), Hájek (2015), and the references therein.

3. Although this is the canonical way to understand conditional probabilities, one can mount a case for taking conditional probabilities to be primitive instead; see Hájek (2003).

4. It follows from the definition of conditional probability that Conditionalization won't

As has been noted elsewhere, there are a couple of different ways to understand Conditionalization.⁵ One important question is whether we're understanding Conditionalization as an *interval* rule or a *sequential* rule. On an *interval* understanding, cr is the subject's credence at any time, cr^{+E} is the subject's credence at any later time, and E is the cumulative evidence the subject receives between those two times. On this understanding the rule tells us, for any interval of time, how our credences at the beginning and end of that interval should be related, given the evidence received in the interim.⁶

On a *sequential* understanding, cr is the subject's credences after she's received all of her evidence but E , cr^{+E} is the subject's credences upon receiving E , and E is the new piece of evidence she's received. On this understanding, the rule tells us how to update whenever we get a new piece of evidence.⁷

Another important question is whether we're understanding Conditionalization as *narrow scope* or *wide scope*.⁸ Let O be the operator that expresses epistemic obligation, A the proposition that the subject has credences cr and gets evidence E , and C the proposition that the subject adopts new credences $cr^{+E}(\cdot) = cr(\cdot | E)$. On the *narrow scope* understanding, Conditionalization takes the form $A \rightarrow O(C)$. On the *wide scope* understanding, Conditionalization takes the form $O(A \rightarrow C)$.

The differences between and relative merits of these understandings have been discussed elsewhere.⁹ For much of the discussion to come, the difference between these understandings won't matter. So to simplify things, unless I say otherwise, I'll assume we're adopting an interval wide scope understanding of Conditionalization. In the few places where this makes a difference to the dialectic, I'll note how things play out if we adopt a different understanding of

impose substantive constraints on subjects without probabilistic credence functions. (If cr is non-probabilistic, $cr(\cdot | E)$ will be undefined, and Conditionalization won't impose any constraints on the subject's beliefs.) But one might wonder whether it would be better to understand the rule as simply requiring one's new credences $cr^{+E}(\cdot)$ to equal the ratio $\frac{cr(\cdot \wedge E)}{cr(E)}$. For that ratio will still be well-defined even if cr isn't probabilistic (assuming $cr(\cdot \wedge E)$ and $cr(E)$ have numerical values, and $cr(E) \neq 0$). Although this is a natural thought, this non-standard interpretation of Conditionalization yields implausible (and arguably unintelligible) verdicts. For example, suppose a subject has non-probabilistic credences $cr(A \wedge E) = 0.6$ and $cr(E) = 0.3$, and gets E as evidence. According to this non-standard rule, the subject's new credence in A should be $cr^{+E}(A) = \frac{cr(A \wedge E)}{cr(E)} = 2$. But it's not clear what it would even mean to adopt a credence of 2 in a proposition. (I thank an anonymous referee for encouraging me to address this point.)

5. See Schoenfield (in press) and Meacham (2016) for discussions of issues regarding how to understand the rule. The two distinctions discussed here are explored in greater detail in Meacham (2016).

6. See Lewis (2010) for a formulation of this kind.

7. See Howson and Urbach (1993) for a formulation of this kind.

8. Recent discussions of the difference between narrow and wide scope norms were initiated by Broome (1999).

9. See Meacham (2016).

Conditionalization.

Probabilism and Conditionalization are not the only norms that Bayesians have adopted. In addition to these two core Bayesian norms, there are further auxiliary norms that some Bayesians accept. For example, many Bayesians endorse something like Lewis's (1980) Principal Principle (a norm tying a subject's beliefs about a proposition to her beliefs about the chance of that proposition), and some Bayesians have endorsed rules like van Fraassen's (1984) Reflection Principle or a version of the Principle of Indifference (see Howson & Urbach 1993). Although I won't discuss these auxiliary norms in detail, it's important to keep them in mind, since interesting issues will arise regarding how these norms interact with updating rules like Conditionalization.

3. Complaints about Conditionalization

Despite the popularity of Conditionalization, a number of complaints have been raised about the rule. Many of these complaints serve as potential motivations for adopting something like Ur-Prior Conditionalization, roughly:

Ur-Prior Conditionalization: If a subject has ur-priors up and current evidence E , her credences cr should be $cr(\cdot) = up(\cdot \mid E)$, if defined.

For the purposes of this section I'll treat the term "ur-priors" as a black box, one which might be cashed out in a number of different ways. (We'll return to look at different ways of understanding ur-priors and Ur-Prior Conditionalization in Section 4.)

In what follows, I'll consider eight complaints about Conditionalization that might move one to adopt something like Ur-Prior Conditionalization. In each case I'll consider a couple of possible replies to these complaints, and sketch how Ur-Prior Conditionalization might help. (These will be brief sketches only; we'll return to a more detailed exploration of how different versions of Ur-Prior Conditionalization deal with these complaints in Section 5.)

Two comments about these complaints before we proceed. First, I'm not trying to provide an exhaustive list of the concerns one might have about Conditionalization. For there are many concerns about Conditionalization that don't motivate a move to Ur-Prior Conditionalization (e.g., concerns regarding uncertain evidence), and I won't discuss such concerns here. Second, I expect few will take *all* of these complaints to provide compelling reasons to adopt Ur-Prior Conditionalization. In some cases, one might not be moved by the complaint about Conditionalization; in other cases, one might be moved by the complaint, but prefer a response that doesn't appeal to Ur-Prior Conditionalization. But although I suspect different people will be moved by different complaints, I expect most will take at least some of these complaints to provide a reason to consider

Ur-Prior Conditionalization.

3.1. Error Correction: Updating Errors

One complaint about Conditionalization is that it doesn't properly handle cases in which subjects have updated incorrectly. Consider a subject who has credences cr_0 at t_0 , who gets evidence E_1 at t_1 , and who updates incorrectly to get credences $cr_1(\cdot) \neq cr_0(\cdot \mid E_1)$ at t_1 . If she then gets some new evidence E_2 at t_2 , what should her credences at t_2 be? Should her t_2 credences be those she'd get by conditionalizing cr_1 on E_2 , $cr_2(\cdot) = cr_1(\cdot \mid E_2)$? Or should her t_2 credences be those she'd get by conditionalizing cr_0 on the evidence she's gotten since t_0 , $cr_2(\cdot) = cr_0(\cdot \mid E_1 \wedge E_2)$? The first answer seems funny: if a subject makes a mistake while updating her credences, one would think that the right thing to do is to correct the mistake, not to ignore it. The second answer seems better: it tells the subject to correct her error, and adopt the credences she would have adopted if she'd updated correctly all along. So one might like one's updating rule to tell them that after making an updating error, a subject should adopt the credences that she would have come to have if she hadn't made that error.

But no understanding of Conditionalization yields this result, though the way in which it fails to yield this result depends on which understanding we adopt. First consider the sequential narrow scope understanding of Conditionalization. This rule tells us that upon receiving new evidence, subjects should adopt credences equal to their credences right before they got their new evidence, conditional on that new evidence. So in the case above, this rule tells us that the subject's t_2 credences should be those she'd get by conditionalizing cr_1 on E_2 , $cr_2(\cdot) = cr_1(\cdot \mid E_2)$. But this isn't the prescription we want: for it tells the subject she shouldn't bother correcting any errors she's made.

Next consider the interval narrow scope understanding of Conditionalization. This rule says that for any earlier credence function cr , the subject's new credences cr^+ should be equal to cr conditional on the cumulative evidence she's received in the interim. So in the case above, this rule tells us that the subject's t_2 credences should be equal to both her t_0 credences conditional on $E_1 \wedge E_2$ and her t_1 credences conditional on E_2 . But these prescriptions are inconsistent.

Now consider the wide scope understandings of Conditionalization, whether sequential or interval. These rules take the form $O(A \rightarrow C)$, or equivalently $O(\neg A \vee C)$, and so only require subjects to satisfy a disjunction. Because of this, these rules aren't committed to the problematic prescriptions of the narrow scope understandings. For example, instead of requiring the subject to adopt the wrong t_2 credences or requiring her to adopt two incompatible sets of t_2 credences, these rules will merely require her to adopt those t_2 credences *or* to have had different earlier credences and evidence. But these wide scope understandings still won't yield the desired results: they don't tell the subject to correct the

errors she's made. For example, given the sequential wide scope understanding, Conditionalization tells the subject to either not have had credences cr_1 and have gotten evidence E_2 , or to adopt new credences $cr_2(\cdot) = cr_1(\cdot \mid E_2)$. So this rule effectively tells her to either not to have made a mistake, or to adopt new credences which effectively ignore the mistake. Neither option is what we're looking for: what we're looking for are prescriptions which tell the subject to adopt t_2 credences which *correct* the mistake she's made.¹⁰

One might try to defend Conditionalization (on any of these understandings) by suggesting that these rules should be thought of as norms that describe the beliefs of ideal subjects. After all, it's well-known that Conditionalization only makes prescriptions to subjects who are at least somewhat idealized. For example, Conditionalization only makes well-defined prescriptions for subjects with probabilistic credences.¹¹ With this in mind, one might propose to understand Conditionalization as only making prescriptions to subjects who are ideal in another respect: they have always updated correctly. Or, to put it another way, one might understand Conditionalization as only directly telling us how ideal subjects should update their beliefs, and only indirectly (if at all) telling us anything about how non-ideal subjects should update their beliefs. And since ideal subjects never fail to update correctly, the difficulties described above will never arise.

Working out rules which describe the behavior of ideal subjects is an interesting and worthwhile project. And it's reasonable to understand Conditionalization as such a rule.¹² But even granting this, one might prefer a rule that requires less idealization and makes plausible prescriptions regarding less idealized subjects than Conditionalization does. Indeed, the verdicts described above—that the subject should correct her mistake by adopting $cr_2(\cdot) = cr_0(\cdot \mid E_1 \wedge E_2)$ —require a rule that applies to less idealized subjects. So insofar as we want a rule which captures these verdicts, we'll be unhappy with a rule like Conditionalization.

In light of Conditionalization's inability to yield the desired verdicts in these cases, one might pursue a more radical option: replacing Conditionalization with a rule whose prescriptions aren't affected by updating errors, such as a version of Ur-Prior Conditionalization. For example, one might adopt a rule whose pre-

10. Unlike the narrow scope formulation, there's a sense in which the wide scope formulation gives the right answer—the subject *shouldn't* have made the mistake. The problem with the wide scope formulation is that it doesn't tell us enough: given that she made a mistake, what should she do now? (I thank an anonymous referee for encouraging me to address this point.)

11. Since conditional probabilities are only defined for probability functions, $cr(\cdot \mid E)$ is undefined if cr is non-probabilistic. So if a subject doesn't have probabilistic credences, Conditionalization won't yield a well-defined prescription.

12. See Meacham (2016) for a defense of this way of understanding Conditionalization. For a discussion of the dangers of applying idealized norms to non-ideal subjects, see Staffel (2016).

scriptions are determined by an array of facts associated with a subject that encode what she should believe given any particular batch of evidence, e.g., some kind of ‘ur-priors’. Such a rule would avoid the above complaint because it would yield the same prescriptions regardless of whether the subject has made any updating errors. And it could yield these prescriptions without weakening the normative scope of the framework.

3.2. Error Correction: Auxiliary Norm Violations

A second complaint about Conditionalization is that it doesn’t properly handle cases in which subjects have violated auxiliary Bayesian norms, such as Lewis’s (1980) Principal Principle. For example, consider a subject whose credences at t_0 , cr_0 , violate the Principal Principle, and who gets some new evidence E_1 at t_1 . Conditionalizing on a credence function that violates the Principal Principle will typically yield another credence function that violates the Principal Principle. So what should her credences at t_1 be? Should her t_1 credences be those she’d get by conditionalizing cr_0 on E_1 , $cr_1(\cdot) = cr_0(\cdot \mid E_1)$, even though these credences will typically violate the Principal Principle? Or should she consider what her t_0 credences should have been in order to satisfy the Principal Principle, and then conditionalize those corrected credences on E_1 to generate her t_1 credences?¹³ The first answer seems funny: if a subject’s credences violate the Principal Principle, one would think the right thing to do is to correct this violation, not ignore it. The second answer seems better: it tells the subject to correct her error, and adopt the credences she would have adopted if she’d successfully abided by the Principal Principle all along. So one might like one’s updating rule to tell them that, given an auxiliary norm violation, a subject should adopt the credences that she would have come to have if she hadn’t violated those auxiliary norms.

Conditionalization won’t yield that result, though the way in which it fails to do so depends on which understanding of the rule we adopt. First, consider the narrow scope understandings, whether sequential or interval. Given these

13. Here’s a concrete example of a case in which this conflict arises, using Lewis’s (1980) framework. Let cr_0 stand for a subject’s (probabilistic) initial credence function, held at t_0 , and let $ch_{TH_t}(A)$ stand for the chance of A given chance theory T and history up to t , H_t . Lewis’s Principal Principle requires that $cr_0(A \mid TH_t) = ch_{TH_t}(A)$. Now suppose that $ch_{TH_0}(A) = ch_{TH_1}(A) = 0.5$, and that the subject’s credences at t_0 are such that $cr_0(TH_0) = 1$, $cr_0(ATH_0) = 0.25$, $cr_0(TH_1) = 0.5$, and $cr_0(ATH_1) = 0.125$. Since $cr_0(A \mid TH_0) = 0.25 \neq ch_{TH_0}(A)$, the subject violates the Principal Principle at t_0 . Now suppose that at t_1 she gets evidence TH_1 . If she updates on TH_1 via Conditionalization, her new credences will be such that $cr_1(A) = cr_1(A \mid TH_1) = cr_0(A \mid TH_1) = 0.25 \neq ch_{TH_1}(A)$, and she’ll continue to violate the Principal Principle at t_1 . On the other hand, if she modifies her credences at t_1 so that $cr_1(A \mid TH_1) = 0.5$, she’ll satisfy the Principal Principle, but will have changed her credences in a way which violates Conditionalization. (I thank an anonymous referee for encouraging me to provide a concrete example here.)

understandings, Conditionalization and the Principal Principle will yield conflicting prescriptions. Conditionalization will tell the subject in the case above to adopt the t_1 credences she'd get by conditionalizing her t_0 credences on E_1 . But the Principal Principle will tell her not to adopt those credences, since they violate the Principal Principle.

Next, consider the wide scope understandings, whether sequential or interval. Unlike the narrow scope understandings of the rule, the wide scope understandings won't yield prescriptions which directly conflict with the Principal Principle. For instead of requiring the subject above to adopt the Principal Principle-violating t_1 credences, they'll merely require her to adopt those t_1 credences *or* to have had different earlier credences and evidence. But these wide scope understandings don't yield the desired prescriptions. For these rules effectively tell the subject to either adopt new credences which effectively ignore the mistake, or to not have made a mistake. Neither option is what we're looking for. What we want is a rule which, holding fixed the mistakes the subject made in the past, prescribes future credences which *correct* those mistakes.

(As before, one might try to defend Conditionalization by urging us to think of the rule as a description of how ideal subjects update their beliefs. Since ideal subjects will never violate auxiliary norms like the Principal Principle, the worries discussed above won't arise. But again, while this is a reasonable way to think of Conditionalization, one might prefer a rule that requires less idealization, and makes plausible prescriptions regarding less idealized subjects. And insofar as we want a rule which captures the verdicts described above, we'll be unhappy with a rule like Conditionalization.)

Given Conditionalization's inability to yield the desired results in these cases, one might want to replace Conditionalization with a rule whose prescriptions aren't affected by the mistakes a subject made in the past, such as a version of Ur-Prior Conditionalization. If we adopt a rule whose prescriptions are determined by some 'ur-priors' which tell us what credences a subject ought to adopt (all things considered) given an arbitrary batch of evidence, then we can avoid these problems. Since the prescriptions of auxiliary norms like the Principal Principle would already be built into these facts, the prescriptions of this rule and the auxiliary norms will never conflict.

3.3. *Changing Desiderata of Theory Choice*

A third complaint about Conditionalization is that it doesn't provide a satisfactory way of handling changing desiderata of theory choice.¹⁴ Consider two theories regarding a certain phenomenon, T1 and T2. Suppose none of the available

14. See Levi (1980) for a version of this complaint.

evidence distinguishes between these two theories.¹⁵ And suppose that while T1 is better than T2 with respect to some theoretical virtues (e.g., simplicity), it is worse with respect to others (e.g., unification). It seems like reasonable subjects could disagree as to which of these two theories is more plausible.¹⁶ Which theory a subject finds more plausible hangs on the relative importance she attaches to different desiderata of theory choice. And one reasonable subject could take simplicity or unification to be more important than another.

Sometimes subjects change their attitudes about the relative importance of the desiderata of theory choice. They come to think that certain desiderata of theory choice are more or less important than they were initially inclined to think. And one might think that these changes needn't be irrational. If so, it seems reasonable for a subject who undergoes such a change to readjust their credences in theories like T1 and T2. But Conditionalization won't allow for such changes. These are changes in one's credences without a corresponding change in one's evidence, and this is something that Conditionalization forbids.

Now, for the most part, that's surely the right verdict—changes in one's credences without changes in one's evidence are usually irrational. But one might think there are exceptions, like the one just described above.¹⁷ And insofar as one thinks such changes can be rational, one has a reason to be dissatisfied with Conditionalization.

In light of this worry, one might be motivated to replace Conditionalization with a version of Ur-Prior Conditionalization. In particular, one might adopt a rule according to which the credences one ought to have is determined by something like (i) one's evidence, and (ii) what hypotheses one takes to be inherently plausible, evidence aside. Such a rule would allow for the kinds of Conditionalization-violating belief changes described above, where belief changes result from changes in what the subject finds inherently plausible. But it would still make it the case that most evidence-less belief changes are impermissible, since most evidence-less belief changes won't reflect changes in what the subject finds inherently plausible.

15. For example, suppose that all of the available evidence is compatible with both all of the worlds at which T1 is true and all of the worlds at which T2 is true.

16. This complaint assumes that Evidential Uniqueness—the claim that every batch of evidence fixes a unique rational doxastic state—is false. So one can think of this as a conditional complaint: if Evidential Uniqueness is false, then Conditionalization doesn't provide a satisfactory way of handling changing desiderata of theory choice. (For further discussion of Evidential Uniqueness, see Section 6.)

17. For a discussion of another class of exceptions, see Perez Carballo (in press). For a discussion of a similar phenomenon, involving subjects changing their minds about which theories are most plausible, see Moss (2014).

3.4. Evidence Loss

A fourth complaint about Conditionalization is that it doesn't adequately handle cases of evidence loss, like memory loss.¹⁸ Suppose you know a fair coin toss is going to take place, and you have a credence of 0.5 that it will land heads. Then you learn that the coin toss landed heads, and you increase your credence in heads to 1. Finally, you suffer from memory loss, and you forget the outcome of the coin toss. What should your credence now be that the coin landed heads?

It seems plausible that your credence in heads should no longer be 1. After all, you don't remember how the coin landed, and so it seems unreasonable to be certain the coin landed heads. Instead, it seems your credence in heads should return to something like 0.5.

But Conditionalization won't allow this kind of belief change. Conditionalization makes certainties permanent: if you come to have a credence of 0 or 1 in a proposition, Conditionalization will require you to always have a credence of 0 or 1 in that proposition. Thus, once you've adopted a credence of 1 in heads, Conditionalization will require you to continue having a credence of 1 in heads, even if you forget the evidence that justified this certainty in the first place. Insofar as one takes this to be implausible, one has a reason to be dissatisfied with Conditionalization.^{19,20}

One response to this worry is to maintain that Conditionalization is a description of how ideal subjects update their beliefs, and that ideal subjects never lose

18. For versions of this complaint, see Skyrms (1983), Talbott (1991), Williamson (2000), Arntzenius (2003), Meacham (2008), Titelbaum (2013), and Hedden (2015).

19. The way in which Conditionalization fails to yield the desired results depends, in part, on the notion of evidence one employs. Suppose one adopts the account of evidence suggested by Howson and Urbach (1993), where a subject's evidence is the logically strongest proposition in which she has a credence of 1. On this account, Conditionalization will entail that evidence loss—having evidence now which is strictly weaker than evidence one had in the past—is irrational. This is because, on this account, evidence loss entails that one has shifted from having a credence of 1 in some proposition to not having a credence of 1 in that proposition, which Conditionalization forbids. Suppose instead one adopts the account of evidence suggested by Lewis (2010), where a subject's evidence corresponds to the set of possibilities compatible with her perceptual experiences and memories. On this account, nothing Conditionalization says forbids losing evidence, but Conditionalization won't allow subjects in such cases to change their beliefs in the manner one might want (e.g., in the case above, reverting to a credence of 0.5 in heads and tails). So regardless of which account we adopt, Conditionalization won't permit the kinds of plausible belief changes described in the case above.

20. This is one way to motivate evidence loss worries, but there are other ways as well. In particular, some might want to allow for cases in which subjects can rationally *reject* evidence. (E.g., perhaps the sight of a moving shadow makes a subject certain that there's someone behind her, but when she turns around, she sees that the shadow comes from a tree shaking in the wind.) Again, Conditionalization won't be able to yield the desired verdicts in these cases. (I thank an anonymous referee for raising this point and suggesting this example.)

evidence. If so, then these evidence loss worries are not a problem for Conditionalization. But while it might be reasonable to understand Conditionalization as a rule describing the belief updates of this kind of ideal subject, one might still prefer a rule that requires less idealization, and makes pleasing prescriptions regarding less idealized subjects. And insofar as we want a rule that yields the kinds of plausible verdicts described above, we have a reason to be unhappy with Conditionalization.

Another response is to argue that we should *accept* Conditionalization's prescriptions in cases of evidence loss. Meacham (2010) and Schwarz (2012) have both argued that the verdict that Conditionalization's prescriptions in cases of evidence loss are implausible stems from implicit internalist or evidentialist intuitions. And, they've argued, these views are incompatible with diachronic norms, like Conditionalization. So insofar as one takes there to be diachronic norms, one should reject these intuitions, and thus reject these kinds of complaints about Conditionalization.

Those dissatisfied with these responses might concede that this is a worry for Conditionalization. And thus they might take this to be a reason to replace Conditionalization with a rule that can take evidence loss into account, such as a version of Ur-Prior Conditionalization. Forgotten evidence shows up in Conditionalization's prescriptions for two reasons: (i) Conditionalization updates using the subject's cumulative evidence, which is not affected by evidence losses, and (ii) Conditionalization updates on a subject's prior credences, credences which take her prior evidence into account, and thus implicitly imports evidence that may have been lost into its prescriptions. And if one adopts a rule which diverges from Conditionalization in each of these two respects, then one can avoid these worries.²¹ First, the rule would update using a subject's current evidence, instead of her cumulative evidence, which allows it to update using a form of evidence that's sensitive to evidence loss. Second, the rule would update on something that encodes how a subject ought to respond to evidence without taking any of her prior evidence into account, which ensures that there is no prior evidence being implicitly imported.

3.5. *Self-Locating Beliefs*

A fifth complaint about Conditionalization, related to the complaint about evidence loss, is that Conditionalization is unable to provide a satisfactory account of how to update self-locating beliefs, such as beliefs about one's temporal location or identity.²² Consider a subject who is looking at a clock that she knows

21. A number of people have suggested accommodating evidence loss in this way, including Skyrms (1983), Meacham (2008), Titelbaum (2013), and Hedden (2015).

22. That is, Conditionalization cannot provide a satisfactory account of how to update what Lewis (1979) called irreducibly *de se* beliefs. For some discussions of this complaint, see

to be accurate. The clock currently reads 9 AM, and so the subject adopts a credence of 1 that it's 9 AM. A minute later the clock reads 9:01 AM. At this point it seems the subject should adopt a credence of 1 that it's 9:01 AM, and decrease her credence that it's 9 AM to 0. But, as we've seen, Conditionalization can't allow this kind of belief change. Conditionalization makes certainties permanent, so once the subject has adopted a credence of 1 that it's 9 AM, it will require her to continue to be certain that it's 9 AM forever.

One might take this worry to be a reason to replace Conditionalization with a different updating rule. There are a number of different kinds of rules one might employ, but one option is familiar from the preceding sections: one could adopt a version of Ur-Prior Conditionalization. In particular, one could adopt a rule whose prescriptions are determined by some facts ('ur-priors') that encode how a subject should respond to evidence, which are independent of her prior credences and evidence. Such a rule would be able to yield the kinds of temporally-shifting prescriptions we want in the clock-watching case described above.²³

3.6. *Learning about New Theories*

A sixth complaint about Conditionalization is that it doesn't provide a satisfactory account of what subjects should believe when they learn about new theories they hadn't considered before.²⁴ Let T , U and V be three mutually exclusive theories that could explain some phenomenon. And suppose a subject who satisfies Probabilism has a credence of 0.4 that theory T provides the correct explanation of that phenomenon, and a credence of 0.6 that theory U provides the correct explanation. Finally, suppose the subject then discovers that theory V , a theory she was not previously aware of, would also explain the phenomenon. If V is a plausible and attractive theory, it would seem reasonable for her to now assign some non-zero credence to V . But Conditionalization won't allow this. Since the subject's prior credence in $T \vee U$ was 1, Conditionalization will require her to continue to be certain that one of these two theories is correct. And since V is incompatible with T and U , it follows that her credence in V must be 0.

One response to this complaint is to again appeal to the idea that Conditionalization is a description of how ideal subjects update their beliefs. If ideal sub-

Elga (2000), Arntzenius (2003), Halpern (2004), Meacham (2008), Kim (2009), Briggs (2010), Meacham (2010), Bradley (2011), Moss (2012), Schwarz (2012), and Titelbaum (2013).

23. Even if one approves of this general approach, there are different functions from evidence and these facts to prescriptions that one might employ. This paper implicitly assumes one would want to adopt something like what Meacham (2008) calls "centered conditionalization". But there are other choices one might consider, such as the alternatives discussed by Halpern (2004) and Meacham (2008).

24. See Glymour (1980) and Garber (1983) for classic discussions of this worry, and Perez Carballo (in press) for a more recent take on this issue.

jects are aware of every possible theory, then the worry that they could discover some new theory won't arise. Of course, it's not clear how the subject described in the case above fails to be ideal. Being unaware of certain theories doesn't make one violate Probabilism or Conditionalization. But one could, perhaps, appeal to some further auxiliary norm—perhaps a norm that requires subjects to be 'modally omniscient'—in order to maintain that subjects who are unaware of certain theories fail to be ideal.

In any case, if we adopt this response, then cases in which subjects learn about new theories become cases of error correction. And as we've seen, Conditionalization (arguably) fails to provide a satisfactory account of error correction. If that's right, then this response does little to vindicate Conditionalization.

A second response is to embrace the thought that cases in which subjects learn about new theories should be treated as cases of error correction, and appeal to a version of Ur-Prior Conditionalization to deal with cases of error correction, in the manner described earlier. One might go about this in one of two ways. First, one might take there to be auxiliary norms which require rational subjects to be aware of every possibility, and take cases in which subjects learn about new theories to be cases in which they violated this norm. That would make these cases instances of the kind of auxiliary norm error correction discussed in Section 3.2. Second, one might take there to be some array of facts ('ur-priors') that encode how a subject should respond to evidence, and take cases in which subjects learn about new theories to be cases where subjects haven't been generating their credences from these facts in the right way. That would make these cases instances of the kind of updating error correction discussed in Section 3.1 (though in this case the updating rule being violated is a version of Ur-Prior Conditionalization, not Conditionalization).

That said, this response does face one potential demerit. Some will want an account that not only tells a subject what credences to adopt when she learns about new theories, but allows subjects who do so to be rationally blameless. That is, some will want an account on which the subject described above never does anything wrong, because both her initial certainty in $T \vee U$ and her subsequent uncertainty in $T \vee U$ are rational. And this second response—which reduces cases of learning about new theories to cases of error correction—won't yield this result.

A third response is to treat learning about new theories as an expansion of a subject's conceptual space, and supplement Conditionalization, which tells you how to update when you get evidence, with a second rule, which tells you how to update when you change your conceptual space. On this approach, we wouldn't take the subject described above to have a prior credence of o in V . Instead, we'd take the subject to not have a prior credence in V ; before she learned about V , V

wasn't a part of the possibility space over which her credences were defined.²⁵

A fourth response is to again handle these cases by appealing to Ur-Prior Conditionalization, but not by treating these as cases of error correction. Instead, one might treat these as cases where learning about new theories changes one's ur-priors. This way of treating these cases would allow subjects to change their beliefs upon learning about new theories in a way that doesn't imply any rational deficiency on their part.

3.7. *The Problem of Old Evidence and Confirmation*

A seventh complaint about Conditionalization is that it's subject to the problem of old evidence.²⁶ The literature on this issue has recognized at least two distinct 'problems of old evidence'—a synchronic problem and a diachronic problem.²⁷ And only the synchronic problem of old evidence raises a new complaint about Conditionalization. But it's worth briefly discussing both problems.

The diachronic version of the worry is that Conditionalization won't yield the desired prescriptions in cases where you learn that some old evidence—that is, evidence you've already received—is entailed by a theory. So, for example, suppose you have probabilistic credences, and you learn that some theory H (such as the General Theory of Relativity) entails some already observed behavior E (e.g., Mercury's perihelion).²⁸ Since E has already been observed, your credence in E is 1, i.e., $cr(E) = 1$. Thus discovering that H entails E can't raise your credence in H , according to Conditionalization, since if $cr(E) = 1$ then $cr^{+E}(H) = cr(H | E) = cr(H)$. But it seems like discovering that H entails E *should* raise your credence in H . Certainly, Einstein took the discovery that the General Theory of Relativity explains Mercury's perihelion—behavior that was well-known to the scientific community—to be a powerful mark in its favor.²⁹

There are a couple of ways to flesh out the diachronic version of the worry,

25. See Garber (1983) and Perez Carballo (in press) for discussions of such approaches. Although commentators have generally focused on what rational constraints apply to cases in which subjects expand their conceptual space, similar questions arise about cases in which subjects shrink their conceptual space.

26. For the classic presentation of this worry, see Glymour (1980). Later discussions of this worry include Garber (1983), Skyrms (1983), Howson (1984), Eells (1985), Earman (1992), Maher (1996), Christensen (1999), and Eells and Fitelson (2000).

27. I borrow this terminology from Christensen (1999). For a more detailed taxonomy of old evidence problems, see Eells (1985).

28. This example comes from Glymour (1980).

29. This way of spelling out the diachronic problem focuses on what Einstein *discovered*—that H entails E —and how this discovery increased Einstein's credence in H , *contra* Conditionalization's prescriptions. Another way of spelling out the diachronic problem is to focus on the worry that E can't raise your credence in H (according to Conditionalization), even though E seems like it should be evidence for H . (I thank an anonymous referee for flagging this alternative.)

but none of them yield new complaints about Conditionalization. If we understand the case as one where you didn't realize that H entails E because theory H hadn't occurred to you yet, we get an instance of the "learning about new theories" worry discussed in Section 3.6. If we understand the case as one where you failed to notice the relationship between H and E when you originally received E , and so you failed to update in the appropriate way, we get an instance of the updating error correction worry discussed in Section 3.1. A third way to flesh out the diachronic worry is as a complaint about the logical omniscience requirement, something entailed by typical formulations of Probabilism. But this is a complaint about Probabilism, not Conditionalization. Since none of these ways of understanding the diachronic worry raise new concerns for Conditionalization, I'll put the diachronic worry aside.

The synchronic version of the worry is that Conditionalization, when paired with a natural claim about the relation between updating and confirmation or evidential support, will yield unsatisfactory verdicts. Suppose we want to assess whether evidence E supports theory H . A natural thought is to look at whether the credence in H a given subject would adopt after receiving E is greater than her current credence in H . That is, it's natural to maintain that, for a subject with credences cr , E confirms H (or E supports H) iff $cr^{+E}(H) > cr(H)$. But if we adopt Conditionalization, then this account of confirmation seems to yield the wrong results in cases in which E is old evidence. If a subject already has E as evidence, then $cr(E) = 1$, and thus $cr^{+E}(H) = cr(H | E) = cr(H)$. So it would seem to follow that no evidence a subject already has can support a hypothesis, which seems incorrect. This is the synchronic problem of old evidence: given Conditionalization and a natural claim about the relation between updating and evidential support, how do we make sense of the fact that evidence can still support a hypothesis after it's been updated on?³⁰

One response to this worry is to argue that we should characterize confirmation using the counterfactual credences a subject would adopt if she did not have

30. One could also formulate versions of the accounts of confirmation I discuss here in terms of conditional probabilities directly, and try to avoid making any claims about rational updating. These accounts will run into the same problems as the ones I consider in the text. And in order to adopt the analog of the fourth proposal I consider, which avoids these problems, one still needs to appeal to ur-priors.

If we set up things this way, one might wonder: do we still have a complaint about *Conditionalization*? We do, though the complaint is less direct: it becomes a complaint about the core Bayesian framework (Probabilism and Conditionalization) rather than a complaint about Conditionalization *per se*. In particular, if we set things up this way, then we can see this as the complaint that the core Bayesian framework doesn't have the resources to provide an adequate account of confirmation. And we can see the fourth proposal I describe below as a way of modifying the core Bayesian account (by replacing Conditionalization with Ur-Prior Conditionalization) in a way that imbues it with the resources we need to construct such an account.

E as evidence (call this $cr_{cf(\neg E)}$) instead of her actual credences, cr .³¹ So, on this account, E supports H for a subject iff $cr_{cf(\neg E)}^{+E}(H) > cr_{cf(\neg E)}(H)$. This response faces a number of obstacles, however.³² For one, $cr_{cf(\neg E)}$ seems underdetermined, since there are a number of different ways to evaluate this counterfactual. For another, there are natural ways of evaluating such counterfactuals that seem to yield the wrong results. To borrow an example from Maher (1996), consider an author who knows her work is popular, and who believes that the popularity of her work supports the claim that her work is important, but who would not believe that popularity is evidence for importance if she did not know her work was popular. The counterfactual account entails that, given the author's beliefs, popularity is not evidence for importance. But this is the wrong result.

Another response is to argue that we should assess claims about what E supports using the subject's credences before they got E as evidence.³³ Doing so would allow old evidence to support hypotheses, since a subject's credences cr^* before they got E as evidence could be such that $cr^{*+E}(H) > cr^*(H)$. But this proposal also faces worries.³⁴ For one, this proposal entails that the order in which a subject receives evidence can bear on what hypotheses that evidence supports. Consider the hypothesis $H = E \vee F$, and consider a subject who initially has some middling credence in H . If that subject receives E and then F as evidence, this proposal will entail that E but not F supports H . Whereas if the subject receives F and then E as evidence, this proposal will entail that F but not E supports H . But these seem like the wrong results: both E and F should support H , regardless of the order in which they were received. Another worry for this proposal is that questions of support often seem to presuppose some subset of the subject's evidence, with the evidence being presupposed varying from context to context. And since this proposal makes no appeal to context, it can't explain this phenomenon.³⁵ A final worry for this proposal is that it can't handle subjects who come into existence with non-trivial (i.e., non-tautological) evidence. Since there is no credence function the subject had before she received such evidence, this proposal can't assess what such evidence supports. And there are a number of accounts of evidence which allow subjects to come into existence with non-trivial evidence. For example, the popular account endorsed by Howson and Urbach (1993), in which a subject's evidence is the logically strongest proposition in which she has a credence of 1, allows subjects to come into existence with non-trivial evidence. Likewise, the account sketched by Lewis (2010), in which

31. For defenses of this suggestion, see Garber (1983) and Howson (1984).

32. See Christensen (1999).

33. See Eells (1985) and Earman (1992) for proponents.

34. See Christensen (1999).

35. See Maher (1996) for a discussion of the role of context, and Eells and Fitelson (2000) for an account which notes the importance of relativizing assessments of support to what background evidence is being presupposed.

a subject's evidence corresponds to the set of possibilities in which her perceptual experience and memory are the same as they actually are, allows subjects to come into existence with non-trivial evidence.

A third response follows the second in appealing to historical credences, but does so in a more nuanced manner. Following Eells and Fitelson (2000), one might argue that we should assess claims about whether E supports H using a credence function that we construct in the following way.³⁶ First, we select some credence function cr_0 the subject had before she received any evidence relevant to the hypothesis H .³⁷ Second, take the context in which the question of confirmation is being assessed to pick out some (possibly empty) subset B of the evidence relevant to H that the subject has received—this is the background evidence that the question is presupposing—and update cr_0 on B . Then E supports H (for this subject, in this context) iff $cr_0^{+B \wedge E}(H) > cr_0^{+B}(H)$.³⁸

This response almost gets us what we want; it avoids the order-dependence and contextual worries facing the second response, and it yields plausible verdicts in almost every case. But unfortunately it's still subject to a version of the third worry facing the second response. Namely, if a subject came into existence with some evidence relevant to H , then there will be no pre-relevant-evidence credence function cr_0 for us to use to assess claims regarding what evidence supports H . So, like the second response, it has difficulties handling old evidence if a subject has had such evidence from the moment she came into existence.

These difficulties might motivate a different kind of reply to these worries. One might adopt something like the third response, but make one key change: replace the appeal to cr_0 —some early credence function the subject has before she's received any evidence relevant to the hypothesis—with an appeal to some 'ur-prior' function up that encodes how a subject ought to respond to evidence.³⁹ Thus, on this account, E supports H (for a subject, in a context) iff $up^{+B \wedge E}(H) >$

36. Eells and Fitelson distinguish between two kinds of questions one might ask about confirmation: a historical question and a non-historical question. It's their answer to the (more difficult) non-historical question that we're concerned with here.

37. This requires a way of assessing whether a piece of evidence is relevant to a hypothesis. The most natural way to do this is to evaluate whether the evidence is probabilistically relevant, but this is something that can change as the subject's credences change. The most straightforward way to deal with this is to assess relevance with respect to a subject's initial credence function, though Eells and Fitelson (2000) are not entirely comfortable with this move.

38. I'm using a different notation than Eells and Fitelson, but this is essentially their proposal for how to assess qualitative non-historical confirmation. See also Maher (1996) for another account which recognizes the importance of appealing to the context in which the confirmational question is being asked in order to pick out the background evidence the question is presupposing.

39. Skyrms (1983) also offers an account of confirmation that appeals to something like ur-priors, though his account lacks the flexibility of the Eells and Fitelson (2000)-inspired view presented in the text.

$up^{+B}(H)$. Now, something like up^{+B} isn't well-defined given Conditionalization. Conditionalization tells us how to go from a credence function the subject has at one time, and the evidence she's received since then, to the credence function she should adopt now. But it doesn't know what to do if you give it some up function that doesn't correspond to any credence function the subject has ever actually had. So to finish off this proposal, we need to replace Conditionalization with a rule that can make sense of updating on up , i.e., a version of Ur-Prior Conditionalization. Once we do this, we have a proposal that yields the desired verdicts regarding confirmation.⁴⁰

3.8. Epistemic Continuity

A final complaint about Conditionalization is that it relies on some notion of personal identity or "epistemic continuity".⁴¹ Conditionalization tells us what a subject's current credences should be given what her prior credences were. This requires a way of determining which time slices of doxastic subjects are time slices of the *same* subject. One might appeal to the ordinary notion of personal identity here, or (following Meacham 2010) one might appeal to an *epistemic continuity* relation which tracks the notion of 'same subject' that's relevant to diachronic credence constraints like Conditionalization. But either way, providing a precise characterization of Conditionalization requires spelling out what the relevant relation is like.

Meacham (2010) takes considerations regarding epistemic continuity to motivate modifying Conditionalization. For example, as standardly presented, Conditionalization can't handle cases of fusion—cases in which one person stage is the successor of two person stages with different credences. Thus he recommends replacing Conditionalization with a different kind of diachronic credence constraint.

Hedden (2015), however, has recently appealed to considerations regarding epistemic continuity to argue for a more radical conclusion: that we should re-

40. Christensen (1999) argues against something like this proposal—the idea that one can appeal to a "pure initial probability function" to resolve these problems. Christensen does so by posing a dilemma: if we appeal to such a function and don't take *any* other evidence into account, we'll often fail to track some of the relevant evidential relations we think should obtain, whereas if we appeal to such a function and take *all* of our other evidence except for E into account, we'll get incorrect results in cases like the one where $H = E \vee F$ (since if our other evidence includes F , this proposal will tell us that E doesn't support H). But as the fourth response shows, this is a false dilemma; there are other options regarding what evidence to take into account besides the two Christensen considers. In particular, following Maher (1996) and Eells and Fitelson (2000), we can take all and only the contextually determined background evidence into account. Doing so avoids the worries Christensen raises.

41. See Hedden (2015).

ject *any* kind of diachronic credence constraint.⁴² Hedden argues that we should be skeptical of the possibility of providing a satisfactory account of epistemic continuity. And since any diachronic credence constraint like Conditionalization has to appeal to something like epistemic continuity, it follows that we should be skeptical of any rule of this kind.

One response to Hedden's worry is to meet the challenge head on, and to develop an account of epistemic continuity to pair with diachronic credence constraints. But those persuaded by Hedden's arguments are unlikely to be satisfied with this response. A more radical response is to follow Hedden and give up on diachronic credence constraints. The most natural way to do this is to replace Conditionalization with a synchronic rule. For example, one might adopt a version of Ur-Prior Conditionalization whose prescriptions are determined wholly by features of the subject's current state, like her current evidence and what hypotheses she currently finds inherently plausible. Such a rule won't require a notion of epistemic continuity. And thus it will avoid the sorts of worries Hedden raises.

4. Ur-Prior Conditionalization

In the last section we saw a number of complaints about Conditionalization. And we saw that one way of addressing these complaints was to replace Conditionalization with a different kind of rule, a rule that appeals to how a subject ought to respond to evidence, or to what hypotheses a subject finds inherently plausible, or something like that—i.e., a rule that makes prescriptions by appealing to something like ur-priors. Thus replacing Conditionalization with something like Ur-Prior Conditionalization seems like a natural way to respond to these complaints.

In Section 3, I sketched Ur-Prior Conditionalization as a rule that prescribes credences equal to a subject's 'ur-priors' conditional on her current evidence. But this sketch leaves open a number of details regarding both the logical form of the rule and the notion of 'ur-priors' it employs. And whether Ur-Prior Conditionalization helps to address the complaints considered in the last section depends on how these details are filled in.

So how should we understand Ur-Prior Conditionalization? In what follows I'll consider six versions of Ur-Prior Conditionalization, which vary along two dimensions. First, they vary with respect to what kind of thing they take the ur-prior function to be. I'll consider three options here: (i) the ur-prior function

42. See Moss (in press) for another proponent of this conclusion. Meacham (2010) also considers (and is agnostic about) the possibility of giving up diachronic credence constraints. But both Moss and Meacham are motivated by different kinds of considerations than Hedden. For a defense of diachronic credence constraints, see Carr (2016).

is a subject's initial credence function, (ii) an ur-prior function can be any function which bears the right relations to a subject's credences and evidence over time, and (iii) the ur-prior function is a function representing the subject's *evidential standards*.⁴³ Second, they vary with respect to whether the rule requires ur-priors to satisfy some further normative constraints. Thus I'll consider both (a) *normatively bare* accounts of ur-priors, which don't require ur-prior functions to satisfy any further normative constraints, and (b) *normatively loaded* accounts of ur-priors, which do require ur-prior functions to satisfy further normative constraints.

In the rest of this section, I'll spell out these six versions of Ur-Prior Conditionalization precisely, and lay out the commitments and demerits that each incurs. In the next section, I'll assess how well each of these versions does with respect to the complaints discussed in Section 3.

4.1. Initial Credence Versions

The first two versions of Ur-Prior Conditionalization I'll consider take a subject's ur-priors to be her initial credences.⁴⁴

Let's start with a normatively bare initial credence account, which imposes no further constraints on what one's initial credences must be like. To formulate the account precisely, we need to introduce some vocabulary.

First, I'll employ some restricted variables. I'll take the variable s to range over doxastic subjects, the variable t to range over times, the variable f to range over functions from (at least some) propositions to real numbers, and the variable A to range over propositions.

Second, we need an obligation operator, O . Following various authors in the literature, I'll take this operator to be subject and time indexed, allowing obligations to vary with subject and time.⁴⁵ So $O_{s,t}(A)$ holds *iff* for s , at t , it's obligatory that A . In some cases, when one or both of the indices don't appear in the scope of the operator and don't play a substantive role in characterizing the content of the norm, I'll simplify by omitting the index and the universal quantifier ranging over it. Thus, for example, I might write $\forall s O_s(\exists t B(s, t))$ as shorthand for $\forall s \forall t' O_{s,t'}(\exists t B(s, t))$.

43. It might seem like I've omitted a fourth understanding of ur-priors suggested in the literature: that a subject's ur-priors are the credences she ought to have if she were to have no evidence (see Meacham 2008). But given Ur-Prior Conditionalization, all of the understandings of ur-priors I discuss will be such that a subject with no evidence should have credences equal to her ur-priors. So the understandings of ur-priors I consider are all just more fine-grained versions of this proposal.

44. For a discussion of Ur-Prior Conditionalization that employs something like this understanding of ur-priors, see Skyrms (1983).

45. See Feldman (1986) for a sustained argument in favor of indexing obligation operators to time and subject.

Third, we need a predicate which characterizes a subject's evidence. Let the *evidence* predicate $E(s, t, A)$ represent a 3-place relation between a subject s , a time t , and a proposition A . $E(s, t, A)$ holds *iff* A is the content of subject s 's evidence at t .

Fourth, we need a predicate which characterizes a subject's initial credences. Let the *initial credence* predicate $IC(s, f)$ to represent a 2-place relation between a subject s , and a function f . $IC(s, f)$ holds *iff* f is s 's initial credence function.

Fifth, we need a predicate that tracks whether conditionalizing a subject's credence on their evidence yields a well-defined function. Let the *well-defined* predicate $W(f, A)$ represent a 2-place relation between a function f and a proposition A . $W(f, A)$ holds *iff* $f(\cdot \mid A)$ is well-defined.⁴⁶

Sixth, we need a predicate which characterizes a subject's credences. Let the *credence* predicate $C(s, t, f)$ represent a 3-place relation between a subject s , a time t and a function f . $C(s, t, f)$ holds *iff* f is s 's credence function at t .

With this in hand, we can formulate the bare initial credence version of Ur-Prior Conditionalization as follows:^{47,48}

Ur-Prior Conditionalization (bare initial credence):

$$\forall s \forall f \forall t \forall A O_{s,t} \left((E(s, t, A) \wedge IC(s, f) \wedge W(f, A)) \rightarrow C(s, t, f(\cdot \mid A)) \right)$$

Roughly, this rule says that it ought to be the case that if a subject's evidence is A , her initial credences were f , and $f(\cdot \mid A)$ is well-defined, then her credences

46. There are three ways in which $f(\cdot \mid A)$ might fail to be well-defined. First, f might fail to be a probability function, making conditional probabilities like $f(\cdot \mid A)$ undefined. Second, f might be a probability function that doesn't have A in the algebra of propositions over which it's defined, making $f(A)$, and thus $f(\cdot \mid A)$, undefined. Third, f might be a probability function that assigns 0 to A , making $f(\cdot \mid A) = \frac{f(\cdot \wedge A)}{f(A)}$ undefined.

47. To avoid notational clutter, I employ the convention of leaving the parentheses indicating the scope of the quantifiers implicit, and taking the scope of the quantifiers to be everything to their right. (The parentheses around $E \wedge IC \wedge W$ are to ensure that the formula is read as having the form $(E \wedge IC \wedge W) \rightarrow C$, not $E \wedge IC \wedge (W \rightarrow C)$.) (I thank an anonymous referee for flagging these potential confusions.)

48. One question that arises here is whether we should insert W into the antecedent or consequent of the conditional. If we insert $W(f, A)$ as part of the consequent of the conditional, then subjects whose initial credences f are such that $f(\cdot \mid A)$ is undefined will violate the rule. If we insert $W(f, A)$ as part of the antecedent of the conditional, then subjects whose initial credences f are such that $f(\cdot \mid A)$ is undefined will trivially satisfy the rule (since the antecedent of the conditional will be false, and thus the conditional true). I pursue the latter option in the text because when we look at the loaded versions of the rule, there are strong reasons to want W in the antecedent (e.g., if W appears in the consequent, then some of the loaded initial credence formulations discussed in Section 4.1.1 become impossible to satisfy). And since we'd like to have the bare and loaded formulations be as similar as possible, I've put W in the antecedent of the bare formulation as well.

are $f(\cdot \mid A)$.⁴⁹

Now let's look at a normatively loaded initial credence account, which does impose further constraints on what one's initial credences can be like. Let $\forall_R$ and $\exists_R$ be restricted quantifiers that range over only the rational f s; that is, only those f s which, when plugged into the rule, yield credences that satisfy whatever further normative constraints one takes to constrain belief (e.g., the Principal Principle, Reflection, etc.). At a first pass, we might try to formulate the loaded initial credence version of Ur-Prior Conditionalization by taking the bare initial credence version and replacing $\forall f$ with $\forall_R f$, as follows:

$$\forall s \forall_R f \forall t \forall A O_{s,t} \left((E(s, t, A) \wedge IC(s, f) \wedge W(f, A)) \rightarrow C(s, t, f(\cdot \mid A)) \right)$$

But this rule won't deliver the desired prescriptions to subjects whose initial credences are irrational. Subjects with irrational initial credences will automatically satisfy this norm—for such subjects, $IC(s, f)$ won't hold for any rational f , making the antecedent of the conditional false, and thus the conditional true. And we'd like the rule to tell irrational subjects to adopt credences which correct their errors, not to tell them they don't need to do anything at all (because they automatically satisfy the norm).

At a second pass, one might instead require subjects to just update in accordance with *some* rational function; that is, one might adopt a rule like this:

$$\forall s \forall t O_{s,t} \left((\forall f \exists_R f' \forall A (E(s, t, A) \wedge IC(s, f) \wedge W(f', A)) \rightarrow C(s, t, f'(\cdot \mid A))) \right)$$

But on this version of the rule the f' the subject updates on needn't bear any resemblance to the f representing her initial credences; it's as viable for her to use a rational function which is radically different from her initial credences as one which is a slight correction of her initial credences.

In order to get a normatively loaded account in which a subject's initial credences play a role, we need the rule to prescribe updates based on the subject's

49. If we take Conditionalization, plug a subject's initial credences in for cr , and take E to be the total cumulative evidence the subject's received, then we get the following as a special case of Conditionalization:

IC-Conditionalization: If a subject has initial credences ic , and E is her total cumulative evidence, then her credences cr should be $cr(\cdot) = ic(\cdot \mid E)$, if defined.

The bare initial credence version of Ur-Prior Conditionalization closely resembles IC-Conditionalization. Both rules tell subjects to adopt credence equal to their initial credences conditional on their evidence. The main difference between the two rules comes from the notion of evidence they appeal to. IC-Conditionalization tells subjects to conditionalize on their total cumulative evidence—the conjunction of all of the evidence they've ever received. The bare initial credence version of Ur-Prior Conditionalization, by contrast, only tells subjects to conditionalize on their current evidence.

initial credences. That is, we need the rule to require subjects to update using something like the closest rational approximation, or closest rational fit, to their initial credences (or *one* of the closest rational fits, if there are ties). As it turns out, formulating such a rule in a satisfactory manner is difficult. This is because of what I'll call the *Quantifier Problem*: there are three compelling desiderata regarding where to place the quantifiers when formulating the rule that are jointly unsatisfiable. Thus every way of formulating such a rule will incur some costs. In the following section, I'll look more carefully at how the Quantifier Problem arises, and what the options for addressing it are.

4.1.1. The Quantifier Problem

Let $\exists_{R(f)}$ be a restricted quantifier that ranges over functions that are the best rational fits to f . Here is one natural way to formulate a normatively loaded initial credence version of Ur-Prior Conditionalization:

Ur-Prior Conditionalization (loaded initial credence (v1)):

$$\forall s \forall t O_{s,t} \left(\forall f \exists_{R(f)} f' \forall A (E(s, t, A) \wedge IC(s, f) \wedge W(f', A)) \rightarrow C(s, t, f'(\cdot | A)) \right)$$

Roughly, this rule says that at every time there ought to be some best rational fit f' to a subject's initial credences such that, if her evidence at that time is A and $f'(\cdot | A)$ is well-defined, then her credences at that time are $f'(\cdot | A)$.

The drawback of this rule is that, in cases in which there are ties for best rational fit, the rule doesn't require the kind of diachronic stability that one might like. This rule requires only that a subject's credences at a time line up with some best rational fit. But it doesn't require them to line up with the *same* best rational fit at different times. So this rule would allow one to arbitrarily flip back and forth between different best rational fits.

The problem arises because $\forall t$ comes before $\exists_{R(f)} f'$. Thus the rule requires that, at all times, there be some best rational fit that the subject's credences line up with. What we want, however, is that there be some best rational fit that the subject's credences line up with at all times.

We can address this problem by moving the quantifier over times, as follows:

Ur-Prior Conditionalization (loaded initial credence (v2)):

$$\forall s \forall t O_{s,t} \left(\forall f \exists_{R(f)} f' \forall t' \forall A (E(s, t', A) \wedge IC(s, f) \wedge W(f', A)) \rightarrow C(s, t', f'(\cdot | A)) \right)$$

Roughly, this rule says that at every time t , there ought to be some best rational fit f' to a subject's initial credences such that, at all times t' , if her evidence at t' is A and $f'(\cdot | A)$ is well-defined, then her credences at t' are $f'(\cdot | A)$.

The drawback of this rule is that it can't handle the kinds of error correction discussed in Sections 3.1 and 3.2.⁵⁰ The reasons are the same for both updating errors and auxiliary norm errors, so to simplify things I'll focus on updating errors. This rule directs a subject to line up her credences with some particular best rational fit at all times. But if she fails to do so at any time, the constraint becomes impossible to satisfy—she'll now fail to satisfy the rule at all times, no matter what credences she adopts. Thus this rule won't tell subjects who have made updating errors to adopt credences which *correct* the errors they've made. Instead, it will tell them that there's nothing to be done (because no matter what credences they adopt, they'll fail to satisfy the norm).

The problem arises because the quantifier $\forall t'$, which binds the times at which the subject is directed to adopt her credences, appears inside the obligation operator. To provide the desired prescriptions to subjects who have made errors, the time indexing O must be the same as the time at which the subject is directed to adopt her credences. (If these two times are the same, then whether the subject satisfies $O_{s,t}$ is determined by her credences at t . So even if she fails to satisfy the rule at t_1 because she has the wrong credences at t_1 , she can still satisfy the rule at t_2 by adopting the right credences at t_2 .) But in order for the same quantifier to bind both the time indexing O and the time at which she's directed to adopt her credences, it must appear outside of the obligation operator.

We can address this problem by moving all of the quantifiers outside of the scope of the obligation operator, as follows:

Ur-Prior Conditionalization (loaded initial credence (v3)):

$$\forall s \forall f \exists_{R(f)} f' \forall t \forall A O_{s,t} \left((E(s, t, A) \wedge IC(s, f) \wedge W(f', A)) \rightarrow C(s, t, f'(\cdot | A)) \right)$$

Roughly, this rule says that there is some best rational fit f' to a subject's initial credences such that, at all times, it ought to be the case that if her evidence at that time is A and $f'(\cdot | A)$ is well-defined, then her credences at that time are $f'(\cdot | A)$.

The drawback of this rule is that, in cases in which there are ties for best rational fit, the rule requires us to essentially privilege one of these ties over the others.⁵¹ The problem arises because $\exists_{R(f)} f'$ appears outside of the scope of the

50. By and large, I'm reserving assessments of how these different versions of Ur-Prior Conditionalization deal with the complaints raised in Section 3 until Section 5. But these issues pop up early here, because the question of whether these rules can handle error correction comes up in trying to figure out how to best *formulate* the loaded initial credence version of Ur-Prior Conditionalization.

51. I'm assuming here that there are no 'epistemic dilemmas'—that epistemic rationality can't impose inconsistent prescriptions. If it could, then this rule wouldn't necessarily require us to privilege one of the best fits—for it could be that a subject is simply required to line up their credences with *every* best fit. (I thank an anonymous referee for flagging this assumption.)

obligation operator. Let's think of the obligation operator as analogous to the necessity operator, where $O(P)$ is true *iff* P is true at all of the best accessible worlds. Because $\exists_{R(f)}f'$ appears outside of O , the rule requires there to be some f' that the subject's credences line up with at all of the best accessible worlds. Thus it requires there to be some particular privileged f' that the subject's credences must line up with. Whereas what we want is a rule where $\exists_{R(f)}f'$ appears inside of O , and thus merely requires that at each best accessible world there be some f' that the subject's credences line up with. So it doesn't matter which of the f 's the subject's credences line up with, so long as it lines up with one of them.

This 'privileging' feature of the rule is unpalatable, because it effectively requires one to deny that there are ties for best rational fit. For while there may be ties for best fit in one sense, there will always be a tie-breaker that determines which of these ties actually imposes normative constraints. And this is an uncomfortable commitment, for it's hard to see what these tie-breakers could be.

We can address this problem by moving the $\exists_{R(f)}f'$ quantifier inside the scope of the obligation operator, but this just returns us to one of the first two versions of the rule. More abstractly, we can see the difficulty with formulating a satisfactory loaded initial credence account as arising from three desiderata that are jointly unsatisfiable:⁵²

1. We want $\exists_{R(f)}f'$ to come before $\forall t$ (to get the appropriate diachronic coherence in cases of ties).
2. We want $\forall t$ to be outside of O (to allow for error corrections).
3. We want $\exists_{R(f)}f'$ to be inside of O (in order to not privilege any best rational fit in cases of ties).

Each pair of these desiderata entails the negation of the third. If $\exists_{R(f)}f'$ comes before $\forall t$ (1), and $\forall t$ is outside of O (2), then $\exists_{R(f)}f'$ must be outside of O ($\neg 3$). If $\exists_{R(f)}f'$ comes before $\forall t$ (1), and $\exists_{R(f)}f'$ is inside O (3), then $\forall t$ must be inside of O ($\neg 2$). And if $\forall t$ is outside of O (2), and $\exists_{R(f)}f'$ is inside O (3), then $\exists_{R(f)}f'$ must come after $\forall t$ ($\neg 1$). This is the Quantifier Problem.

I'm not sure what the best response to the Quantifier Problem is. I'm inclined to think that we should give up either the first desideratum (as (v1) does) or the third desideratum (as (v3) does). So in what follows, I'll restrict my attention to these two versions of the loaded initial credence account. But I'll leave it open which of these two versions we should favor. In any case, the Quantifier Problem will be one of the demerits we need to take into account when assessing the overall attractiveness of these views.⁵³

52. Where " $\forall t$ " is the quantifier binding the times at which the subject is directed to adopt her credences.

53. It's worth noting that proponents of Evidential Uniqueness (the view that every batch of

4.1.2. Worries

Now let's turn to canvass the demerits facing the bare and loaded initial credence accounts. I'll flag four worries for these accounts here, two which apply to the bare account, and two which apply to the loaded account.

The first worry specific to the bare account is that it won't apply to subjects who don't have initial credences. So, for example, consider immortal subjects who have existed forever, subjects existing along closed time-like paths, or subjects who were created in open temporal intervals. None of these subjects will have initial credence functions, and so the bare account won't impose constraints on their beliefs. But, it seems, the way in which they update their beliefs should still be subject to rational constraints.

(By contrast, the loaded account has a built-in response to these concerns. The loaded account tells subjects to line up their credences with some best rational fit to their initial credences. For subjects without pre-existing initial credences, this procedure simply becomes very permissive: since the subject has no initial credences, every rational function is an equally good fit. Thus the loaded account has no difficulty providing substantive prescriptions to subjects without initial credences.)

A second worry specific to the bare account is that it loses its normative force when applied to subjects with non-probabilistic initial credences f . Since f needs to be probabilistic in order to yield well-defined conditional probabilities, $f(\cdot \mid A)$ will be undefined and thus $W(f, A)$ will be false. This makes the antecedent of the conditional false, and thus the conditional itself trivially true. So the bare initial credence account won't impose any constraints on what the beliefs of such subjects should be like. (This worry doesn't arise for the loaded initial credence account because it's not the subject's initial credences f , but the rational fixed-up f' , which needs to be probabilistic in order for the rule to provide well-defined prescriptions. And given Probabilism, subjects are rationally required to have probabilistic credences.)

A third worry specific to the loaded initial credence account, at least as it's been described so far, is that it's incomplete. In order to complete this proposal, one needs to provide a detailed account of how to assess what the best rational fit to a function is. And this is a non-trivial task—one that will depend on what one takes the rational constraints, and their relative importance, to be.

A fourth worry specific to the loaded account is that it faces the Quantifier Problem. As a result, however we formulate it, it seems to incur some further demerit. It either fails to provide the desired diachronic coherence in cases of

evidence determines a unique rational doxastic state) have a straightforward way of responding to the Quantifier Problem. Since the worries facing the (v3) formulation won't apply to them, they can simply adopt the (v3) formulation (see Section 6).

ties, fails to allow for substantive error corrections, or requires us to privilege some best rational fit in cases of ties.

4.2. Functional Versions

Next, let's consider a pair of accounts which adopt a "functional role" understanding of ur-priors.⁵⁴ In some ways this is the trickiest understanding of ur-priors, so let's start by getting a feel for the difference between functional and non-functional understandings.

One could understand Ur-Prior Conditionalization's appeal to ur-priors in two kinds of ways. First, one could take 'ur-priors' to pick out some particular feature of the subject—say, her initial credences. So, on this understanding of the Ur-Prior Conditionalization, the norm requires a subject's credences to equal her initial credences conditional on her evidence. Alternatively, one could understand the norm's appeal to 'ur-priors' functionally. On this understanding, the norm just requires that there exist some function—call it the *ur-prior function*—that plays a certain role with respect to a subject's beliefs. In particular, the norm requires that there exist some function, *up*, such that her credences at every time are equal to *up* conditional on her evidence at that time.

On this functional understanding of the norm, there isn't some particular 'ur-prior' function that a subject is required to line up her beliefs with. Rather, it's just required that there exist some function that her beliefs line up with at every time. And the term 'ur-priors' picks out any function that plays this role. (So a subject may have no ur-prior function, if no function lines up with her beliefs over time in the right way. And a subject may have many ur-prior functions, if many functions line up with her beliefs over time in the right way.⁵⁵)

Let's look at how to formulate this kind of functional understanding of Ur-Prior Conditionalization more precisely. Let's start with a normatively bare functional account, which imposes no further normative constraints on what *up* must be like. We can formulate the bare functional version of Ur-Prior Conditionalization as follows:

Ur-Prior Conditionalization (bare functional):

$$\forall s O_s (\exists f \forall t \forall A (E(s, t, A) \rightarrow (W(f, A) \wedge C(s, t, f(\cdot | A))))$$

Roughly, this rule says that for each subject, there ought to be a function *f* such that whenever her evidence is *A*, she has (well-defined) credences $f(\cdot | A)$.⁵⁶

54. For discussions regarding pairing something like Ur-Prior Conditionalization with this understanding of ur-priors, see Meacham (2007) and Titelbaum (2013).

55. For example, this will happen if there's a world to which a subject assigns a 0 credence at all times.

56. Again, we need to be careful about choosing where to insert *W*. In Footnote 48, we saw

This version of Ur-Prior Conditionalization is normatively bare, in that it doesn't impose any further normative constraints on f other than that it be such that it can line up with the subject's credences and evidence at various times. Now let's turn to look at a normatively loaded version of the functional account, one which imposes further constraints on what the f the subject is told to update on should be like. Imposing such constraints is much easier here than it was for the initial credence accounts. Here all we have to do is require f to be rational, i.e., replace $\forall f$ with $\forall_R f$. Thus we can formulate a loaded functional version of Ur-Prior Conditionalization as follows:

Ur-Prior Conditionalization (loaded functional):

$$\forall s O_s \left(\exists_R f \forall t \forall A (E(s, t, A) \rightarrow (W(f, A) \wedge C(s, t, f(\cdot | A)))) \right)$$

Roughly, this rule says that for each subject, there ought to be a rational function f such that whenever her evidence is A , she has (well-defined) credences $f(\cdot | A)$.

With respect to the commitments they incur, these functional versions of Ur-Prior Conditionalization are the most attractive of the six versions we'll consider. For these functional accounts have no more commitments than the standard Bayesian framework. They require subjects to have credences, and to have evidence—albeit evidence at a time instead of cumulative evidence—but these kinds of commitments are familiar in the standard Bayesian framework. The loaded version appeals to some further normative constraints on credences, but these constraints are again familiar to the Bayesian. So in this respect, these functional accounts offer an appealing way to understand Ur-Prior Conditionalization.

4.3. Evidential Standards Versions

Finally, let's consider a pair of accounts which take a subject's ur-priors to be what I'll call *evidential standards*.⁵⁷ People have used the term 'evidential stan-

that for the initial credence versions, we want W to appear in the antecedent of the conditional. But if we insert $W(f, A)$ as part of the antecedent of the conditional here, we get a rule that is too easy to satisfy, for there will always be some f that will make $W(f, A)$ false, and thus the conditional true. Alternatively, if we insert $W(f, A)$ as a separate clause inside O , alongside the conditional, we get a rule that is too hard to satisfy, since satisfying the rule will require $W(f, A)$ to be true for all A ; which no probabilistic f can do (when $A = \perp$, for example). So the right place to put W is in the consequent of the conditional.

(In Footnote 48 I noted that we have a *prima facie* reason to put W in the antecedent of the bare initial credence formulation, because putting W in the consequent of the loaded initial credence formulations is problematic, and we'd like the bare and loaded formulations to be similar. But this rationale doesn't apply here, because there's nothing problematic about putting W in the consequent of the loaded functional formulation.)

57. For discussions regarding pairing something like Ur-Prior Conditionalization with this understanding of ur-priors, see Meacham (2007). For proponents of something like this

dards' to pick out different kinds of things.⁵⁸ So the 'evidential standards' understanding of ur-priors isn't a specific proposal, but rather a cluster of related proposals, with the following common feature: on all of these views, a subject's evidential standards supervene on her mental state at a time. Thus, in contrast to the functional and initial credence understandings of ur-priors, evidential standards are the kinds of things which can (potentially) change over time.

One can develop this kind of proposal in a number of ways. One rough thought is that evidential standards represent what a subject finds inherently plausible, independently of any evidence she's received. Another rough thought is that evidential standards should be understood as analogous to moral standards—something which the subject is, in some sense, committed to or an advocate of.⁵⁹ And each of these rough thoughts is compatible with a number of different concrete proposals. For example, one might take evidential standards to be determined by a subject's dispositions to adopt new beliefs, a subject's beliefs about what beliefs she should adopt, a subject's intentions about what beliefs to adopt, some mixture of these factors, or even some primitive propositional attitude. We'll return to look at some of these proposals more carefully below.

In order to formulate such rules, we need to introduce a predicate which characterizes a subject's evidential standards. I'll take $ES(s, t, f)$ to represent a 3-place relation between a subject s , a time t , and a function f from propositions to real numbers. (Unlike IC , ES is relativized to times, since evidential standards, like any other feature of a subject's current psychological state, might change over time.) $ES(s, t, f)$ holds *iff* f correctly represents the relevant features of s 's mental state at t .

With this in hand, we can formulate the desired evidential standards versions of Ur-Prior Conditionalization by simply taking the corresponding initial credence formulation of the rule and substituting $ES(s, t, f)$ for $IC(s, f)$. This will yield a normatively bare evidential standards account similar to the bare initial credence account from Section 4.1:

Ur-Prior Conditionalization (bare evidential standards):

$$\forall s \forall f \forall t \forall A O_{s,t} \left((E(s, t, A) \wedge ES(s, t, f) \wedge W(f, A)) \rightarrow C(s, t, f(\cdot | A)) \right)$$

And it will yield three versions of the loaded evidential standards account that mirror the three versions of the loaded initial credences account. As before, I'll

understanding of ur-priors, independent of Ur-Prior Conditionalization, see Schoenfield (2014) and Titelbaum (2016).

58. For example, see Schoenfield (2014). While many of these uses are compatible with how I'll use the term, some are not. For example, some have used the term 'evidential standards' to refer to a subject's initial credences, which is incompatible with how I'm using the term here.

59. See Titelbaum (2016).

restrict my attention to the first and the third versions of this account in what follows:

Ur-Prior Conditionalization (loaded evidential standards (v1)):

$$\forall s \forall t O_{s,t} \left(\forall f \exists_{R(f)} f' \forall A (E(s, t, A) \wedge ES(s, t, f) \wedge W(f', A)) \rightarrow C(s, t, f'(\cdot | A)) \right)$$

Ur-Prior Conditionalization (loaded evidential standards (v3)):

$$\forall s \forall f \exists_{R(f)} f' \forall t \forall A O_{s,t} \left((E(s, t, A) \wedge ES(s, t, f) \wedge W(f', A)) \rightarrow C(s, t, f'(\cdot | A)) \right)$$

The (v1) formulation says, roughly, that at every time there ought to be some best rational fit f' to a subject's evidential standards such that, if her evidence at that time is A and $f'(\cdot | A)$ is well-defined, then her credences at that time are $f'(\cdot | A)$. The (v3) formulation says, roughly, that there exists some best rational fit f' to a subject's evidential standards such that, at all times, it ought to be the case that if her evidence at that time is A and $f'(\cdot | A)$ is well-defined, then her credences at that time are $f'(\cdot | A)$.

As we saw above, there are a number of different ways to cash out evidential standards. Let's briefly look at some of these possibilities, before turning to look at some worries facing evidential standards accounts.

One possibility is that evidential standards encode a subject's doxastic dispositions. In particular, a subject's evidential standards might represent the credences the subject is disposed to adopt given various batches of evidence (those credences being equal to her evidential standards conditional on that evidence). This understanding of evidential standards trivializes Ur-Prior Conditionalization if we adopt a normatively bare account; for then the rule merely says that subjects should adopt the credences they're disposed to adopt. But the rule remains non-trivial if we adopt a normatively loaded account, on which the rule says that subjects should adopt the credences suggested by the closest rational fit to their doxastic dispositions.⁶⁰

60. One might also consider a natural variant of this view, according to which a subject's evidential standards represent the credences she's disposed to adopt given no evidence, but say nothing about her credal dispositions in other evidential situations. This view is attractive because it doesn't require evidential standards to encode as much information. Furthermore, on this understanding of evidential standards, the normatively bare account no longer trivializes Ur-Prior Conditionalization.

The main worry facing this proposal is that there are a number of accounts of evidence according to which it's difficult, if not impossible, for a subject to have no evidence. For example, suppose we follow Lewis (1996) in understanding the content of a subject's evidence as the set of centered worlds centered on individuals-at-a-time whose perceptual experiences and memories are the same as her own. On this account of evidence, subjects always have non-trivial evidence. So given this account of evidence, the proposed understanding of evidential

Another possibility is that a subject's evidential standards encode a subject's normative beliefs.⁶¹ In particular, a subject's evidential standards might represent the credences she believes she should adopt given various batches of evidence (where her evidential standards conditional on the evidence yield those credences). Or, if a subject is unsure about what her credences should be given some evidence E , we can take the expectation of the credences the subject believes she should adopt given E , and require her evidential standards to yield that expectation when conditionalized on E .⁶²

A third possibility is that a subject's evidential standards encode her doxastic intentions. In particular, a subject's evidential standards might represent the credences the subject intends to adopt given various batches of evidence (those credences being equal to her evidential standards conditional on that evidence).

There are also some more open-ended possibilities. For example, one might take a subject's evidential standards to encode some mixture of the first three possibilities. Or one might take a subject's evidential standards to represent some primitive propositional attitude that's not characterizable in terms of a subject's other attitudes.

Now let's turn to look at some demerits of these evidential standards versions of Ur-Prior Conditionalization. These evidential standards accounts employ a more substantive notion of ur-priors than the functional or initial credence versions. As we'll see in Section 5, this allows them to handle the broadest range of complaints about Conditionalization. But it also leaves them open to a broader range of worries. I'll flag five worries for these accounts here: one which applies equally to both bare and loaded accounts, two which only apply to the bare account, and two which only apply to the loaded account.

The first worry, which applies to both the bare and loaded accounts, arises from the fact that evidential standards can change over time. Ur-Prior Conditionalization is supposed to tell us how a subject should change her beliefs over time. And if evidential standards can change, then part of that story is telling us how a subject's evidential standards should change over time. But the story

standards doesn't make sense.

(One might consider similar variants of the other understandings of evidential standards we'll look at below, variants on which evidential standards only represent information about cases in which subjects have no evidence. But all of them will face the same worries regarding accounts of evidence according to which it's difficult, if not impossible, for a subject to have no evidence.)

61. A related, though slightly different, possibility is that a subject's evidential standards encode the rules of belief formation that the subject believes to be most truth conducive (see Schoenfield 2014).

62. That is, letting es stand for the subject's evidential standards, cr stand for the subject's credences, and O stand for the obligation operator, and letting i range over the credence functions the subject thinks might be the ones she should adopt: $es(A) = \sum_i cr(\langle O(cr_i(A) = x) \rangle) \cdot x$.

provided here doesn't say anything about the norms that govern evidential standards changes.⁶³

The most straightforward response to this worry is to reject the claim that there are further norms governing changes in evidential standards. For one may find it plausible that one's credences should line up with one's evidential standards, and (on a loaded approach) that there are synchronic constraints on these evidential standards, but yet not be pulled by the thought that there are further diachronic constraints on these evidential standards.⁶⁴ A more concessive response is to grant that there may be further norms governing evidential standards changes, but to note that this doesn't take away from the substance of the evidential standards versions of Ur-Prior Conditionalization. We can get a reasonable grasp on what evidential standards are, and these updating rules tie a subject's credences to her evidential standards. The fact that there may be further parts to the story—stories regarding how her evidential standards change—doesn't show that there's anything wrong with *this* part of the story.

The second worry, specific to the bare account, is that it won't yield substantive prescriptions for doxastic subjects who don't have evidential standards.⁶⁵ And given some understandings of evidential standards (e.g., as representing a subject's normative beliefs), this seems like a real possibility (e.g., there can be subjects who don't have normative beliefs of this kind). One reply would be to argue that this result is unproblematic, because subjects without evidential standards plausibly shouldn't be bound by updating rules like Ur-Prior Conditionalization. The viability of this reply depends on the notion of evidential standards one adopts. E.g., suppose one took evidential standards to be some primitive propositional attitude with no clear relation to belief. It would seem *prima facie* strange to maintain that the beliefs of some subject shouldn't be constrained by updating rules because, though she has beliefs, she doesn't have this belief-independent primitive propositional attitude. (Just as it would be strange to say that a subject's beliefs shouldn't have to satisfy Conditionalization because she doesn't have some *prima facie* unrelated propositional attitude, like fears. Presumably, whether a subject has fears or not has no bearing on whether we should

63. For a discussion of some related issues, regarding how to make rational assessments of subjects who change their minds or have second thoughts, see Moss (2014).

64. For discussion, see Meacham (2010), Hedden (2015), Carr (2016) and Moss (in press).

65. On the bare account, if a subject doesn't have evidential standards, then the antecedent of the rule's conditional will be automatically false (since there will be no f which makes $ES(s, t, f)$ true). And if the antecedent is automatically false, then the conditional will be automatically true, and so such subjects will automatically satisfy the rule, regardless of what credences they adopt.

The loaded account will not be subject to this worry because it tells subjects to line up their credences with some best rational fit to their evidential standards. And for subjects without evidential standards, this procedure simply becomes very permissive: since the subject has no evidential standards, every rational function is an equally good fit.

apply rational constraints to her beliefs!) A second reply would be to argue that doxastic subjects must have evidential standards. Again, the plausibility of this reply depends on the notion of evidential standards one adopts. E.g., it may be tenable to maintain that doxastic subjects must have doxastic dispositions, but it's far less plausible to maintain that doxastic subjects must have doxastic intentions. A third reply would be to concede that subjects without evidential standards aren't bound by Ur-Prior Conditionalization, but are bound by some 'back-up' rule instead—perhaps some functional version of Ur-Prior Conditionalization. But it seems *ad hoc* to bring in a different updating rule for such subjects.

A third (and related) worry, specific to the bare account, is that it won't make substantive prescriptions to subjects whose evidential standards aren't probabilistic.⁶⁶ For example, a subject's normative beliefs might just be a grab bag of rules encoding what sorts of things the subject thinks she should believe, something much less detailed than a probability function. The three potential replies to the first worry are also available here—one might maintain that this result is unproblematic, maintain that evidential standards must be probabilistic, or adopt a back-up rule for subjects with non-probabilistic evidential standards. And, as above, the plausibility of these replies depends on the notion of evidential standards we adopt.

A fourth worry, specific to the loaded account, is that, as it stands, the account described here is incomplete. For a complete description of this account would need to include a metric for assessing how close a fit one function is to another which we could use to assess what the best rational fits to a subject's evidential standards are.

A fifth worry, specific to the loaded account, is the Quantifier Problem. As we saw in Section 4.1, it seems that all of the natural ways of formulating this rule have some undesirable features. They either won't yield the kind of diachronic coherence we'd like (in cases of ties), won't allow for updating error corrections, or will require us to privilege one best rational fit over the others (in cases of ties).

66. On the bare account, if a subject's evidential standards f aren't probabilistic, then the antecedent of the rule's conditional will be automatically false (since if f is non-probabilistic, then $W(f, A)$ will be false). And if the antecedent is automatically false, then such subjects will automatically satisfy the rule, regardless of what credences they adopt.

The loaded account will not be subject to this worry because while a subject's evidential standards may not be probabilistic, the best rational fit to their evidential standards will be. And it's these best rational fits that the rule employs to make prescriptions.

5. Assessing the Complaints

In Section 4 we looked at several ways in which one might develop Ur-Prior Conditionalization, and we examined the costs of adopting each of these views. But a cost/benefit assessment of these views can't focus solely on the costs; it also has to look at the benefits. So let's now turn to explore some of the benefits of adopting each view. In particular, let's consider in more detail how each of these views does at helping us address the slew of complaints about Conditionalization discussed in Section 3.

5.1. Error Correction: Updating Errors

One complaint about Conditionalization was that it can't properly handle cases in which subjects have updated incorrectly. We'd like an updating rule's prescriptions to lead subjects to correct previous updating errors. And Conditionalization doesn't do this.

The functional versions of Ur-Prior Conditionalization are subject to the same complaint. These versions of the rule require a subject's credences at all times to line up with some f . And if she adopts credences at a pair of times which don't line up with some f , then these rule's prescriptions become impossible to satisfy—there's no credence function she can now adopt which will allow her to satisfy these rules. Thus these functional versions of the rule won't make the desired prescriptions to subjects who have made updating errors.

The initial credence and evidential standards versions of Ur-Prior Conditionalization, on the other hand, are not subject to this complaint. These norms tell a subject that her current credences should be equal to her ur-prior (or some best rational fit) conditional on her current evidence. Thus they direct her to *correct* her errors—to adopt the same credences as she would have had if she hadn't made any updating errors.

5.2. Error Correction: Auxiliary Norm Violations

Another complaint about Conditionalization is that it can't adequately handle cases in which subjects have violated auxiliary Bayesian norms, like the Principal Principle. We'd like an updating rule's prescriptions to lead subjects to correct violations of auxiliary norms, or at least to not conflict with such corrections. But Conditionalization will typically direct a subject whose credences violate an auxiliary norm to continue to adopt credences which violate that norm.

The normatively bare versions of Ur-Prior Conditionalization can't avoid this complaint either. In order to direct subjects to satisfy auxiliary norms, the rule's prescriptions need to take these auxiliary norms into account. And only the normatively loaded versions of Ur-Prior Conditionalization make prescriptions

which take auxiliary norms into account.

But adopting a normatively loaded version of Ur-Prior Conditionalization isn't enough. Consider the loaded functional account. This rule requires a subject to have credences which line up with some rational f at all times. As a result, if at any time she adopts credences which don't line up with f —she adopts credences that violate the Principal Principle, for instance—then the rule's prescriptions become impossible for her to satisfy. So the loaded functional version of the rule won't yield the desired prescriptions for subjects who have violated auxiliary norms.

By contrast, the loaded initial credence and evidential standards versions of Ur-Prior Conditionalization can avoid this complaint. These rules tell a subject to adopt credences equal to some best rational fit to her ur-priors conditional on her current evidence. Thus these rules effectively tell her to adopt the credences she would have had if she hadn't violated any auxiliary norms.

5.3. Changing Desiderata of Theory Choice

A third complaint about Conditionalization is that it can't accommodate cases of rational belief change due to changes in a subject's assessment of the relative importance of different desiderata of theory choice.

The functional and initial credence versions of Ur-Prior Conditionalization don't help with this complaint. In order to avoid this complaint, the rule needs to allow the ur-priors that generate the prescriptions to change over time. And on both the functional and initial credence accounts these ur-priors are static. (The functional accounts require a subject's credences to line up with the same f at all times, and so don't allow for changing f s. And the initial credence accounts require a subject to line up her credences with her initial credences, which don't change.)

The evidential standards versions of Ur-Prior Conditionalization, however, can address this complaint. This is because a subject's evidential standards can change over time. A subject may find certain things inherently plausible which, over time, come to seem less inherently plausible. Thus the evidential standards version of the rule can accommodate this kind of belief change.

5.4. Evidence Loss

A fourth complaint about Conditionalization is that it can't adequately handle cases of evidence loss. If a subject learns that a fair coin toss landed heads, and then loses this evidence, it seems her credence in heads should revert to 0.5. But Conditionalization won't allow subjects to go back to their earlier credences in this way.

Every version of Ur-Prior Conditionalization can handle this complaint in

most cases. All of these rules tell subjects to set their credences equal to their ur-priors conditional on their current evidence. It follows that if a subject loses evidence, and so comes to have the same evidence as she had at some earlier time, she will be prescribed the same credences as she was at that earlier time.

That said, if one adopts an account of evidence that allows subjects to have (non-trivial) evidence as soon as they come into existence,⁶⁷ then there will be some cases of evidence loss which initial credence versions of Ur-Prior Conditionalization can't handle. Consider a case in which a subject comes into existence with evidence *E* (e.g., a robot who is turned on for the first time and has *E* programmed into her memory banks). If she then loses that evidence (has her memory banks deleted), the initial credence versions of the rule will prescribe her the same credences as before. But in many cases this seems like the wrong prescription—it should at least be possible for her credences to reflect this evidence loss in some way.

5.5. Self-Locating Beliefs

A fifth complaint about Conditionalization is that it can't handle changes in self-locating beliefs, such as beliefs about what time it is.

The initial credence versions of Ur-Prior Conditionalization won't help with this complaint. Start with the normatively bare initial credence approach. A subject's initial credences can assign a value of 1 to some self-locating belief, such as that it's 9 AM. If such a subject gets new evidence that doesn't rule out the possibility that it's 9 AM, then the bare initial credence approach will require her to continue to have a credence of 1 that it's 9 AM. And if a subject gets new evidence that does rule out the possibility that it's 9 AM, then the bare initial credence approach will fail to impose any constraints on her beliefs at all.⁶⁸ Either way, the bare initial credence approach won't provide a satisfactory account of how the subject should update her self-locating beliefs.

The normatively loaded initial credence approach faces similar problems. It seems an ideally rational subject can have initial credences which assign a value of 1 to some self-locating belief, such as that it's 9 AM. (E.g., consider a rational robot who's just been switched on for the first time, and who has been programmed to start with the correct belief about what time it is.) And the loaded initial credence version of Ur-Prior Conditionalization will yield the same implausible prescriptions as the bare initial credence approach given these kinds of

67. E.g., Lewis's (1996) account of evidence or Howson and Urbach's (1993) account of evidence.

68. In this case, conditionalizing the subject's initial credences on her evidence will fail to yield a well-defined result. So the well-defined clause will make the antecedent of the rule's conditional false, and thus the conditional itself true, regardless of what future credences the subject adopts.

initial credences.

Likewise, it's not clear the bare evidential standards version of Ur-Prior Conditionalization helps with this complaint. After all, a (rather strange) subject might have evidential standards which assign a value of 1 to some self-locating belief, such as that it's 9 AM. And given such evidential standards, the bare evidential standards approach will run into the same problems as the initial credence approaches.

The loaded evidential standards approach, on the other hand, is better positioned to handle this complaint. For one could plausibly *deny* that rational evidential standards will (say) assign a value of 1 to the belief that it's 9 AM. After all, one would not expect a rational robot's evidential standards—say, her beliefs about what beliefs she ought to have, given various batches of evidence—to require her, given *any* batch of evidence, to assign a credence of 1 to it being a particular time. For example, one would not expect a rational robot to believe that her credence that it's 9 AM should remain 1 given the evidence that an accurate clock reads 10 AM. So the loaded evidential standards approach can maintain that rational evidential standards are such that they always yield plausible prescriptions regarding self-locating beliefs.

Finally, both the bare and loaded functional versions of Ur-Prior Conditionalization can handle this complaint. For on the functional accounts, a subject's credences can still satisfy the rule even if they're certain they're watching an accurate clock, and so become certain that it's t_1 at t_1 , and t_2 at t_2 . For all the functional accounts require is that there exist some function which, when conditionalized on her evidence at each of those times, yields the appropriate credences. And there will be plenty of functions which satisfy these constraints.⁶⁹

5.6. *Learning About New Theories*

A sixth complaint about Conditionalization is that it doesn't provide satisfactory prescriptions in cases in which subjects learn about new theories. In particular, there seem to be cases where a subject who learns about some new theory V should come to assign V a non-zero credence, even though she hasn't acquired any new evidence. And Conditionalization won't make such prescriptions.

The functional versions of Ur-Prior Conditionalization can't handle this complaint. If a subject didn't assign a non-zero credence to V before, these rules will forbid her from assigning a non-zero credence to V now (given the same evidence). The initial credence versions of Ur-Prior Conditionalization can't handle

69. Although the functional and evidential standards versions of Ur-Prior Conditionalization can accommodate self-locating beliefs, the rules themselves don't require as much diachronic continuity as some might like. As a result, proponents of such views may want to impose further constraints on rational ur-priors that ensure such diachronic continuity. See Meacham (2008) for a discussion of these issues.

this complaint either. For if a subject's initial credences fail to assign a non-zero credence to V , then these rules will forbid her from assigning a non-zero credence to V at any point in the future.⁷⁰

On the other hand, the evidential standards versions of Ur-Prior Conditionalization can handle this complaint in a couple of ways. One way is to treat this as a case of updating error correction—as a case where, given the subject's evidential standards, her initial failure to assign V a non-zero credence was irrational. Given an auxiliary norm that (say) requires subjects to be 'modally omniscient', one could treat this as a case where the subject needs to correct for past violations of auxiliary norms. And loaded evidential standards accounts will direct such a subject to adopt credences which correct these errors; i.e., to adopt a non-zero credence in V . Alternatively, one could treat this as a case where the subject satisfies all of the relevant auxiliary norms, but has just failed to update correctly, and both bare and loaded evidential standards accounts will direct such a subject to adopt credences which correct these errors. Of course, both of these options entail that only subjects who are have been irrational in some way can learn about new theories.

Another way evidential standards versions of Ur-Prior Conditionalization can handle this complaint is to treat this as a case of changing evidential standards. This allows us to maintain that subjects who learn about new theories needn't be rationally deficient. Assuming V is compatible with the subject's prior evidence, a subject's prior evidential standards could fail to yield a non-zero credence in V for one of two reasons. First, the subject's evidential standards might have assigned a zero value to V . Second, the subject's evidential standards might have been defined over an algebra of propositions that doesn't include V as a member, and so might have failed to assign any value to V . Either way, if we take learning about a new theory like V to correspond to a change in a subject's evidential standards—perhaps a boost in the value assigned to V , or the addition of V to the algebra of propositions—then we can maintain that a subject who learns about some new theory V and comes to assign it a non-zero

70. With one exception: if the subject adopts a loaded initial credence version of the rule, and takes there to be an auxiliary norm which requires her to adopt a non-zero initial credence in V given her initial evidence, then we can treat this as a case of auxiliary norm error correction, and the rule can allow her to adopt a non-zero credence in V even though her actual initial credence in V was 0.

credence needn't be in any way irrational.^{71,72}

5.7. *The Problem of Old Evidence and Confirmation*

A seventh complaint about Conditionalization is that, when paired with a natural claim about the relation between rational updating and confirmation, it has a hard time making sense of how evidence can still support a hypothesis after it's been updated on. Some sophisticated accounts of confirmation, such as the one proposed by Eells and Fitelson (2000), can accommodate much of what we'd like to say here, by appealing to credence functions the subject had before she received the relevant evidence E . But these accounts still run into difficulties in cases in which subjects have E from the very start, and so have no 'pre- E ' credence function to use when assessing claims regarding the evidential relevance of E .

The initial credence versions of Ur-Prior Conditionalization are subject to similar worries. In cases in which subjects have E from the start, E will be encoded in their initial credences. Thus these accounts don't provide a 'pre- E ' function to employ to assess claims about confirmation.

For the functional versions of Ur-Prior Conditionalization, scenarios in which subjects start with E aren't necessarily problematic. Recall that the functional accounts require there to exist some function up such that one's credences are equal to up conditional on one's evidence at every time. So suppose there's a subject who starts with E as evidence, but who later gets amnesia at some time t , loses all of her evidence, and comes to have a credence $cr_t(E) = 0.5$. In order for $cr_t(E) = up(E \mid \top) = 0.5$, as the functional accounts require, it must be the case that $up(E) = 0.5$. And if $up(E) = 0.5$, then there can be hypotheses that E can

71. If we pursue the second way in which one's evidential standards might change—via an expansion or contraction of the algebra of propositions over which one's evidential standards are defined—then this response becomes similar to the third response discussed in Section 3.6, which appeals to expansions of the space over which one's credences are defined and the project of working out the rational constraints on such expansions. And if one pursues the project of working out the rational constraints on what kinds of changes in evidential standards are permissible (see Section 4.3), and takes these possible changes to include expansions and contractions of the space over which one's evidential standards are defined, then these two projects naturally coincide.

72. Note that this second way of handling the complaint, which allows us to maintain that subjects who learn about new theories needn't be rationally deficient, is unavailable to proponents of Evidential Uniqueness (the view that every batch of evidence picks out a unique rational doxastic state). Proponents of Evidential Uniqueness must maintain that if a subject has the same evidence at two different times, then they're rationally required to have the same doxastic state at both times. But if epistemic standards can change, then the epistemic standards versions of Ur-Prior Conditionalization will allow a subject with the same evidence at two different times (before and after they learn about the new theories) to rationally maintain different doxastic states. (See Section 6.)

confirm (i.e., hypotheses A such that $up(A | E) > up(A)$).

While the functional accounts don't have problems providing 'pre- E ' functions, in some cases they'll provide too many 'pre- E ' functions. Consider, for example, a case in which a subject has E and a credence of $cr(E) = 1$ at all times. In this case, if any function satisfies the constraint the functional accounts impose, then an infinite number of functions will satisfy this constraint—for every $x \in (0, 1]$, there will be some function $up(E) = x$ that yields the subject's credences when conditionalized on her evidence at every time. Thus in these cases there won't be a fact of the matter as to whether E confirms other hypotheses or not. For there won't be a unique up function to appeal to—just a set of them that satisfy the relevant constraints—and the up functions in this set won't agree about whether E can confirm anything.

The evidential standards versions of Ur-Prior Conditionalization are not subject to either of these worries, however. For these accounts do provide a 'pre- E ' function to employ; namely, the subject's evidential standards es . So even if a subject starts with E , and has E at all future times, it can still be the case that $es(E) \neq 1$, and thus that E can confirm other hypotheses.⁷³

5.8. Epistemic Continuity

A final complaint about Conditionalization is that it relies on something like a personal identity or epistemic continuity relation, and one might be skeptical of the prospects of providing a satisfactory account of this relation.

The functional and initial credence versions of Ur-Prior Conditionalization won't help with this worry, since they both rely on something like personal identity or epistemic continuity. The functional approaches consider whether there exists some function up such that a subject's credences at different times can be generated by conditionalizing up on her total evidence at that time. And the initial credence approaches consider whether the subject's credences are equal to her initial credences conditional on her total evidence. In both cases, these approaches require a way of determining whether time slices of subjects at different times belong to the same subject.

The evidential standards versions of Ur-Prior Conditionalization, on the other hand, do help with this worry. The evidential standards approaches consider whether a subject's credences are equal to her current evidential standards conditional on her total evidence. And this doesn't require a way of identifying subjects over time. Thus the evidential standards approaches avoid any complaints

73. What if a subject's evidential standards assign 1 to E ? Well, in that case it seems she *shouldn't* take E to confirm anything. E.g., if a subject believes that, even given no evidence at all, she should be certain that E is true, then it doesn't seem E should have any confirmational power. For such a subject, E will have the same epistemic status as a tautology, and we wouldn't expect a tautology to confirm anything.

regarding how to spell out the notion of epistemic continuity.

5.9. Summing Up

Thus each version of Ur-Prior Conditionalization has a different profile with respect to how it handles these complaints about Conditionalization. Putting these results together, we can map the ability of these accounts to handle these complaints in a single chart (see Figure 1).

Figure 1. Ur-Prior Conditionalization: Assessing The Complaints

	Error Correction: Updating Errors	Error Correction: Auxiliary Norms	Changing Desiderata of Theory Choice ^d	Memory Loss _b	Self-Locating Beliefs	New Theories	POE and Confirmation ^e	Epistemic Continuity
Bare Initial Credence	☺	☹	☹	☹	☹	☹	☹	☹
Loaded Initial Credence	☺ ^d	☺ ^d	☹	☹	☹	☹	☹	☹
Bare Functional Loaded	☹	☹	☹	☹	☺	☹	☹	☹
Bare Evidential Standards	☹	☹	☺	☺	☹	☺ ^e	☺	☺
Loaded Evidential Standards	☹ ^f	☺ ^f	☺	☺	☺	☺ ^e	☺	☺

a. If one adopts Evidential Uniqueness, and thus denies the possibility of multiple reasonable stances regarding desiderata of theory choice, then one can simply ignore this complaint.

b. If one adopts an account of evidence according to which subjects never come into existence with evidence, then all of these accounts will be able to handle all cases of evidence loss, and this will become a column of ☺s.

c. If one adopts an account of evidence according to which subjects never come into existence with evidence, then all of these accounts can provide accounts of confirmation that yield the desired results in every case, and this will become a column of ☺s.

d. This assumes one adopts the (v1) or (v3) formulations of the loaded initial credence account given in Section 4.1. If one adopts the (v2) formulation, then the account won't be able to handle updating error correction in the desired way, and this will become a ☹.

e. If one adopts Evidential Uniqueness, then one can't avail oneself of the second way of dealing with this complaint described in Section 5.6, and this arguably becomes a ☺.

f. This assumes one adopts the (v1) or (v3) formulations of the loaded initial credence account given in Section 4.3. If one adopts the (v2) formulation, then the account won't be able to handle updating error correction in the desired way, and this will become a ☹.

6. Evidential Uniqueness

Let *Evidential Uniqueness* be the thesis that for any total body of evidence *E*, there's a unique doxastic state that it's rationally obligatory for any subject with total evidence *E* to adopt. Some proponents of Ur-Prior Conditionalization, such as Carnap (1950), Williamson (2000) and Hedden (2015), adopt Evidential Uniqueness, and take it to be a vital part of the epistemic package they endorse. How does adopting Evidential Uniqueness affect the preceding discussion?

First, adopting Evidential Uniqueness adds a few wrinkles regarding the complaints about Conditionalization discussed in Section 3. For example, proponents of Evidential Uniqueness won't take the changing desiderata of theory choice complaint to be a reason to reject Conditionalization (see Footnote 16). But by and large, the discussion in Section 3 goes just as before.

Second, adopting Evidential Uniqueness adds a few wrinkles regarding the assessments of the different versions of Ur-Prior Conditionalization discussed in Section 5. For example, proponents of Evidential Uniqueness can't avail themselves of any way of dealing with these complaints that appeals to rational changes in epistemic standards, such as the replies sketched in Sections 3.3 and 3.6. For different epistemic standards will yield different prescriptions given the same evidence, which Evidential Uniqueness forbids. But again, by and large, the discussion in Section 5 will go just as before.

Third, adopting Evidential Uniqueness substantially changes the kinds of worries facing some of the versions of Ur-Prior Conditionalization discussed in Section 4. In particular, adopting Evidential Uniqueness allows one to avoid some of the worries facing the loaded initial credence and loaded evidential standards approaches.

Let's start with the loaded initial credence approach. Section 4.1 raised two challenges for the loaded initial credence approach: specifying the best rational fit (or fits) to an arbitrary credence function, and resolving the Quantifier Problem.

If we adopt Evidential Uniqueness, then it might seem like the first challenge is easy to meet: the best rational fit to any credence function is the one rational credence function. But this benefit is illusory. Saying "the best rational fit to *cr* is the one rational credence function" only addresses the challenge if one also provides a detailed account of what the one rational credence function is. Otherwise, one is just asserting that there exists a best rational fit, not telling us what it is. And if *that* were all we needed to do, then opponents of Evidential Uniqueness would be off the hook too—in every case they could just assert that there exists some best rational fit (or fits). For proponents of Evidential Uniqueness to meet this challenge, they need to provide a detailed account of what the one best rational fit is. And that task is no easier than the task of specifying best rational

fits facing opponents of Evidential Uniqueness.

That said, adopting Evidential Uniqueness does help with the second challenge—the Quantifier Problem. In particular, the proponent of Evidential Uniqueness can just adopt the loaded initial credence (v3) formulation. As we saw in Section 4.1.1, the worry for the (v3) formulation was that it effectively privileges one best rational fit over the others in cases of ties. But if we adopt Evidential Uniqueness then there are no ties, and this worry goes away.

Things are much the same for the loaded evidential standards approach. Section 4.3 raised three worries for the loaded evidential standards approach: providing an account of how to assess best rational fits, resolving the Quantifier Problem, and spelling out what norms (if any) govern changes in evidential standards. The first two challenges are affected in just the same way as before. And adopting Evidential Uniqueness also helps with the third challenge: since there's only one rational set of evidential standards, the evidential standards one should have *can't* change over time.

For the most part, adopting Evidential Uniqueness has little effect on the preceding discussion—proponents of Evidential Uniqueness won't recognize some of the complaints about Conditionalization discussed in Section 3, and can't avail themselves of some of the replies to complaints discussed in Section 5, but most of the preceding discussion goes just as before. There's one notable exception, however: adopting Evidential Uniqueness does a fair bit to alleviate some of the worries facing the loaded initial credence and evidential standards approaches. This provides some *prima facie* support for a package of views which adopts Evidential Uniqueness and either the loaded initial credence or evidential standards versions of Ur-Prior Conditionalization. But, of course, the ultimate appeal of such a package will depend on one's assessment of the plausibility of Evidential Uniqueness, and the case for adopting one of these versions of Ur-Prior Conditionalization.⁷⁴

7. Conclusion

Each version of Ur-Prior Conditionalization is accompanied by a web of pros and cons. And some of these pros and cons are conditional on other issues, such as Evidential Uniqueness or one's account of evidence. So assessing these different versions of Ur-Prior Conditionalization is a complicated task.

But two of these accounts stand out. The first is the bare functional account, when paired with an account of evidence which entails that subjects who have just come into existence have no (non-trivial) evidence. This view avoids the sub-

74. For some recent discussion on the merits of Evidential Uniqueness, see Kelly (2013), Horowitz (2013), Meacham (2013), White (2013), Schoenfield (2014), and Titelbaum and Kopec (2016).

stantial commitments or demerits of the other views; indeed, it incurs no further commitments beyond those of the standard Bayesian approach. And yet it offers a way to deal with complaints regarding evidence loss and self-locating beliefs, and (given the assumption about evidence) is not subject to worries regarding confirmation. Thus the bare functional account *dominates* Conditionalization: it is a strictly more attractive updating rule.

The second account that stands out is the loaded evidential standards account. This view can address all of the complaints about Conditionalization we've considered, and is, I think, the most promising version of Ur-Prior Conditionalization. Of course, the cost of adopting this account is incurring some substantive commitments: providing an account of evidential standards, providing an account of how to assess best rational fits, and deciding which of the three natural formulations of the rule, and its accompanying demerits, to accept. But if one is willing to accept the commitments of such a view, then it provides one with a promising neo-Bayesian account which avoids many of the worries facing its predecessor.

Acknowledgments

I'd like to thank Phil Bricker, Lisa Cassell, Maya Eddon, Nina Emery, Alejandro Perez-Carballo, Mike Titelbaum, and two anonymous referees for helpful comments and discussion.

References

- Arntzenius, Frank (2003). Some Problems for Conditionalization and Reflection. *Journal of Philosophy*, 100(7), 356–370.
- Bradley, Darren (2011). Self-Location is No Problem for Conditionalization. *Synthese*, 182(3), 393–411.
- Briggs, Rachael (2010). Putting a Value on Beauty. In Tamar Szabo Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 3, 3–34) Oxford University Press.
- Broome, John (1999). Normative Requirements. *Ratio*, 12(4), 398–419.
- Carnap, Rudolf (1950). *The Logical Foundations of Probability*. University of Chicago Press.
- Carr, Jennifer Rose (2016). Don't Stop Believing. *Canadian Journal of Philosophy*, 45(5-6), 744–766.
- Christensen, David (1999). Measuring Confirmation. *Journal of Philosophy*, 96(9), 437–461.
- Earman, John (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirma-*

- tion Theory. MIT Press.
- Easwaran, Kenny (2014). Regularity and Hyperreal Credences. *Philosophical Review*, 123(1), 1–41.
- Eells, Ellery (1985). Problems of Old Evidence. *Pacific Philosophical Quarterly*, 66(3), 283.
- Eells, Ellery and Branden Fitelson (2000). Measuring Confirmation and Evidence. *Journal of Philosophy*, 97(12), 663–672.
- Elga, Adam (2000). Self-Locating Belief and the Sleeping Beauty Problem. *Analysis*, 60(2), 143–147.
- Feldman, Fred (1986). *Doing the Best We Can: An Essay in Informal Deontic Logic*. D. Reidel.
- Garber, Daniel (1983). Old Evidence and Logical Omniscience in Bayesian Confirmation Theory. In John Earman (Ed.), *Testing Scientific Theories: Minnesota Studies in the Philosophy of Science* (Vol. 10, 99–131). Minnesota University Press.
- Glymour, Clark (1980). *Theory and Evidence*. Princeton University Press.
- Hájek, Alan (2003). What Conditional Probability Could Not Be. *Synthese*, 137(3), 273–323.
- Hájek, Alan (2015). Staying Regular. Manuscript in preparation.
- Halpern, Joseph (2004). Sleeping Beauty Reconsidered: Conditioning and Reflection in Asynchronous Systems. In Tamar Szabo Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 1, 111–142). Oxford University Press.
- Hedden, Brian (2015). Time-Slice Rationality. *Mind*, 124(494), 449–491.
- Horowitz, Sophie (2013). Immoderately Rational. *Philosophical Studies*, 167, 41–56.
- Howson, Colin (1984). Bayesianism and Support by Novel Facts. *British Journal for the Philosophy of Science*, 35(3), 245–251.
- Howson, Colin and Peter Urbach (1993). *Scientific Reasoning: The Bayesian Approach*. Open Court.
- Kelly, Thomas (2013). Evidence Can Be Permissive. In Matthias Steup and John Turri (Eds.), *Contemporary Debates in Epistemology* (298–311). Blackwell.
- Kim, Namjoong (2009). Sleeping Beauty and Shifted Jeffrey Conditionalization. *Synthese*, 168(2), 295–312.
- Levi, Isaac (1980). *The Enterprise of Knowledge an Essay on Knowledge, Credal Probability, and Chance*. MIT Press.
- Lewis, David (1979). Attitudes de Dicto and de Se. *Philosophical Review*, 88(4), 513–543.
- Lewis, David (1980). A Subjectivist's Guide to Objective Chance. In Richard C. Jeffrey (Ed.), *Studies in Inductive Logic and Probability* (Vol. 2, 83–132). University of California Press.

- Lewis, David (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Lewis, David (2010). Why Conditionalize. In Antony Eagle (Ed.), *Philosophy of Probability: Contemporary Readings* (403–407). Routledge.
- Maher, Patrick (1996). Subjective and Objective Confirmation. *Philosophy of Science*, 63(2), 149–174.
- Meacham, Christopher J. G. (2007). *Chance and the Dynamics of de Se Beliefs* (Doctoral dissertation). Retrieved from <http://dx.doi.org/doi:10.7282/T3oZ73N5>
- Meacham, Christopher J. G. (2008). Sleeping Beauty and the Dynamics of de Se Beliefs. *Philosophical Studies*, 138(2), 245–269.
- Meacham, Christopher J. G. (2010). Unravelling the Tangled Web: Continuity, Internalism, Non-Uniqueness and Self-Locating Beliefs. In Tamar Szabo Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 3, 86–125). Oxford University Press.
- Meacham, Christopher J. G. (2013). Impermissive Bayesianism. *Erkenntnis*, 79(6), 1–33.
- Meacham, Christopher J. G. (2016). Understanding Conditionalization. *Canadian Journal of Philosophy*, 45(5-6), 767–797.
- Moss, Sarah (2012). Updating as communication. *Philosophy and Phenomenological Research*, 85(2), 225–248.
- Moss, Sarah (2014). Credal Dilemmas. *Noûs*, 48(3), 665–683.
- Moss, Sarah (in press). Time-Slice Epistemology and Action Under Indeterminacy. In John Hawthorne and Tamar Gendler (Eds.), *Oxford Studies in Epistemology* (Vol. 5). Oxford University Press.
- Perez Carballo, Alejandro (in press). New Boundary Lines. In Kristoffer Ahlstrom-Vij and Jeff Dunn (Eds.), *Epistemic Consequentialism*. Oxford University Press.
- Schoenfield, Miriam (2014). Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief. *Noûs*, 48(2), 193–218.
- Schoenfield, Miriam (in press). Conditionalization Does Not (in General) Maximize Expected Accuracy. *Mind*.
- Schwarz, Wolfgang (2012). Changing Minds in a Changing World. *Philosophical Studies*, 159(2), 219–239.
- Skyrms, Brian (1983). Three Ways to Give a Probability Function a Memory. In John Earman (Ed.), *Testing Scientific Theories, Minnesota Studies in the Philosophy of Science* (Vol. 10, 157–161). Minnesota University Press.
- Staffel, Julia (2016). Should I Pretend I’m Perfect? Manuscript in preparation.
- Talbott, W. J. (1991). Two Principles of Bayesian Epistemology. *Philosophical Studies*, 62(2), 135–150.
- Titelbaum, Michael G. (2013). *Quitting Certainties: A Bayesian Framework Modeling*

- Degrees of Belief*. Oxford University Press.
- Titelbaum, Michael G. (2016). Fundamentals of Bayesian Epistemology. Manuscript in preparation.
- Titelbaum, Michael G. and Matthew Kopec (2016). Plausible Permissivism. Manuscript in preparation.
- van Fraassen, Bas C. (1984). Belief and the Will. *Journal of Philosophy*, 81, 235–256.
- Weatherson, Brian (2005). Should We Respond to Evil with Indifference? *Philosophy and Phenomenological Research*, 70(3), 613–635.
- White, Roger (2013). Evidence Cannot Be Permissive. In Matthias Steup and John Turri (Eds.), *Contemporary Debates in Epistemology* (312–324). Blackwell.
- Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford University Press.