

ARTICLE

Commutativity, Normativity, and Holism: Lange Revisited

Lisa Cassell

Department of Philosophy, University of Maryland, Baltimore County, Baltimore, MD, USA
Email: lisarcassell@gmail.com

Abstract

Lange (2000) famously argues that although Jeffrey Conditionalization is non-commutative over evidence, it's not defective in virtue of this feature. Since reversing the order of the evidence in a sequence of updates that don't commute does not reverse the order of the experiences that underwrite these revisions, the conditions required to generate commutativity failure at the level of experience will fail to hold in cases where we get commutativity failure at the level of evidence. If our interest in commutativity is, fundamentally, an interest in the order-invariance of information, an updating sequence that does not violate such a principle at the more fundamental level of experiential information should not be deemed defective. This paper claims that Lange's argument fails as a general defense of the Jeffrey framework. Lange's argument entails that the inputs to the Jeffrey framework differ from those of classical Bayesian Conditionalization in a way that makes them defective. Therefore, either the Jeffrey framework is defective in virtue of not commuting its inputs, or else it is defective in virtue of commuting the wrong kinds of ones.

Keywords: Commutativity; Jeffrey Conditionalization; Bayesian updating; experience; evidence; confirmation holism; Bayes factors

1. Introduction

The non-commutativity of Jeffrey Conditionalization is typically regarded as a mark against it. The reason isn't difficult to grasp. Consistency seems to require that identical pieces of evidence be treated the same no matter the order in which they are received. A non-commutative updating rule seems to flout this requirement.

Lange (2000) famously argues that although Jeffrey Conditionalization is non-commutative over weighted evidence partitions, the Jeffrey framework isn't defective in virtue of this feature. This is because reversing the order of the evidence in a sequence of non-commutative updates does not reverse the order of the experiences that underwrite these revisions. If our interest in commutativity is, fundamentally, an interest in the order-invariance of information, an updating sequence that does not violate such a principle at the more fundamental level of *experiential* information should not be deemed defective.

This paper claims that Lange's argument fails as a general defense of the Jeffrey framework. Lange's argument entails that the inputs to the Jeffrey framework differ from those of classical Bayesian Conditionalization in a way that makes them defective. This means that either the Jeffrey framework is defective in virtue of not commuting its inputs, or else it is defective in virtue of commuting the wrong kinds of ones.

The main contribution this paper makes, then, is to spell out what Lange's argument teaches us about the normative structure of the Bayesian updating framework, and, in particular, about what it teaches us about the relation between classical Bayesian Conditionalization and Jeffrey Conditionalization. It's widely believed that the question of when an agent *has* the evidence that she updates on falls beyond the purview of a Bayesian theory of rational belief revision. Most would say that it's no

criticism of the Bayesian framework that it allows a coherent agent to update on bad information—on information that leads her astray. But Lange’s defense of the Jeffrey framework entails something much stronger. If we take the commutative property to be a constraint on the adequacy of the inputs to the updating process, as Lange’s argument assumes that we should, then we are left with a framework that not only does not entail *some* account of evidence, but which is incompatible with *any* account of evidence. The Jeffrey framework is defective in virtue of what the commutative property implies about its inputs in a way that the classical Bayesian framework is not.

Here’s how this discussion will go. In section 2, I outline Lange’s argument. In section 3 and section 4, I consider the account of experience that Lange’s argument assumes. I argue that the account of experience we would need to defend Jeffrey Conditionalization against the charge that it is defective in virtue of being non-commutative is a stronger account of experience than the one that Jeffrey Conditionalization entails. Therefore, the norm that isn’t defective in virtue of being non-commutative is actually a different norm than Jeffrey Conditionalization. In section 5, I argue that this different norm faces a problem that classical Bayesian Conditionalization avoids: it entails that the inputs to the Jeffrey framework cannot be adequately governed by any substantive constraint. The Jeffrey framework is still defective by virtue of the desideratum that its inputs commute in a way that the classical Bayesian framework is not. Finally, in section 6, I briefly consider how the history surrounding this problem supports these results, and how these results, in turn, help to unify the ideas in some of these earlier discussions of the problem.

2. Lange’s argument

Standard Bayesianism assumes that an agent’s degrees of confidence, or credences, in the propositions she entertains can be represented as an assignment of real numbers in the unit interval to those propositions. It further assumes that the following synchronic norm governs this assignment:

Probabilism: An agent’s credences should obey the probability axioms.

More importantly, for our purposes, most Bayesians also take an agent’s credences to satisfy a diachronic constraint. This constraint tells us how we should update our beliefs when we are presented with some new evidence:

Bayesian Conditionalization: When B represents the total information you have acquired, your new degree of belief in A, for any A, should be $p'(A) = p(A|B)$.

Since our evidence always ought to get a credence of one on this updating framework, a commitment that this framework requires us to take on is that we should always be certain of our evidence. But, to many, it seems as though the evidence we get from our sensory experience is often uncertain. Indeed, to some it seems as though the experiences that underwrite our belief revisions almost never leave us certain of anything at all.

To accommodate this intuitive idea, Jeffrey (1965) proposed an alternative to Conditionalization that generalizes this model. Jeffrey suggests that our evidence takes the form of a weighted partition of propositions: an n -tuple of propositions $\langle B_1, \dots, B_n \rangle$ that partitions the agent’s credal state, and that are assigned an n -tuple of credences:

Jeffrey Conditionalization: When the total information you have acquired can be represented as a change over the partition $\{B_i\}$ from $p(B_i)$ to $p'(B_i)$, your new degree of belief in A, for any A, should be $p'(A) = \sum_i p(A|B_i) p'(B_i)$.

It’s easy to see that, in many cases, the order in which two Jeffrey updates happen will determine the credence distribution one ends up with. Consider the case where I see a raven and this leads to me directly changing my credence in the proposition that *the raven is black* to .9, i.e., $p(RB) = .9$. And then after a second glance, I come to directly change my credence in this proposition to .7. After I

update on these two pieces of evidence, I am left with a credence of .7 in RB. But had I made these two updates in reverse order, I would have been left with a credence of .9 in RB. Therefore, the order in which these two pieces of evidence are received makes a difference to my post-observational credence distribution. Jeffrey Conditionalization is non-commutative over weighted evidence partitions.

In some cases, this feature of the framework seems exactly right: the probative value of my most recent evidence seems to swamp the value of the evidence I've gotten before. If I glance an object from afar, and then glance it again at a much closer distance, the information I get from the second glance should make my first glance irrelevant. But in just as many cases, this seems like the wrong result. Two consecutive glances taken from the very same vantage point shouldn't necessarily lead us to discount the deliverances of the first.

An assumption made by those who criticize Jeffrey Conditionalization for being non-commutative is that the elements for which the commutative property ought to hold are weighted evidence partitions.¹ In his paper, Lange calls into question these criticisms by pointing out that updates that don't commute in this way will be underwritten by different *experiences*, in the original sequence and its permutation. If experiences are different in two sequences of updates where the order of the weighted evidence partitions have been reversed, then the conditions required to generate commutativity failure at the level of experience will fail to hold in cases where we get commutativity failure at the level of weighted evidence partitions. Therefore, there's a more fundamental type of information that isn't non-commutative under the framework: experiential information.

Why think that experiences will be different in cases where weighted evidence partitions don't commute? Lange claims that the answer to this question lies with the agent's background beliefs. It lies especially with one kind of background belief: the agent's *prior opinion* about the proposition she gets as evidence. It's easy to see why such beliefs will be important. If my prior credence in the proposition that the raven I will see a moment from now will be black is pretty low, an experience of a raven that's not so clear or intense will result in my still having a low credence in the proposition that the raven is black. By contrast, if my prior opinion in the proposition that the raven will be black had been higher, this same mildly informative experience would have resulted in my having a relatively higher credence in this same proposition. *The same experience in the presence of different prior opinions will yield different posterior credence distributions.*

This same reasoning seems to suggest something else. It suggests that *the same posterior credence distribution arrived at from different prior opinions implies that the agent has had different experiences.* As Lange (2000, 398) puts it:

For an experience at twilight to have lowered our confidence in e from 0.99 to 0.8, the bird must have not looked much the way a black bird would be expected to look at twilight, whereas for an experience at twilight to have raised our confidence in e from 0.75 to 0.8, the bird must have looked about the way that any dusky colored object would be expected to look under those conditions. Plainly, these are different experiences.

We can walk through the reasoning in this passage to see how it is supposed to lead to the conclusion that the Jeffrey framework isn't defective in virtue of being non-commutative. First, consider that since any update involves the agent changing her credence in the evidence proposition, a non-commutative update will entail that the agent's priors in the evidence proposition are different in the original sequence and its permutation, provided we begin from the same credence function ($p(e) \neq q'(e)$ and $p(e) \neq q(e)$). If an agent has lowered her credence in e from 0.99 to 0.8, then the further update on the evidence that leaves her with a credence of 0.75 in the original

¹See, for example, Domotor (1980) and Doring (1999). For ease of exposition, in what follows, I will simply talk about weighted evidence partitions, or other elements, "commuting" rather than "commuting under the Jeffrey framework."




Figure 1. Two Glances at a Raven.

sequence will begin from 0.8, whereas the update on the same evidence in its permutation will begin from 0.99, and vice versa (Figure 1).

Second, notice that in order for an update to be non-commutative, we must have had the same evidence in the original sequence and its permutation. Therefore, it must be the case that the agent’s posterior credence in the evidence proposition after the first update in the first sequence is identical to her posterior credence in the evidence proposition after the second update in the second sequence, and vice versa ($q(e) = r'(e)$ and $q'(e) = r(e)$).

By Lange’s reasoning in the previous passage, the experiences that trigger the two updates in question must have been different in the original sequence and its permutation ($\xi_1 \neq \xi_4$ and $\xi_2 \neq \xi_3$) in order to offset the difference in prior opinions to yield the same posterior credence: ξ_1 must have “not looked much the way a black bird would be expected to look at twilight,” whereas ξ_4 must have “looked about the way that any dusky colored object would be expected to look under those conditions.” A similar story could be told about ξ_2 and ξ_3 .

And this gets us Lange’s conclusion. Since the conditions required to get a violation of experience commutativity (two experiences whose orders have been reversed) fail to hold in cases where we have reversed the order of the weighted evidence partitions, the non-commutativity of weighted evidence partitions does not entail the non-commutativity of experiences. Since the framework isn’t non-commutative over a more fundamental set of elements, we ought to conclude that it is not defective.²

Let’s consider a final example. Imagine that a whiff of pie leaves you with a credence of .3 in the proposition that *the pie in the oven is rhubarb* (R) at t_1 and then, a moment later, another whiff leaves you with a credence of .7 in the same proposition at t_2 . Now imagine that you’d gotten the previous evidence in reverse order (Figure 2).

Here, again, we have a case where the fact that weighted evidence partitions don’t commute under the framework doesn’t seem like a defect of the framework, since the experiences underlying these partitions in the original case and its permutation aren’t likely to be identical. The fact that we have raised our credence slightly at t_1 from $p(R) = .1$ to $q(R) = .3$ in the first sequence, in response to the evidence gotten in between t_0 and t_1 , suggests that the experience we’ve had is something like a decisive whiff of rhubarb pie. By contrast, where we decrease our credence from $q'(R) = .7$ to $r'(R) = .3$ in the second sequence, in response to having gotten the same weighted evidence partition in between t_1 and t_2 , that must have been because we had a “disconfirming” experience—maybe a whiff of lemon (which

²One might ask: what about those non-commutative updates that don’t involve a change in the agent’s prior opinions about her evidence? What about cases where our updates involve two different evidence partitions, and still the updates do not commute? Similar reasoning applies to these sorts of cases. The fact that our priors and posteriors in the propositions in question will be different in the original sequence and its permutation are all that are needed for Lange’s justification to go through. And this will be the case even where we have a sequence of updates on two different evidence partitions. (For a more detailed discussion of the conditions under which an update will fail to commute, see Diaconis and Zabell (1982, section 3.3), as well as Wagner’s (2003) discussion and defense of what he calls the Uniformity Principle.)




Figure 2. The Smell of Pie.

is not an ingredient in rhubarb pie). A similar story could be told about the second update in the first sequence, and the first update in the second sequence.

Though weighted evidence partitions don't commute under Jeffrey Conditionalization, then, that's altogether appropriate, since they don't supervene on the same experiences in the original case and its permutation. If we begin both sequences from the same initial, p , and if $p(R) \neq q'(R)$ and $p(R) \neq q(R)$, but $q(R) = r'(R)$ and $q'(R) = r(R)$, then the experience that triggered the first update in the first sequence must have been different than the experience that triggered the second update in the second sequence, and vice versa. Therefore the experiences underlying these sequences aren't the same. Therefore there's no reason we would want it to be the case that $r = r'$. Although Jeffrey Conditionalization isn't commutative over weighted evidence partitions, then, examples like this one suggest that it's not defective in virtue of this feature of it.

3. A partial account of experience

The crucial move in Lange's argument is his claim that Jeffrey Conditionalization is not defective because it doesn't fail to commute experiences. A natural question to ask is whether this defensive move commits Lange to some positive thesis about commutativity. It seems as though it must. In particular, it seems as though it must commit Lange to the claim that experiences ought to commute, and that they do commute. Without this assumption, his reasoning that Jeffrey Conditionalization is not defective because it doesn't fail to commute experiences seems to make no sense. Lange's argument seems to be committed to the following principle:

Lange's Assumption: The elements that ought to commute under Jeffrey Conditionalization are experiences and Jeffrey Conditionalization is commutative over experiences.

But maybe this is too quick. Maybe not everyone committed to Lange's argument needs to be committed to Lange's Assumption. Maybe the reason that Jeffrey Conditionalization does not fail to commute experiences is that experiences aren't the sorts of things that *could* fail to commute. One might maintain a view about the identity conditions for experience that make it impossible to get the same two experiences in reverse order. Consider a view according to which each experience's identity is determined by its qualitative character along with the totality of the agent's background beliefs, which are, in turn, determined by all of that agent's prior history. On this sort of story, it will be impossible to get the same two experiences in reverse order. If my experiences are individuated by the totality of what I have learnt up until the time in question, then experiences could not be had in reverse order: no two experiences could have been informed by the same history if they occur at different times in this history. On this sort of story, the conditions required to generate a failure of commutativity—two experiences taken in reverse order—could not possibly be met. Therefore, on this sort of story, the nature of experience itself

ensures that experiences could not fail to commute under the Jeffrey framework. No additional norm is required.³

I think there are a couple of reasons for resisting this sort of story. First, the account just described advocates for a radically holistic view, one whereupon *every* background belief matters for individuating experiences or, at least, determining their probative value. But this seems implausible. My belief that the lighting is tricky might matter for determining the probative value of my experience of seeing a red shoe on the floor. But my belief that Barack Obama was the forty-fourth president of the United States is surely irrelevant here. One might instead opt for a more modest form of holism that says not all background beliefs are relevant for individuating experiences or for determining their probative value. But this more plausible form of holism isn't going to get us the result we need. If I have two experiences, and no background beliefs that are relevant to the identity of these experiences, then I may very well have these experiences again, in reverse order. Adopting a more plausible, moderate form of holism that says that only *relevant* background beliefs individuate experiences will not make it impossible for experiences to fail to commute by making it impossible to get the same two experiences in reverse order. Therefore, adopting a more modest form of holism does not seem to preclude the need for Lange's Assumption—though, as we will see in section 5, it may inform the norm that this assumption justifies.

One might still object that a more moderate form of holism will make it impossible to get, in reverse order, the particular pairs of experiences featured in the last section: pairs of experiences that both bear on at least some of the same evidence propositions. Therefore, it may be that a more moderate form of holism will render a commutative norm unnecessary for precisely those experiences that are in danger of not commuting. However, I think there's a simpler reason why anyone, including Lange, should still endorse what I've called "Lange's Assumption." It seems close to a truism that we would expect identical information to be treated the same by our updating framework. The idea that information, understood in the most fundamental sense, should commute is surely more intuitive than any particular account of experience that would preclude the need for such a norm. I want to suggest, then, that Lange's Assumption is the weakest assumption that we need in order to be committed to Lange's argument.

It looks like we have reason to accept Lange's Assumption. Does the Jeffrey framework vindicate Lange's Assumption? Let's focus on the second conjunct: that Jeffrey Conditionalization is commutative over experiences. Some remarks Lange (2000, 400–402) makes near the end of the paper about the account of experience his argument assumes suggest how he hopes to secure this part of the assumption. Recall that Lange's conclusion is that experiences must have been different in two sequences of updates that feature different priors, but identical posterior credence distributions. This implies that had both prior and posterior credence distributions been the same, the experiences involved could have also been inferred to have been the same. We get something like this idea near the end of Lange's discussion:

Consider two agents who undergo sensory experiences, where neither agent is left with a full belief in some proposition that captures all that the agent learned from her sensory experience. I am inclined to suggest that the two agents are undergoing the same sensory experience exactly when it is the case that *had* the two agents begun with the *same* prior probability distribution, then they *would* as a result of their actual sensory experiences have imposed exactly the *same* constraints on that distribution, and this agreement would have resulted no matter what the two agents' common prior probability distribution had been. (2000, 401; emphasis in original)

This passage makes it sound as though Lange is putting forth necessary and sufficient conditions for experience identity. In the next section, we will consider whether or not this passage should

³Thanks to an anonymous referee for raising this objection.

indeed be read in this way. For now, we can say that, at the very least, this passage says that if two experiences are identical, then they will have the same *impact* on an agent's credence distribution.⁴ While there are a number of ways this claim might be read, one interpretation of it seems natural. The impact, or degree of change, that some experience induces in an agent's credence distribution is often identified with a *Bayes factor*. Formally, two updates have the same Bayes factor, \mathcal{BF} , just in case the ratios of the new-to-old odds of the elements of each of the evidence partitions, taken pairwise, are identical. More formally still, for experiences, ξ and ξ^* , and credence functions, $p(\cdot)$ and $q(\cdot)$ defined over every proposition, X , in some σ -algebra, \mathcal{A} (and where $p_{\xi}(\cdot)$ comes from $p(\cdot)$ by updating on the weighted evidence partition induced by ξ , and $q_{\xi^*}(\cdot)$ comes from $q(\cdot)$ by updating on the weighted evidence partition induced by ξ^*), the claim that two updates have the same Bayes factors amounts to the following:

$$\mathcal{BF}_{p,p_{\xi}} = \mathcal{BF}_{q,q_{\xi^*}}$$

$$\frac{p_{\xi}(E_{i_1})/p_{\xi}(E_{i_2})}{p(E_{i_1})/p(E_{i_2})} = \frac{q_{\xi^*}(E_{i_1})/q_{\xi^*}(E_{i_2})}{q(E_{i_1})/q(E_{i_2})},$$

for all elements, i_1, i_2 , of the evidence partition, $\{E_i\}$.

With this in mind, an account of the impact of an experience can be formulated in the following way (for ease of exposition, let us simply formulate experience identity as $\xi = \xi^*$):

An Account of the Impact of an Experience:

$$\forall p \forall q \left[\forall \xi \forall \xi^* (\xi = \xi^* \rightarrow (\mathcal{BF}_{p,p_{\xi}} = \mathcal{BF}_{q,q_{\xi^*}})) \right]$$

This way of understanding the passage from above gets us the connection to experience commutativity we are after. This is because a series of results confirm that two updates will commute just in case they yield the same Bayes factors in the original case and its permutation.⁵ Since *An Account of the Impact of an Experience* says that a necessary condition for experience identity is a property that is necessary and sufficient for experiences to commute, it vindicates the second conjunct of Lange's Assumption. Therefore, *An Account of the Impact of an Experience* is a necessary condition for the account of experience Lange's argument relies upon.

4. A complete account of experience

Is *An Account of the Impact of an Experience* also a *sufficient* condition for the account of experience Lange's argument relies upon? Although Lange certainly makes it sound as though he is putting forth necessary and sufficient conditions in the passage that points us toward this account, at the

⁴I assume here that Lange does not mean to say that two experiences must be had against the same priors in order to be identical. Especially in light of his emphasis in the passage above, it seems clear that Lange's view is that two experiences are identical, not in virtue of *actually* having been had by two agents with the same priors who update to the same posteriors, but in virtue of the fact that *had* they been had by two agents with the same priors, these agents *would* have reached the same posterior distribution. Given this, it seems to me that the most plausible and charitable reading of this passage is one that individuates experiences based, at least in part, upon the impact that these experiences have on an agent's credence distribution. As we will see in section 6, this reading is well-supported by the literature. (Thanks to an anonymous referee for pressing me to be clearer about this.)

⁵See Field (1978) for the left-to-right half of the biconditional. See Wagner (2002) for a theorem that establishes both directions of the biconditional for commutativity on updates. Wagner's theorem also "generalizes out" Field's nominalism in establishing the result that Bayes factor identities individuate commutative updates involving countable evidence partitions as well. We will consider these two important discussions more carefully in sections 5 and 6.

very end of the paper, he says the following: “Of course, a fuller account of the sameness of sensory experiences would be very welcome. But I cannot offer one at present” (2000, 402).

How do we reconcile these two seemingly inconsistent passages? Plausibly, Lange makes the latter statement because he believed that a necessary condition for experience identity is some sort of *qualitative* constraint—one that we might think of as distinguishing experiences with different phenomenal characters. In this section, I argue that not only is such a constraint plausible, but that some such constraint on experience identity looks like it’s required to secure the conclusion that Jeffrey Conditionalization is not defective. This means that the norm that Lange’s argument defends as not defective is actually a stronger norm than Jeffrey Conditionalization, for it must include a constraint that Jeffrey Conditionalization does not entail.

To begin to see this, recall again the following assumption that Lange’s argument relies upon:

Lange’s Assumption: The elements that ought to commute under Jeffrey Conditionalization are experiences and Jeffrey Conditionalization is commutative over experiences.

Let’s now focus on the first conjunct: that the elements that ought to commute under Jeffrey Conditionalization are experiences. Lange is careful in his discussion to distinguish the claim that: (1) Jeffrey Conditionalization is *formally* non-commutative over weighted evidence partitions, from the claim that (2) Jeffrey Conditionalization is *defective* in virtue of being non-commutative over weighted evidence partitions (2000, 393, 397). The first is a mathematical claim; the second is a normative claim. And, of course, it is the second claim that Lange’s argument targets.

In order for experience commutativity to save Jeffrey Conditionalization from being defective, then, we must assume that, unlike arithmetic operands, the identity conditions for experience make experience commutativity a good thing. The most obvious thing to say here is that the reason why experiences ought to commute is that they are *phenomenally indistinguishable*. In other words, what saves Jeffrey Conditionalization from being defective is that it treats phenomenally indistinguishable experiences the same, as we would want. But maybe this is too quick. Maybe we don’t actually need to identify experiences with anything phenomenal or qualitative to do justice to the idea that treating experiences consistently makes the Jeffrey framework nondefective. Maybe there’s something about the degree of an impact on a credence distribution that makes *this* the thing that ought to be treated consistently, and so that ought to commute.

There are several reasons to reject this suggestion. First, it does not seem to be what Lange had in mind; never once in the examples he provides throughout the paper does he talk about the agent’s doxastic behavior. Instead, he talks about “sensory experiences”—or sometimes just “experiences”—being different in two sequences of updates where the weighted evidence partitions have been reversed (2000, 394–95, 397). Sometimes he talks about particular sensory experiences, like “the way that any dusky colored object would be expected to look under those conditions” (2000, 398). And, of course, there is the passage we were left with in the last section, wherein Lange acknowledges he has failed to offer a complete account of experience in providing an account of how identical experiences impact an agent’s credence distribution.

Second, regardless of Lange’s intentions again, this is clearly not what he *should* have had in mind. To whatever extent we think that there is something intuitively good about belief revisions commuting, it’s difficult to imagine that this is not because we are assuming that the experiences that ground these revisions are phenomenally identical, i.e., that the feeling that induces these changes is the same. As Bradley (2005, 6) puts it, “it does not follow from the fact that your Bayes factors represent what you have gleaned from observation, that they have the kind of objectivity which obligates others to modify their beliefs using them as constraints on their posteriors.”

To see Bradley’s point more clearly, consider the lack of rhetorical force Lange’s argument would have if it did not include a description of the experiences that prompted the belief revisions he describes. To claim that Jeffrey Conditionalization is not defective because reversing the order of the weighted evidence partitions involved in two updates does not entail that the magnitudes of these revisions are

identical comes pretty close to claiming that Jeffrey Conditionalization is not defective in virtue of being non-commutative because its updates don't commute. It comes pretty close to being just a description of the fact that Jeffrey Conditionalization is non-commutative, rather than a defense of it.

These considerations suggest that Lange's Assumption—that Jeffrey Conditionalization is not defective in virtue of commuting experiences—requires that we think of experiences as individuated by something more than merely Jeffrey Conditionalization's commutative property. Plausibly, this "something more" is an experience's phenomenal character.⁶ But then this means that experience commutativity won't be trivially preserved by the Jeffrey framework. If we want to ensure that experiences, individuated in a way that would make Jeffrey Conditionalization not defective in virtue of commuting such experiences (as required by the first conjunct of Lange's Assumption) actually do commute (as required by the second conjunct of Lange's Assumption), we need an additional norm to secure that experiences with the same phenomenal characters actually commute under the framework. The norm that isn't defective in virtue of being non-commutative isn't Jeffrey Conditionalization, then, but an updating rule that comprises two constraints. The first constraint is Jeffrey Conditionalization. The second constraint secures Lange's Assumption by ensuring that phenomenally indistinguishable experiences yield updates with the same Bayes factors:

Experience-Commuting Jeffrey Conditionalization (ECJC):

1. When an experience, ξ , directly changes your credences over a partition $\{E_i\}$ from $p(E_i)$ to $p'(E_i)$, your new degree of belief in A, for any A, should be $p'(A) = \sum_i p(A|E_i) p'(E_i)$.
2. For any two experiences, ξ and ξ^* , that are identical with respect to phenomenal character, and for any credence distributions, $p(\cdot)$, $q(\cdot)$, the following should hold: $\mathcal{BF}(p, p_\xi) = \mathcal{BF}(q, q_{\xi^*})$.⁷

Lange's argument entails that Jeffrey Conditionalization alone does not commute the elements that ought to commute. We need, in addition, ECJC's second constraint.

5. Commutativity, normativity, and holism

5.a The normativity problem

Lange's argument entails that the norm that isn't defective in virtue of being non-commutative isn't actually Jeffrey Conditionalization, but the stronger norm that I've called "ECJC." A natural thought is that by trading in Jeffrey Conditionalization for ECJC, we can avoid the conclusion that the Jeffrey framework—the framework that permits us to update on uncertain evidence—is defective. However, this natural thought is mistaken. Lange's argument assumes that the commutative property is a constraint on the adequacy of the inputs to the updating process, and that this constraint picks out experiences, individuated qualitatively. In this section, I argue that this assumption entails that the Jeffrey framework is still defective, in virtue of the desideratum that its inputs commute, in a way that the classical Bayesian framework is not.

To begin to see the problem, consider Joan who gets a whiff of rhubarb pie. Suppose that, in response to this experience, Joan updates in a way that leaves her credence in the proposition that

⁶Strictly speaking, the structure of my objection does not require that experience identity be individuated by phenomenal character. To get the conclusion that the norm that isn't defective in virtue of being non-commutative is a stronger norm than Jeffrey Conditionalization it suffices that (1) experiences be partly individuated by some feature that makes commutativity on experiences plausibly a norm, and that (2) this feature is not entailed by Jeffrey Conditionalization. While this could be any number of things, an experience's phenomenal character seems like the most natural candidate.

⁷Where, again, $p_\xi(\cdot)$ is the function that comes from $p(\cdot)$ by updating when ξ directly changes the agent's credences along the partition $\{E_i\}$, and $q_{\xi^*}(\cdot)$ is the function that comes from $q(\cdot)$ by updating when ξ^* directly changes the agent's credences along the same partition, $\{E_i\}$.

the pie in the oven is cherry at $p(C) = .7$. For suppose that Joan is a small child who has not learned to properly distinguish the smell of rhubarb from the smell of cherry. Now suppose that later on in life when Joan has become capable of making such distinctions, she has the same phenomenal experience and, in response, changes her credence in the proposition that *the pie in the oven is rhubarb* to $p(R) = .7$. This is a violation of ECJC. But, intuitively, it seems rational. Intuitively, we wouldn't want an agent to be beholden to how she has updated in the past, if her past update was on unreliable evidence.

Can we supplement ECJC in a way that would allow it to yield the right result in these kinds of cases? I think that Lange's argument entails a problem for ECJC on this front that Bayesian Conditionalization manages to avoid. The case I've just described dramatizes a feature of all Bayesian updating rules, namely, that they lack an account of evidence. While the Bayesian framework tells us what to do once we have evidence—once we have the inputs to the framework—it does not tell us what constraints some information must meet in order to count as evidence in the first place. It does not tell us which inputs it is rational to have. Granted that none of Bayesianism's updating rules entail an account of evidence—granted that we don't think these rules fall short whenever the inputs to their updates lead the agent astray—I think we would surely want these updating rules to be *compatible* with an account of evidence. There should be *some* account of evidence to which Bayesian updating rules could, in principle, be paired. It would be a bad result—a defect—if a framework that tells us how to revise our beliefs when we get evidence were incompatible, in principle, with any account of evidence more substantive than one that merely tells us that we should revise our beliefs when we change our credences in some propositions noninferentially. This would imply, contra the example above, that Bayesian updating is not only necessary, but sufficient, for epistemic rationality.

One might argue that it is impossible to raise a problem for the inputs to the Jeffrey framework that the inputs to the classical Bayesian framework will not also face. Since an update on the classical Bayesian framework is generally taken to be a special case of an update by the Jeffrey framework, any problem for the latter will be a problem for the former. Interestingly, however, taking the commutative property to be a constraint on our inputs entails that updates on certain evidence are *not* a special case of Jeffrey updates. This is because Bayes factors are undefined whenever one of the evidence propositions receives a value of one. If we assume with Lange, then, that the commutative property picks out the inputs to our updating framework, we are left with the result that the classical Bayesian updating framework and the Jeffrey updating framework have different inputs.⁸ The inputs to the Jeffrey framework are experiences, while those to the classical Bayesian updating framework must be propositions.

It's clear that propositional inputs are amenable to the sorts of constraints on evidence that traditional epistemology has to offer. They are amenable to constraints that are grounded in internalist and externalist theories of justification. We might, for instance, take some proposition to be evidence iff it supervenes upon our current mental states. Or we might take some proposition to be evidence iff it was formed by a reliable process.⁹ In each of these cases, Bayesian Conditionalization will say that, when we get some information that meets the constraint on evidence in question, we should update our beliefs on this information.

Can we impose these sorts of internalist and externalist constraints on the way that experience gives rise to a Bayes factor? Can we get a modified version of ECJC that commutes norm-governed experiences instead? To see why this poses a problem, consider what might seem like a good candidate for the sort of constraint that we are after. Suppose that I have some experience that directly changes my credences along a partition. And suppose that I know the objective probability of

⁸See Field (1978, 365) for this point. See also Lange (2000, 397) for something close to this concession.

⁹For an example of the first kind of account of evidence, in the context of Bayesian epistemology, see Hedden (2015). For an example of the second kind of account, see Dunn (2010).

my having had this experience, conditional on each of the members of the evidence partition being true. In this case, we might think that the Bayes factor we ought to adopt in response to having had this experience should reflect these objective likelihoods. The inputs to the framework should reflect the objective chance of my having had the experience in question, conditional on it being veridical.¹⁰

To see this constraint in action, suppose that Joan has an experience that changes her credences over the partition $\{R, \neg R\}$, where R is again the proposition that *the pie in the oven is rhubarb*. Suppose that the objective probability of Joan having had this experience, conditional on R being true is .8, and the objective probability of her having had this experience, conditional on $\neg R$ being true is .2. Then we might plausibly take the Bayes factor that Joan's experience ought to give rise to to be $.8/.2=4$. This constraint looks like it avoids the problem identified above by making the Bayes factor that experience ought to give rise to a function of this experience, as well as a normative constraint.

However, I think this conclusion is too quick. First, notice that the constraint we've described does not always seem to get us the right answer about how we should revise. The fact that I am not very likely to get a strong and decisive whiff of rhubarb pie, even when the pie in the oven *actually is* rhubarb, because rhubarb does not typically have a strong smell, does not mean that my credence in the proposition that the pie is rhubarb should not increase significantly if I do happen to get a strong and decisive whiff of it. Where the intensity of my experience does not line up with the objective chance of having had the experience, our constraint ignores the former in a way that seems objectionable.

This example points to a more general worry about our constraint, one that illustrates why no other constraint is likely to succeed. Like Christensen (1992, 554), we might think of experience as having both a magnitude and a direction. According to Christensen, the direction of an experience corresponds to the propositional content of the experience, while the magnitude of an experience corresponds to the strength and clarity of the experience. In light of this picture of things, it's easy to say where the previous account goes wrong. The magnitude of an experience—its clarity and strength—isn't the sort of thing to which an objective probability can be ascribed. To return to an earlier example, "the way a black bird would be expected to look at twilight" isn't the sort of thing that has an objective probability.

The point generalizes to constraints that don't involve objective probabilities. The clarity and strength of an experience isn't the sort of thing that we can have reasons for. It isn't the sort of thing that could have been formed by a reliable process. We can, perhaps, impose constraints on propositions *about* these things. For instance, we can perhaps ascribe an objective probability to the proposition that the black bird looks a certain way.¹¹ But we've already established that weighted propositions aren't the inputs to the Jeffrey framework: they aren't the things that ought to commute.

The previous constraints fail because they require we take into account the character of experience, in a way that seems impossible to do. But perhaps a constraint that determines Bayes factors in some different way will do better. To test this hypothesis, we might again appeal to some more familiar norms. There are well-known ways that reliabilism, for instance, is able to provide norms for degrees of belief. Some have taken degrees of belief to be justified iff they are well-calibrated: if the degree of belief assigned to some proposition corresponds to the frequency with which the process producing this belief produces true beliefs.¹² Since weighted evidence partitions are, roughly speaking, just degrees of belief for evidence, one might reason that an account of justification for degrees of belief can provide an account of justification for the Bayes factors that represent a change in these degrees of belief.

¹⁰Thanks to an anonymous referee for this proposal.

¹¹I say "perhaps" here because it isn't obvious to me that a proposition could express the clarity or strength of the experience that it is about. It's difficult to see how we could fill out the "certain way" that a black bird looks—that which corresponds to how strongly and clearly it looks like a black bird—in a way that would allow us to ascribe an objective probability to this proposition. I take this to be a different way of expressing the same concern I've just raised.

¹²See van Fraassen (1984) and, more recently, Dunn (2015).

But this proposal again misses its mark. Any constraint that justifies a Bayes factor, qua input, will need to do so directly. It will need to justify, not a weighted partition, but the *relation between* the values we assign this partition, before and after we have had some experience. It will need to justify the *magnitude of the change* in these values. But it's difficult to see how some process could be described as directly producing, not a belief, but a relation between beliefs. To see this most clearly, consider that the degree of reliability of this process could not possibly correspond to the value we assign this relation, for while the value of this relation must be a magnitude, the degree of reliability of this process must be a probability.

Again, the point generalizes. The magnitude of a change isn't the sort of thing that we can have reasons for; it isn't the sort of thing that can supervene on a mental state. This is, again, in contrast to the proposition *that* some change has a certain magnitude, which might be amenable to these sorts of constraints. Such a proposition could perhaps be assigned a value that corresponds to the reliability of the process that produced it. But this would, again, assume that the inputs to the Jeffrey framework are weighted propositions, an assumption we have already rejected. Unsurprisingly, then, governing the impact of an experience directly will be impossible for the very same reason that governing it by appealing to the character of experience looks to be impossible. The magnitude of a belief change, like the magnitude of the experience that produces it, isn't the sort of thing that can be targeted by our norms.

Summing up, then, where we take the commutative property to be a constraint on the adequacy of the inputs to our updating framework, we are left to conclude that the inputs to the regular Bayesian framework are propositions, whereas the inputs to the Jeffrey framework are experiences. This formal distinction in inputs results in a normative distinction in these frameworks. While the regular Bayesian framework is capable of having its inputs governed by the sorts of substantive constraints described above, the Jeffrey framework is not. There are, then, two important upshots of this section. The first is that the Bayesian framework and the Jeffrey framework turn out to be fundamentally different updating norms. The second is that the Jeffrey framework is still defective in a way that the Bayesian framework is not, even once we reinterpret it in the way that Lange's argument suggests that we should.

5.b The holism problem

The normative problem for ECJC is only the beginning of its troubles. Even if we could get a modified version of ECJC that commutes norm-governed experiences, such a norm would still yield weird results in cases where background beliefs other than the agent's prior opinions about her evidence seem relevant to how she ought to update. Earlier we saw that endorsing a moderate form of holism about the role that background beliefs play does not make a commutative norm redundant, since it does not make it impossible to get the same experience twice. However, endorsing a moderate form of holism does mean that our commutative norm may need to govern a more holistic sort of experience.

Garber (1980) famously argues that mapping phenomenal experiences to Bayes factors yields the wrong result in cases where we have the same phenomenal experience over and over again. To modify Garber's example, say that I keep having the same blurry visual experience of a black raven. If each blurry experience yields the same Bayes factor, then having it enough times will lead me to hold the proposition that the raven is black with something close to certainty, which seems absurd. Notice that this is a different problem than the one that Lange's argument resolves. Whereas the lesson of Lange's argument is that the same weighted evidence partitions don't always correspond to the same phenomenal experiences, the lesson of Garber's argument is that the same phenomenal experiences don't always correspond to the same Bayes factors.

Weisberg (2009) takes Garber's objection from confirmation holism in a different direction. Rather than showing that we get bad results when we assume that the impact of phenomenally identical experiences is additive, he notices that we get bad results when we assume that the impact

of phenomenally identical experiences is commutative. To illustrate with his own example, say that I update on the proposition that the lighting is red tinted, and then I update on the proposition that the sock is red. If the phenomenal experiences that underwrite these updates are the same, then they should commute. But this means that my background belief about the tricky lighting cannot influence the probative value of my experience of the red sock in the way that it should if some moderate version of confirmation holism is correct. For whether it should influence the probative value of my experience of a red sock depends upon whether it comes before or after the experience, contra the assumption that the order of these experiences does not make a difference to the agent's final credence distribution.

Wagner (2002) offers a suggestion that overcomes these worries: it is *considered* experiences—experiences in the light of all relevant background beliefs—that should be mapped to Bayes factors. His suggestion is that we replace the second condition of ECJC with something like the following:

- 2*. For any two considered experiences ξ and ξ^* , and for any credence distributions $p(\cdot)$, $q(\cdot)$, the following should hold: $\mathcal{BF}(p, p_\xi) = \mathcal{BF}(q, q_{\xi^*})$

On Wagner's interpretation of ECJC, Garber's worry disappears. Since considered experiences will be different in each case where we have the same mildly informative phenomenal experience over and over again—since the memory of our previous experiences will make each successive experience less probative—what this will yield are updates whose Bayes factors decrease over time. Continually updating on the same phenomenal experience won't, then, lead us to be almost certain of some proposition.

In a similar way, in the case that Weisberg describes, the fact that the agent's considered experiences will be different in the original case and its permutation—since the agent's belief about the tricky lighting will be a relevant background belief when it comes before the experience of the red sock, but not after—means that the Bayes factors in the original case and its permutation won't be identical, as desired in order to accommodate confirmation holism. Having the same phenomenal experiences in reverse order, then, needn't lead us to ignore our background beliefs.

But while Wagner's proposal helps with these problems, it comes at a cost. Mapping considered experiences to Bayes factors seems much more problematic than mapping simple phenomenal experiences to Bayes factors. Any rule that does the former would have to encode information about which background beliefs are relevant to the probative value of an experience—alone and in combination with other background beliefs—and of how relevant these background beliefs will be. In the last section, we saw the problems that arise when we try to impose substantive constraints on our experiential inputs. We now see that even if we could resolve the normativity problem, we would still be left with the holism problem: the problem of how to map considered experiences to Bayes factors prior to imposing any further substantive constraints.¹³

6. Carnap and Field

I've argued for two claims. First, I've argued that the norm that Lange's argument targets is not Jeffrey Conditionalization, but the norm that I've called ECJC. Second, I've argued that, like Jeffrey Conditionalization, ECJC is defective. ECJC is defective in virtue of the nature of its inputs and, in particular, in virtue of the problems that these inputs raise. Before closing, I want to briefly consider

¹³For a comprehensive discussion of the difficulties involved in providing this sort of account, see Christensen (1992). Whether or not this idea requires that we reject condition 2 of ECJC and adopt 2* instead depends upon whether one thinks that background beliefs actually affect the phenomenal character of some experience. If they do, then there's no need to appeal to a modified version of ECJC. In this case, what Wagner's argument entails is that we don't actually have the same phenomenal experiences in the cases where Garber and Weisberg claim that we do. But the main point in either case is the same: giving content to the second constraint of ECJC in a nonarbitrary way is an impossible task.

the history of previous attempts to formulate a norm like ECJC. This history lends some support to my results. My results, in turn, help to unify some of the ideas in these earlier discussions.

The first worries about a norm that tells us how to get from some particular phenomenal experience to a posterior credence distribution come from Carnap's 1957 correspondence with Jeffrey (reprinted in Jeffrey [1975]), where Carnap tells Jeffrey that he himself had attempted to formulate such a norm. His correspondence expresses two worries about its prospects. Carnap's first worry is one that our own discussion has circumvented. It arises as a result of assuming that experience ought to be represented, not as a Bayes factor, but as a *probability* attached to some experientially affected sentence which indicates "the subjective certainty of the sentence on the basis of the observational experience" (1975, 42). Carnap noticed that since probabilities, unlike Bayes factors, aren't defined in terms of the agent's prior and posterior probabilities, this made it impossible to represent experience in the right way: as the impact on an agent's credence distribution. Even before Lange, then, there were worries about how to represent the inputs, though these worries took for granted that these inputs were experiences. Call this *the representation problem*.

Carnap's second concern is closely related to the one that we were left with at the end of the last section. Even if we were to represent experience as a Bayes factor so that we could describe the Jeffrey framework as representing the agent's evidence as a function of her experiences and prior expectations, we still would not have a rule for how to map a *particular* experience to a particular Bayes factor in a way that takes into account how features like "the clarity of the observation (or the feeling of certainty connected with it or something similar)" (1975, 42) should determine the latter. Carnap complained to Jeffrey that his framework did not include the latter sort of rule:

You emphasize correctly that your [posterior probability distribution] is behavioristically determinable. But this concerns only the *factual* question of the actual belief of [the agent] in [her evidence]. But [the agent] desires to have a rule which tells him what is the *rational* degree of belief. (1975, 44; emphasis in original)

As we saw in the last section, the latter sort of rule is not even compatible with a certain understanding of the Jeffrey framework. This is *the normative problem*.

Carnap's frustration that he could not adequately address the representation problem and the normative problem led him to abandon the project of crafting a rule like ECJC.¹⁴ The first discussion that actually addresses both the representation problem and the normative problem is Field (1978). Field distinguishes his approach from Carnap's attempt at such an account based upon the responses he offers to the representation problem and the normative problem (1978, 363–64). Field's account resolves the representation problem by assuming, like Lange, that experience should be represented by a Bayes factor. Field's response to the normative problem is equally decisive. He claims that it is not a legitimate problem at all:

Carnap's criticism of Jeffrey is that there are cases where a person ought *if he is rational* to come to attach a high probability q to a directly affected sentence E , but that nothing in Jeffrey's constraints requires him to do so. I do not think that this way of putting the problem makes the problem very persuasive. In any case, it did not persuade Jeffrey. (1978, 364; emphasis in original)

Having identified experience with a Bayes factor, rather than a probability, the normative problem for Field is close to the problem we were left with at the end of the last section. It is the problem that there are cases where a person ought, if he is rational, to come to adopt a certain Bayes factor on the basis of an experience, but that nothing in Jeffrey's constraints requires him to do this.

¹⁴See Jeffrey (1975, 41). For other expressions of skepticism, especially concerning the normative problem, see Christensen (1992) and Schwarz (2018).

As the previous passage suggests, Field rejects the normative problem. He argues that the task of providing an account of how experience figures into the updating process should be conceived of as a purely descriptive one, as “the problem of giving a complete psychological theory for a Bayesian agent” (1978, 364). Unlike Carnap, Field believed that a norm for evidence does not belong as part of the framework.

Both Field and Carnap took the representation problem and the normative problem to be independent of each other. Our discussion establishes a relation between them by arguing that it is precisely how we are forced to represent the inputs to the Jeffrey framework—as experiences formalized as Bayes factors—that gives rise to the version of the normative problem we have discussed. Whether or not we side with Field or with Carnap on the question of whether a substantive norm for evidence belongs as part of the Jeffrey framework, I’ve argued that the Jeffrey framework, like the classical framework, should at least be compatible with such a norm. And the fact that we must represent our inputs as experiences, or Bayes factors, entails that it is not. No such norm is going to be able to directly govern these inputs since no such norm is going to be able to directly govern these inputs qua magnitudes. I think, then, we cannot escape the conclusion that the Jeffrey framework is defective. Either it is defective in virtue of not commuting its inputs, or else it is defective in virtue of commuting the wrong kinds of inputs.

Acknowledgments. Thanks to Chris Meacham, Alejandro Pérez Carballo, and several anonymous referees for very helpful comments and suggestions.

Lisa Cassell is an assistant professor at the University of Maryland, Baltimore County.

References

- Bradley, Richard. 2005. “Radical Probabilism and Bayesian Conditionalization.” *Philosophy of Science*, 72(2): 342–64.
- Christensen, David. 1992. “Confirmation Holism and Bayesian Epistemology.” *Philosophy of Science*, 59(4): 540–57.
- Diaconis, Persi, and Sandy Zabell. 1982. “Updating Subjective Probability.” *Journal of the American Statistical Association* 77 (380): 822–30.
- Domotor, Zoltan. 1980. “Probability Kinematics and Representation of Belief Change.” *Philosophy of Science*, 47: 384–403.
- Doring, Frank. 1999. “Why Bayesian Psychology is Incomplete.” *Philosophy of Science (Proceedings)*, 66: 379–89.
- Dunn, Jeffrey. “Bayesian Epistemology and Having Evidence” (PhD diss., University of Massachusetts, Amherst, 2010).
- Dunn, Jeffrey. 2015. “Reliability for Degrees of Belief.” *Philosophical Studies*, 172(7): 1929–52.
- Field, Hartry. 1978. “A Note on Jeffrey Conditionalization.” *Philosophy of Science*, 45(3): 361–67.
- Garber, Daniel. 1980. “Field and Jeffrey Conditionalization.” *Philosophy of Science*, 47(1): 142–45.
- Hedden, Brian. 2015. *Reasons without Persons: Rationality, Identity, and Time*. Oxford, UK: Oxford University Press.
- Jeffrey, Richard. 1965. *The Logic of Decision*. Chicago: Chicago University Press.
- Jeffrey, Richard. 1975. “Carnap’s Empiricism.” In *Studies in the Philosophy of Science. Induction, Probability, and Confirmation* Vol. 6, edited by G. Maxwell and R. M. Anderson, Jr., 37–49. Minneapolis: University of Minnesota Press.
- Lange, Marc. 2000. “Is Jeffrey Conditionalization Defective in Virtue of Being Non-Commutative.” *Synthese*, 123(3): 393–403.
- Schwartz, Wolfgang. 2018. “Imaginary Foundations.” *Ergo*, 29: 764–89.
- van Fraassen, Bas. 1984. “Belief and the Will.” *Journal of Philosophy*, 81: 235–56.
- Wagner, Carl. 2002. “Probability Kinematics and Commutativity.” *Philosophy of Science*, 69: 266–78.
- Wagner, Carl. 2003. “Commuting Probability Revisions: The Uniformity Rule.” *Erkenntnis* 59 (3): 349–64.
- Weisberg, Jonathan. 2009. “Commutativity or Holism? A Dilemma for Conditionalizers.” *British Journal for the Philosophy of Science*, 60(4): 793–812.