

Jeffrey Conditioning, Rigidity, and the Defeasible Red Jelly Bean

Abstract: Jonathan Weisberg has argued that Jeffrey Conditioning is inherently "anti-holistic" By this he means, *inter alia*, that JC does not allow us to take proper account of after-the-fact defeaters for our beliefs. His central example concerns the discovery that the lighting in a room is red-tinted and the relationship of that discovery to the belief that a jelly bean in the room is red. Weisberg's argument that the rigidity required for JC blocks the defeating role of the red-tinted light rests on the strong assumption that all posteriors within the distribution in this example are rigid on a partition over the proposition that the jelly bean is actually red. But individual JC updates of propositions do not require such a broad rigidity assumption. Jeffrey conditionalizers should consider the advantages of a modest project of targeted updating focused on particular propositions rather than seeking to update the entire distribution using one obvious partition. Although Weisberg's example fails to show JC to be irrelevant or useless, other problems he raises for JC (the commutativity and inputs problems) remain and actually become more pressing when we recognize the important role of background information.

Introduction

In "Commutativity or Holism? A Dilemma for Conditionalizers," Jonathan Weisberg (2009) argues *inter alia* that Jeffrey Conditioning (hereafter JC) is inherently "anti-holistic," by which he means that JC does not allow us to take proper account of after-the-fact defeaters for our contingent beliefs.

Weisberg (2009, pp. 797-98) defines "holism" in strongly anti-foundationalist terms: Holism, on his definition, requires that *all* propositions, including appearance propositions, be sensitive to background evidence in the sense that there always exists some defeater proposition F that, if we knew it to be true, would defeat our justification for any proposition, including a proposition such as "I am having a red-jelly-bean-like experience."

Since they reject Weisberg's thoroughgoing holism anyway, classical foundationalists might seem to have little interest in a critique of JC as anti-holistic. But Weisberg's central example uses a proposition that classical foundationalists would not regard as indefeasible--

E The jelly bean is red.

Weisberg's argument shows that, under certain rigidity assumptions about the transition induced by a red-jelly-bean-like experience, our justification for E when we seem to see a red jelly bean cannot later be defeated by another proposition

F The lighting of the room is tinted red.

This is a highly counterintuitive result. It seems that the discovery, while looking at a red-appearing jelly bean, that the room's lighting is tinted red should indeed undermine our confidence in the actual redness of the jelly bean. Even foundationalists should be interested in this result if they are interested in making use of JC for purposes of analysis or for updating where it would be onerous to speak always in terms of foundational

appearance propositions. Foundationalists and anti-foundationalists share an interest in evaluating Weisberg's critique of JC.

I will argue that Weisberg's critique can be answered, though the answer requires that we take a rather modest approach to the use of JC as an update rule. This modest approach does differ from the more common approach in the literature and sets aside some of its assumptions about the possibility of using an obvious partition of directly affected propositions to update an entire distribution.

Weisberg's argument

Weisberg (2009, 798, 806-7) asks us to imagine a situation in which, first, we are outside of a room containing a jelly bean. We have not been told what color the jelly bean is. This situation is represented by the prior probability distribution designated p . We then enter the room, look at the jelly bean, and have an experience as if of a red jelly bean. The distribution induced by this new experience is designated q . Weisberg points out that, intuitively, F and E should be probabilistically independent in p but dependent in q .

Suppose that we stipulate that, in q , conditioning on F would lower the probability of E to its prior probability by negating the effect of the red appearance of the jelly bean. In that case, F is negatively relevant to E in q . This seems correct, broadly speaking.^{1, 2} If we discovered the lighting to be tinted red, we would have a full explanation of our experience without assuming that the jelly bean is actually red. Under red-tinted light we would expect to have a red-jelly-bean experience even if the jelly bean were not red. It is also intuitively correct to say that F and E are independent in p . After all, why should we think any differently about the actual color of the jelly bean prior to any experiential evidence merely because we are told that the lighting in the room is tinted red? Weisberg captures this intuition by referring to F as an "undercutting defeater" of E rather than a "rebutting defeater" of E, where the latter would be a proposition that is negatively relevant to E from the outset.

Weisberg makes the further, apparently plausible, assumption that the partition $\{E, \sim E\}$ should be used for JC in the shift from p to q , and he formalizes this assumption by supposing that "the transition from p to q is rigid with respect to the partition $\{E, \sim E\}$ " (2009, 807).³ This means, of course, that all posterior probabilities of all other propositions must be rigid on this partition, and in particular that $p(F|E)$ and $p(F|\sim E)$ must equal $q(F|E)$ and $q(F|\sim E)$ respectively.

But these conditions are not all jointly satisfiable. For if the posteriors of F on E are rigid from p to q , it is not possible for F and E to be independent in p and negatively relevant to each other in q . Hence, F cannot act as an undercutting defeater as defined. Weisberg concludes that JC is inherently incapable of taking into account the proper epistemic role of a defeating proposition such as F.

¹ Of course, from both a foundationalist and a holist perspective, no experience could really make one certain of F, so imagining our "finding out that" F and treating this as being "given" F at probability 1 is an idealization.

² Weisberg's argument also applies if F is negatively relevant to E in q but does not completely remove the support for E in q .

³ I am silently changing Weisberg's notation for negation.

The Jeffrey conditionalizer's response

In *The Logic of Decision* (1965, pp. 168ff), Richard Jeffrey discusses in detail what he calls the "important special case" in which some proposition B and its negation provide the proper partition through which the impact of some passage of experience is propagated to an entire probability distribution. However, Jeffrey makes it clear that he does not take it that this will always be the situation in real life. In some cases the proper partition might be more complex, involving all of B_1 - B_n where $n > 2$. The use of the simple binary partition $\{B, \sim B\}$ in Jeffrey Conditioning is, of course, just a special case, and a partition can have any number of members E_1 - E_n as long as the members of the partition basis are mutually exclusive and jointly exhaustive and the posteriors are rigid relative to all members of the set. One way that n could be greater than two would be if two or more propositions and, necessarily, their negations, were directly affected in the transition from the old to the new distribution.⁴ In that case, the posterior probabilities of one of the directly affected propositions would not be rigid *vis a vis* a partition over another directly affected proposition.

Whether or not a proposition is directly affected by a passage of experience can, of course, be a difficult thing to determine, and I do not propose here to provide a formula for determining it. Direct and indirect effects will depend (at least) on the concepts that a subject has and to that extent on background knowledge; it seems plausible to go farther and to say that the "directness" or otherwise of the impact of an experience depends on background knowledge in a great many ways. For purposes of this example, Weisberg regards it as unproblematic to regard E as one of the propositions (with its negation) directly affected, and we may treat that as a point granted on both sides. But we should not try to update some other proposition (such as F) using JC and a partition over E if F is *also* a directly affected proposition. This possibility (that F is also directly affected) introduces both a complication and a possible resource for a response to Weisberg's argument.

A good case can be made that, if E is a proposition directly affected by the experience, then at least for subjects with ordinary background knowledge (which Weisberg assumes) F is also a proposition directly affected by the experience. That is to say, E and F in this example are on a par as far as being directly affected is concerned-- F should be considered to be just as "directly affected" as E . The reason is not far to seek: E and F are both possible explanations of the observer's red-jelly-bean-like experience when he enters the room, and they are in competition with each other as explanations.

The fact that both E and F can explain the experience renders both of them more probable in q than they were in p , but, more to the point, neither is more probable (speaking informally) "because" the other is more probable. Each is more probable simply because it would explain the experience that induces the shift to a new distribution. This is important, because we can think in an epistemically valuable way about indirect epistemic effects in terms of the channeling of the epistemic force of some

⁴ Note that here and in what follows, in phrases such as "two or more propositions" and "another directly affected proposition," I am not speaking of the negation of a proposition as "another" directly affected proposition. So, for example, E and F as described above are *different* "directly affected propositions" from one another as that concept is being used here. It is necessarily true that if E is directly affected, its negation is as well. What is less intuitive is that E and F are both directly affected.

experience to one proposition by another. Suppose, for example, that we think of a proposition like

S The jelly bean is strawberry flavored.

It makes sense to say that the probability of S goes up *because* the probability of E goes up. In q the subject has more reason to think that the jelly bean in the room is strawberry flavored *because* he has more reason to think that it is red. Therefore, S can be usefully thought of as indirectly affected *by way of* a proposition like E. But this is not the case for the relationship between E and F. Rather, if we obtain additional reason to believe that one of them is the true explanation for the red appearance, we have less reason to believe the other.

Conversely, $q(F|\sim E) > p(F|\sim E)$. Once we have observed the red appearance of the jelly bean (but not before), if we were to learn only that the jelly bean is not really red, the probability of F would be significantly raised, since some explanation of the red appearance is needed. All of this, of course, is what lies behind the negative relevance of E and F to each other in q and the failure of rigidity between them in the shift. It makes sense to think of the same epistemic facts as showing us that, if E can be regarded without further argument as "directly affected" by the passage of experience, so can a competing explanation of the experience, F.

The fact that F and E are competing explanations of the passage of experience that induces the shift is somewhat obscured by the fact that in normal epistemic situations $p(F)$ is a good deal lower than $p(E)$. But on careful consideration it is obvious that the posterior probabilities of F on a partition of E are not rigid, contrary to Weisberg's rigidity assumption in the example. The Jeffrey conditionalizer should not grant a rigidity assumption for the partition $\{E, \sim E\}$ that applies across the entire probability distribution.⁵

How does this help?

It might be asked whether this response is really of any help in analyzing the epistemic situation. Does this response show us a plausible set of circumstances in which the intuitively correct negative relation--the undercutting defeater relation--between E and F arises?

The answer to this question becomes more evident when we consider the oddness of even attempting to update F using JC and E. Precisely because in ordinary circumstances it does not seem that E is in any sense a premise for F nor that a change in the probability of F occurs "because of" E, there is something very awkward about trying to use JC to assimilate E as uncertain evidence about F. The sort of proposition that we would typically be interested in updating in this type of scenario using a partition on E is one we discussed above,

⁵ Once we realize that the posteriors for F on E are not rigid, we can see that there are other posteriors for propositions influenced epistemically by F that are also not rigid on $\{E, \sim E\}$, because they are indirectly affected by the experience via F--for example, the posteriors for

T Someone has been tampering with the light in the room.

S The jelly bean in the room is strawberry flavored.

There is every reason to think that the posteriors of S on $\{E, \sim E\}$ are rigid from p to q , nor does the possible action of F as an "undercutting defeater" conflict with *this* rigidity assumption in any way, since no assumption about the rigidity of the posteriors of F is required for the application of JC to S using $\{E, \sim E\}$. Something similar is true, given normal background information, of other propositions, e.g.,

J My daughter, who is fond of red jelly beans, will like the jelly bean in the room.

With some unusual background information that makes me more likely to win the lottery if the jelly bean is really red, E would be relevant to

L I will win the lottery.

There seems no reason to doubt that, if the only thing we know in addition to our normal background knowledge is that I am more likely to win the lottery if the jelly bean in the room is red, a partition on E could be used to update L using JC. In fact, there are plausibly a great many propositions across the entire distribution that have posteriors that are rigid on $\{E, \sim E\}$.

What happens if, in the transition from p to q , we update the probabilities of these propositions using the JC formula and a partition on E? We will find exactly what we should intuitively find--namely, a negative relationship between E and F. In this transition, F is free to be, as intuitively it should be, independent of both E and S in p and negatively relevant to both E and S in q . In such a shift, E and F individually along with more complex propositions such as $(E \& F)$ are affected directly by the experience. Nor does the negative relationship between F and E in q arise by any Jeffrey update; indeed, it cannot, since their posteriors are not rigid *vis a vis* each other, so the Jeffrey formula cannot be used between them. Rather, the negative relationship between them arises immediately from the passage of experience. The propositions that it seems *ought* to be updatable using JC on $\{E, \sim E\}$ are updatable, while F undercuts our justification where it seems that this should be the case.

A formally similar response is available to the non-foundationalist (the thoroughgoing holist) for propositions the holist treats as defeaters for experiential beliefs. Suppose one believes that even a proposition like

E' The jelly bean appears red

is defeasible by

F' I appear to have just gotten a call from the neurology lab informing me that I am a poor judge of my own color experiences.

(This is Weisberg's example [pp. 798-99].) If we take F' to be independent of E' in p and negatively relevant to it in q , the posteriors of F' on E' will not be rigid, but this will make

no difference to the possibility of using JC to update other salient propositions in the distribution on a partition over E' , including E itself.⁶

In general, the posteriors for any after-the-fact defeater F on any proposition E that F undermines will not be rigid. But this is not a problem for the Jeffrey conditionalizer unless he makes the strong assumption that the posteriors of all other propositions in the distribution are rigid on $\{E, \sim E\}$.

The Jelly Bean and the Holy Grail

It will be helpful to be explicit here on a point that has been implicit heretofore: I do not see any necessity for us to discover and be able to access a single, complete partition through which some experience affects the rest of an entire distribution. If we assume the number of distinct propositions in the distribution to be finite, it seems that there must *be* a partition (which I will call the Holy Grail partition) that could be used to propagate the experience throughout the rest of the distribution.⁷ That partition, however, may be of little epistemic use in assimilating experience--a conclusion Weisberg's argument does support. We might previously have assumed that there is some fairly obvious and relatively simple partition of directly affected propositions, perhaps consisting of intuitively identifiable observation sentences related to some experience (and their negations), and that this partition can be used to propagate the effect of the experience through the entire distribution. Weisberg's example strongly suggests that this is not the case, that a larger number of propositions than previously suspected are directly affected by experience (or at least as much so as ordinary observation sentences), and that these after-the-fact defeater propositions will be un-obvious and not much like "The jelly bean is red."

What examples like Weisberg's show, then, is that there is no route from p to q by way of a Jeffrey update on a single, relatively straightforward and readily accessible partition. But if there are many propositions, including un-obvious defeater propositions, directly affected by an experience like that of a red jelly bean, a partition that includes all of them will become unwieldy and Byzantine and hence, inconvenient to use for actual Jeffrey updates. This is the case, first of all, because a complete partition must include all Boolean combinations for the propositions involved. For two propositions, such as E and

⁶ It could be argued that the holist should regard F' as a rebutting defeater of E' , in Weisberg's terminology, rather than an undercutting defeater (2009, p. 807). If we can give any meaning to the notion of my being a poor judge of my color experiences (and a foundationalist will have serious doubts about this), it would seem to have some such consequence as that I may "judge" that I am having any of a very large number of different color experiences on any given occasion. This would seem to reduce the probability before the fact of my judging that I am having a red-jelly-bean experience when I enter the room, since on ordinary background information (without the proposed defeater), my expectations before the fact about jelly bean appearances will be related instead to my information about the actual colors of jelly beans, which are more limited in number. Thus, it would seem that F' , if it is taken in the way that the holist wants it to be taken, should be negatively relevant to E' even in p . I am waiving this point, however, for purposes of the main discussion, since Weisberg apparently intends F' to be treated as an undercutting defeater of E' .

⁷ If the thoroughgoing holist maintains that for *any* partition proposed there is *always* some defeater proposition that has non-rigid posteriors over that partition (see Weisberg 2009, p. 797), it seems that the holist may be envisaging a distribution with an infinite number of distinct propositions, which is psychologically implausible. It would be interesting to see strong holists address this question.

F, a complete partition would be $\{(E \ \& \ F), (E \ \& \ \sim F), (\sim E \ \& \ F), (\sim E \ \& \ \sim F)\}$. For three propositions and their negations, each member of the partition will be a conjunction of three items, and the partition will have eight members, and so forth.

Even when we consider only possible causal defeaters of the general type Weisberg brings up, we can imagine others besides F. A moment's imaginative thought suggests something similar to F', above:

N I am suffering from neurological damage that causes me to see objects as red regardless of whether they are red or not,

as well as

P Someone has put a very realistic picture of a red jelly bean in the room where I will see it,

and the more Cartesian

D A powerful deceiver is causing me to perceive objects as red regardless of whether they are red or not.

None of these is antecedently probable, but then again, we do not usually find that random rooms we enter have red-tinted lights. Most rooms in our present world don't. Nothing in Weisberg's example encourages us to limit our consideration only to antecedently probable defeater propositions. Presumably the idea is to consider propositions which, on our background knowledge, would be sufficient alternative causes of the red-jelly-bean appearance even in the absence of a real red jelly bean. There are many more such possibilities that one could dream up. But even if our partition for transmitting the effect of the experience to the rest of the distribution contained only E, F, N, P, and D and their negations, the complete partition would have a total of thirty-two members, which already merits the adjective "unwieldy."

It is also important to bear in mind that a Holy Grail partition--that is to say, a partition that can be used to transmit the effect of the experience to the entire distribution, not just to some of its propositions--would have to include *all* propositions (and their negations) that are in the distribution but that cannot be updated, because of a failure of rigidity, on the other propositions in the Holy Grail partition. Every proposition in the distribution either is not in the Holy Grail partition (for a new experience) and does have rigid posteriors based on the Holy Grail partition, or it belongs in the partition in order to make the partition complete. This would mean, however, that we would have to include not only defeaters but also propositions (with their negations) that express a level of detail irrelevant to a given proposition we actually want to update. Hence, something like

A The jelly bean is approximately the same shade of red as the dress I seem to remember my aunt's wearing to my wedding

could belong in the Holy Grail, because it is, let us say, “directly affected” by the experience, because of my apparent memories about my aunt's dress and the shade of red the jelly bean appears to be. It cannot then be updated using a partition on E, for even if I knew E to be true, the specific shade of red experienced gives me *additional* reason to believe A. The posteriors of A would therefore not be rigid, though this rigidity failure has nothing to do with defeaters. The relation between the jelly bean's specific shade and my aunt's dress is under many circumstances irrelevant to S, the proposition that the jelly bean is strawberry flavored. Yet A would have to be included in the Holy Grail partition even though its level of detail is irrelevant to S.⁸

The Holy Grail partition might even have to include propositions that are *entirely* epistemically irrelevant to some higher-level proposition of interest we wish to update. We can see this if we consider propositions that are epistemically affected by the fact of my having a red-jelly-bean experience, regardless of its cause. Suppose that I always have terrible cravings for jelly beans whenever I have an experience as if of a red jelly bean, even if that experience occurs in a dream, when I look at a picture, or when I vividly imagine a red jelly bean. Suppose that I know that I usually yield to such cravings, eat far too many jelly beans, and gain several pounds. Suppose further that, because I have easy access to jelly beans whenever I want to go and purchase them, it makes no difference to how much weight I subsequently gain whether I am initially in the presence of a real jelly bean when I have the experience. With this background knowledge in place, I might want to update a proposition like

W I am going to gain several pounds in the next few weeks.

A relevant proposition for transmitting the effect of my experience to W would be

C I will soon experience intense cravings for jelly beans.

C seems like a plausible candidate for inclusion in the Holy Grail partition because of its relevance to W and its increased probability in q , given my background evidence. But propositions like E and F, which relate to the *actual* cause of my experience, are epistemically irrelevant to W. Yet E and F belong in the Holy Grail partition, for E is relevant to updating other propositions like S, and, as we have already seen, F does not have rigid posteriors on a partition over E and hence must be included in the Holy Grail.

From a purely practical epistemic perspective, the existence of the Holy Grail partition would have little point even if we had access to it, since we are at any given time not trying to think about updating all the myriad propositions in our entire distribution at once. The JC formula for updating some particular proposition requires that all the members of *some* partition be used, but for the reasons just discussed, it would be both onerous and unnecessary to do a Jeffrey update using the JC formula and

⁸Since we are now discussing a partition constructed using more than one proposition and hence made up of members which are conjunctions, it would be a good idea to define the concept of “including” a proposition in a partition. For purposes of this discussion, a proposition or a negation of a proposition is included in a partition just in case either it is a member of the partition or is equivalent to the disjunction of all the members of a partition in which it appears as a conjunct. My thanks to a reviewer for *Philosophical Studies* for raising this question and to Timothy McGrew for helping me to refine the definition.

plugging in all the many members of the Holy Grail partition. A simpler partition is what we are looking for in any specific application of JC.

For all these reasons, it seems misguided to make the search for a complete partition reflecting the effects of some experience the Holy Grail of a quest to legitimize Jeffrey Conditioning.

A better approach is a more modest one: JC should be thought of in terms of a three-part relation among an experience, a target proposition H we would like to update, and some relevant uncertain evidence E that we would like to use to update H . This approach is fully sufficient for filling in what is needed for the JC formula itself, and we then need only focus on whether the posteriors of H over a simple partition on E are rigid, which can easily be the case even if a simple partition on E is not the Holy Grail partition for the entire distribution.⁹

Targeted updating permits us simply to ignore propositions such as those discussed above that are entirely irrelevant or that contain a level of detail that is irrelevant to the target proposition. For purposes of targeted updating, we focus on the correct level of detail for a partition proposition for some H without being worried by the fact that another (less or more detailed) version of the same proposition cannot be updated on this partition. Similarly, we can ignore causal propositions if we are trying to update a proposition like W to which the specific cause of our experience is, on our background knowledge, irrelevant.

Targeted updating using JC appears to have received no explicit attention from seminal authors, though it is widely understood that multiple propositions may be directly affected by a passage of experience (see Jeffrey 1965, p. 168-9, Field 1978, p. 362, Christensen 1992, pp. 551-2, Zynda 1995, pp. 77ff). This possibility in itself entails some degree of targeting, if only in a negative sense. If A and B (and their negations) are directly affected by some passage of experience, we must at least be able to do a sufficient amount of targeting to avoid the blunder of attempting to update A on a partition over B or *vice versa*, and the possibility of several directly affected propositions should lead to caution against strong rigidity assumptions about some one proposition E and the rest of the distribution.

But even though Jeffrey, Field, et. al. allow for the possibility of multiple directly affected propositions, it does not follow that they are interested in targeted updating in the strong sense proposed here. To the contrary, Jeffrey was interested in propagating the effect of an experience through *some* partition to the rest of the distribution.

[T]he problem is this. Given that a passage of experience has led the agent to change his degrees of belief in certain propositions B_1, B_2, \dots, B_n from their original values...to new values..., how should these changes be

⁹ It is not necessary, for purposes of using the JC formula, that the uncertain evidence used as a partition be in some sense "as close as possible" to the experience. As long as $P(H|E)$ and $P(H|\sim E)$ are rigid in the transition induced by the experience, the uncertain evidence used as a partition can be much "higher up" in one's evidential framework than a simple sentence like "The jelly bean is red." This point about "closeness" to the experience has been emphasized by James Hawthorne (personal communication). See also Hawthorne's discussion of a basis (2004, p. 93). This point is crucial to the argument in McGrew and McGrew (2008).

propagated over the rest of the structure of his beliefs (Jeffrey, 1965, p. 168)?

Lyle Zynda expressly defines JC in terms of a shift from one entire distribution to another using a single, though possibly complex, partition (Zynda 1995, p. 75).¹⁰

It seems fair to say that many writers on the subject have a more ambitious project in mind than targeted updating, or at least that they identify JC with a more ambitious project--namely, to show how an entire distribution can be updated in one fell swoop if only we could find *the correct* partition to reflect the entire effect of the passage of experience. If defeater examples like Weisberg's do show the difficulty and even the practical pointlessness of trying to find such a partition, Jeffrey conditionalizers will have to lower their sights from such an ambitious project. This may seem to be a high price to pay, but the benefits are also great.

The ability to leave irrelevant or unnecessarily detailed propositions, discussed above, out of consideration is one such benefit. Another benefit is this: If Jeffrey conditionalizers are willing to restrict themselves to asking whether some particular H can be updated using some E and the JC formula, it is difficult to see why they should be worried at all by the mere possibility of defeaters. The strong holist argues that there will always be some defeater F that will not have rigid posteriors on any ordinary partition we might wish to use; this seems like a problem for Jeffrey Conditioning only if we are bound and determined to find a single partition for the transition from p to q . The ability to be unfazed by the possibility of defeaters is a great advantage of the more modest project.

Someone might argue that targeted updating complexifies the use of JC; now, apparently, we must find not only the right partition for every update but also the right partition for updating every target proposition.¹¹ I respond, to the contrary, that simplification is one of the great virtues of targeted updating. Consider the fact that the JC formula itself involves focusing on a specific proposition whose new probability we are trying to calculate. When we try to update some H using the JC formula and some partition, we are already reaching for an accessible partition that can be used *to update H*. This does not become more difficult because we are *not* also, simultaneously, trying to find a partition that can be used to update the entire distribution. On the contrary, such a search could be a distraction from a real and interesting epistemological question before us--the relation of the new experience to H. To be sure, there are many other propositions besides H that we wish to update, but each application of the JC formula involves considering them individually in any event. Moreover, a simple and plausible partition that is usable for H may well be usable (as in the case of $\{E, \sim E\}$) for a great many other propositions in the distribution.

Another objection to targeted updating might be that, for probabilistic consistency, JC can be used to update H on a partition over E only if H is independent of every defeater F in q , a claim that might seem to introduce difficulty in deciding whether

¹⁰ Zynda's interest is in seeing how the notion of confirmation is affected by the possibility that one proposition in the partition made up of directly affected propositions is positively relevant to some H while another is negatively relevant--in other words, where one portion of the partition seems to confirm H while another seems to disconfirm it (Zynda 1995, p. 77).

¹¹ This objection was raised by a reviewer.

we can do even a modest and focused update.¹² But this claim is false if taken to mean that H must be *absolutely* independent of F in q and is non-threatening if it is not taken in that strong sense. As long as a target proposition such as S is independent of F in q *modulo* $\{E, \sim E\}$, that is to say, independent of F conditional on E and independent of F conditional on $\sim E$, the probabilistic relevance of F to S presents no difficulty for a targeted update of S using the partition $\{E, \sim E\}$. We may express this relation by saying that E screens off the influence of F on S.

A target proposition such as

S The jelly bean is strawberry flavored

is, it is true, not *unconditionally* independent of F. If we had strong evidence for

F The lighting of the room is tinted red

this would on normal background be negatively relevant to S in q . In q , F makes it less probable that the jelly bean really is red, and on many sets of normal background evidence, if I have less reason to believe that the jelly bean is really red, I have less reason to believe that the jelly bean is strawberry flavored. So F may be regarded as a defeater of S for reasons closely related to its being a defeater of E. However, the relevance of F to S is exhausted by its relevance to E. The entire reason why F is relevant to S in q is that it makes E less probable. It is worth noting that a similar point can be made about all of N, P, and D above, on normal background information.

Nor is it necessary to think of every possible defeater F and to ask whether it is relevant to the target proposition H only by way of its relevance to E. It is enough to ask whether the new experience that induces the shift is relevant to H only by way of its relevance to E, which is just another way of saying that we need to ask whether the posteriors of H on $\{E, \sim E\}$ are rigid. If they are, JC can be used; the insight by means of which we discern this rigidity entails that any undercutting defeater that might be "out there" somewhere, *even if we have not thought of it*, is in fact independent of H in q *modulo* $\{E, \sim E\}$.

Naturally, this response raises the entire question of how we can know that the posteriors of H are rigid for some partition. But that question has always been present for JC; the possibility of undercutting defeaters does not in itself introduce any new or special problems for JC in connection with targeted updates. As already noted, a proposition like F, N, D, or P which merely provides an alternative to an obvious causal explanation like E but has no other relevance, on the subject's background, to an inferred proposition like S will be screened by a causally relevant proposition like E (and its negation) from S. Other propositions in the distribution can cause rigidity to fail for S and the partition $\{E, \sim E\}$ *only* if they relate the experiential evidence to S in a probabilistically relevant way that is independent of E. It does not matter whether, on the subject's background, such a proposition happens to be a defeater for E or not.

In fact, it might be impossible to update S on $\{E, \sim E\}$ because of a rigidity failure caused by a proposition that, so far from being a defeater for E, is positively relevant to

¹² This objection was raised by a reviewer.

E. Suppose that my new experience--let us say, an experience as if of a round, red object that looks like what I seem to remember a jelly bean looks like--also directly affects

R The jelly bean is round.

I might have statistical background evidence indicating that most round jelly beans are actually red and that all round jelly beans are actually strawberry flavored, even if they are some color other than red. On this background, R is positively relevant to both E and S in q . If I discovered, in q , that E was false (suppose I were to discover that the jelly bean is really white but that the lighting in the room makes it appear red), the experience would still be positively relevant to S, because the experience also raises the probability of R, which raises the probability of S even if the jelly bean is not really red. Hence, S cannot under these circumstances be updated using $\{E, \sim E\}$. But this is because of my background knowledge *about the target proposition S* and its relationship to the experience I am having. It has nothing to do with defeaters for E. This statistical information about roundness and strawberry flavor is also unusual, and on more common background evidence, S could be updated on $\{E, \sim E\}$.

Whatever else we may say about the difficulties in ascertaining when the posteriors are rigid for JC, we ought at least to grant the subject access to his own background knowledge about a target proposition. There seems to be no point otherwise in talking about differences in background information. The analysis here suggests that a subject attempting to ascertain if rigidity holds should make use of background knowledge about the target proposition and its relation to the experience to discover whether an obvious proposition like E, and its negation, screen off the experience from the target proposition. Such screening off would be a sufficient condition for rigidity to hold and for a JC update to be probabilistically permissible (that is, consistent with the probability calculus) using the obvious updating proposition for that target.¹³

I have given some example cases in which, on normal as opposed to unusual background knowledge, those independence conditions do seem to hold. For example, they seem to hold for normal background knowledge (as opposed to the unusual statistical background knowledge about roundness and flavor just discussed) and the updating of S on E. They seem to hold for normal background knowledge and the updating of J, the proposition about my daughter's enjoyment of the jelly bean, on E. And even in unusual cases, the sufficient independence conditions could be discerned to hold for some relevant proposition. For example, in the case of my jelly bean cravings, it makes sense to think that W, the proposition about my weight gain, could be updated on a simple partition over C, the proposition about my impending cravings. If such independence conditions are ascertainable at all (which seems a rather modest assumption, and is certainly a much more modest requirement than that the subject have access to a Holy Grail partition), targeted Jeffrey updates can be carried out. In any event, the possibility of defeaters *qua* defeaters is not the issue for purposes of the targeted update.

Conclusion: Remaining problems for JC

¹³ For a full discussion of this proposal concerning rigidity, see McGrew (2010).

How serious a problem is it if there is no clear and simple JC partition for propagating an experience throughout an entire distribution? The question does not admit of a single answer, because it depends entirely on what a given philosopher wanted from JC in the first place. If one had hoped that JC could be used to assimilate experience to whole distributions in a rule-governed way via obvious observation sentences, thereby minimizing the role of epistemic intuition about the effect of experience on non-certain propositions, the realization that propositions like F are also directly affected by experience will be a bitter pill to swallow. If, on the other hand, one regarded the JC formula only as a useful epistemic tool for particular propositional updates, then the fact that a great many salient and interesting propositions can be updated using an obvious partition and Jeffrey's rule will be quite sufficient.

Nothing in this discussion implies that all the problems Weisberg raises for JC are answered by the argument given here about after-the-fact defeaters. In fact, both the inputs problem and the commutativity problem remain as vexed as ever partly because of what we have seen--the complex direct effects of experience.¹⁴

Weisberg is certainly right to suspect that a solution to these remaining problems like Field's (1978, p. 366), which involves using the same Bayes factor to model the impact of an experience regardless of background knowledge, is epistemologically incorrect. The defense of JC against the charge that it is intrinsically incapable of dealing with after-the-fact defeaters depends on recognizing that E and F are competing explanations of the red appearance. But epistemologically this also means that, if I have strong experiential reason to believe F *before* experiencing the appearance of a red jelly bean, my Bayes factor for the red-jelly-bean experience should register little or no confirmation for E, since I will already have a good reason to expect the red appearance even if the jelly bean is not really red. Bayes factors for sensory evidence must be sensitive to background evidence--hence, for some pairs of experiences, order-dependent--for reasons directly related to the fact that multiple propositions may be equally directly affected by sensory evidence. Field is seeking an input parameter that "can be viewed as in some sense a function of sensory stimulation alone" (1978, p. 364); such a thing is apparently not to be found if observation sentences are not the only things directly affected by experience.

Foundationalists and thoroughgoing holists will give diverging, indeed, antithetical advice to Jeffrey conditionalizers about how to solve all of their difficulties. The central argument here has shown that we need not violate epistemic intuitions about after-the-fact defeaters when we use the JC formula. If the correct answer to the defeater challenge makes the inputs problem and the commutativity problem more pressing than ever, Jeffrey conditionalizers will have a choice to make between the different camps of their advisers and well-wishers.¹⁵

¹⁴ Someone who wished to use the JC formula could avoid the commutativity problem by treating JC as an analytical tool while assuming an underlying structure of true Bayesian conditioning on certainties. That approach would, of course, be contrary to thoroughgoing holism in any event, thus confirming Weisberg's point that commutativity and thoroughgoing holism as he defines it are incompatible.

¹⁵ I wish to thank James Hawthorne, Timothy McGrew, and an anonymous reviewer for *Philosophical Studies* for comments on this paper.

References

- Christensen, David (1992). "Confirmational Holism and Bayesian Epistemology," *Philosophy of Science* 59:540-557.
- Field, Hartry (1978). "A Note on Jeffrey Conditionalization." *Philosophy of Science* 45:361-367.
- Hawthorne, James (2004). "Three Models of Sequential Belief Updating on Uncertain Evidence." *Journal of Philosophical Logic* 33:89-123.
- Jeffrey, Richard (1965). *The Logic of Decision*. Chicago: University of Chicago Press.
- McGrew, Lydia (2010). "Probability Kinematics and Probability Dynamics." *Journal of Philosophical Research* 35:89-105.
- McGrew, Timothy and McGrew, Lydia (2008). "Foundationalism, Probability, and Mutual Support." *Erkenntnis* 68:55-77.
- Weisberg, Jonathan (2009). "Commutativity or Holism? A Dilemma for Jeffrey Conditionalizers." *British Journal for the Philosophy of Science* 60:793-812.
- Zynda, Lyle (1995). "Old Evidence and New Theories." *Philosophical Studies* 77:67-95.