

Commutativity or Holism? A Dilemma for Conditionalizers*

Jonathan Weisberg
University of Toronto

Abstract Conditionalization and Jeffrey Conditionalization cannot simultaneously satisfy two widely held desiderata on rules for empirical learning. The first desideratum is confirmational holism, which says that the evidential import of an experience is always sensitive to our background assumptions. The second desideratum is commutativity, which says that the order in which one acquires evidence shouldn't affect what conclusions one draws, provided the same total evidence is gathered in the end. (Jeffrey) Conditionalization cannot satisfy either of these desiderata without violating the other. This is a surprising problem, and I offer a diagnosis of its source. I argue that (Jeffrey) Conditionalization is inherently anti-holistic in a way that is just exacerbated by the requirement of commutativity. The dilemma is thus a superficial manifestation of (Jeffrey) Conditionalization's fundamentally anti-holistic nature.

1 Introduction

If something looks red, whether or not you should believe that it is red depends on what you think about the quality of the lighting, the reliability of your vision, etc. In general, a belief's empirical justification is sensitive to background belief. Call this general view *confirmational holism*. Respect for confirmational holism is a widely held desideratum on rules for updating our beliefs in response to empirical input. Another commonly held desideratum is *commutativity*, the view that the order in which information is learned should not matter to the conclusions we ultimately draw, provided the same total information is collected. It shouldn't matter whether I find the murder weapon in the maid's room first and then hear testimony about her alibi, or the other way around. Either way my ultimate attitude about her guilt should be one of guarded suspicion.

Commutativity and holism both have strong intuitive pull, and it will be an unhappy state of affairs if our epistemology cannot satisfy them both. My

* I am extremely grateful to Frank Arntzenius, David Christensen, Phil Kremer, Chris Meacham, Jennifer Nagel, Jim Pryor, Adam Sennet, Carl Wagner, Ralph Wedgwood, and an anonymous referee for *B/PS* for their generous feedback and patient discussion. Earlier versions of this paper were given at the 2006 Formal Epistemology Workshop and the University of Southern California, and I owe much to those audiences for their helpful suggestions and criticism.

aim here is to consider whether Bayesian epistemology can satisfy them both. In particular, I want to present the worry that Jeffrey Conditionalization has built-in limitations in this regard. The worry that Jeffrey Conditionalization is not compatible with commutativity has a long history (Field 1978; Domotor 1980; Skyrms 1986; van Fraassen 1989; Doring 1999), though recent work seems to have resolved that problem (Lange 2000; Wagner 2002). Compatibility with holism is a concern that has received much less attention, and to my knowledge has only been seriously explored by Christensen (1992). Christensen concludes that existing attempts to incorporate holistic considerations into Bayesianism fail, though he leaves open the possibility that a more sophisticated approach might be made to work. The surprising result to be argued for here is that the recent work on Jeffrey Conditionalization and commutativity points to a necessary tension with holism. Thus we are threatened with an unappealing dilemma: if we endorse Jeffrey Conditionalization, we must choose between commutativity and holism. The purpose of this paper is to lay out the argument for this dilemma and propose a diagnosis. Ultimately, I will be arguing that Jeffrey Conditionalization is inherently anti-holistic, and that considerations of commutativity just serve to make the problem more apparent.

Actually, a form of the same dilemma arises for Jeffrey Conditionalization's predecessor, Strict Conditionalization. Since many already see Strict Conditionalization as having unrealistically strong foundationalist commitments, I expect that many readers will not be surprised to find that it runs afoul of holism. Thus the really interesting case is Jeffrey Conditionalization. Still, it is instructive to see how the problem arises in the simpler case. So, after taking a moment to clarify our two desiderata in §2, I present the Strict Conditionalization version of the dilemma in §3. I then lay out the dilemma as it arises for Jeffrey Conditionalization in §4, and offer my diagnosis of the problem in §5. I close in §6 with some accompanying conclusions.

2 Clarifying Commutativity and Holism

Clarifying the desideratum of commutativity, especially as it applies to cases of Jeffrey Conditionalization, turns out to be quite tricky. To avoid hanging our discussion on contentious assumptions about what commutativity comes to in general, we will end up making use of a fairly minimal assumption about how commutativity applies in a particular case. But, for the purposes of later discussion, it will be useful to have considered some of the candidates for general formulations of commutativity, before coming to the minimal assumption that will be our primary resource in constructing the dilemma.

Roughly, commutativity says that the order in which you acquire information

shouldn't affect what conclusions you draw, provided that you glean the same total information in the end. But how do we understand 'information' here, and what does it correspond to in a Bayesian model? The natural thing to say is that the information an agent acquires is her propositional evidence, so that our requirement becomes:

Commutativity on Propositions Any rule for updating degrees of belief on propositional evidence should be such that the result of updating credences on one proposition and then another is the same as the result of updating on the same two propositions in the opposite order.

This formulation may be adequate for discussions of the classical rule of Strict Conditionalization, but Jeffrey Conditionalization is expressly designed to free us from a propositional conception of evidence.¹ It is intended to handle cases where the agent has an informative experience, but does not acquire any propositional evidence. In such cases, this requirement will not apply.

How should we understand commutativity in the cases that motivate Jeffrey Conditionalization? In these cases, an experience directly informs the agent's credences over a partition, and the rule then uses those new credences, together with her prior credences, to settle what the rest of her new credences ought to be. Since it is the distribution over the partition that determines her new credences as a function of her old credences, it is natural to regard that distribution as the information the agent acquires. It will be useful to have some terminology in this connection. Let's call the partition over which the new credences determined directly by experience are distributed, the *input partition*, and the credences distributed over it, the *input values*. The input partition and the input values together will be called the *input distribution*. The proposed interpretation of the commutativity desideratum is then:

Commutativity on Input Distributions Any rule for updating degrees of belief on input distributions should be such that the result of updating credences on one input distribution and then another is the same as the result of updating on the same two distributions in reverse order.

As sensible as it appears though, there is a good reason to reject commutativity on input distributions.

As Lange (2000) points out, having the same two perceptual experiences in opposing orders won't typically yield the same input distributions in opposing orders. Suppose you are looking at a jellybean in dim light. A first glance moves

¹ See p. 7 for a statement of Strict Conditionalization, and p. 9 for a statement of Jeffrey Conditionalization.

your credence in its redness from .1 to .8, and a second glance moves you from .8 to point .9. Now imagine that things had happened differently, so that those input values are reversed: first you move from .1 to .9 and then from .9 to .8. In that case, your first glance must have been much more clearly red than your second glance was in the first scenario, and your second glance must have been much less clearly red than the first glance was in the first scenario. So reversing input values does not correspond to reversing the order of experiences. But the look of the jellybean — your perceptual experience of it — is the source of your information regarding the jellybean's color. So a natural alternative to commutativity on input distributions is:

Commutativity on Experiences Any rule for updating degrees of belief on experiences should be such that the result of updating credences on one experience and then another should be the same as the result of updating on the same two experiences in reverse order.

This is, to a first approximation, the kind of commutativity that we will be interested in. But this statement has some obvious problems that need to be addressed.

Some will question whether it makes sense to talk about reversing the order of two experiences. Maybe experiences have their ordering essentially, so that it is metaphysically impossible for two experiences to have happened in the opposite order. Can we avoid this concern by talking about qualitatively identical experiences happening in reverse order? Some will worry that it is also impossible to have two qualitatively indistinguishable experiences in opposing orders, since an experience's place in the flow of consciousness colors the quality of the experience. This might be the case if, for example, experiences are always "bundled" with a kind of meta-information about when, or in what phenomenological context, they occur. And if that is so, shouldn't order matter to what conclusions we draw? Come to think of it, it seems that the order of experience often *does* matter. If a traffic light looks red and then green, your conclusions about whether the intersection will be clear when you drive through will be very different than if you had had those experiences in reverse order.

Examples like the traffic-light example notwithstanding, there are still those cases where the order of experience should not affect your ultimate conclusions about certain questions. When sitting at your desk, it should not matter whether you see or hear the rain on your window first when judging the weather. The order in which your perceptions of the rain come — visual and then aural or aural and then visual — may affect the qualities of those experiences, so that the two visual experiences are qualitatively different, and likewise for the aural ones. But it seems

clear enough that whatever differences the order makes, they are not differences that should affect your ultimate confidence that it is raining. Fortunately for us, we can sidestep the complications gestured at in the previous paragraph by working with only these kinds of cases. For our purposes, all that we will need is the following very weak assumption:

Occasional, Partial Commutativity on Experiences In *some cases* where an agent has two subsequent experiences, she could have had two qualitatively very similar experiences, but in reverse order, such that the bearing of those experiences on her opinions about *some hypotheses* should be the same in either case.

As long as we can say that there are cases of this kind, then we will have the materials we need to generate our dilemma.

Officially then, occasional, partial commutativity on experiences is the claim that we will be concerned with. For brevity, I will talk unqualifiedly about “commutativity on experiences”. But it is important to bear in mind that this is both a partial and occasional requirement, and that exactly when and to what extent it applies has not been specified.² All we have said is that it applies sometimes and to some extent. The examples that I will discuss in generating our dilemma will, however, be ones where it does apply. Also for brevity, I will talk about reversing the order of two experiences, but this can be understood in terms of qualitatively very similar types occurring in opposite orders.

Our second desideratum, confirmational holism, derives from the Quinean dictum that a theory can never be tested against experience in isolation, but only given a body of background assumptions (Quine 1953). Thus the verdict of experience for a theory depends on what background assumptions we hold. The kind of holism that concerns us here claims that no empirical proposition³ is immune to this kind of background sensitivity:

Holism For any experience and any proposition, there is a “defeater” proposition, such that your degree of belief in the first proposition, upon having the experience, should depend on your degree of belief in the defeater proposition.

It is plausible that most of our beliefs are background-dependent in this way, since

² Indeed, we might suspect that the general class of cases where it applies cannot be specified in any non-circular way. But as long as we can point to particular cases where it does apply, we can proceed with our discussion.

³ A priori beliefs may be an exception and are not crucial to our discussion, so we will set them to the side.

many of our beliefs are based on observation, and will thus depend on our beliefs about the accuracy of our observational abilities. But the kind of holism we are after is more thorough than that. Not only does your belief that the jellybean is red depend on your background beliefs, your belief that there *appears* to be a red jellybean does too. If you believe that you are a poor judge of your own experiences, the holist maintains that it is unreasonable for you to form any strong judgments about them. The kind of holist we envision points out that there is always a possibility of error when moving from an experience to a belief, and concludes that you should be reluctant to draw conclusions from your experiences when you think that you are prone to such errors.

While the kind of commutativity we need to generate our dilemma is quite weak, the kind of holism we need is quite strong. Nevertheless, this kind of holism is widely held.⁴ Moreover, we only need this strong kind of holism to apply in some cases, namely those where commutativity on experiences applies. So for the purposes of most of the paper, we can distill our two requirements down to their application in a single case, the jellybean case:

The Jellybean Case. You have standard background beliefs when you see what appears to be a red jellybean on a table. You then notice that the light-fixtures in the room appear to be red-tinted.

Our commutativity assumption will be that the order in which you catch sight of the jellybean and of the light-fixtures should not affect your final credence in the jellybean's redness. Either way, your credence in its redness should, in the end, be the same as it was to start with. If someone prefers to talk in terms of appearance propositions, then we just alter the case to talk about appearances: the first experience is as of a red jellybean, and the second is as of a phone-call from the local neurology lab informing you that you are a poor judge of your own color experiences. We then assume that your credence that the jellybean looks red on a given glance should not depend on the order of the experiences. No matter the order, your final degree of belief that the jellybean looked red to you should be the same as it was to begin with.

Doubtless, some readers will be unwilling to grant these assumptions, because they take issue with commutativity or with holism,⁵ or with their joint application to a single case. Nevertheless, they must acknowledge that it is a problem for the

⁴ See (Christensen 1992) for some discussion of the prevalence of holism and its connections with Bayesian epistemology.

⁵ Those who hold strongly externalist positions about empirical justification, for example, may be inclined to say that your beliefs about the jellybean's redness should not depend on your assessment your vision's reliability in such cases — because it depends only on your vision's actual reliability, for example.

many philosophers who have endorsed our two desiderata in the literature. They might even take our dilemma as a new argument supporting of their rejection of holism or commutativity.

3 The Dilemma for Strict Conditionalization

Suppose your degrees of belief are represented by the probability function p . The classic Bayesian rule for updating your degrees of belief in response to new evidence is Strict Conditionalization:

Strict Conditionalization When you acquire new evidence E , your new degree of belief function should be q , where $q(\cdot) = p(\cdot|E)$.⁶

Strict Conditionalization is well known to be commutative on propositions: if p is updated by Strict Conditionalization on the propositions E and F , the order in which those updates are done will not affect the final results. In either case, the resulting probability function will be $p(\cdot|EF)$.

Is Strict Conditionalization commutative on experiences? That all depends on what propositions we conditionalize on in response to our experiences. We could, of course, insist that each experience be associated with a unique proposition — the proposition that you should conditionalize on when you have an experience of that type. Then commutativity on propositions would carry over to commutativity on experiences. But this is where we run into the conflict with holism. If an experience is always treated by conditionalizing on the same proposition, then that experience always yields absolute certainty in that proposition, *regardless of what the agent's background beliefs are*. If the experience \mathcal{E} is always appropriately treated by conditionalizing on the proposition E , then the agent will be obliged to be certain of E no matter whether she has a background belief that, intuitively, should undercut the support that \mathcal{E} lends to E .⁷ This is precisely the sort of immunity to background belief that holism rejects.

Let's apply this point to our jellybean case, to see the problem in action. Let \mathcal{E} be the red visual appearance of the jellybean, let E be the proposition that the jellybean is red, and let F be the proposition that the lighting is red-tinted. If \mathcal{E} is always treated by conditionalizing on E , then an agent who already believes F will have to become certain of E when she has experience \mathcal{E} , which is a mistake. It's immediately tempting to say that the appropriate response to \mathcal{E} is to conditionalize on some *other* proposition besides E , like $E' = \textit{The jellybean appears red}$. But then

⁶ $p(\cdot|E)$ is the probability function that assigns $p(H|E)$ to each proposition H .

⁷ Throughout I'll be using the script letters \mathcal{E} and \mathcal{F} for sensory experiences (types, not tokens), and plain capital letters like E and F for propositions.

we just use a different proposition to play the role of F , namely $F' = I \text{ appear to have just gotten a call from the neurology lab informing me that I am a poor judge of my color experiences}$. Recall, the crux of holism is that *any* belief's empirical support can be undercut by some background belief or other. So no matter what proposition we try to associate with \mathcal{E} , whether it is E , E' , or something else, there will be some background belief that would undermine \mathcal{E} 's support for that belief. Holism forbids us associating with each experience a unique proposition to conditionalize on.

If we can't say that an experience should always be accommodated by conditionalizing on the same proposition, can't we just associate the experience with different propositions, depending on your background beliefs? Maybe the appropriate response to experience \mathcal{E} depends on your background beliefs. If you're not confident in F then you should conditionalize on E , but if you are fairly confident that F then you should conditionalize on E' instead. Unless you also happen to be confident that F' , in which case you should conditionalize on some other proposition.

This is where we get snagged on the other horn of our dilemma, since the proposed escape route sacrifices commutativity. Generally speaking, if we do Strict Conditionalization updates on E and then F , we will get different results than if we had updated on F and then some other proposition $E' \neq E$. After all, in the first case the end result will be certainty in E and F , and in the second case it will be certainty in E' and F . It's logically possible, of course, for $p(\cdot|EF)$ to be the same as $p(\cdot|E'F)$, but this can only happen when $p(E|E'F) = 1$. Thus, in order to preserve commutativity, we would have to be certain of E in the end anyway, our belief in the defeater F notwithstanding.

There is one last hole to be plugged in the construction of our dilemma. As we saw earlier, experience commutativity has its limits. It is only true that the order of *some* experiences shouldn't matter to our beliefs about *some* hypotheses. But the argument I just gave assumed that the order of our two experiences shouldn't matter to our confidence in *any* hypothesis, since I assumed that we wanted $p(\cdot|EF) = p(\cdot|E'F)$. But a thorough identity between $p(\cdot|EF)$ and $p(\cdot|E'F)$ isn't really essential to the argument. All we really need to require is that $p(E|EF) = p(E|E'F)$, so that we can get $p(E|E'F) = 1$. Thus we're only assuming that the order of experiences shouldn't matter to our final state of belief in whatever proposition is *prima facie* supported by \mathcal{E} , before the defeater F is discovered. And that's clearly appropriate. Whether we notice the apparent color of the jellybean or the tint of the lighting first, our ultimate confidence in the jellybean's redness should be reserved.

So there's the dilemma: to preserve holism, Strict Conditionalization forces us to give up on commutativity. Since Strict Conditionalization has a notoriously

strong foundationalist character, you might not be too surprised to learn that it runs into troubles with holism. All belief-changes via Strict Conditionalization are based on taking a single proposition as absolutely certain, and this radical foundationalism has led many to prefer Jeffrey's liberalized conditionalization rule instead. The really interesting thing is that Jeffrey's more liberal rule runs into the same dilemma.

4 The Dilemma for Jeffrey Conditionalization

Because Strict Conditionalization has the unacceptable result that evidence always gets credence 1, Jeffrey (1965) offered an alternate rule for updating:

Jeffrey Conditionalization When experience directly⁸ changes your credences over a partition $\{E_i\}$ from $p(E_i)$ to $q(E_i)$, set your new credences to $q(\cdot) = \sum_i p(\cdot|E_i)q(E_i)$.

Jeffrey Conditionalization avoids Strict Conditionalization's strong foundationalist commitment by allowing for uncertain evidence. Experience may not furnish us with any certain information, but it can give us uncertain information in the form of a distribution over a partition, which can then be assimilated into our corpus of beliefs via Jeffrey Conditionalization.

4.1 Jeffrey Conditionalization and Commutativity

A classic complaint about Jeffrey Conditionalization is that it blatantly violates commutativity (see, for example, Domotor (1980)), since it is not commutative on input distributions. This is most easily seen when we consider two subsequent Jeffrey updates on a single partition, $\{E, \bar{E}\}$, first with the input values $x, 1-x$ and then the values $y, 1-y$ ($x \neq y$). The first update leaves E with its input value, x , and the second leaves it with y . Reversing the order of the updates, E gets y first and then x in the end.

Clearly Jeffrey Conditionalization is not commutative *on input distributions*. But earlier we saw that this is a desirable feature because of the considerations raised by Lange (2000): reversing the order of input values does not correspond to reversing the order of experiences. Lange's point saves Jeffrey Conditionalization from the objection that Jeffrey Conditionalization is incorrect because it is not

⁸ The use of 'directly' here is deliberately vague, since there is disagreement about what sorts of 'direct' changes in response to experience Jeffrey Conditionalization should be used to handle. We will address this point momentarily (p. 10), when we consider whether these direct changes are subject to epistemological normativity.

commutative on input distributions, but it leaves us wondering: is Jeffrey Conditionalization appropriately commutative on experiences? This, of course, depends on how partitions and input values are selected for updating on in response to experiences. Thus Lange's point draws our attention to what I'll call *the inputs problem*. This is the problem of specifying what partition, and what values on that partition, we should apply Jeffrey Conditionalization to when we have a given experience. We can think of this as the problem of specifying a supplementary rule for Jeffrey Conditionalization that maps sensory experiences to input distributions, which can then be plugged into Jeffrey Conditionalization.

We had a similar problem with Strict Conditionalization, deciding what propositions to conditionalize on in response to our experiences. But the problem is even more pressing here since, without a solution, Jeffrey Conditionalization is actually vacuous. Unless we say what partition is directly affected by an experience and what values it ought to get, we can do any update we want without violating Jeffrey Conditionalization. To get an arbitrary q from p via Jeffrey Conditionalization, we just do a Jeffrey Conditionalization update on the set of epistemic possibilities, $\{w_i\}$, with the values $q(w_i)$.⁹ In applications, of course, we generally have an intuitive understanding of Jeffrey Conditionalization that rules this sort of thing out. But in many respects intuition is inadequate, and we need a proper theory.

Some people will think that the inputs problem doesn't even ask a legitimate normative question. Some, including Jeffrey himself, think that the input values for Jeffrey Conditionalization are provided by unconscious or involuntary psychological processes that are beyond the purview of normative epistemology. The results of those processes are neither rational nor irrational reactions to experience; they are arational. On this view, it may be a legitimate psychological question how an experience gets its foot in the doxastic door, but there are no shoulds or oughts about it. I happen to think that's wrong — someone who responds to the smell of strawberries by increasing their confidence in Marxism is surely doing something wrong — but this isn't the place to argue the point. Instead I'll simply note that, if you take this view, you've already abandoned holism. According to holism, every degree of belief you adopt in response to an experience should depend on your background beliefs, and so is subject to rational evaluation. So long as we are taking holism seriously — even if only for the sake of argument — we must acknowledge the inputs problem as raising a legitimate question.

The first attempt I know of to face up to the inputs problem is Field's (1978). Field acknowledges the seriousness of the problem and proposes as a solution that experiences be represented by a numerical parameter α , which captures the extent

⁹ I'm simplifying by assuming a countable set of possibilities. If the set is larger, we can still obtain any distribution we like over a given countable partition.

to which an experience supports the input proposition. (For ease of exposition, I'm following Field in sticking to the simple case where the input partition is $\{E, \bar{E}\}$, so we can speak of *the* input proposition, E .) Field suggested that α be the Bayes factor:

$$\beta_{q,p}(E:\bar{E}) = \frac{q(E)/q(\bar{E})}{p(E)/p(\bar{E})},$$

where p describes your prior credences and q describes the posterior credences you ought to have. Once each experience is assigned an α -value, $q(E)$ can be obtained from an experience-prior pair by fixing $p(E)$ in the above equation and plugging the α -value in for $\beta_{q,p}$.¹⁰

One of the virtues Field advertised for his proposal is its respect for commutativity on experiences. Doing a Jeffrey Conditionalization update with $\alpha = x$ and then $\alpha = y$ yields the same result as doing an update with $\alpha = y$ first and then with $\alpha = x$. Since each experience gets a single α -value on Field's proposal, this guarantees us commutativity on experiences.¹¹

But, as Christensen (1992) points out, Field's proposal is strongly anti-holistic. Consider the way an experience fixes the values over its input partition: the experience gives us a number, α , and we solve for $q(E)$ in the equation,

$$\alpha = \frac{q(E)/q(\bar{E})}{p(E)/p(\bar{E})}.$$

So your prior degree of belief in E is the only background belief that ends up affecting what bearing the experience has on E . In fairness, that's more holism than we got from Strict Conditionalization, but it certainly doesn't capture the kind of thoroughgoing defeasibility that holists see for empirical learning. Holists think that the empirical support for your credence in E can be undercut by background beliefs in propositions besides E . If your belief that there is a red jellybean is supported by the appearance that it is so, then background beliefs about the tint of the lighting can undercut that support. If we go instead to the belief that the jellybean merely *looks* red, the support can still be undercut — by the belief that you generally make mistakes about color-appearances. Whatever conclusion

¹⁰ Interestingly, I have not seen it noted that Field's proposal is really only a bare sketch of a solution since it does not actually say anything about what experiences get which α -values, or even which E should go with a given experience. That may be because Field's proposal, bare as it is, ran into an immediate problem, posed in (Garber 1980).

¹¹ There is, of course, the worry that Field's proposal ends up yielding too much commutativity, since the commutativity we get is neither occasional nor partial.

you want to draw from a sensory experience, if you know that you generally make mistakes about this kind of thing then you shouldn't draw the conclusion. Since Field's proposal makes $q(E)$ depend only on $p(E)$ and the experience, it cannot respect these holistic considerations.

4.2 The Tension with Holism

Now, just because Field's proposal is commutative but not holistic doesn't mean the two desiderata are, in principle, incompatible under Jeffrey Conditionalization. Couldn't there be another answer to the inputs problem that satisfies both desiderata? Surprisingly, the answer is no. A recent result by Wagner (2002) shows that Field's proposal is the only one that satisfies commutativity. Wagner shows that, on certain minimal assumptions, if two updates commute then they must yield the same bayes factor in both cases:

Wagner's Theorem Let p be the initial probability function, and consider two possible sequences of Jeffrey Conditionalization updates it might undergo: first to q and then to r , or first to q' and then to r' . Suppose also that the updates happen on the partitions $\mathbf{E} = \{E_i\}$ and $\mathbf{F} = \{F_j\}$ in opposing orders so that, schematically, we have

$$\begin{array}{ccccc} p & \xrightarrow{\mathbf{E}} & q & \xrightarrow{\mathbf{F}} & r \\ p & \xrightarrow{\mathbf{F}} & q' & \xrightarrow{\mathbf{E}} & r' \end{array}$$

Then, if $r = r'$ and

$$(1) \quad \forall i_1 \forall i_2 \exists j : p(E_{i_1} F_j) p(E_{i_2} F_j) > 0$$

$$(2) \quad \forall j_1 \forall j_2 \exists i : p(E_i F_{j_1}) p(E_i F_{j_2}) > 0,$$

we are guaranteed the bayes factor identities

$$(3) \quad \forall i \forall j : \beta_{q,p}(E_i : E_j) = \beta_{r',q'}(E_i : E_j)$$

$$(4) \quad \forall i \forall j : \beta_{r,q}(F_i : F_j) = \beta_{q',p}(F_i : F_j).$$

That is, provided p regards any two elements of \mathbf{E} as probabilistically consistent with some element of \mathbf{F} and vice versa, commutativity implies that the q and r' updates yield the same bayes factors on \mathbf{E} , as must the q' and r updates on \mathbf{F} . Thus, if two experiences are to commute, the α -values they yield must be the same even if the background beliefs those experiences happen against are different.

Let's apply Wagner's theorem to our jellybean example to see the problem in action. As before, let E be the proposition that the jellybean is red and F the proposition that the lighting is red-tinted. If I have a visual experience as of a red jellybean before noticing the tinted lighting, E gets boosted initially; from 1/10 to 9/10, let's say. Now suppose I notice the lighting first, and then look at the jellybean. For these experiences to commute, Wagner's result tells us that the red-jellybean appearance must yield the same Bayes factor in the second scenario as it had in the first: $\beta_{q,p}\{E,\bar{E}\} = \beta_{r',q'}\{E,\bar{E}\} = 81$. But then the appearance as of a red jellybean will have to strongly support E , even if I am already aware that the lighting is tricky. In fact, since my confidence in the jellybean's redness will not be affected when I notice the tricky lighting first, $q'(E)$ will be 1/10 and $r'(E)$ will thus have to be 9/10. So even if I'm aware of the defeater F , the experience will still have the same degree of support for E .

It's important to keep in mind, as always, that the choice of E and F here is not essential. If we prefer to use an appearance proposition to play the role of E , holism says that there will be another proposition that can play the role of F . It is the general structure of the scenario that drives the problem: whatever sorts of propositions we choose for our input partitions, commutativity will require that their experiential support be immune to background beliefs, contrary to holism.

4.3 Loose Ends & Technical Worries

It is not absolutely impossible to escape this consequence of Wagner's theorem. There are gaps in the argument when we move from Wagner's formal result to the philosophical conclusion that Jeffrey Conditionalization makes holism and commutativity incompatible. I regard these as mere technical worries, but some might hope to exploit them as escape routes. So let me dispose of those hopes before moving on to my diagnosis of the problem.

As it is stated, Wagner's result depends on an unreasonable form of commutativity. The theorem assumes that $r = r'$, which amounts to assuming that the order of experience should make no difference at all to our final degrees of belief. As we noted earlier (p. 4), that isn't reasonable since the order of experience will often matter to some hypotheses. But Wagner's theorem doesn't crucially depend

on this strong form of commutativity. If we only require that

$$(5) \quad \forall i \forall j : r(E_i F_j) = r'(E_i F_j),$$

the same result goes through (see the Appendix). This much more minimal requirement is surely reasonable in examples like our jellybean case, and is likely to apply quite generally.

Still, we might criticize the setup assumed in Wagner's theorem for supposing that the two sequences of updates happen on the same partitions, just in opposing orders. Maybe when we reverse the order of the experiences, different partitions should be 'directly affected'. Then the right setup to examine would be:

$$\begin{array}{ccccc} p & \xrightarrow{\mathbf{E}} & q & \xrightarrow{\mathbf{F}} & r \\ p & \xrightarrow{\mathbf{F}'} & q' & \xrightarrow{\mathbf{E}'} & r' \end{array}$$

In fact, because Jeffrey Conditionalization is trivially satisfiable in the way described above (p. 10), it is even guaranteed that we can specify partitions \mathbf{F}' and \mathbf{E}' , and distributions over them, such that r' will match r in whatever ways we like. But we pay a price if we take this way out, since we would trivialize Jeffrey Conditionalization in the process. Of course we can always set up the input partition to get the results we want, but then our input rule is doing all the work, and the inputs problem becomes the only interesting problem on the table. Jeffrey Conditionalization only has substance if we are allowed to assume things such as that \mathbf{E} and \mathbf{E}' are the same, or at least that their contents are very similar. We might allow, for example, that E_i says that the jellybean looks red at t_1 where E'_i says that it looks red at t_2 , in which case a tweaked version of Wagner's result still goes through (again, see the Appendix). But if \mathbf{E} and \mathbf{E}' are not on the same topic in this kind of way, our input rule will be doing so much epistemological heavy lifting that Jeffrey Conditionalization will cease to make any interesting epistemological claim.

In any case, those who are willing to trade Jeffrey Conditionalization's substance for a way out of our dilemma can expect only short-lived relief. In the next section I'll argue that the dilemma is just a symptom of a more basic problem with Jeffrey Conditionalization, one that can be raised without getting into questions of commutativity. So even if we were willing to pay this price for commutativity, we would end up with nothing in the end.

5 Diagnosis

It is puzzling that something as seemingly-sensible as Jeffrey's rule should force a choice between such plausible considerations as holism and commutativity. Even

if we can't avoid the choice, we should try to understand why it is being forced on us. What tacit assumptions do we make in adopting the rule that force us to choose between holism and commutativity? I think there are a couple of enlightening things we can say here. First, we can identify the source of the dilemma as a feature common to both Strict Conditionalization and Jeffrey Conditionalization called *rigidity*. Second, we can identify the problematic implications of rigidity, and then explain how the commutativity/holism dilemma results as a sort of byproduct. Our final diagnosis will be that rigidity entails a kind of anti-holism all by itself, and that commutativity only serves to expose that pre-existing problem by making it more obvious and thoroughgoing.

Strict Conditionalization and Jeffrey Conditionalization are both rigid, meaning that they preserve the conditional probabilities on the evidence. When we apply Strict Conditionalization to evidence E , $q(H|E) = p(H|E)$. Similarly, if we apply Jeffrey Conditionalization to the partition $\{E_i\}$, then $q(H|E_i) = p(H|E_i)$ for each E_i . This feature of the rules, it turns out, enforces a kind of anti-holism, since it ensures that a belief's empirical support cannot be undercut by the later discovery of a fact that should, intuitively, defeat the initial empirical support. Suppose, for example, that the reddish appearance of the jellybean, \mathcal{E} , supports the proposition that there is a red jellybean, E . Intuitively speaking, when we later discover that the lighting is tricky, F , the empirical support for our belief that there is a red jellybean has been undercut, and our confidence in the jellybean's redness should drop. Surprisingly though, rigidity will prevent this from happening. Our later discovery about the tricky lighting cannot return our confidence in the jellybean's redness to its initial, lower value. In fact, it can't reduce it at all. The reason is not immediately obvious, and we will bring it out rigorously in a moment. But the basic problem is that a probabilistic correlation between F and E needs to be introduced when the experience \mathcal{E} is had. Initially, F has no probabilistic bearing on E , but it should have a negative bearing on E once E has been boosted on the basis of \mathcal{E} . Rigidity, however, prevents any such correlation from being introduced when \mathcal{E} has its effects. Because conditional probabilities *on* E are held fixed, conditional probabilities *of* E cannot be amended. In particular, $q(E|F)$ cannot be set to match $p(E)$, as we are about to show.

In the lingo of defeaters, the experience \mathcal{E} provides a defeasible reason for believing E , while F is an *undercutting defeater* for that reason. That is, F undermines the support for E without being evidence against E directly. What we are about to see is that rigidity makes such undercutting defeaters for non-doxastic reasons impossible. If \mathcal{E} is a defeasible non-doxastic reason for believing E , then F cannot later defeat that reason without being a reason against believing E to begin with. That is a consequence of the following fact:

Rigidity Preserves Independence Suppose that the transition from p to q is rigid with respect to the partition $\{E, \bar{E}\}$. Then the following two conditions are incompatible:

$$(6) \quad p(E|F) = p(E)$$

$$(7) \quad q(E|F) < q(E).$$

In particular, (6) entails $q(E|F) = q(E)$.¹²

Conditions (6) and (7) must be jointly satisfied for \mathcal{E} to be a defeasible reason for E with F as an undercutting defeater. Condition (7) follows from the fact that F defeats the support \mathcal{E} provides, and condition (6) follows from the fact that F does this by *undercutting* that support, as opposed to being a *rebutting* (or “opposing”) defeater.

Consider how these conditions apply in the jellybean example. Condition (7) ought to hold because of the way holism applies in the example: having increased your confidence that the jellybean is red based on its looking that way, you should be prepared to lower your credence in the jellybean’s redness if you learn that the lighting is deceptive. And condition (6) ought to hold because, before you enter the room, the color of the lighting has no bearing on the jellybean’s actual color. What the result tells us is that, as long as updates are rigid, these two conditions cannot be jointly satisfied. As long as updates are rigid, the empirical support lent to E cannot be undercut. If E ’s probability is going to drop as a result of discovering F , it must be because F was evidence against E in its own right. We can gloss this point as follows: *rigidity implies the absence of after-the-fact, undercutting defeaters for non-doxastic reasons*.

I propose that this fact is the ultimate source of our troubles, since it alone is an anti-holistic consequence of rigidity. Moreover, we can see from there how requiring commutativity leads to thoroughgoing anti-holism. We have just seen that rigidity prevents the introduction of a probabilistic correlation between F and E , and thus prevents E ’s support from being removed when F is discovered later. Now suppose we time-reverse the scenario: F is discovered first, and then the experience \mathcal{E} is had. To ensure commutativity, we must have it that \mathcal{E} ends up supporting E , since we were unable to undercut that support in the first scenario. Thus \mathcal{E} will end up supporting E even in the face of F , and we have

¹² *Proof.* Assume the transition is rigid as stated and that (6) holds. From (6) and the symmetry of independence, $p(F|E) = p(F)$. By rigidity then, $q(F|E) = q(F)$. So, by the symmetry of independence again, $q(E|F) = q(E)$.

unbridled holism. The source of the problem was that we couldn't, on account of rigidity, introduce an after-the-fact undercutting defeater, and commutativity just exacerbates the problem by making it time-symmetric; we can't have the undercutting defeater do its work before the fact either. The bottom line is that rigidity is inherently anti-holistic; commutativity just serves to expose the problem in an especially egregious way, as Wagner's result illustrates.

6 Morals & Connections

What I think all this teaches us is that conditionalization-based epistemologies are inherently anti-holistic in spirit. Cartesian foundationalism sought certainties to serve as epistemic foundations, and Jeffrey did free Bayesianism from that radical foundationalist commitment. But replacing Strict Conditionalization with Jeffrey Conditionalization just yields foundationalism with uncertain foundations, and so leaves us with an anti-holistic epistemology nonetheless. Because Jeffrey Conditionalization is rigid, a Bayesianism based on Jeffrey Conditionalization still needs foundational propositions, in the form of partitions that we can assign credences to on the basis of experience alone. If we think there are no such propositions, then we've learned that conditionalization rules were simply the wrong way to think about empirical learning.

We could, of course, go the other way and say so much the worse for holism. The literature on foundationalism and holism is vast, so arguing for either option goes well beyond the scope of this paper. But I do want to draw one connection with another literature before quitting.

The completeness of conditionalization rules has been in question for some time. Do Strict Conditionalization and Jeffrey Conditionalization handle every possible case of empirical learning? Van Fraassen's Judy Benjamin problem ([van Fraassen 1981](#)) is a classic case alleged to show that they do not. In Judy Benjamin's case, what she supposedly learns is not any one proposition, nor even a distribution over a partition. Instead, she learns a conditional probability — the probability that she is in Red Team's headquarters given that she is in Red Team's territory at all. Other alleged examples of this kind of thing are proposed by Bradley ([2005](#)). It is debatable whether these really are cases where the evidence cannot be represented in a conditionalization-friendly way. For example, maybe what Judy learns when she radios her HQ and gets a report of a conditional probability is a proposition about HQ's credences, or about the objective chances.¹³

But the problem we've uncovered here can circumvent any quibbles about the particulars of these cases and go straight to the point. Assuming holism, we've seen that the evidence is never representable in a conditionalization-friendly way. Unless

¹³ See ([Grove & Halpern 1997](#)) for an analysis along these lines.

of course (as Bradley notes), we are willing to trivialize Jeffrey Conditionalization by using a super-fine-grained partition. One interesting connection between our problem and other literatures, then, is that we have here an argument for the incompleteness of Jeffrey Conditionalization, one that does not depend on the contentious particulars of cases like Judy Benjamin's.

Appendix

We want to refine the framework of Wagner's result to accommodate the fact that the input partitions might vary with the experience-ordering, since different times will be salient. The setup is now:

$$\begin{array}{ccccc} p & \xrightarrow{\mathbf{E}} & q & \xrightarrow{\mathbf{F}} & r \\ p & \xrightarrow{\mathbf{F}'} & q' & \xrightarrow{\mathbf{E}'} & r' \end{array}$$

where the contents of the \mathbf{E} and \mathbf{E}' partitions are similar, except that they are about different times. For example, E_i might be "there is a red jellybean at t_0 " while E'_i is "there is a red jellybean at t_1 ". The same goes for the partitions \mathbf{F} and \mathbf{F}' .

We also don't want to assume that $r = r'$ as Wagner does, since we aren't assuming that the order of experience doesn't make any difference. All we have to assume is that the order shouldn't make a difference to the agent's credences over the input partitions. More precisely, let's assume that

$$(8) \quad \forall i \forall j : r(E_i F_j) = r'(E'_i F'_j),$$

which is reasonable because we are considering cases where, intuitively, the order of things doesn't matter to the evidence at hand, and so the agent's final state of belief shouldn't reflect a difference. Similarly, we'll want to assume

$$(9) \quad \forall i \forall j : p(E_i F_j) = p(E'_i F'_j),$$

which says that the agent doesn't regard the order as important from the outset. We can then prove bayes-factor identities equivalent to Wagner's, as follows.

Proof. From the rigidity of Jeffrey Conditionalization, we have the identities

$$(10) \quad q(E_i F_j) = q(E_i) p(F_j | E_i)$$

$$(11) \quad q(E_i F_j) = q(F_j) r(E_i | F_j)$$

$$(12) \quad q'(E_i'F_j') = q'(F_j')p(E_i'|F_j')$$

$$(13) \quad q'(E_i'F_j') = q'(E_i')r'(F_j'|E_i')$$

for all i, j . Conjoining (10) with (11) and (12) with (13) yields

$$(14) \quad \forall i \forall j : q(E_i)p(F_j|E_i) = q(F_j)r(E_i|F_j)$$

$$(15) \quad \forall i \forall j : q'(E_i')r'(F_j'|E_i') = q'(F_j')p(E_i'|F_j')$$

Now if we let $i = i_1$ and then i_2 in (14), and solve for $q(E_{i_1})$ and $q(E_{i_2})$, we get:

$$(16) \quad q(E_{i_1}) = q(F_j)r(E_{i_1}|F_j)/p(F_j|E_{i_1})$$

$$(17) \quad q(E_{i_2}) = q(F_j)r(E_{i_2}|F_j)/p(F_j|E_{i_2}).$$

Then for $\beta_{q,p}(E_{i_1} : E_{i_2})$ we have:

$$(18) \quad \beta_{q,p}(E_{i_1} : E_{i_2}) = \frac{q(E_{i_1})/q(E_{i_2})}{p(E_{i_1})/p(E_{i_2})}$$

$$(19) \quad = \frac{q(E_{i_1}) p(E_{i_2})}{q(E_{i_2}) p(E_{i_1})}$$

$$(20) \quad = \frac{q(F_j)r(E_{i_1}|F_j)/p(F_j|E_{i_1}) p(E_{i_2})}{q(F_j)r(E_{i_2}|F_j)/p(F_j|E_{i_2}) p(E_{i_1})}$$

$$(21) \quad = \frac{q(F_j)r(E_{i_1}|F_j) p(F_j|E_{i_2}) p(E_{i_2})}{p(F_j|E_{i_1}) q(F_j)r(E_{i_2}|F_j) p(E_{i_1})}$$

$$(22) \quad = \frac{q(F_j)r(E_{i_1}|F_j) p(F_jE_{i_2})}{p(F_jE_{i_1}) q(F_j)r(E_{i_2}|F_j)}$$

$$(23) \quad = \frac{r(E_{i_1}|F_j) p(F_jE_{i_2})}{p(F_jE_{i_1}) r(E_{i_2}|F_j)}$$

$$(24) \quad = \frac{p(F_jE_{i_2}) r(E_{i_1}|F_j)}{p(F_jE_{i_1}) r(E_{i_2}|F_j)}$$

$$(25) \quad = \frac{p(F_jE_{i_2}) r(E_{i_1}F_j)}{p(F_jE_{i_1}) r(E_{i_2}F_j)}.$$

That is,

$$(26) \quad \beta_{q,p}(E_{i_1} : E_{i_2}) = \frac{p(E_{i_2} F_j) r(E_{i_1} F_j)}{p(E_{i_1} F_j) r(E_{i_2} F_j)}$$

Similar substitutions in (15) and parallel reasoning yield

$$(27) \quad \beta_{r',q'}(E'_{i_1} : E'_{i_2}) = \frac{p(E'_{i_2} F'_j) r'(E'_{i_1} F'_j)}{p(E'_{i_1} F'_j) r'(E'_{i_2} F'_j)}$$

From (8) and (9) we know that (26) and (27) must be identical, which gives us a bayes factor identity like the one in (3):

$$(28) \quad \forall i \forall j : \beta_{q,p}(E_i : E_j) = \beta_{r',q'}(E'_i : E'_j).$$

Similarly, substituting $j = j_1, j_2$ in (14) and (15) and following analogous reasoning gives us a bayes factor identity like in (4):

$$(29) \quad \forall i \forall j : \beta_{r,q}(F_i : F_j) = \beta_{q',p}(F'_i : F'_j).$$

So the worries raised so far don't eliminate the problem. Even if we allow that input partitions can vary with the order of experience, the same result follows so long as we can make assumptions like (8) and (9).

References

- Bradley, Richard. 2005. Probability kinematics and bayesian conditioning. *Philosophy of Science* 72: 342–364.
- Christensen, David. 1992. Confirmational holism and bayesian epistemology. *Philosophy of Science* 59: 540–557.
- Domotor, Zoltan. 1980. Probability kinematics and representation of belief change. *Philosophy of Science* 47(3): 384–403.
- Doring, Frank. 1999. Why bayesian psychology is incomplete. *Philosophy of Science* 66(Proceedings): S379–S389.
- Field, Hartry. 1978. A note on jeffrey conditionalization. *Philosophy of Science* 45: 361–367.
- Garber, Daniel. 1980. Field and jeffrey conditionalization. *Philosophy of Science* 47(142-145).
- Grove, Adam & Joseph Y. Halpern. 1997. Probability update: Conditioning vs. cross-entropy. In *Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence*.

- Jeffrey, Richard C. 1965. *The Logic of Decision*. Chicago: University of Chicago Press.
- Lange, Marc. 2000. Is jeffrey conditionalization defective by virtue of being non-commutative? remarks on the sameness of sensory experience. *Synthese* 123(3): 393–403.
- Quine, Willard V. O. 1953. Two dogmas of empiricism. In *From A Logical Point of View*. Cambridge: Harvard University Press.
- Skyrms, Brian. 1986. *Choice and Chance*. Belmont: Wadsworth, 3rd edn.
- van Fraassen, Bas. 1981. A problem for relative information minimizers. *British Journal for the Philosophy of Science* 32: 375–379.
- van Fraassen, Bas. 1989. *Laws and Symmetry*. Oxford: Clarendon Press.
- Wagner, Carl. 2002. Probability kinematics and commutativity. *Philosophy of Science* 69: 266–278.