

Understanding Conditionalization

Christopher J. G. Meacham

Department of Philosophy, University of Massachusetts, Amherst, MA, USA.

ABSTRACT

At the heart of the Bayesianism is a rule, Conditionalization, which tells us how to update our beliefs. Typical formulations of this rule are underspecified. This paper considers how, exactly, this rule should be formulated. It focuses on three issues: when a subject's evidence is received, whether the rule prescribes sequential or interval updates, and whether the rule is narrow or wide scope. After examining these issues, it argues that there are two distinct and equally viable versions of Conditionalization to choose from. And which version we choose has interesting ramifications, bearing on issues such as whether Conditionalization can handle continuous evidence, and whether Jeffrey Conditionalization is really a generalization of Conditionalization.

ARTICLE HISTORY Received 31 October 2015; Accepted 5 November 2015

KEYWORDS Conditionalization; Bayesianism; belief revision; belief updating; formal epistemology

1. Introduction

At the heart of the Bayesian account of rationality is a rule – Conditionalization – which tells us how to update our beliefs in light of evidence. At a first pass, one might characterize this rule as follows:

Conditionalization: If a subject with credences cr gets evidence E , she should adopt new credences cr_E such that $cr_E(\cdot) = cr(\cdot|E)$, if defined. This formulation of the rule is adequate for most purposes, but it leaves open a number of questions. And, by answering these questions in different ways, we get different versions of Conditionalization.

In this paper, I'll explore these questions. My focus here will be on questions regarding the logical form of Conditionalization. There are, of course, many other interesting questions to ask about the rule, such as how to understand the notions of 'credence' or 'evidence' the rule employs. But I won't try to address

CONTACT Christopher J. G. Meacham  cmeacham@philos.umass.edu

© 2016 Canadian Journal of Philosophy

those kinds of questions here. Instead, I'll restrict myself to questions regarding logical form.¹

The rest of this paper will proceed as follows. In Section 2, I'll sketch some background. In the following sections, I'll consider three questions left open by formulations of Conditionalization like the one given above. In Section 3, I'll consider the time of evidence question. In Section 4, I'll consider the sequential vs. interval updating question. In Section 6, I'll consider the narrow vs. wide scope question. For each question, I'll present and assess some plausible answers. For two of these three questions, I'll argue that one of the answers is better than the rest. But for one of these questions – the sequential vs. interval updating Question – I'll suggest that there are two viable answers. Thus I'll suggest that, at the end of the day, there are two viable versions of Conditionalization for us to choose from.

My examination of these three questions does not, of course, guarantee that there aren't further open questions regarding the logical form of Conditionalization that haven't been answered. In order to do this, we would need to provide a logically precise formulation of Conditionalization. So to ensure that no more details of formulation are left unspecified, I'll conclude in Section 7 by presenting two formal characterizations of Conditionalization, one corresponding to each of the two viable versions of Conditionalization mentioned above. (Readers who would like to be forewarned about what formulations of Conditionalization I'll endorse can skip ahead and skim Section 7 before reading the rest of the paper.)

2. Background

Let a subject's *credences* be an assignment of real numbers to propositions representing the subject's confidence in those propositions, where an assignment of 0 indicates that the subject is virtually certain the proposition is false, and an assignment of 1 indicates that the subject is virtually certain the proposition is true.² Note that we are not assuming that a subject's credences assign numbers to every proposition; there may be some propositions in which a subject doesn't have a credence.

One popular normative constraint on credences is *Probabilism*, a constraint on what a subject's credences should be like at a time:

Probabilism: A subject's credences should be probabilistic.³ A second popular normative constraint on credences is *Conditionalization*, a constraint on how a subject's credences should change over time in light of evidence. At a first pass, we can characterize Conditionalization as follows:

Conditionalization: If a subject with credences cr gets evidence E , she should adopt new credences cr_E such that $cr_E(\cdot) = cr(\cdot|E)$, if defined.

I've appended the '(v0)' to this formulation to highlight that this is only a first pass approximation; we'll consider more precise formulations of Conditionalization in the sections to come.⁴ Intuitively, Conditionalization tells us that upon receiving evidence E , we should assign E a credence of 1 and renormalize; that is, shift all of our credence to E , and distribute that credence among the propositions entailing E in a way that keeps the ratios between them the same.

One common complaint about Conditionalization is that it requires us to adopt a credence of 1 in our evidence. In some situations, it's been suggested, it seems like we get 'uncertain' evidence – evidence to which we should assign a credence of less than 1. For example, if a subject sees her friend through the window, and the lighting outside is poor, then it might seem like she should assign the proposition that her friend is outside a value less than 1. Worries of this kind motivate another popular normative constraint on credences, *Jeffrey Conditionalization*.

Like Conditionalization, Jeffrey Conditionalization is a constraint on how a subject's credences should change over time in light of evidence. But on this picture, one's evidence isn't a proposition E ; instead, it's a weighted partition of propositions $S = \{(E_1, x_1), (E_2, x_2), \dots\}$ (where $E_1 - E_n$ are mutually exclusive and jointly exhaustive propositions, and the weights $x_1 - x_n$ are real numbers in the $[0, 1]$ -interval that sum to 1). Given such evidence, Jeffrey Conditionalization tells us to update as follows:

Jeffrey Conditionalization (v0): If a subject with credences cr gets evidence partition $S = \{(E_1, x_1), (E_2, x_2), \dots\}$ she should adopt new credences cr_S such that:

$$cr_S(\cdot) = \sum_i x_i \cdot cr(\cdot|E_i), \text{ if defined.}$$

Intuitively, Jeffrey Conditionalization tells us that upon receiving evidence partition S , we should assign each E_i a credence of x_i , and then renormalize; that is, move our credence in the E_i 's to the indicated amount, and then distribute the credence assigned to each E_i among the propositions entailing E_i in a way that keeps the ratios between them the same.

Jeffrey Conditionalization is generally taken to be a generalization of Conditionalization that yields Conditionalization as a special case. This is because when we plug in simple evidence partitions of the form $S = \{(E, 1), (\neg E, 0)\}$, we get:

$$\begin{aligned} cr_S(\cdot) &= \sum_i x_i \cdot cr(\cdot|E_i), \\ &= 1 \cdot cr(\cdot|E) + 0 \cdot cr(\cdot|\neg E), \\ &= cr(\cdot|E), \end{aligned}$$

which is just what Conditionalization prescribes.

Call a view *Bayesian* if it takes both Probabilism and Conditionalization to impose normative constraints on credences. For the purposes of this paper, I'll be assessing the question of how to understand Conditionalization under the assumption that some Bayesian view is correct.

Although the discussion in this paper is couched in terms of questions about Conditionalization, the same questions can be raised regarding Jeffrey Conditionalization. Likewise, the same answers to these questions can be offered, and the same considerations for and against these answers obtain. (With one exception in Section 4.1 – but I'll flag this difference when we come to it.) So, although the following discussion will focus on Conditionalization, most of the conclusions of this paper will apply to both updating rules.

3. The time of evidence question

Conditionalization (v_0) makes reference to three events: the subject having credences cr , the subject receiving evidence E , and the subject adopting new credences cr_E .⁵ How are the times of these events – call them $t(cr)$, $t(E)$ and $t(cr_E)$ – related?⁶ Certain constraints are clear: the time at which a subject adopts cr_E should not be before the time at which she receives E as evidence, and she should adopt cr_E at some point after she has cr . But are there other constraints on the timing of these events?

Let's focus here on the relation between $t(E)$ and $t(cr_E)$. (We'll consider the relation between these times and $t(cr)$ at the end of Section 3.1.)

Q3. The Time of Evidence Question: How is the time at which the subject receives her evidence related to the time at which she should adopt her new credences?⁷

Answers to this question fall into two camps. First, one might hold that subjects should adopt their new credences after they've received their new evidence, so that $t(E) < t(cr_E)$. Because time is *dense* – given any two distinct times there will be some time in-between – it follows that if $t(E) < t(cr_E)$, there are times in-between $t(E)$ and $t(cr_E)$. That is, there will be a temporal gap between $t(E)$ and $t(cr_E)$:

Answer 1 (Posterior). The subject should adopt her new credences some finite amount of time after she gets her evidence. We can make this understanding explicit by adding times to our formulation, following the usual convention of using time indices to reflect the differences between these times (e.g. t_1 is one unit of time after t_0 , according to some linear measure):

Conditionalization (v1.1): If a subject with credences cr gets evidence E at t_0 , she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined.

Second, one might hold that subjects should receive E and adopt cr_E simultaneously:

Answer 2 (Concurrent). The subject should adopt her new credences at the same time as she gets her evidence. Again, we can make this understanding explicit by inserting the appropriate times into our formulation of the rule:

Conditionalization (v1.2): If a subject with credences cr gets evidence E at t_r she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined.

3.1. *The time of evidence question: assessing the answers*

Now, let's turn to assess each of these answers.

The posterior answer, that the subject should adopt cr_E some amount of time after receiving E , leads to implausible verdicts. Consider a subject who receives evidence E at t_0 . Given the posterior answer the subject shouldn't adopt cr_E until some later time t_r . And, if she adopts cr_E at t_0 she's epistemically deficient for having 'jumped the gun'. But this seems like an odd verdict. After all, a subject who adopts cr_E at t_0 is in a strictly better position, epistemically speaking, than the subject who doesn't – she's taken all of her evidence into account. And, it's hard to see why it's epistemically irrational to take all of one's evidence into account as soon as one has it.⁸

The posterior answer is also in tension with the popular account of evidence endorsed by Howson and Urbach (1993), where a subject's evidence is the strongest proposition in which they have a credence of 1. Suppose a subject receives evidence E at t_0 . What should her credences be at t_0 ? Since she isn't required to adopt her new credences until t_r , the most natural option is to maintain that her credences at t_0 are still cr . But, since cr generally won't assign 1 to E , this option is incompatible with Howson and Urbach's picture of evidence. Of course, there are other options one might try, but none of them are satisfactory. For example, one might maintain that at t_0 a subject should assign a credence of 1 to E (as Howson and Urbach's account requires), but should otherwise assign the same values as cr does. But since this credence function will generally violate the probability axioms, this option's incompatible with Probabilism. Alternatively, one might maintain that at t_0 a subject should adopt cr_E . But this is to abandon the posterior answer for the concurrent answer, as this entails that one should receive E and adopt cr_E simultaneously. So if we want to hold on to the posterior answer, it seems we must say that the subject's credences at t_0 should be cr . And this is incompatible with Howson and Urbach's account of evidence.

Now consider the concurrent answer, that subjects should adopt cr_E at the same time as they receive E . One worry for the concurrent answer is that it places too strong a demand on subjects like us. In particular, one might worry that this answer is incompatible with the principle that ought implies can. If

subjects like us aren't capable of updating instantaneously, then how can we be obligated to do it?

Of course, this is not a fair criticism of the concurrent answer if the posterior answer also has this problem. And, as given, the posterior answer doesn't take the cognitive capacities of subjects into account either. It states that a subject's new credences should be adopted some amount of time after the subject gets her evidence, but it doesn't say anything about what this later time is, or whether it's possible for the subject to adopt those credences at that time.

But one might naturally think that one can modify the posterior answer so that it does take the cognitive capabilities of subjects into consideration. And if, by modifying the posterior answer, we can get a plausible formulation of Conditionalization that avoids these kinds of ought-implies-can worries, then we have a reason to favor the posterior answer over the concurrent answer. Let's see whether this is true: by modifying the posterior answer, can we get a plausible version of Conditionalization that avoids these kinds of ought-implies-can worries?

Here's a natural way to modify the posterior answer to take the updating time lag of cognitively limited subjects into account:

Answer 1a (Posterior-a). The subject should adopt her new credences *as soon as possible* after receiving her evidence.

Conditionalization (v1.1a): If a subject with credences cr gets evidence E at t_0 , she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$ (if defined), where t_1 is the earliest time following t_0 at which the subject is capable of adopting cr_E . This formulation of Conditionalization is sensitive to the temporal limitations of subjects with limited cognitive capacities. But it yields inconsistent prescriptions. Suppose a subject with prior credences cr gets evidence E at t_0 and evidence F at $t_{1/2}$. And suppose the subject is cognitively limited in such a way as to not be able to update her credences until t_1 . But at t_1 , she's capable of changing her credences in any way she likes. Because the subject gets E at t_0 , Conditionalization (v1.1a) requires her to adopt new credences equal to $cr(\cdot|E)$ at t_1 since t_1 is the first time at which she's capable of updating on E . Likewise, because the subject gets F at $t_{1/2}$, Conditionalization (v1.1a) requires her to adopt new credences equal to $cr(\cdot|F)$ at t_1 since her credences at $t_{1/2}$ will still be cr , and t_1 is the first time at which she's capable of updating on F . But these two prescriptions will usually be inconsistent.

The problem here is that Conditionalization (v1.1a) doesn't take into account the possibility that a subject might receive further evidence besides E between t_0 and t_1 . So, we might repair our formulation of Conditionalization by adding a clause which rules out this possibility:

Conditionalization (v1.1b): If a subject with credences cr gets evidence E at t_0 , she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$ (if defined),

where t_j is the earliest time following t_0 at which the subject is capable of adopting cr_{E_j} and the subject doesn't get any other evidence between t_0 and t_j . This rule avoids making inconsistent prescriptions in the kinds of 'multiple evidence cases' described above by simply falling silent in such cases. But an adequate formulation of Conditionalization should give us some guidance in these cases – it shouldn't just fall silent. So this rule is too weak.

We can get around this problem by adding a further clause which tells us what credences to adopt in these cases:

Conditionalization (v1.1c): If a subject with credences cr gets evidence E at t_0 , she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$ (if defined), where t_j is the earliest time following t_0 at which the subject is capable of adopting cr_{E_j} and assuming she hasn't gotten any other evidence between t_0 and t_j . If she has gotten other evidence, she should update on the conjunction of all of the evidence $E_1 - E_n$ she's received up to t_j ; i.e. she should adopt credences $cr_{E_1 \dots E_n}(\cdot) = cr(\cdot|E_1 \wedge \dots \wedge E_n)$.

This formulation avoids the inconsistent prescriptions worry, and offers prescriptions in multiple evidence cases. But it avoids these worries by sliding back toward the concurrent answer to the time of evidence question. For, like the concurrent answer, Conditionalization (v1.1c) effectively takes into account all of the evidence the subject receives up until the time at which she should adopt her new credences. And in doing so, Conditionalization (v1.1c) runs afoul of the same kinds of ought-implies-can worries that prompted these modifications of Conditionalization in the first place. Consider a multiple evidence case like the one described above, but with the following modification: at t_j the subject isn't capable of changing her credences in any way one likes. Instead, she can only update on one of the pieces of evidence she's received, and so can only adopt new credences equal to either $cr(\cdot|E)$ or $cr(\cdot|F)$. Since Conditionalization (v1.1c) requires the subject to adopt $cr(\cdot|E \wedge F)$, such a subject isn't capable of satisfying the prescriptions the rule makes.

One might continue to finesse the formulation of Conditionalization, but there are some more general reasons why this strategy won't work. First, the intuitive idea we're trying to capture is that subjects should update on as much as they can, as soon as they can. But in order to capture this idea, a rule would need to be considerably more complex than the formulations offered above, having to provide verdicts regarding choices between updating on different batches of evidence, choices between updating sooner on less evidence versus updating later on more, and so on. And it's hard to see how any such rule could remain similar enough to the formulations given above to plausibly be identified with Conditionalization, the norm that people have been talking about in the Bayesian literature.

Second, it's not clear that subjects *should* update on their evidence as soon as they can. For example, suppose a subject who gets evidence E also violates

some other epistemic norm, such as the Principal Principle – roughly, the claim that your credences should line up with what you think the chances are.⁹ And suppose she is only capable of doing one thing at a time – either updating on *E*, or altering her credences so that they satisfy the Principal Principle. In this case, it's not clear that she should update as soon as she can – it could well be that she should alter her credences to satisfy the Principal Principle first. So even the general idea that subjects should conditionalize on their evidence 'as soon as they can' seems too simplistic to yield the desired prescriptions, given the variety of situations and cognitive limitations facing imperfect subjects.

To sum up, in order to make allowances for the restricted capacities of cognitively limited subjects, we need a very complex updating rule. And no plausible understanding of Conditionalization is this complex.

The moral is that it's a mistake to think of Conditionalization as a norm which is supposed to provide guidance to cognitively limited subjects. Instead, following Christensen (2004), we should think of Conditionalization as an ideal toward which to aim, a description of optimal performance in the epistemic realm.¹⁰

To borrow Christensen's analogy, we should think of Bayesian norms as like the norms describing perfect tournament chess play. Better tournament chess players can make better moves in less time, and in the limit, ideal tournament chess players would make perfect moves in no time. But even though this is a good description of perfect tournament chess play, we wouldn't expect any actual subject to be able to live up to this ideal. These 'ideal performance' norms aren't the kinds of norms to which ought implies can generally applies. And once we properly understand Conditionalization as an ideal performance norm, we can see that ought-implies-can worries regarding it are misplaced.¹¹

A different worry that one might raise for the concurrent answer is that it's incompatible with plausible 'procedural' requirements on rational belief. Following Simon (1976), let us distinguish between two kinds of rational requirements: *substantive* rational requirements, which concern the rationality of the result, and *procedural* rational requirements, which concern the rationality of the process by which one obtains those results. Thus, when assessing whether a subject is epistemically rational, we might not only want her to have the right beliefs (i.e. to satisfy the relevant substantive requirements), but also to come to have those beliefs in the right way (i.e. to satisfy the relevant procedural requirements).

Conditionalization imposes a substantive requirement on rational updating: it requires subjects to come to have the right beliefs given their evidence and prior beliefs. Conditionalization doesn't care about how subjects came to have those beliefs – it doesn't care whether their beliefs were formed by reflecting on their evidence or whether their beliefs were formed by random quantum mechanical fluctuations. Now, one might think there are further requirements on rational updating beyond those imposed by Conditionalization. For example, one might also take there to be procedural requirements which require that a

subject's evidence be the cause of her changing her beliefs in this way, or require that she come to have her new beliefs by reasoning about her evidence. And, one might worry that the concurrent answer to the time of evidence question is incompatible with such procedural requirements. For if a subject needs to update on her evidence instantaneously, then it doesn't seem like she *can* come to have those credences in the right way – via some causal process initiated by the receipt of her evidence, or by reasoning in light of her evidence. For any such process requires a non-zero amount of time.¹²

But this worry is misplaced. For these substantive and procedural requirements aren't actually in conflict. There isn't anything logically incoherent about a subject who instantaneously comes to have the right beliefs for the right reasons. It's true that subjects like *us* can't do this. But that's only a reason to think that these can't be rational requirements if we're assuming something like ought implies can. And, as we've seen, once Conditionalization is understood in the right way – as a description of optimal performance – this kind of reasoning is not compelling. For ideal performance norms aren't the kinds of norms to which ought implies can generally applies.¹³

Thus, all things considered, the concurrent answer is the best way to understand Conditionalization. The posterior answer yields implausible prescriptions, by effectively requiring subjects not to take all of their evidence into account, and is in tension with some popular accounts of evidence. And while the concurrent answer appears to face ought-implies-can worries and to conflict with procedural requirements, further reflection makes it clear that these worries are unreasonable (in the first case) and mistaken (in the second).¹⁴

Adopting the concurrent answer also settles the relationship between $t(cr)$, $t(E)$ and $t(cr_E)$. As we noted earlier, we want $t(cr) < t(cr_E)$. If we adopt the concurrent answer, so that $t(cr_E) = t(E)$, then it follows that $t(cr) < t(E)$ as well. Thus, the right relationship between these three times is: $t(cr) < t(E) = t(cr_E)$.

Building this into our formulation gets us the following rule:

Conditionalization (v1): If a subject with credences cr at t_0 gets evidence E at t_1 , she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined.

4. The sequential vs. interval updating question

Conditionalization (v1) says, roughly, that if a subject has credences cr at $t(cr)$, and gets evidence E at $t(E)$, then she should adopt cr_E at $t(cr_E) = t(E)$. But suppose a subject receives evidence F at some time between $t(cr)$ and $t(E)$. Then, as written, Conditionalization will still tell the subject to adopt credences $cr_E(\cdot) = cr(\cdot|E)$ at $t(E)$. But one might worry that this is the wrong prescription – after all, the subject also received F as evidence, and cr_E doesn't seem to take that into account!

There are two natural ways to reply to this worry. The first is to place a further constraint on $t(cr)$. Namely, require $t(cr)$ to be a time such that the subject *doesn't*

receive any evidence between $t(cr)$ and $t(E)$. The second is to place a further constraint on the content of E . Namely, require that E incorporate (i.e. entail) any other evidence the subject gets between $t(cr)$ and $t(E)$.

These two replies lead to two different ways of thinking about Conditionalization. The first reply depicts Conditionalization as a rule which tells us how to update whenever we get a new piece of evidence. On this conception, $t(cr)$ is some time before the subject gets E but after she's received any of her other evidence, and E is the evidence the subject receives at $t(E)$. We might call this the 'sequential updating' picture, since on this picture we determine what a subject's credence should be, given some earlier credence function, by sequentially applying the rule to each of the pieces of evidence the subject receives.¹⁵

The second reply depicts Conditionalization as a rule which tells us, for any interval of time, how our credences at the endpoints of that interval should be related given the evidence received in the interim. On this conception, $t(cr)$ and $t(E)$ can be any times we like, and E is the cumulative evidence the subject receives during this interval.¹⁶ We might call this the 'interval updating' picture, since on this picture the rule tells us how to update over arbitrary intervals.

So which of these two pictures of Conditionalization should we adopt?

Q3. The Sequential vs. Interval Updating Question: Does the rule tell subjects how to update whenever they get a piece of evidence? Or does it tell them how to update over arbitrary intervals, given the cumulative evidence they've received during that interval?¹⁷ The first answer to this question takes the rule to be telling us how to update our credences whenever we get a piece of evidence:

Answer 1 (Sequential). The rule tells a subject how to update when she gets a piece of evidence. Thus $t(E)$ is the time after $t(cr)$ at which they next get evidence, and E is the evidence they get at $t(E)$. If we adopt the sequential answer, we need to add a clause to our formulation of the rule that states that the subject hasn't received any other evidence after $t(cr)$ prior to receiving E :

Conditionalization (v2.1-): If a subject with credences cr at t_0 gets evidence E at t_1 (and no evidence between t_0 and t_1), then she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$ if defined. This formulation puts constraints on what a subject's credences should be at $t(E)$ given her credences at $t(cr)$. But it doesn't place any constraints on what her credences between $t(cr)$ and $t(E)$ should be. Since any deviations from cr during this period would amount to evidence-less belief changes – something we presumably want the rule to forbid – we'll also want to add a clause requiring the subject's credences between $t(cr)$ and $t(E)$ to remain the same:

Conditionalization (v2.1): If a subject with credences cr at t_0 gets evidence E at t_1 (and no evidence between t_0 and t_1), then her credences should remain cr between t_0 and t_1 and she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined. The second answer to this question takes the rule to be telling us

how to update over any interval of time. We choose a time interval, plug in our credences at the start of that interval and the cumulative evidence we receive during that interval, and the rule tells us what our credence should be at the end of this interval.

Answer 2 (Interval). The rule tells a subject how to update over an arbitrary interval, given the evidence received during that interval. Thus $t(cr)$ and $t(E)$ are the endpoints of an arbitrary interval, cr is the credence function of the subject at the start of the interval, and E is the cumulative evidence she receives during that interval. We can make this understanding of the rule explicit by adding the appropriate clause to our formulation, as follows:

Conditionalization (v2.2): If a subject with credences cr at t_0 gets cumulative evidence E in the $[t_0, t_1]$ interval, then she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined.

4.1. *The sequential vs. interval updating question: assessing the answers*

Both of these answers to the sequential vs. interval updating question have uncomfortable consequences.

If we adopt the sequential answer, and think of Conditionalization as telling subjects how to update whenever they get a piece of evidence, then the rule can't accommodate cases in which subjects get evidence continuously. This understanding of Conditionalization requires a 'no evidence between t_0 and t_1 ' clause. Thus, the rule requires a pair of distinct times t_0 and t_1 between which the subject doesn't get any evidence. But if the subject continuously receives evidence, then there is no such pair of times – between *any* two distinct times, there will be times in-between at which the subject gets evidence. So if we adopt a sequential understanding of Conditionalization, the rule simply goes silent in cases in which subjects receive continuous evidence.

But we also face some uncomfortable consequences if we adopt the interval answer, and think of Conditionalization as telling subjects how to update over arbitrary intervals, given the cumulative evidence they receive during that interval. In particular, the interval answer is incompatible with the conjunction of two widely held claims:

- (1) Conditionalization is a special case of Jeffrey Conditionalization.
- (2) The formalism of Jeffrey Conditionalization is neutral with respect to whether evidence is 'credence-dependent'.

First, there's the claim that Conditionalization is a special case of Jeffrey Conditionalization. The thought here is that Jeffrey Conditionalization is just a generalization of Conditionalization, one that returns Conditionalization in the special case in which all of one's evidence partition takes the form $\{(1, E), (0, \neg E)\}$

. This understanding of the relationship between Conditionalization and Jeffrey Conditionalization is a standard part of the Bayesian lore.¹⁸

Second, there's the claim that the formalism of Jeffrey Conditionalization is neutral with respect to whether evidence is 'credence-dependent'. On one way of thinking about evidence, the evidence partition a subject receives is determined by factors that are independent of her credences. For example, one might take a subject's evidence partition to be determined solely by her sensory information. On another way of thinking about evidence, the evidence partition a subject receives is determined by factors that are dependent on her credences. For example, one might take the evidence partition a subject receives to be a function of both her sensory information and her beliefs about what this kind of sensory information suggests. Thus, if the subject sees a figure through the window who looks like her friend, the proposition that her friend is outside might be assigned a large value by her evidence partition if she believes her friend is coming over to visit, but a small value if she believes her friend is out of the country. But while various considerations have convinced many people to favor the second picture of evidence over the first, the formalism of Jeffrey Conditionalization itself has generally been taken to be neutral with respect to these two ways of thinking about evidence.¹⁹

We can see that these two claims and the interval answer are inconsistent as follows. Suppose that the interval answer is true: we should understand Conditionalization as a rule which tells us how to update over arbitrary intervals, given the cumulative evidence we've received during that interval. Given the first claim – that Conditionalization is a special case of Jeffrey Conditionalization – Jeffrey Conditionalization should also be understood this way; as a rule which tells us how to update over arbitrary intervals, given the cumulative evidence we've received during that interval. Given the second claim – that the formalism of Jeffrey Conditionalization is neutral with respect to whether evidence is credence-dependent – it follows that this interval understanding of Jeffrey Conditionalization should be consistent with credence-independent evidence.

But it isn't. If we adopt this interval understanding of Jeffrey Conditionalization, then Jeffrey Conditionalization must yield (on pain of inconsistency) the same prescriptions regardless of what intervals we choose to update on. (For instance, updating on the evidence received during the $[t_0, t_1]$ interval and then updating on the evidence received during the $[t_1, t_2]$ interval had better yield the same result as updating all at once on the evidence received during the $[t_0, t_2]$ interval.) But as I show in Appendix A, if evidence is credence-independent, then Jeffrey Conditionalization's prescriptions *will* depend on what intervals we choose to update on. So the interval understanding of Jeffrey Conditionalization is not consistent with credence-independent evidence.

Thus the interval answer and these two claims are jointly incompatible. Given this answer, Conditionalization is an interval rule. Given the first claim, it follows that Jeffrey Conditionalization is also an interval rule. But given the second claim,

it follows that Jeffrey Conditionalization can't be an interval rule – for if evidence is credence-independent, Jeffrey Conditionalization will yield inconsistent verdicts depending on what intervals we choose to update on.

To sum up, there are compelling reasons in favor of both answers. The sequential answer, which adopts a sequential-updating understanding of Conditionalization, allows us to retain the familiar picture of how Conditionalization and Jeffrey Conditionalization are related, without having to accept the surprising claim that the formalism of Jeffrey Conditionalization requires evidence to be credence-dependent. The interval answer, which adopts an interval-updating understanding of Conditionalization, allows us to apply Conditionalization to cases in which subjects get evidence continuously. And neither of these reasons clearly trumps the other.

So the sequential vs. interval updating question leaves us with two viable ways to understand Conditionalization. There's the understanding suggested by the sequential answer:

Sequential Conditionalization (v2): If a subject with credences cr at t_0 gets evidence E at t_1 (and no evidence between t_0 and t_1), then her credences should remain cr between t_0 and t_1 , and she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined. And there's the understanding suggested by the interval answer:

Interval Conditionalization (v2): If a subject with credences cr at t_0 gets cumulative evidence E in the $[t_0, t_1]$ interval, then she should adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined.

5. Interlude: deontic logic

In preparation for the final two sections, let's pause to say a bit more about the deontic operators that Conditionalization employs. This will provide us with the tools to discuss some natural answers to the third question we'll consider in Section 6, and it will provide us with some of the vocabulary we'll need to spell out the rule precisely in Section 7.

In standard deontic logic, the permission and obligation operators mirror the possibility and necessity operators in modal logic.²⁰ We begin with a set of worlds and an accessibility relation over these worlds, where the accessible worlds are intuitively the 'best' worlds that one can get to from that world. A proposition A is then permissible at w iff it's true at some world accessible to w , and obligatory at w iff it's true at all worlds accessible to w .

On the standard approach, deontic claims are true or false at a world. Thus the same deontic claims apply to every subject at a world, and at every time at that world. But it's natural to want to allow for different subjects at a world to have different obligations, and to allow a subject at different times to have

different obligations.²¹ And we can allow for such variations by taking deontic operators and accessibility relations to be subject and time-indexed.²²

The move to subject and time-indexed deontic operators makes it natural to modify our formulations of Conditionalization in order to make these indices explicit. One natural way to do this is to understand the rule as making prescriptions that are indexed to the subject whose credences we're considering, and to the time at which cr_E should be adopted, as follows:²³

Sequential Conditionalization (v2): If a subject s with credences cr at t_0 gets evidence E at t_1 (and no evidence between t_0 and t_1), then her credences should $_{s,t_1}$ remain cr between t_0 and t_1 and she should $_{s,t_1}$ adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined.

Interval Conditionalization (v2): If a subject s with credences cr at t_0 gets cumulative evidence E in the $[t_0, t_1]$ interval, then she should $_{s,t_1}$ adopt new credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined. I will understand Conditionalization along these lines from now on.²⁴

In order to avoid cluttering the text with subscripts, I'll leave these subject and time-indices implicit when tracking them isn't important, or context makes it clear what they are.

6. The narrow vs. wide scope question

Both versions of Conditionalization (v2^{*}) are conditional norms. That is, they have the form: 'If a subject ..., then she should ...'. Let A be the antecedent clause, C the consequent clause, and O the obligation operator. Then this phrase can be understood to have one of two logical forms: $A \rightarrow O(C)$ or $O(A \rightarrow C)$.²⁵ In which of these two ways should Conditionalization be understood?

Q5. The narrow vs. wide scope question: Does the obligation operator apply to the consequent, or to the entire conditional?²⁶ One way of answering this question takes the obligation operator to take narrow scope, applying to just the consequent:

Answer 1 (Narrow). The rule has the form $A \rightarrow O(C)$. On this understanding, the rule tells us that if a subject actually satisfies A , then she satisfies C at all of the best worlds.

The other way of answering this question takes the obligation operator to have wide scope, applying to the entire conditional:

Answer 2 (Wide). The rule has the form $O(A \rightarrow C)$. On this understanding, the rule tells us that at all of the best worlds, the subject satisfies the conditional $A \rightarrow C$.²⁷

There's a literature on whether we should understand fundamental rational requirements as narrow or wide scope norms in general.²⁸ And one might think that the question of whether Conditionalization should be understood as narrow or wide scope should be deferred to this literature, and to the general

question of whether rational requirements should be understood as narrow or wide scope. I'm inclined to think that this is a mistake. Different considerations come into play for different norms, and we shouldn't expect there to be one general answer to how all norms should be understood.²⁹ In any case, in what follows I'll only consider whether one particular rational requirement – Conditionalization – should be understood as a narrow or wide scope norm.

6.1. *The narrow vs. wide scope question: assessing the answers*

At least at first glance, both answers seem reasonable. The narrow answer is natural when we're thinking of Conditionalization primarily as an updating rule, a constraint on what one's posterior credences should be given one's prior credences. The wide answer is natural when we're thinking of Conditionalization primarily as a diachronic credence constraint, a constraint on how one's prior and posterior credences should line up.

Some people have argued that we should understand all fundamental rational requirements as narrow scope norms.³⁰ Perhaps the most popular argument against wide scope rational requirements is that such rational requirements make symmetrical prescriptions, and it's argued that this symmetry is implausible. Thus, in the case of Conditionalization, the wide scope version of the rule $O(A \rightarrow C) \equiv O(\neg A \vee C)$ requires that subjects either satisfy $\neg A$ or satisfy C , and the rule treats these two options symmetrically: the rule gives us no reason to prefer satisfying C over satisfying $\neg A$. But, it's argued, satisfying C should be preferred to satisfying $\neg A$. If, somehow, we were given a choice between adopting the appropriate new credences given our prior credences and evidence, or adopting whatever new credences we like and having whatever prior credences and evidence would yield these new credences, we should prefer the former option to the latter. And it's argued that the wide scope formulations of rational requirements cannot accommodate this fact. Insofar as these kinds of arguments are compelling, and insofar as we want Conditionalization to be a fundamental norm, this gives us a reason to favor the narrow answer.³¹

That said, there also seem to be considerations that tell against the narrow answer.³² The narrow understanding of the rule takes the form $A \rightarrow O(C)$. Thus, it tells us that if a subject actually satisfies A , then they will satisfy C at all of the best worlds. But this claim seems strange, since it requires the subject to satisfy C at all of the best worlds, regardless of whether the subject satisfies A at those worlds. And it's hard to see why the subject should have to satisfy C at best worlds in which they don't satisfy A .

In light of this worry, one might modify this answer by adding A to the consequent as well:

Answer 1a (Narrow-a). The rule has the form $A \rightarrow O(A \wedge C)$. This formulation of the rule gets around the above worry, since it tells us that a subject who actually satisfies A should satisfy both A and C at all of the best worlds.

But this isn't what we want the rule to say. If $A \wedge C$ is true at all of the best worlds, then A is true at all of the best worlds. So the above norm entails that $A \rightarrow O(A)$. But we don't want A 's actually obtaining to entail that A ought to obtain. In the case of Conditionalization, this would mean that the subject's actually having credences cr and getting evidence E entails that she should have credences cr and get evidence E . And this is implausible.

In light of this, one might consider a second way of salvaging the narrow answer, inspired by the ethics literature, that leaves the form of the rule the same: $A \rightarrow O(C)$.³³ The idea is to avoid the difficulties sketched above by adding certain background assumptions. We will go through this in more detail in a moment, but here's a preliminary sketch of how the story will go.

One can avoid the first worry – that the rule implausibly requires C to be true at all best worlds, even those where A is false – by adding assumptions which entail that if A is currently true, then A is also true at all of the best worlds. Now, this way of avoiding the first worry leads to the second worry: if the current truth of A entails that A is true at all of the best worlds, then the current truth of A entails that A is obligatory. One can blunt this second worry by distinguishing between 'trivial' and 'non-trivial' oughts, and showing that A only entails that A is obligatory in a trivial sense.

Let's go through this in more detail. Standard deontic logic employs a notion of accessibility, call it *O-accessibility*, that determines what things are permissible and obligatory. Let's add to this another notion of accessibility, call it *C-accessibility*, that determines what things we have the capacity to influence.³⁴ So while *O-accessibility* intuitively picks out the normatively 'best' worlds, *C-accessibility* intuitively picks out the worlds that one is capable of getting to.

With these two notions of accessibility in hand, we can impose a constraint that implements the idea that 'ought implies can':

Ought Implies Can: For any world w , the worlds that are *O-accessible* _{s,t} from w are a subset of the worlds that are *C-accessible* _{s,t} from w . It follows from this that a proposition can only be obligatory for a subject if she's capable of bringing about a state of affairs in which that proposition is true.

Given Ought Implies Can, there will be many propositions that are strictly speaking obligatory, but only in an uninteresting sense. For example, take the proposition that $1+1=2$. This proposition will be true at every world, and *a fortiori* at all of the best worlds. Thus the proposition that $1+1=2$ will be obligatory. But it is only obligatory in a trivial sense. Since it's inevitable, nothing a subject can do has any bearing on whether $1+1=2$ turns out to be true or not. I'll call oughts of this kind *trivial oughts*. A proposition A is *trivially obligatory* iff it's true at both every *O-accessible* world (making it obligatory) and every *C-accessible* world (making it inevitable).

Contrast this with the interesting obligations that subjects have, obligations towards propositions which are not inevitable. These are the obligations we

normally talk about, such as the obligation to save a drowning baby, or to adopt the appropriate credences. I'll call these *non-trivial oughts*. A proposition A is *non-trivially obligatory* iff it's true at every O -accessible world (making it obligatory) but not true at every C -accessible world (making the subject capable of falsifying it).

Introducing the notion of C -accessibility also allows us to impose a constraint that implements the idea that the past is immutable:³⁵

The past is immutable: For any world w , the only worlds that are C -accessible _{s,t} to w are worlds that agree with w about the state of the world before t . It follows from this that we can't change what's in the past. (Note that if we didn't take the accessibility relation to be relativized to times, we couldn't formulate a constraint like this.) Finally, let's suppose that the posterior answer to the time of evidence question is true. I argued in Section 3.1 that we should adopt the concurrent answer, and maintain that subjects should get E and adopt cr_E simultaneously. But now let's suppose that we adopt the posterior answer instead, and hold that subjects get E some amount of time before they adopt cr_E .³⁶

With all this in hand, let's turn back to the Narrow vs. Wide Scope question. Suppose we adopt the narrow answer to this question, and maintain that the rule has the form $A \rightarrow O(C)$. We now have the resources to get around the first worry raised for the narrow answer: that the rule requires C to be true at all of the best worlds, regardless of whether A is true at those worlds.

Given the posterior answer to the time of evidence question, A concerns solely facts that occur at t_0 – the subject's having credences cr at t_0 and the subject getting evidence E at t_0 . So the truth value of A will be fixed by the state of the world before t_1 . Given the past is immutable, it follows that A will be true at all of the C -accessible _{s,t_1} worlds, and thus (given Ought Implies Can) at all of the C -accessible _{s,t_1} worlds. Thus, given this set-up, the first worry no longer arises: if A is actually true, then A will be true at all of the best _{s,t_1} worlds too, and the rule won't have the awkward consequence of potentially requiring C to be true at best _{s,t_1} worlds where A is not.

Now, this way of avoiding the first worry leads to the second worry raised above for answer 1a: that if A is true at all of the best worlds, then A is obligatory, which seems implausible. But we now have the resources to explain away this worry too. For while it's true that on this account A is obligatory, it's only *trivially* obligatory, since it's true at all of the C -accessible worlds. So it won't be the case that A is obligatory in any sense that should bother us.

Thus if we adopt this set-up, we can lay out a way of understanding the narrow answer that avoids the worries raised above. Furthermore, this set-up fits nicely with the motivations for the posterior answer to the time of evidence question. Since the posterior answer is largely motivated by ought-implies-can worries, it's a natural fit with this way of fleshing out the narrow answer, which also appeals to ought-implies-can-like considerations.

This is, I think, the closest one can get to a viable narrow scope understanding of Conditionalization. But ultimately, even this account does not give us a compelling reason to adopt the narrow answer. First, given this account, we have no reason to favor a narrow scope over a wide scope understanding of Conditionalization. As I show in Appendix C, given the assumptions we've been making, the narrow and wide scope formulations of Conditionalization are equivalent.³⁷ And this equivalence undercuts the arguments that have been offered in favor of narrow over wide scope understandings of rational requirements. For example, the argument that wide scope understandings yield implausibly symmetric prescriptions while narrow scope understandings do not won't work since both understandings will yield the same prescriptions.³⁸

Second, this account inherits the demerits of the posterior answer to the time of evidence question. So even this version of the narrow answer carries some unpleasant baggage.

Third, all of the versions of the narrow answer we've considered, including this one, yield implausible verdicts in cases in which the subject's t_0 credences are irrational.³⁹ Consider a variant of an example we discussed in section 3.1. Suppose a subject both violates the Principal Principle and receives evidence E . The subject can conditionalize on E , or adopt credences that satisfy the Principal Principle, but it's impossible for her to do both (since conditionalizing on E yields credences that violate the Principal Principle). The narrow answer yields the undesirable result that this subject must always update on E . ($A \rightarrow O(C)$ will be true *iff* either A is false or at all best worlds C is true; and since A is true, it follows that at all best worlds C must be true.) By contrast, the wide answer yields the desired result that the subject needn't update on E . ($O(A \rightarrow C)$ will be true *iff* at all best worlds either A is false or C is true; and since A commits the subject to irrational credences, A will be false at all best worlds. So it doesn't follow that C must be true at all (or even any) best worlds.)⁴⁰

So, when all is said and done, we're left with little motivation to adopt the narrow answer. In the best case scenario for the narrow answer, the narrow and wide answers are equivalent. And even in this best case scenario, the narrow answer has unpleasant consequences, as it inherits the demerits of the posterior answer and yields implausible verdicts in cases in which a subject's starting credences are irrational.

So the wide scope understanding of Conditionalization is the better of the two options. Making the wide scope understanding explicit gives us the following formulations:

Sequential Conditionalization (v3): It should _{s, t_1} be the case that if a subject s with credences cr at t_0 gets evidence E at t_1 (and no evidence between t_0 and t_1), then her credences will remain cr between t_0 and t_1 are cr , and she will adopt credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined.

Interval Conditionalization (v3): It should _{s, t_1} be the case that if a subject s with credences cr at t_0 gets cumulative evidence E in the $[t_0, t_1]$ interval, then she will adopt credences cr_E at t_1 such that $cr_E(\cdot) = cr(\cdot|E)$, if defined.

7. Precise formulations

In the previous sections, we considered three questions about how Conditionalization should be understood. And if we adopt the answers I suggested, we end up with two viable formulations of Conditionalization: the Sequential and Interval (v3) formulations. But one might reasonably worry about whether we've addressed all of the open questions regarding how to formulate Conditionalization. After all, the formulations of Conditionalization I've provided are in English. And until we've spelled out these rules in logical form, it's reasonable to worry that there may be further ambiguities, and thus further questions, that we've missed. So let's conclude by providing formal characterizations of these two rules.

To start, let's introduce the non-logical vocabulary that we need to formulate these rules.

First, we need a predicate that characterizes what a subject's credences are. Let the *credence* predicate $C(s, t, f)$ represent a 3-place relation holding between a subject s , a time t , and a function f which assigns real numbers to (at least some) propositions. Intuitively, $C(s, t, f)$ holds *iff* s 's credence function at t is f .

Second, we need a predicate that ensures that the conditional credence function to be prescribed is well-defined. Let the *well-defined* predicate $W(f, E)$ represent a 2-place relation between a function f that assigns real numbers to propositions, and a proposition E . Intuitively, $W(f, E)$ holds *iff* $f(\cdot|E)$ is well-defined; i.e. *iff* f is a probability function which assigns a non-zero value to E .

Third, we need a predicate that characterizes what evidence a subject gets, where this will depend on what version of Conditionalization we're considering. For use in the sequential formulation of Conditionalization, let the *sequential evidence* predicate $SE(s, t_-, t, E)$ represent a 4-place relation between a subject s , a time t_- , a time t after t_- , and a proposition E . Intuitively, $SE(s, t_-, t, E)$ holds *iff* either (i) s doesn't receive any evidence in the $[t_-, t]$ interval, and receives E as evidence at t , or (ii) s doesn't receive any evidence in the $[t_-, t]$ interval, and E is the trivial proposition Ω . (The second clause is needed to ensure that the rule not only requires subjects to change their credences in the right way when they get evidence, but to also not change their credences when they don't get evidence.)

For use in the interval formulation of Conditionalization, let the *interval evidence* predicate $IE(s, t_-, t, E)$ represent a 4-place relation between a subject s , a time t_- , a time t after t_- , and a proposition E . Intuitively, $IE(s, t_-, t, E)$ holds *iff* E is the cumulative evidence s receives during the interval $[t_-, t]$.

Finally, with all of this in hand, we can formulate our two versions of Conditionalization in logical form as follows:

Sequential Conditionalization (final):

$$\forall E, \forall f, \forall s, \forall t, \forall t- < t, O_{s,t} \left(C(s, t-, f) \wedge SE(s, t-, t, E) \wedge W(f, E) \rightarrow C(s, t, f(.|E)) \right)$$

Interval Conditionalization (final):

$$\forall E, \forall f, \forall s, \forall t, \forall t- < t, O_{s,t} \left(C(s, t-, f) \wedge IE(s, t-, t, E) \wedge W(f, E) \rightarrow C(s, t, f(.|E)) \right)$$

These, I maintain, are the two best ways to understand Conditionalization.⁴¹

Notes

- 1 Of course, questions of how to understand the concepts the rule employs and questions of how to formulate the rule aren't entirely distinct. For example, different choices regarding how to think of credences can have a bearing on the logical form of the predicate one uses to represent (say) having a certain credence in a proposition. Likewise, different choices regarding how to formulate the rule can make certain views regarding how to think about credences more or less attractive. But, the discussion that follows will be focused primarily on questions of logical form.
- 2 The 'virtually' caveat is required because of the potential gap between having a credence of 1 in something and being certain of it. (E.g. one should have a credence of 1 that a countably infinite number of fair coin tosses will land tails at least once, but one shouldn't be certain of this.)
- 3 That is, a subject's credences should assign values to propositions that form a Boolean algebra – a set of propositions closed under conjunction and negation – and the values it assigns should satisfy the (finite) probability axioms: (i) $cr(\cdot) \geq 0$, (ii) $cr(\Omega) = 1$ (where Ω is the trivially true proposition), (iii) if A and B are mutually exclusive, then $cr(A) + p(B) = cr(A \vee B)$.
- 4 Of course, there are a number of different formulations of Conditionalization in the literature to choose from. I've chosen this one because it leaves all three of the questions I'll be exploring open. In order to figure out how to best formulate Conditionalization, we need to assess the different possible answers to these questions. And, a formulation of Conditionalization which leaves all three of these questions open, like the one given above, provides us with an ideal place to start.
- 5 In the next section, we'll be considering whether we should be thinking of E as something you get at a particular time, or as the cumulative evidence one receives over some interval of time. In the latter case, I'll understand ' $t(E)$ ' as the final time over which one gets E ; i.e. the future endpoint of this interval.
- 6 I'm assuming in the text that there's a time $t(E)$ at which one gets E as evidence and a time $t(cr_E)$ at which one should first adopt cr_E . But, this needn't be the case. It could be that the temporal interval during which one has E as evidence is an open interval, so that there is no initial time at which E is received. Likewise, it could be that the temporal interval during which one should have credences cr_E in an open interval, so that there is no first time at which cr_E should be adopted. Although these possibilities complicate things, they don't end up changing the dialectic in any interesting way (see footnote 14). So, I'll put these possibilities aside.
- 7 The formulations of Conditionalization given in the literature vary widely with respect to this question: some assume the posterior answer (e.g. Earman 1992,

Howson and Urbach 2006), some assume the concurrent answer (e.g. (e.g. Lewis 2010 and Strevens 2015), and some are silent on the matter (e.g. Easwaran 2011 and Weisberg 2011).

- 8 I'm assuming here that the time at which a subject 'gets' evidence is the time at which the evidence becomes relevant to our epistemic evaluation of the subject. But I do not take this to be controversial, since this assumption is shared by virtually everyone. (Consider: if this were not the case, then the oft repeated truism that a subject's beliefs should take all of her evidence into account (the so-called 'Principle of Total Evidence') would be untenable. If, for example, one characterized 'receiving evidence' such that the time at which a subject received visual evidence was one minute (or one year, or one century) before the light struck her eyes, it would be implausible to say that she should take all of her evidence into account.)
- 9 See Lewis (1980).
- 10 Other proponents of this kind of stance regarding epistemic norms include Feldman (2001) and Wolterstorff (2010).
- 11 One might reasonably want to hear more about what, exactly, it means to say that Conditionalization is 'an ideal at which to aim' or an ideal performance norm', and about how this understanding of Conditionalization interacts with things like a subject's cognitive capabilities and ought-implies-can. I describe one natural way of spelling out these notions and their interaction, using the framework of Kratzer (1991), in Appendix B.
- 12 Of course, this kind of worry won't arise given certain natural pictures of what receiving and updating on evidence is like. For example, suppose one adopts an account of evidence (such as Howson and Urbach's 1993 or Williamson's 2000) according to which receiving evidence *E* is, at least in part, a matter of coming to believe *E*. And, suppose one takes such belief changes to be global and concurrent. That is, just as stepping on to a trampoline changes the elevation of both the place one's standing and the surrounding area, getting *E* as evidence changes both one's credence in *E* and one's credence in the 'surrounding' propositions. And just as the full change in elevation of the place one steps doesn't happen *before* any of the other changes in elevation take place – the changes in elevation of different parts of the trampoline are (roughly) concurrent – getting *E* as evidence doesn't take place before these other belief changes take place, these belief changes are concurrent. On this picture, Conditionalization is naturally thought of as describing what the shape of these global concurrent belief changes should be like. And, the worries regarding potential conflicts between instantaneous updating and procedural norms described in the text won't arise on such a picture.
- 13 In making this reply, I assume that in order for a set of norms to provide a coherent standard of ideal performance, or a coherent ideal for us to aim at, it only has to be logically possible to satisfy them. But one might wonder why these ideals shouldn't also have to be *metaphysically* possible to live up to. And if they do, and if instantaneous causation is metaphysically impossible, then it seems that these natural procedural norms and the concurrent understanding of Conditionalization *are* in conflict after all. (Having the right beliefs (those prescribed by Conditionalization) in the right way (via some causal process initiated by the receipt of evidence) seems to require instantaneous causation, since one's beliefs must instantly change in light of one's evidence. Thus, if instantaneous causation is metaphysically impossible, then so is jointly satisfying Conditionalization and these procedural norms. And if jointly satisfying Conditionalization and these procedural norms has to be metaphysically possible in order for these norms

to be jointly true, then they can't be jointly true.)⁶ So why do I require these ideals to be logically possible, but not metaphysically possible? Here is why. For something to usefully serve as an ideal at which to aim, it needs to be something which we can approach by degrees; something for which we can discern paths of states that lead to it such that each state along the path gets closer to satisfying the ideal. Thus, metaphysically impossible norms can serve as useful ideals at which to aim: one can make sense of moving toward or away from the ideal of updating instantly by updating more or less quickly. But it's hard to see how logically impossible ideals could serve as useful ideals at which to aim: it's not clear how one could move closer or father away from the ideal of being a round square, or a married bachelor.

- 14 As noted in footnote 6, I've simplified this discussion by assuming that there are particular times $t(E)$ and $t(cr_E)$ at which a subject receives her evidence and should adopt her new credences. But this needn't be the case (e.g. if the period during which a subject should have cr_E is an open interval). Introducing these possibilities complicates the dialectic in two ways, but these complications end up effectively canceling each other out. The first complication is that these possibilities leave us with *three* natural ways to group answers to the question of how the receipt of E and the adoption of cr_E are related:

- (1) E is received before cr_E is adopted, and there is a gap between the receipt of E and cr_E .
- (2) E is received at the same time as cr_E is adopted, and thus there's no gap between them.
- (3) E is received before cr_E is adopted, but there is no gap between the receipt of E and cr_E (e.g. E 's received at some time t , and cr_E is adopted at the open endpoint of an interval starting at t).

(1) and (2) correspond to the posterior and concurrent answers considered in the text, while (3) is a possibility which only comes into view once we drop the simplifying assumption. The second complication is that these possibilities allow us to see that the worries raised for the different answers track slightly different issues. The worries regarding total evidence and fit with Howson-and Urbach-like pictures of evidence arise for any view on which E is received before cr_E is adopted (thus applying to (1) and (3)). The worries regarding ought implies can and substantive and procedural requirements arise for any view on which there's no gap between the receipt of E and the adoption of cr_E (thus applying to (2) and (3)). But together these complications allow us to see that no interesting positions are left out by simplifying and ignoring (3). For (3) is strictly less appealing than (1) and (2), as it is subject to the worries raised for both.

- 15 Of course, it's a well known feature of Conditionalization that conditionalizing on E and then on F yields the same result as conditionalizing on the conjunct $E \wedge F$. So one could obtain the result either way. But, on this understanding of the rule, what's really going on 'under the hood' is a sequence of updates; updating on the conjunction is merely a convenient calculational short cut.

- 16 The notion of 'cumulative evidence' is best thought of as one of the basic notions that this interval updating understanding of conditionalization employs. (I.e. whereas the 'sequential updating' picture takes as basic a notion of getting a piece of evidence at a particular time, the 'interval updating' picture takes as basic a notion of getting cumulative evidence over an interval.) Of course, when discussing Conditionalization, we have the option of understanding the *cumulative evidence* a subject receives during an interval as the logically weakest proposition that entails all of the evidence she receives during that interval (or, equivalently, the conjunction of all of the evidence she receives during that interval). This would allow both sequential and interval updating understandings of Conditionalization to take as basic the notion of getting a piece of evidence at a particular time. But when we turn to discuss Jeffrey Conditionalization, in Section 4.1, we lose the option of providing a reductive understanding of the notion of 'cumulative evidence', since Jeffrey Conditionalization deals with weighted evidence partitions to which the notions of entailment and conjunction don't apply.
- 17 The formulations of Conditionalization given in the literature vary widely with respect to this question: some assume the sequential answer (e.g. Earman 1992; Howson and Urbach 2006), others assume the interval answer (e.g. Lewis 2010), and yet others are silent on the matter (e.g. Earman 1992; Weisberg 2011; Strevens 2015).
- 18 Though see Christensen (1992) for a notable dissent: '[Jeffrey Conditionalization] is thus not simply an elegant generalization of [Conditionalization], a pure improvement which merely removes some gratuitous idealization. It removes idealization, but at a price. The additional cases covered by the liberal model are not covered in the same way; and consequently, the account as a whole must be given a different philosophical interpretation' (547).
- 19 A number of people, including Levi (1967), Carnap (in Jeffrey 1975), Field (1978), Christensen (1992) and Lange (2000), have maintained that what evidence partition a subject receives should depend on what the subject's credences are. But they've argued for this under the assumption that it's a substantive question left open by the formalism. And thus they've appealed to various intuitive and epistemic considerations to make their case. If the formalism of Jeffrey Conditionalization itself required evidence to be credence-dependent, then the kinds of considerations that Christensen and others have offered in support of this claim would be superfluous. No interesting discussion needs to take place to establish that evidence is credence-dependent if the formalism itself entails it.
- 20 For discussions of standard deontic logic, see Aqvist (2002) and McNamara (2010).
- 21 See Feldman (1986) for some reasons for wanting to have this kind of subject and time sensitivity in the context of ethics.
- 22 A different (and more powerful) way to allow for such variations is to take deontic operators to range over centered worlds instead of worlds. But since this approach raises complications orthogonal to the issues at hand, I employ the more traditional approach of subject and time indexing the deontic operators in the text.
- 23 Recognizing the time-indexed nature of the obligation operator could be seen as raising a fourth question, in addition to the three questions discussed in the text: a 'Time of Obligation' question, regarding what times the obligation operators should be indexed to. One answer is that they're indexed to the initial time we're considering (t_0), another is that they're indexed to the final time we're considering (t_1), and a third is that they're indexed to all times – the norm requires a subject to satisfy this constraint at every time, regardless of the times the prescription

involves. I tentatively favor the second and third answers over the first, since according to the first answer, it will never be the case that a subject's obligated to adopt cr_E at the time at which she's supposed to adopt it – by the time t_1 rolls around, she'll only have obligations to adopt some further credences in the future. And I adopt the second answer (instead of the third) in the text because the most plausible version of the narrow scope answer, which I'll be arguing against in section 6, requires this time to be t_1 . (Looking ahead: we need the time to be t_1 in order to ensure that the antecedent is true at all C-accessible worlds and yet the consequent is not.) So adopting the second answer allows me to stack the deck in favor of my opponents.

- 24 In linguistics, the standard framework for formalizing claims involving modals is the framework developed by Kratzer (1991). Although the discussion in the text assumes we're using standard deontic logic, everything I say can be translated into Kratzer's more sophisticated framework. Indeed, moving to Kratzer's framework offers some benefits, for it provides us with the tools to spell out several things that it's difficult to flesh out using standard deontic logic. Since working out these details takes a bit of time, I've relegated my discussion of how to set things up using Kratzer's framework to Appendix B.
- 25 See Broome (1999).
- 26 In the literature, Conditionalization is generally presented in a way that is neutral with respect to this question (such as in Earman 1992; Howson and Urbach 2006; Lewis 2010; Easwaran 2011; Weisberg 2011; Strevens 2015).
- 27 Broome (2007) shows that changing a norm from narrow to wide scope or vice versa won't change whether the actual world is one of the best worlds. Given this, one might worry about whether there's anything substantive at stake here. But these two answers are logically distinct. (E.g. if A is false and at all of the best worlds A is true and C is false, then the narrow scope conditional will be true but the wide scope one false; while if A is false and at all of the best worlds A and C are false, then the narrow scope conditional will be false but the wide scope conditional true.) And, as we'll see, these two answers will have importantly different implications.
- 28 For some of the recent literature bearing on this issue, see Broome (1999), Schroeder (2004), Broome (2007), Kolodny (2007), Bedke (2009), Brunero (2010), Way (2011), Brunero (2012), Lord (2013), Shpall (2013), and Titelbaum (forthcoming). (The reason it's *fundamental* rational requirements that are of interest is because one might be able to derive some wide scope norms from narrow scope norms or vice versa. So the interesting claims aren't whether there are any wide/narrow scope rational requirements, but rather whether there are any fundamental (i.e. non-derivative) wide/narrow scope rational requirements.)
- 29 For example, in the next section, certain questions regarding how we understand Conditionalization (such as how we answer the Time of Evidence question) will bear on the plausibility of wide vs. narrow scope understandings of the rule. We wouldn't expect the same dialectic to play out for other norms – the shape of this dialectic is particular to Conditionalization.
- 30 For examples of arguments in favor of narrow scope understandings of fundamental rational requirements, see Schroeder (2004), Kolodny (2009). It's worth noting that the terms 'narrow scope' and 'wide scope' are used somewhat equivocally in this literature; see Titelbaum (forthcoming). But everyone calls norms of the form $A \rightarrow O(C)$ narrow scope, and norms of the form $O(A \rightarrow C)$ wide scope, so we can skirt these complications here.

- 31 For discussions of these kinds of symmetry arguments, see Way (2011), Brunero (2012), Lord (2013) and Shpall (2013).
- 32 In the text, I focus on worries regarding whether the narrow answer has the right normative profile. For further kinds of arguments against narrow scope understandings of rational requirements, see Brunero (2010) and Shpall (2013).
- 33 In particular, this proposal follows many of the ideas regarding how to set up deontic structure laid out in Feldman (1986), chapter 2.
- 34 The notion of 'C-accessibility' is the same as the notion of accessibility described in Feldman (1986), chapter 2.
- 35 Though see Feldman (1986), Section 2.1.1 for some reasons against adopting such a constraint.
- 36 The reason we need to assume the posterior answer is that it ensures there's a temporal gap between the events described in the antecedent *A* (having credences cr and receiving evidence *E*) and the events described in the consequent *C* (adopting credences cr_E). And this temporal gap is crucial to getting the maneuver described in the text to work. (Briefly: we need *A* to be in the past to ensure that *A* is true at all best worlds but only trivially obligatory, and we need *C* to not be in the past to ensure that *C* isn't trivially obligatory.)
- 37 It's important to not confuse this result with the result shown by Broome (2007) and discussed in footnote 27. Broome's result is both broader and weaker. Broome shows that, in general, switching narrow and wide scope formulations of a rule won't change whether the actual world is one of the best worlds. This result yields the stronger conclusion that the narrow and wide scope formulations will logically entail each other, but only in a narrower class of cases (those in which the norm is Conditionalization, and the assumptions we've been making hold).
- 38 In a similar fashion, this equivalence undercuts Kolodny's (2007) 'Problem of Conflict' for wide scope understandings of Conditionalization, the worry that '[s]ome requirements of formal coherence not only are not explained by a concern for the true and the good, but moreover would forbid what that concern requires'. (Kolodny 2007, 231) This kind of argument cannot support the narrow over wide scope understandings of Conditionalization, since the two understandings make the same prescriptions, and thus 'forbid' the same things. Likewise, it undercuts Kolodny's (2007) 'Problem of Normativity' for wide scope understandings of Conditionalization, the worry that there's no plausible explanation for how wide scope rational requirements could have normative force. Since the narrow and wide scope understandings of Conditionalization are equivalent, any considerations that one could use to justify one of these sets of prescriptions could also be used to justify the other.
- 39 See Brunero (2010) for a more general version of this worry for narrow scope norms.
- 40 This assumes that we don't adopt the additional assumptions sketched above. If we do adopt those assumptions, then the wide answer will yield the same undesirable results as the narrow answer (as one would expect, since given these assumptions the narrow and wide answers are equivalent).
- 41 I'd like to thank the May 2014 UMass Brown Bag Presentation group and the audience of the 2015 Belief, Rationality and Action over Time conference for helpful comments and discussion. In addition, I'd like to thank Jennifer Carr for flagging the third worry for the narrow answer discussed in Section 6.1, Brian Hedden for discussion about the ways of thinking about evidence mentioned in footnote 12, Sarah Moss and Ralph Wedgwood for discussion about Kratzer semantics and the issues discussed in Appendix B, Miriam Schoenfeld for suggesting the names

for the answers to the Time of Evidence Question, and for pushing me to address the issues discussed in footnote 13, and Michael Titelbaum for pushing me to get clearer on what one might mean by things like ‘an ideal at which to aim’ (which I now try to do in Appendix B). Finally, I owe special thanks to Lisa Cassell, Maya Eddon, and Alejandro Perez-Carballo, for detailed comments on the entire paper, which led to more substantive improvements than I could reasonably list.

- 42 I think these are plausible assumptions for the ordering base corresponding to the notions of epistemic obligation and permission. But not everyone would agree. For example, Christensen (2007) argues that there are inconsistent epistemic ideals.
- 43 For simplicity, I talk here as if we were taking the modal base and ordering source to be functions from worlds to set of propositions, and to not be indexed to anything. But I think our final account will want these functions to be indexed to subjects and times (or to be functions from centered worlds); see Section 5.
- 44 To see that these two thoughts are equivalent: given the assumption about the modal base, the formulations of Conditionalization given in the text will entail that at every ideal world subjects conditionalize (in the manner specified by that formulation), and thus that the ordering source must entail that subjects conditionalize. Going the other way, if the ordering source entails that subjects conditionalize (in the manner corresponding to some formulation), then it follows that subjects conditionalize at every ideal world, and thus (given the assumption about the modal base) that subjects *should* conditionalize, which is just what the corresponding formulation of Conditionalization asserts.
- 45 Of course, none of the formulations of Conditionalization we’ll consider specify what this range of weaker propositions is. But this isn’t something we should expect from Conditionalization – spelling out how to best decompose the proposition that subjects conditionalize into these weaker claims is a task which requires a lot more information than a simple rule like Conditionalization could provide.
- 46 The wide scope rule requires that $\neg A \vee C$ be true at all O-accessible worlds. If $O(A)$ is true, then the first part of that disjunct ($\neg A$) will be false at all O-accessible worlds, and thus the second part of the disjunct (C) must be true at all O-accessible worlds. Thus $O(C)$ is true.

Notes on contributors

Christopher J. G. Meacham is an Associate Professor at the University of Massachusetts, Amherst. His primary research interests are in formal epistemology, decision theory, and the philosophy of physics.

References

- Aqvist, L. 2002. "Deontic Logic". In *Handbook of Philosophical Logic*, edited by D. M. Gabbay and F. Geunther. Dordrecht, Netherlands: Kluwer Academic Publishers.
- Bedke, M. S. 2009. "The Iffiest Oughts: A Guise of Reasons account of Endgiven Conditionals." *Ethics* 119 (4):672–698.
- Broome, J. 1999. "Normative Requirements." *Ratio* 12 (4):398–419.
- Broome, J. 2007. "Wide or Narrow Scope?" *Mind* 116 (462):359–370.

- Brunero, J. 2010. "The Scope of Rational Requirements." *Philosophical Quarterly* 60 (238):28–49.
- Brunero, J. 2012. "Instrumental Rationality, Symmetry and Scope." *Philosophical Studies* 157 (1):125–140.
- Christensen, D. 1992. "Confirmational Holism and Bayesian Epistemology." *Philosophy of Science* 59 (4):540–557.
- Christensen, D. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief*. New York, NY: Oxford University Press.
- Christensen, D. 2007. "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals." *Oxford Studies in Epistemology* 2:3–31.
- Earman, J. 1992. *Bayes or Bust?* Cambridge, MA: MIT Press
- Easwaran, K. 2011. "Bayesianism I: Introduction and Arguments in Favor." *Philosophy Compass* 6 (5):312–320.
- Feldman, F. 1986. *Doing the Best We Can: An Essay in Informal Deontic Logic*. Dordrecht, Netherlands: D. Reidel Publishing Company.
- Feldman, R. 2001. "Voluntary Belief and Epistemic Evaluation". In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by M. Steup. 77–92. Oxford: Oxford University Press.
- Field, H. 1978. "A Note on Jeffrey Conditionalization." *Philosophy of Science* 45 (3):361–367.
- Hacquard, V. 2011. "Modality". In *Semantics: An International Handbook of Natural Language Meaning*, edited by C. Maienborn, K. von Heusinger, and P. Portner. 1484–1515. Berlin: Mouton de Gruyter.
- Howson, C., and P. Urbach. 1993. *Scientific Reasoning: The Bayesian Approach*. Chicago: Open Court.
- Howson, C., and P. Urbach. 2006. *Scientific Reasoning: The Bayesian Approach*, Chicago: 3rd ed. Chicago: Open Court.
- Jeffrey, R. C. 1975. "Carnap's Empiricism." *Minnesota Studies in the Philosophy of Science* 6:37–49.
- Kolodny, N. 2007. "How does Coherence Matter?" *Proceedings of the Aristotelian Society* 107 (1pt3):229–263.
- Kratzer, A. 1991. "Modality". In *Semantics: An International Handbook of Contemporary Research*, edited by D. W. Arnim von Stechow. 639–650. Berlin: W. de Gruyter.
- Lange, M. 2000. "Is Jeffrey Conditionalization Defective by Virtue of being Non-commutative? Remarks on the Sameness of Sensory Experiences." *Synthese* 123 (3):393–403.
- Levi, I. 1967. "Probability Kinematics." *British Journal for the Philosophy of Science* 18 (3):197–209.
- Lewis, D. 1980. "A Subjectivist's Guide to Objective Chance". In *Studies in Inductive Logic and Probability*, edited by R. C. Jeffrey. Vol. 2, 83–132. Berkeley: University of California Press.
- Lewis, D. 2010. "Why Conditionalize". In *Philosophy of Probability: Contemporary Readings*, edited by Antony Eagle, 403–407. London: Routledge.
- Lord, E. 2013. "The Real Symmetry Problem(s) for Wide-Scope accounts of Rationality." *Philosophical Studies* 3:1–22.
- McNamara, P. 2010. "Deontic logic". In *Stanford Encyclopedia of Philosophy*, edited by E. Zalta. Published April 21, 2010. Accessed September 1, 2015. <http://plato.stanford.edu/entries/logic-deontic/>
- Schroeder, M. 2004. "The Scope of Instrumental Reason." *Philosophical Perspectives* 18 (1):337–364.
- Shpall, S. 2013. "Wide and Narrow Scope." *Philosophical Studies* 163 (3):717–736.

- Simon, H. A. 1976. "From Substantive to Procedural Rationality". In *Method and Appraisal in Economics*, edited by S. Latsis 129–148. Cambridge, MA: Cambridge University Press.
- Strevens, M. 2015. *Notes on Bayesian Confirmation Theory*. Unpublished.
- Swanson, E. 2008. "Modality in Language." *Philosophy Compass* 3 (6):1193–1207.
- Titelbaum, M. G. 2015. "How to Derive a Narrow-Scope Requirement from Wide-Scope Requirements." *Philosophical Studies* 172 (2): 535–542.
- Way, J. 2011. "The Symmetry of Rational Requirements." *Philosophical Studies* 155 (2):227–239.
- Weisberg, J. 2011. "Varieties of Bayesianism". In *Handbook of the History of Logic*, edited by D. Gabbay, S. Hartmann, and J. Woods. Vol. 10, 477. Elsevier.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wolterstorff, N. 2010. *Practices of Belief*. Cambridge: Cambridge University Press.

Appendix A

In Section 4.1, I stated that if evidence is credence-independent and we adopt an interval understanding of Jeffrey Conditionalization, then Jeffrey Conditionalization's prescriptions will depend on what intervals we choose to update on. Let's see why this is the case. If evidence is credence-independent, Jeffrey Conditionalization can only yield the same prescriptions regardless of what intervals we update on if the following claim is true:

Cumulative Evidence: For any sequence of evidence partitions $S_1 - S_n$, there is a 'cumulative' evidence partition S such that, for any credence function, consecutively Jeffrey Conditionalizing on $S_1 - S_n$ will yield the same results as Jeffrey Conditionalizing on S . We see that Cumulative Evidence is false by constructing a counterexample.

For simplicity, let's focus on a simple case where the subject's credences are only defined over four possibilities, each corresponding to one of the possible truth values for a pair of propositions A and B . Let S_1 be the evidence partition $\{(A, 1/2), (\neg A, 1/2)\}$, and S_2 the

cr_0	B	$\neg B$
A	1/10	2/10
$\neg A$	3/10	4/10

evidence partition $\{(B, 1/2), (\neg B, 1/2)\}$.

To start, consider a subject whose initial credences cr_0 over these possibilities are:

cr_1	B	$\neg B$
A	1/6	2/6
$\neg A$	3/14	4/14

Jeffrey Conditionalizing cr_0 on S_1 gives us cr_1 :

cr_2	B	$\neg B$
A	7/32	7/26
$\neg A$	9/32	6/26

And Jeffrey Conditionalizing cr_1 on S_2 gives us cr_2 :

What evidence partition S would get us from cr_0 to cr_2 in one step? Well, recall that Jeffrey Conditionalization keeps the ratios of credences within each element of an evidence partition the same. So unless S is maximally fine-grained – unless S is an evidence partition which puts each of these four propositions in a different element – some of the ratios between the credences of cr_0 and the credences of cr_2 will be the same. Since none of the ratios between the credences of cr_0 and the credences of cr_2 are the same, S must be an evidence partition which puts each of these four propositions in a different element of the partition.

If one updates on such a partition, one will adopt the credences it assigns to each of these elements. So the only partition that will get us directly from cr_0 to cr_2 is the one that assigns each of these values directly. Thus, if we want an evidence partition S to yield the cr_0 -to- cr_2 transition in one step, then S must be

cr_0^*	B	$\neg B$
A	1/4	1/4
$\neg A$	1/4	1/4

$\{(A \wedge B, 7/32), (A \wedge \neg B, 7/26), (\neg A \wedge B, 9/32), (\neg A \wedge \neg B, 6/26)\}$.

Now consider a subject who starts out with a different initial credence function, cr_0^* : Jeffrey Conditionalizing cr_0^* on S_1 gives us cr_0^* again. And Jeffrey Conditionalizing cr_0^* on S_2 gives us cr_0^* again. So updating cr_0^* on S_1 and then S_2 will leave the subject's credences unchanged. But Jeffrey Conditionalizing cr_0^* on S will give us cr_2 , which changes the subject's credences quite a bit.

So there is no evidence partition which, for any credence function, will yield the same results as consecutively updating on S_1 and S_2 . For the only evidence partition that will yield the same results as updating on S_1 and S_2 if we start with cr_0 is S , and S will not yield the same results as updating on S_1 and S_2 if we start with cr_0^* . Thus Cumulative Evidence is false.

Appendix B

Although the proposed understanding of Conditionalization presented and defended in the text is characterized using standard deontic logic, one can also characterize this proposal using Kratzer (1991) account of the semantics of modals. And doing so provides us with a natural way to flesh out some more details regarding this proposal, details that it's hard to provide using standard deontic logic.

On Kratzer's framework, the truth values of deontic modals are determined by two functions from worlds to sets of propositions. First, a *modal base*, which yields sets of propositions representing constraints on the range of possibilities – we use the modal base to pick out the relevant worlds by selecting all and only those worlds at which these propositions are true. Second, an *ordering source*, which yields sets of propositions representing the 'ideal' – we use the ordering source to provide a partial ordering over worlds, one that ranks worlds according to how close they are to this ideal (i.e. according to how many of these propositions are true). For simplicity, let's adopt the Limit Assumption (that in any set of worlds there's always a highest ranked subset, instead of a infinite ascending sequence of more highly ranked worlds). Then the thought is that one should A iff A is true at all of the highest ranked accessible worlds. (For a more detailed presentation of Kratzer's framework, see Kratzer (1991), Swanson (2008) and Hacquard (2011)).

In what follows, I'll assume that the ordering source yields the same set of propositions at every world (i.e. that there's a single uniform ideal), and that all of the propositions in the ordering source can be jointly satisfied (i.e. that the ideal is consistent).⁴²

Given this framework, we can accept the formulations of Conditionalization provided in the text, as long as we understand them as implicitly requiring the modal base to be broad enough to admit at least one ideal world – a world at which all of the propositions in the ordering source are true. The formulations given in the text then assert that, given such a modal base, the highest ranked worlds (and thus the worlds a subject should aim for) will be worlds at which subjects satisfy that formulation of Conditionalization.⁴³ This is equivalent to thinking of these formulations of Conditionalization as providing partial descriptions of the ordering source. In particular, we can think of these formulations as requiring the propositions in the ordering source to entail that subjects conditionalize (in the manner specified by that formulation).⁴⁴

I didn't say that the ordering source *consists* of the proposition that subjects conditionalize for two reasons. First, if one thinks that there are other epistemic norms that bind subjects (e.g. the Principal Principle) then propositions concerning these norms will appear in the ordering source as well; so the ordering source won't *just* consist of propositions regarding Conditionalization. Second, if one wants a detailed ranking which takes into account things like partially satisfying the requirement to conditionalize, or approximately conditionalizing to a greater or lesser degree, then one won't want the claim that subjects conditionalize to appear as a single proposition in the ordering source. Instead, one will want to use a batch of weaker propositions that together entail that subjects conditionalize – this allows one to assess the magnitude of deviations from this ideal, by seeing how many of these weaker propositions are violated.⁴⁵

This framework allows us to give the claim that Conditionalization is 'an ideal at which to aim' a precise meaning. At the ideal worlds, in which all of the propositions in the ordering source are satisfied, subjects will conditionalize. And given an appropriate decomposition of the claim that subjects conditionalize into weaker propositions, the ordering source will tell us how to move closer or farther from this ideal – i.e. it will tell us in what direction to aim.

Likewise, this framework provides a way to spell out the relationship between a subject's cognitive limitations and the rule's ability to provide guidance and satisfy ought-implies-can. A norm can only provide a subject with guidance and satisfy ought-implies-can if that subject is cognitively capable of 'getting to' the highest ranked worlds picked out by the modal base. And since Conditionalization, as I'm understanding it, implicitly requires us to work with a modal base that admits at least one ideal world, this norm will only provide guidance and satisfy ought-implies-can for subjects whose cognitive capabilities are powerful enough to allow them to get to ideal worlds. So the norm won't provide guidance or satisfy ought-implies-can for more cognitively limited subjects like ourselves. (If we can complete the difficult task of working out what exactly to put in the ordering base, we can construct a norm that provides guidance to *any* subject: we simply require the modal base to admit all and only worlds that the subject is cognitively capable of getting to, and then direct her towards the subset of those worlds that are highest ranked. This norm would have the same 'normative heart' as Conditionalization, since it makes prescriptions based on the same ordering source. But, unlike Conditionalization, it would not always tell subjects to conditionalize, since many subjects aren't capable of doing so. (If we distinguish between 'evaluative' and 'guidance' norms, then we can think of Conditionalization as the evaluative norm suggested by this ordering source, and this other norm as Conditionalization's 'guidance counterpart' – the guidance norm suggested by the same ordering source.))

Appendix 3

Appendix C

Suppose that Ought Implies Can, that the Past is Immutable, and that the posterior answer to the Time of Evidence Question is correct. Then the narrow and wide scope formulations of Conditionalization will be equivalent.

Let's see why this is so. To begin, recall that since the time of evaluation is t_1 and The Past is Immutable, it follows that anything true before t_1 will be true at all *C-accessible* _{s, t_1} worlds, and *a fortiori* (given that Ought Implies Can) at all *O-accessible* _{s, t_1} worlds. Given the posterior answer to the Time of Evidence question, A only describes events at t_0 , so it follows that if A is true then $O_{s, t_1}(A)$ is true, and likewise if $\neg A$ is true then $O_{s, t_1}(\neg A)$ is true. For legibility, let's suppress the s and t_1 indices on the O operators. We can now see that the narrow scope answer ($A \rightarrow O(C)$) entails the wide scope answer ($O(A \rightarrow C)$). At the actual world, either A or $\neg A$ is true. If A is true, then it follows from the narrow scope rule that $O(C)$ is true. Thus $O(\neg A \vee C) \equiv O(A \rightarrow C)$ is true. On the other hand, if $\neg A$ is true, then $O(\neg A)$ is true. Thus $O(\neg A \vee C) \equiv O(A \rightarrow C)$ is true. Either way, if the narrow scope rule is true, then the wide scope rule is true as well.

Likewise, given the above, the wide scope answer ($O(A \rightarrow C)$) entails the narrow scope answer ($A \rightarrow O(C)$). At the actual world, either A or $\neg A$ is true. If A is true, then $O(A)$ is true. If $O(A)$ is true, it follows from the wide scope rule that $O(C)$ is true as well.⁴⁶ Thus $\neg A \vee O(C) \equiv A \rightarrow O(C)$ is true. On the other hand, if $\neg A$ is true, then it follows immediately that $\neg A \vee O(C) \equiv A \rightarrow O(C)$ is true. Either way, if the wide scope rule is true, then the narrow scope rule is true as well. So, given these assumptions, the narrow scope and wide scope understandings are equivalent.