Holistic Conditionalization and Underminable Perceptual Learning

Abstract

Seeing a red hat can (i) increase my credence in *the hat is red*, and (ii) introduce a negative dependence between that proposition and potential undermining defeaters such as *the light is red*. The rigidity of Jeffrey Conditionalization makes this awkward, as rigidity preserves independence. The picture is less awkward given 'Holistic Conditionalization', or so it is claimed. I defend Jeffrey Conditionalization's consistency with underminable perceptual learning and its superiority to Holistic Conditionalization, arguing that the latter is merely a special case of the former, is itself rigid, and is committed to implausible accounts of perceptual confirmation and of undermining defeat.

1 Introductory Matters

What do we expect from a theory of perceptual learning? Here's a plausible thought: a complete theory of the epistemology of perceptual learning would specify how having some particular experience affects the beliefs of rational agents. More carefully, it would provide a rule of the form: $(P(\cdot), \mathcal{E}) \mapsto P^{+}(\cdot)$, where $P(\cdot)$ is the agent's prior credence function, $\mathcal{E}$ is the experience, and $P^{+}(\cdot)$ is the posterior credence function that an agent with $P(\cdot)$ ought to adopt upon having experience $\mathcal{E}$. Bayesian Conditionalization (specifically: Jeffrey Conditionalization), on the other hand, specifies how a change in a handful of attitudes ought to affect an agent's other attitudes: it's a rule of the form $(P(\cdot), \{< e_i, \omega_i >\}) \mapsto P^{+}(\cdot)$, where the e_i are propositions that partition the agent's prior probability space and the ω_i are the revised weights of the e_i. Experiences are not weighted partitions — $\mathcal{E}$ and $\{< e_i, \omega_i >\}$ are very different sorts of things — so Bayesian Conditionalization is not a complete theory of the epistemology of perceptual learning.

In what sense, then, is Bayesianism a theory of perceptual learning at all? The idea seems to be that the initial or immediate effect of experience $\mathcal{E}$ is to spark revisions to a small number of credences, which lead to other revisions that are mediated by an update rule. Bayesianism is then a theory of the mediate effects of experience: it takes as its input a prior credence function together with the immediate effects of $\mathcal{E}$ — weighted partition $\{< e_i, \omega_i >\}$ — and it produces a posterior credence function as output via Jeffrey Conditionalization:

Jeffrey Conditionalization: $P^+(\cdot) = \sum_i P(\cdot \mid e_i) \cdot \omega_i$

In what follows it will be important to clearly distinguish the credence revisions that proceed via the various forms of Conditionalization from those that provide the weighted partition to be conditionalized upon, so for convenience I'll introduce some terminology. The effects of experience that are not modeled or regulated by Conditionalization I'll call *exogenous* revisions (as in *exogenous to the model*), and the revisions that are modeled and so proceed by Conditionalization I'll call *endogenous* revisions.[1] Hence the general Bayesian picture of perceptual learning is a two-stage process that involves both types of revision:

$$\underbrace{(P(\cdot),\mathcal{E}) \mapsto \rlap{$\overbrace{\phantom{(P(\cdot),\{< e_i,\omega_i >\}) \mapsto P^+(\cdot)}}^{\text{Endogenous revision}}$}(P(\cdot),\{< e_i,\omega_i >\})}_{\text{Exogenous revision}} \mapsto P^+(\cdot)$$

We can now state more carefully how Bayesianism is an *incomplete* theory of perceptual learning. Whether the posterior credence function adopted is rationally appropriate for an agent who has experience $\mathcal{E}$ will depend not only upon the adequacy of Jeffrey Conditionalization, but also upon whether conditionalizing on $\{< e_i, \omega_i >\}$ was the appropriate response to $\mathcal{E}$. Bayesianism is silent on that question, so Bayesianism doesn't completely determine whether the posterior credence function adopted is rationally appropriate.

Familiar objections to Bayesianism focus on putative problems *inside* the model, problems that arise either from the demand for probabilistically coherent credences (e.g. the problem of logical omniscience) or from the demand that all modeled credence revisions proceed via Conditionalization (e.g. the problem of old evidence). In a more recent line of criticism, Jonathan Weisberg (2009; 2014) argues that Jeffrey Conditionalization is inconsistent with common intu-

[1] This terminology originates in Howson and Urbach (1993).

itions about the defeasibility of perceptual learning, and in particular with the vulnerability of perceptual learning to undermining defeat.

Suppose that I have a visual experience as of a red hat. Plausibly, that experience won't just affect my beliefs about the color of the hat or my beliefs about my own experiences, it also affects which propositions function as defeaters for those beliefs. Before I have my experience as of the hat I would regard *I'm hallucinating* as evidentially independent of *the hat is red* — neither confirming nor disconfirming it — an independence expressed formally as $P(red \mid hallucinating) = P(red)$. After my experience as of the hat's redness I become much more confident that the hat is in fact red, but at this point I no longer think that those propositions are independent. After all, my high confidence is based on the experience, and learning that I was hallucinating is a good reason to doubt that my experience is an appropriate basis for my belief, so $P^+(red \mid hallucinating) < P^+(red)$. But that loss of independence is impossible, Weisberg argues, because Jeffrey Conditionalization is 'rigid' with respect to the elements of the update partition:[2]

Rigidity: For any endogenously revised A and any exogenously revised partition element e_i, $P(A \mid e_i) = P^+(A \mid e_i)$

Rigidity says that conditionalizing on partition $\{e_i\}$ can't change my credence in any other proposition conditional on some e_i. That's problematic because rigidity is independence preserving:[3]

RIP: If the transition from $P(\cdot)$ to $P^+(\cdot)$ is rigid on the partition $\{e_i\}$ and $P(A \mid e_i) = P(A)$ for all $e_i \in \{e_i\}$, then $P^+(A \mid e_i) = P^+(A)$ for every $e_i \in \{e_i\}$

Hence if *the hat is red* and *I'm hallucinating* are evidentially independent, and then I conditionalize on a partition including *the hat is red* as an element, those

[2]Proof: Let e_1 be one the $e_i \in \{e_i\}$. As elements of a partition the e_i are pairwise inconsistent, so for any $e_j \in \{e_i\}$ such that $e_j \neq e_1$, $P(A\&e_1 \mid e_j) = 0$, so $P(A\&e_1 \mid e_j) \cdot P^+(e_j) = 0$. By Jeffrey Conditionalization, $P^+(A\&e_1) = \sum_i P(A\&e_1 \mid e_i) \cdot P^+(e_i)$, but whenever some $e_j \neq e_1$ is the value of e_i, the resulting summand equals 0. Hence $P^+(A\&e_1) = P(A\&e_1 \mid e_1) \cdot P^+(e_1)$, so ${}^{P^+(A\&e_1)}/{}_{P^+(e_1)} = P(A\&e_1 \mid e_1) = P(A \mid e_1)$. By the definition of conditional probability ${}^{P^+(A\&e_1)}/{}_{P^+(e_1)} = P^+(A \mid e_1)$, so for any partition element e_1 and any proposition A whose credence is determined by conditionalizing on weighted partition $\{e_i\}, P(A \mid e_1) = P^+(A \mid e_1)$.

[3]Proof: By the total probability theorem and the definition of conditional probability, $P^+(A) = \sum_i P^+(A \mid e_i) \cdot P^+(e_i)$. The rigidity of the transition ensures that $P(A \mid e_i) = P^+(A \mid e_i)$, so $P^+(A) = \sum_i P(A \mid e_i) \cdot P^+(e_i)$. The prior independence of A and each e_i means that $P(A \mid e_i) = P(A)$, so this becomes $P^+(A) = \sum_i P(A) \cdot P^+(e_i)$. $\{e_i\}$ forms a partition, so the $P^+(e_i)$ sum to 1, so $P^+(A) = P(A) \cdot 1 = P(A)$. Finally, by prior independence $P(A) = P(A \mid e_i)$, which by rigidity is equal to $P^+(A \mid e_i)$, so $P^+(A) = P^+(A \mid e_i)$.

propositions must remain independent. That's inconsistent with the compelling story that I just told about the functioning of undermining defeaters, and so Weisberg concludes that Jeffrey Conditionalization should be rejected.

In response to Weisberg's Puzzle, Gallow (2014) argues that Jeffrey Conditionalization must be rejected in favor of an alternative update rule that he calls 'Holistic Conditionalization':

> [Weisberg's puzzle shows that] neither Conditionalization nor Jeffrey Conditionalization. . . is capable of accommodating the confirmation holist's claim that beliefs acquired directly from experience can suffer undermining defeat. I will diagnose this failure as stemming from the fact that neither of these rules give any advice about how to rationally respond to experiences in which our evidence is theory-dependent, and I will propose a novel updating procedure which does tell us how to respond to these experiences. (Gallow, 2014, 493-4)

My purpose in this essay is to defend the superiority of Jeffrey Conditionalization over Holistic Conditionalization. My argument proceeds in three steps. First, I argue against both Gallow and Weisberg that Jeffrey Conditionalization is perfectly consistent with perceptual learning that is vulnerable to undermining defeat. Second, I show that Holistic Conditionalization is a special case of Jeffrey Conditionalization, rather than an alternative to it. Finally, I argue that there are independent reasons to prefer Jeffrey Conditionalization.

2 Jeffrey Conditionalization and Undermining Defeat

Jeffery Conditionalization is consistent with perceptual learning that is vulnerable to undermining defeat. It's not that Jeffrey Conditionalization isn't rigid, or that rigidity doesn't preserve independence; it is, and it does. But the only independence that Rigidity preserves is between individual partition elements and propositions not in the partition. As a result, constructing an instance of Weisberg's puzzle requires careful attention to partition selection: in order to preserve the independence of *the hat is red* and *I'm hallucinating* (as the puzzle requires), exactly one of those propositions must be a partition element.

How are partition elements selected? One appealing thought is that partition elements are propositions directly affected by experience. For example, an experience as of a hat might directly affect *the hat is red* and *the hat is dirty* and no other propositions, in which case the input partition would include those two propositions as elements. Clearly the propositions directly affected by experience should be among those exogenously revised, and hence they must appear in the input partition, in some sense of 'appear in'. But there is good reason to doubt that they must always appear as *elements* of that partition, a reason independent of Weisberg's puzzle. The problem is that partition elements must be pairwise inconsistent and exhaustive of the prior probability space, and *the hat is red*, *the hat is dirty* are likely to be neither (depending on the details of $P(\cdot)$). Jeffrey Conditionalization takes only weighted partitions as inputs, so in this case the agent is unable to update.

Having identified this potential problem himself, Jeffrey proposed that, in many cases at least, input partitions must be more complicated than a mere set of immediately affected propositions. His proposal was that the partitions contain a set of conjunctions, each conjunct of which is either one of the directly affected propositions or its negation, with every directly affected proposition or its negation appearing exactly once in each conjunctive element. (1983, p. 173) Hence upon having an experience as of the dirty red hat, instead of updating on $\{$*red, dirty*$\}$, which is unlikely to partition the probability space, I should conditionalize on $\{red\ \&dirty, red\ \&\ \neg dirty, \neg red\ \&\ dirty,\ \neg red\ \&\ \neg dirty\}$, which is guaranteed to partition any probability space.

Jeffrey's proposal allows the propositions immediately affected by experience to be included in the input partition without including them as elements of that partition. This in turn allows the posterior credences of those propositions to be determined exogenously and conditionalized upon (indirectly, by conditionalizing upon the partition elements).

Though motivated by a very different set of problems, Jeffrey's proposal can be repurposed a response to Weisberg's puzzle. Rigidity prevents the introduction of a negative correlation between partition elements and other propositions via Jeffrey Conditionalization. Taking the partition elements to be conjunctions doesn't change that: Jeffrey Conditionalization still cannot introduce a negative correlation between a conjunctive element and some other proposition. What it can do, however, is to introduce a negative correlation between the *conjuncts* of those conjunctive elements.

Here's why: $P^+(A\&B)$ and $P^+(\neg A\&B)$ together determine $P^+(B)$ and (trivially) $P^+(A\&B)$. Similarly, $P^+(A\&B)$ and $P^+(A\&\neg B)$ together determine $P^+(A)$. Hence the weights of the conjunctive partition elements determine

$P^+(A|B)$. A and B are independent iff $P^+(A|B) = P^+(A)$, and hence their independence (or lack thereof) is completely determined by the posterior weights of the conjunctive elements of the partition, which are themselves determined exogenously. The upshot is that it's possible to introduce the desired correlation between A and B by exogenously re-weighting the conjunctive elements of the partition.

This approach is not without cost. Conjunctive elements are not the direct effects of experience in any intuitive sense, so on this approach Bayesianism cannot be a theory of the indirect epistemic effects of experience; far more will have to be left out. In particular, much of what's interesting about undermining defeat will be determined exogenously at the point of weighted partition selection rather than endogenously via Conditionalization.[4]

For a response to this objection and further motivation for this approach see Miller (2015). The purpose of this section is merely to demonstrate that Jeffrey Conditionalization is consistent with perceptual learning that is vulnerable to undermining defeat. It should now be obvious that it is, on the condition that both the propositions acquired directly from experience and their potential underminers are taken as conjuncts of the conjunctive elements of the input partition.

3 Holistic Conditionalization

Jeffrey Conditionalization is perfectly consistent with the phenomenon of undermining defeat, and hence Gallow's claim to the contrary is false. Nonetheless his proposed solution to Weisberg's puzzle — the rejection of Jeffrey Conditionalization in favor of Holistic Conditionalization — might be preferable for other reasons.

Both Gallow and Weisberg understand the phenomenon of undermining defeat as arising from the the (putative) theory dependence of perceptual evidence. On this view, the propositional evidence produced by an experience depends upon the agent's background theories, and accounting for this dependence is essential to responding to Weisberg's puzzle. Background theories are propositions, and so they have credences according to $P(\cdot)$. Thus a version of confirmation holism is true — red-hat experience $\mathcal{E}_{RH}$ might produce one weighted partition for an agent with $P(\cdot)$ and another for an agent with $P'(\cdot)$

[4] For example see Christensen (1992) and Weisberg (2014).

— because $P(\cdot)$ and $P'(\cdot)$ might assign different credences to the relevant background theories on which the epistemic significance of $\mathcal{E}_{RH}$ depends.

According to the Holistic Conditionalizer, the problem with Jeffrey Conditionalization is that although it's sensitive to *the fact that* $\mathcal{E}_{RH}$ produced the propositional evidence that it did, it's insensitive to *the reason why* $\mathcal{E}_{RH}$ produced that evidence. Since those dependence facts aren't reflected in $P(\cdot)$, and since they aren't introduced by Jeffrey Conditionalization, those dependence facts won't be reflected in $P^{+}(\cdot)$ either. Finally, since reference to those facts is essential to any solution to Weisberg's puzzle (see below), the Jeffrey Conditionalizer will be unable to solve the puzzle.

The general idea is simple: the propositional evidence generated by experience depends in part on the agent's attitudes towards their background theories, and since agents are not always certain which background theory is true, they are not always in a position determine whether some proposition is evidence. What they are in a position to do, however, is to determine whether some proposition would be evidence, given a particular background theory.

For example, consider background theories t_V = *my experiences are all veridical* and t_M = *my experiences are all misleading.* Had I been sure that t_V, then my red-hat experience $\mathcal{E}_{RH}$ would have produced *the hat is red* as propositional evidence. In that case, after have $\mathcal{E}_{RH}$ I would remain sure that t_V (trivially) and I would have become sure that *the hat is red*, and so I would be sure of their conjunction. I know what to do when I become sure of a proposition: I Strictly Conditionalize, setting $P^{+}(\cdot) = P(\cdot \mid t_V \&$ the hat is red). Alternately, had I been sure that t_M, then $\mathcal{E}_{RH}$ would have produced *the hat is not red* as propositional evidence. In that case I would be sure in both t_M and *the hat is not red*, and by Strictly Conditionalizing on their conjunction I would set $P^{+}(\cdot) = P(\cdot \mid t_M \&$ the hat is not red).

The interesting cases are those in which I'm unsure which of my background theories is true, and hence I'm unsure about the evidence propositions that depend on those theories. For example, I might be unsure between t_V and t_M, but sure that: conditional on t_V my evidence includes *the hat is red*, but conditional on t_M it doesn't. In this case my uncertainty about the background theories translates into uncertainty about whether *the hat is red* is evidence. Gallow proposes two very similar rules for updating on uncertain, theory-dependent evidence. Where t_i is a background theory and e_i is an evidence proposition that depends on t_i, both rules involve calculating $P^{+}(\cdot)$ as a weighted sum of $P(\cdot \mid t_i \& e_i)$, for each t_i/ e_i pair. First:

Holistic Conditionalization: $P^+(\cdot) = \sum_i P(\cdot \mid t_i \& e_i) \cdot P(t_i)$

Holistic Conditionalization offers a response to Weisberg's Puzzle. Recall that the puzzle arises because experience has at least two distinct epistemic effects: it provides propositional evidence, and it introduces negative correlations between that propositional evidence and its potential undermining defeaters. The putative problem for Jeffrey Conditionalization is that although there is no barrier to incorporating newly acquired propositional evidence into the posterior credence function, its rigidity appears to make it impossible to introduce the necessary correlations between propositional evidence and its undermining defeaters. I have proposed that Jeffrey Conditionalizers respond to Weisberg's puzzle by conditionalizing upon conjunctions of the newly acquired propositional evidence and its potential undermining defeater. This solves the puzzle by introducing the needed correlation at the point of input partition selection — the exogenous revision stage — which obviates the need to introduce that correlation via Conditionalization (which is impossible).

Holistic Conditionalization avoids the problem in essentially the same way. Any propositions that need to become correlated with the evidence propositions, including any potential undermining defeaters, are taken to be the background theories: the t_i. After holistically conditionalizing, each e_i becomes certain conditional upon t_i.[5] Assuming $P^+(e_i) < 1$, e_i and t_i will be positively correlated after conditionalizing. This correlation holds regardless of the relationship between $P(e_i \mid t_i)$ and $P(e_i)$,[6] and in particular in holds even if e_i and t_i are independent relative to $P(\cdot)$. Because a *positive* correlation is established between e_i and t_i, a *negative* correlation is established between e_i and $\neg t_i$, meaning that any subsequent increased confidence in $\neg t_i$ means a decreased confidence in e_i. In other words, $\neg t_i$ is now a defeater for e_i.

However, Holistic Conditionalization has a problematic consequence: for each conjunction $t_i \& e_i$, $P^+(t_i \& e_i) = P(t_i)$,[7] which ensures that for each t_i, $P^+(t_i) = P(t_i)$.[8] In other words, according to Holistic Conditionalization, per-

[5] Proof: combining results from footnotes 7 and 8 yields $P^+(e_i \& t_i) = P^+(t_i)$, so $P^+(e_i \& t_i)/P^+(t_i) = 1$ (or undefined), so $P^+(e_i \mid t_i) = 1$ (or undefined).

[6] Assuming $P(e_i \& t_i) > 0$.

[7] Proof: by Holistic Conditionalization, $P^+(t_1 \& e_1) = \sum_i P(t_1 \& e_1 \mid t_i \& e_i) \cdot P(t_i)$. One of the t_i will be t_1 itself, and so one of the summands must be $P(t_1 \& e_1 \mid t_1 \& e_1) \cdot P(t_1) = P(t_1)$. The other summands are calculated using the other t_i, but those values will all be 0: since the background theories form a partition they must be pairwise inconsistent, so for every $t_i \neq t_1$, $P(t_1 \& e_1 \mid t_i \& e_i) \cdot P(t_i) = 0$. The result is that $P^+(t_1 \& e_1)$ is equal to the sum of $P(t_1)$ and a bunch of 0's, so it's equal to $P(t_1)$.

[8] Proof: $P^+(t_i)$ can't be any lower than $P^+(t_i \& e_i)$, and in order to be higher there must be

ceptual experience can't affect credences in background theories. But that's implausible. Suppose I'm sure that either *the lighting is normal* or *the lighting is red* and then I have an experience as of a red hat. That's exactly the sort of experience that one would expect given that the lighting is red, so my experience should make me more confident that the lighting is red, i.e. it should change my credence in a background theory. But by Holistic Conditionalization that's impossible.

Anticipating this objection, Gallow offers a variant of Holistic Conditionalization on which credences in background theories vary according to their degree success in predicting the evidence.

Holistic Conditionalization*: $P^+(\cdot) = \sum_i P(\cdot \mid t_i \& e_i) \cdot P(t_i) \cdot \Delta_i$

Here Δ_i is a *probability ratio*: one measure of a theory's success in predicting the evidence. This value is multiplied by the prior probability of the theory to determine its posterior probability. Understood this way both Strict and Jeffrey Conditionalization have Δ-values. When the evidence is a propositional certainty (as required by Strict Conditionalization), the probability ratio of theory t to evidence e is:

$$\Delta_t = \frac{P(e \mid t)}{P(e \mid \top)}$$

Informally, the denominator establishes a baseline probability of the evidence against which to compare the probability of that evidence conditional on the theory, as represented in the numerator. If the evidence is made more probable by the theory, then $\Delta_t > 1$, and since $P^+(t) = P(t) \cdot \Delta_t$, that means that t is confirmed by the evidence. And since we've stipulated that the background theories form a partition, if one theory receives a credence boost by having a Δ-value greater than 1, that boost must come at the expense of some other theory with a Δ-value less than 1.

When the evidence is a weighted partition rather than a propositional certainty (as permitted by Jeffrey Conditionalization), the probability ratio of theory t to evidence $\{e_j\}$ is:

$$\Delta_t = \sum_j \frac{P(e_j \mid t)}{P(e_j \mid \top)} \cdot \omega_j$$

some $i' \neq i$ such that $P^+(t_i \mid t_{i'} \& e_{i'}) > 0$. But the t_i form a partition, so that's impossible. Hence $P^+(t_i) = P^+(t_i \& e_i)$, which by fn. 7 equals $P(t_i)$.

Here each element of the evidence partition establishes its own baseline against which the predictive success of the theory is measured. As before, if e_j is more probable conditional on t than conditional on $\top$ (i.e. than the unconditional probability of that element), then $P(e_j|t)/P(e_j|\top)$ is greater than 1 and t receives some confirmation. The value of Δ_t, then, is the sum of those fractions (one for each e_j) weighted by the posterior credences of ω_j. Finally, Δ_t will be greater than 1 (thus indicating that t is confirmed by $\{e_j\}$) iff a sufficient number of partition elements are made sufficiently more probable relative to their individual baselines and then weighted sufficiently highly.

Holistic Conditionalization*'s Δ-values are determined by considerations similar to those of Jeffrey Conditionalization: the theory's relative success in predicting the evidence. However, for Holistic Conditionalization* the formal implementation of that approach is complicated by the fact that the background theories are allowed to disagree about what the evidence is. For example, it might be the case that if I were sure that t, then the evidence would be e, but if I were sure that t', then the evidence would be e'. This is important in the present context because the prior probability of the evidence is the baseline against which each theory's predictive success is measured, and hence without a shared body of evidence there's no shared baseline.

Although Holistic Conditionalization* allows background theories to disagree about whether some proposition is part of the evidence, for some other proposition they might agree. For example, suppose that my background theories are *the lighting is normal* and *the lighting is red* and then I have an experience as of a red hat. While my background theories might disagree about whether *the hat is red* is part of my evidence, presumably they will agree that *it appears that the hat is red* is part of my evidence. Presumably it will also be the case that one theory does a much better job at predicting this shared evidence than the other: if the lighting is red, then any hat that I see will appear to be red, whereas normal lighting is consistent with the appearance of a non-red hat. Hence the shared evidence more strongly confirms *the lighting is red* than *the lighting is normal.* Informally, then, the proposal is that we calculate the Δ-value for each theory using only shared evidence and ignoring disputed evidence.

Formally, we begin by establishing the shared baseline against which the predictive success of our background theories can be measured. Let $\{e_j\}$ be the set of propositions accepted as evidence by at least one theory, and let $\{t_i\}$ be a set of background theories (as before the background theories partition the probability space). For any $e_1 \in \{e_j\}$ there will be a non-empty subset of $\{t_i\}$ consisting of theories that regard e_1 as evidence; call it τ_1. Since each $t_i \in \tau_1$ agrees that e_1 is an evidence proposition, we can use the probability of e_1 conditional on τ_1[9] as a common baseline against which to measure each $t_i \in \tau_1$'s success in predicting e_1, i.e. to measure the probability of e_1 conditional on each of the t_i. Finally, the Δ-value for each background theory t is determined by taking the weighted sum of these measurements of t's success:

$$\Delta_i \equiv_{df} \sum\nolimits_j \frac{\delta(e_j \mid t_i)}{P(e_j \mid \tau_j)} \cdot \frac{P(\tau_j)}{\sum_k P(\tau_k)}$$

where:[10]

$$\delta(e_j \mid t_i) \equiv_{df} \begin{cases} P(e_j \mid \tau_j) \text{ if } t_i \notin \tau_j \\ P(e_j \mid t_i) \text{ if } t_i \in \tau_j \end{cases}$$

4 Special Cases

One might be forgiven for thinking that Holistic Conditionalization is a generalization of Jeffrey Conditionalization (just as Jeffrey Conditionalization is a generalization of Strict Conditionalization), and that it's in virtue of this greater generality that Holistic Conditionalization is able to respond to Weisberg's puzzle. We've now seen that the latter point is false: both Jeffrey and Holistic Conditionalization are able to respond to Weisberg's puzzle. In this section I show that the former point is false as well: that both Holistic Conditionalization and Holistic Conditionalization* are special cases of Jeffrey Conditionalization

[9] This is somewhat confusing: how can we define the probability of proposition e_1 conditional on *set of propositions* τ_1? Answer: replace each set τ_i with the disjunction of all the $t_i \in \tau_i$. I've adopted Gallow's notation here, and this appears to be what he has in mind.

[10] If the numerator on the left represents the agent's credence in e_j conditional on t_i, then why '$\delta(e_j \mid t_i)$' rather than '$P(e_j \mid t_i)$'? The point of the Δ-values is to calculate the credence increase or decrease that theories receive its success in predicting each evidence proposition e_j when that theory regards e_j as evidence, and for that predictive success to be irrelevant when that theory does regard e_j as evidence. Hence what's wanted is for:

$$\frac{\delta(e_j \mid t_1)}{P(e_j \mid \tau_j)} \cdot \frac{P(\tau_j)}{\sum_k P(\tau_k)} = \frac{P(\tau_j)}{\sum_k P(\tau_k)}$$

This requires that $\frac{\delta(e_j|t_1)}{P(e_j|\tau_j)} = 1$, which is exactly what we get when $\delta(e_j \mid t_1)$ is replaced with $P(e_j \mid \tau_j)$, in which case:

$$\frac{\delta(e_j \mid t_1)}{P(e_j \mid \tau_j)} \cdot \frac{P(\tau_j)}{\sum_k P(\tau_k)} = \frac{P(e_j \mid \tau_j)}{P(e_j \mid \tau_j)} \cdot \frac{P(\tau_j)}{\sum_k P(\tau_k)} = 1 \cdot \frac{P(\tau_j)}{\sum_k P(\tau_k)}$$

in precisely the same sense that Strict Conditionalization is a special case of Jeffrey Conditionalization.

What exactly does it mean to say that one update rule is a special case of another rule? Here's an initial account: update rules are mappings from elements of an input set to posterior credence functions, and R_S is a special case of R_G iff (i) R_S's inputs are a proper subset of R_G's inputs, and (ii) R_S and R_G map each of of their shared inputs to the same posterior credence function. I'll call this the Strict Account, for reasons that will become apparent below.

Given the Strict Account it's clear why Strict Conditionalization is a special case of Jeffrey Conditionalization, at least on one way of understanding Special Conditionalization. As I understand it, Strict Conditionalization is a rule for updating on new propositional certainties: it's a norm governing how to revise one's credences upon becoming certain in the truth of some evidence proposition.[11] On this understanding, the input to our rule is a kind of doxastic state, together with a prior credence function. To facilitate an important distinction below, call this interpretation 'Strict Conditionalization (dox)'. Jeffrey Conditionalization too is a rule for updating on credence changes, but this time there's no demand for certainty, and credences can take any value in the interval [0,1]. Hence the forms of the two rules are:

Strict Conditionalization (dox): $(P(\cdot), \{< e, 1 >, < \neg e, 0 >\}) \mapsto P^+(\cdot)$

Jeffrey Conditionalization: $(P(\cdot), \{< e_i, \omega_i >\}) \mapsto P^+(\cdot)$

Since $\{< e, 1 >, < \neg e, 0 >\}$ is one of many possible instances of $\{< e_i, \omega_i >\}$, the inputs to Strict Conditionalization (dox) are a proper subset of the inputs to Jeffrey Conditionalization. And since Jeffrey Conditionalizing upon $\{< e, 1 >, < \neg e, 0 >\}$ means setting $P^+(\cdot)$ equal to $(P(\cdot \mid e) \cdot 1 + P(\cdot \mid \neg e) \cdot 0) = P(\cdot \mid e)$ – precisely what Strict Conditionalization (dox) recommends – both rules recommend the same posterior credence function for each shared input. Both conditions of the Strict Account are met, so Strict Conditionalization (dox) is a special case of Jeffrey Conditionalization.

Complicating the picture is a second way of understanding Strict Conditionalization, on which one updates upon propositions rather than propositional certainties. On this understanding Strict Conditionalization is a norm governing how one should revise credences upon obtaining e as evidence, rather than

[11] Authors who understand Strict Conditionalization this way include Jeffrey (1983, 165) and Talbott (2016).

a norm governing how one should revise credences upon becoming certain that e. Understood in this second way, the form of Strict Conditionalization is:

Strict Conditionalization (prop): $(P(\cdot), e) \mapsto P^{+}(\cdot)$

Strict Conditionalization (prop) and Jeffrey Conditionalization have different kinds of evidential inputs – propositions and weighted partitions, respectively – so the possible inputs to the former are not a proper subset of the possible inputs to the latter. Hence according to the Strict Account, Strict Conditionalization (prop) is not a special case of Jeffrey Conditionalization.[12] Nonetheless the near consensus in the literature is that both versions of Strict Conditionalization are special cases of Jeffrey Conditionalization, at least in some sense.[13] If that near-consensus is correct, then the Strict Account is too strict.

The first order of business is to clarify the relationship between the two versions of Strict Conditionalization. The main difference, of course, is that they take different sorts of inputs: propositions, and doxastic states. Nonetheless, there's an intuitive sense in which the rules are the same (there's a reason it passes without comment that they're both referred to as 'Strict Conditionalization'); call that intuitive sameness 'quasi-equivalence'. One plausible explanation for this quasi-equivalence of the two versions of Strict Conditionalization begins by noting the ease of translating between the propositional inputs of Strict Conditionalization (prop) and the doxastic inputs of Strict Conditionalization (dox). In order to translate propositional input $(P(\cdot), e)$ into doxastic input $(P(\cdot), \{< e, 1 >, < \neg e, 0 >\})$, we first determine the content of the doxastic state by identifying it with the propositional evidence (along with its negation). The content of the doxastic state is then weighted as prescribed by Strict Conditionalization (prop) itself: $\omega_e = P^{+}(e|e) = 1$, and $\omega_{\neg e} = P^{+}(\neg e|e) = 0$. Amenability to translation in this way is the first component of the quasi-equivalence of Strict Conditionalization (prop) and Strict Conditionalization (dox). The second component is simply that that both rules determine the same posterior credence function from equivalent possible inputs: Strictly Conditionalizing (prop) on $(P(\cdot), e)$ yields the same posterior credence function as Strictly Conditionalizing (dox) on $(P(\cdot), \{< e, 1 >, < \neg e, 0 >\})$.

[12]Authors who understand Strict Conditionalization this way include Meacham (2016, 768), Van Fraassen (1980, 167-8), and Williamson (2000, 214).

[13]According to Meacham (2016, 778), that Strict Conditionalization is a special case of Jeffrey Conditionalization is 'a standard part of Bayesian Lore'. Nearly every author who comments on the topic seems to agree; see also van Fraassen (1980, 170), Gallow (2014, 495), Hartmann and Sprenger (2011, 620), Jeffrey (2004, 53-5), Titelbaum (ms, 147), Weisberg (2011, 501), and Williamson (2000, 214-16). For an important dissent see Christensen (1992).

With this on the table, we can succinctly state the sense in which Strict Conditionalization (prop) is a special case of Jeffrey Conditionalization: Strict Conditionalization (prop) is quasi-equivalent to Strict Conditionalization (dox), which is itself a special case of Jeffrey Conditionalization according to the Strict Account. I'll have more to say about this translation procedure below, but first I'll show that Holistic Conditionalization is a special case of Jeffrey Conditionalization in this same sense.

Holistic Conditionalization is a rule of the form $(P(\cdot), \{e_i\&t_i\}) \mapsto P^+(\cdot)$: it determines posterior credences from a prior credence function together with a set of background theory/ evidence proposition conjunctions. As with Strict Conditionalization (prop), the evidential inputs to Holistic Conditionalization are propositional rather than doxastic. Hence in order to show that Holistic Conditionalization is a special case of Jeffrey Conditionalization in the same sense as Strict Conditionalization (prop), we first employ our procedure for translating between propositional and doxastic inputs. The doxastic inputs to Jeffrey Conditionalization being represented by weighted partitions of the prior probability space, the immediate goal is to show that each $(P(\cdot), \{e_i\&t_i\})$ input to Holistic Conditionalization translates to a weighted partition.

Using the same translation procedure as before, each of Holistic Conditionalization's possible $(P(\cdot), \{e_i\&t_i\})$ inputs is identified with a partition whose elements are the members of $\{e_i\&t_i\}$, and where the weight ω_i of each $e_i\&t_i$ partition element is equal to $P^+(e_i\&t_i)$, as determined by Holistically Conditionalizing on $(P(\cdot), \{e_i\&t_i\})$. How can we be sure that the resulting $(\{< e_i\&t_i, \omega_i >\})$ actually partitions the posterior probability space? In order for $(P(\cdot), \{e_i\&t_i\})$ to be a possible input to Holistic Conditionalization, the t_i must partition the prior probability space. As we've seen (footnote 7), Holistic Conditionalization ensures that $P^+(e_i\&t_i) = P(t_i)$, so for any possible $(P(\cdot), \{e_i\&t_i\})$ input to that rule, $\sum_i P^+(e_i\&t_i) = 1$. What's more, given that the background theories partition the prior probability space, $P(t_i\&t_j) = 0$, and hence $P[(t_i\&e_i)\&(t_j\&e_j)] = 0$ as well. Holistic Conditionalization cannot raise credences from 0 any more than Strict or Jeffry Conditionalization, so it follows that $P^+[(t_i\&e_i)\&(t_j\&e_j)] = 0$. This shows that each possible input to Holistic Conditionalization maps to a possible input to Jeffrey Conditionalization via the same translation procedure we used to map each possible input of Strict Conditionalization (prop) to a possible input to Strict Conditionalization (dox). However, since Jeffrey Conditionalization lacks Holistic Conditionalization's constraints upon partition weighting, the possible inputs to the latter are a proper subset of the possible inputs to the former. Finally, recall that Holistic Conditionalization says that:

$$P^{+}(\cdot) = \sum_i P(\cdot \mid t_i \& e_i) \cdot P(t_i)$$

$P^{+}(t_i \& e_i) = P(t_i)$ for every $t_i \& e_i$, so by substitution:

$$P^{+}(\cdot) = \sum_i P(\cdot \mid t_i \& e_i) \cdot P^{+}(t_i \& e_i)$$

This is precisely what Jeffrey Conditionalization would advise when updating upon a partition with elements of the form $t_i \& e_i$, which is the form shared by all inputs common to both rules. Hence Holistic Conditionalization is a special case of Jeffrey Conditionalization in precisely the same sense that Strict Conditionalization (prop) is.[14,15]

Holistic Conditionalization* too is a special case of Jeffrey Conditionalization. Holistic Conditionalization* takes inputs of the form $(P(\cdot), \{e_i \& t_i\})$. Any $P(\cdot)$ will be partitioned by an appropriately weighted set of conjunctions of e_i's and t_i's along with their negations as described above. Holistic Conditionalization* weights each $e_i \& t_i$ according to $P(t_i) \cdot \Delta_i$, and Gallow (2014, 517-9) proves that $\sum_i P(e_i \& t_i) \cdot \Delta_i = 1$. Since the t_i are required to be pairwise inconsistent, it follows that the inputs to Holistic Conditionalization* are weighted partitions of the probability space. In other words, each $(P(\cdot), \{e_i \& t_i\})$ input to Holistic Conditionalization* determines a $(P(\cdot), \{< e_i, \omega_i >\})$ input to Jeffrey Conditionalization via our familiar translation procedure. Some possible inputs to Jeffrey Conditionalization are not possible inputs to Holistic Conditionalization* – e.g. any partition such that $\omega_i \neq (P(t_i) \cdot \Delta_i)$ – so the latter are a proper subset of the former. Any weighted partition such that $\omega_i = P(t_i) \cdot \Delta_i$ determines the same $P^{+}(\cdot)$ by either rule, and in that case Jeffrey Conditionalization's $\sum_i P(\cdot \mid t_i \& e_i) \cdot \omega_i$ is equivalent to Holistic Conditionalization*'s $\sum_i P(\cdot \mid t_i \& e_i) \cdot P(t_i) \cdot \Delta_i$. Hence Holistic Conditionalization* is a special case of Jeffrey Conditionalization.

So what's the significance of this result? Importantly, observing that one update rule is a special case of another does not trivialize either, or render

[14]Compare Huber (2014).

[15]Like Jeffrey Conditionalization, Holistic Conditionalization is also a rigid update rule. Proof: we've just seen that Holistic Conditionalization is equivalent to $P^{+}(\cdot) = \sum_i P(\cdot \mid t_i \& e_i) \cdot P^{+}(t_i \& e_i)$. the posterior probability space is partitioned by $\{t_i \& e_i\}$, so $P^{+}(\cdot) = \sum_i P^{+}(\cdot \mid t_i \& e_i) \cdot P^{+}(t_i \& e_i)$ is an instance of the total probability theorem. Combining terms and simplifying yields $P^{+}(\cdot \mid t_i \& e_i) = P(\cdot \mid t_i \& e_i)$.

either of them uninteresting. Indeed, as I discuss below there might be important advantages of the special case over its generalization. What's more, the special case relation that I've described, the one that obtains between Strict Conditionalization (prop) and Jeffrey Conditionalization, has some surprising instantiations. An anonymous referee provides an example. Consider:

Field Conditionalization: $P^{+}(h) = \frac{\sum_i P(h\&e_i)\cdot\alpha_i}{\sum_j P(h\&\neg e_j)\cdot\alpha_j}$

The point of Field's rule is to isolate an 'input parameter' (the α_i) representing the evidential significance of an experience for each evidence proposition e_i, an impact that's independent of the agent's prior credences.[16] A positive value for α_i indicates that evidence proposition e_i is confirmed by the experience, and negative values indicate disconfirmation; it is required that $\sum_i \alpha_i = 0$. As a result, the inputs to Field Conditionalization are not weighted partitions representing doxastic states. Nonetheless, given our translation procedure and the Strict Account, Field Conditionalization is a special case of Jeffrey Conditionalization. The surprising thing is that *Silly* Field Conditionalization is also a special case of Jeffrey Conditionalization:

Silly Field Conditionalization: $P^{+}(h) = \frac{\sum_i P(h\&e_i)\cdot(-\alpha_i)}{\sum_j P(h\&\neg e_j)\cdot(-\alpha_j)}$

Both rules have the general form $(P(\cdot), \{< e_i, \alpha_i >\}) \mapsto P^{+}(\cdot)$, but given the same $(P(\cdot), \{< e_i, \alpha_i >\})$ input the two rules determine very different posterior credence functions.

However, there is another sense in which this is not surprising at all. As we saw, the Strict Account is too strict, as it only allows us to compare rules that share the same kinds of inputs (propositions, doxastic states, etc). As a result, if we want to compare rules such as Strict Conditionalization (prop) and Jeffrey Conditionalization, we must translate between inputs of different types, and in particular we must translate between propositional inputs and doxastic inputs. Doxastic inputs have two components. First are the contents of the doxastic state, which are represented by the partition elements. Second are the the agent's credences in those contents, which are represented by the weights of the partition elements. The obvious candidate for the content of the doxastic state correlated with propositional evidence e is e itself, together with its negation $\neg e$ (in order to form a partition). Determining the credence in e that is correlated with *possessing proposition e as evidence* is left up to the

[16] Field (1978)

propositional evidence rule in question. After all, the role of an update rule is to determine posterior credences from old beliefs and new evidence. When the new evidence is propositional, the posterior credences determined will include the posterior credence in the evidence proposition itself. Rules differ in precisely which posterior credences are determined by evidence proposition e, and in particular they differ in the posterior credence determined for e itself. As a result, a single piece of propositional evidence might be translated into different doxastic inputs by different rules. For example, compare:

Strict Conditionalization (prop): $[P(\cdot), e] \mapsto P^+(\cdot) = P(\cdot|e)$

Contrarian Conditionalization: $[P(\cdot), e] \mapsto P^+(\cdot) = P(\cdot|\neg e)$

Dogmatic Conditionalization: $[P(\cdot), e] \mapsto P^+(\cdot) = P(\cdot|\top)$

The three rules share the same set of possible propositional inputs, each of which can be translated into a possible doxastic input to Jeffrey Conditionalization (and not vice versa). But which particular doxastic input a propositional input translates into depends upon the rule in question. For example, proposition input $(P(\cdot), e)$ translates into doxastic input $(P(\cdot), \{< e, 1 >, < \neg e, 0 >\})$ relative to Strict Conditionalization(prop), $(P(\cdot), \{< e, 0 >, < \neg e, 1 >\})$ relative to Contrarian Conditionalization, and $(P(\cdot), \{< e, P(e) >, < \neg e, P(\neg e) >\})$ relative to Dogmatic Conditionalization. Updating on $(P(\cdot), e)$ by any of these three rules produces the same posterior credences as Jeffrey Conditionalizing upon its doxastic equivalent, whatever that happens to be, and hence each of the three rules is a special case of Jeffrey Conditionalization.

The question remains: is this really how we should be thinking about what it means for one update rule to be a special case of another? In the present context the answer is yes. Holistic Conditionalization(*) is intended as a response to Weisberg's Puzzle, which purports to show that Jeffrey Conditionalization is inconsistent with common intuitions about underminable evidence propositions. That response is essentially to describe how updating via Holistic Conditionalization(*) on some $(P(\cdot), \{e_i, t_i\})$ input can produce a posterior credence function in which $P^+(e_i \& t_i) > P^+(e_i)$, thus allowing $\neg e_i$ to serve as a defeater for e_i (see above). Note that this response depends only on Holistic Conditionalization(*)'s inputs and the posterior credences produced by updating on them via that rule.

In that context it's highly significant that Holistic Conditionalization(*) is in our sense a special case of Jeffrey Conditionalization, i.e. that (i) each possible input ϕ to Holistic Conditionalization(*) translates into some possible input ϕ' to Jeffrey Conditionalization, and (ii) Holistically Conditionalizing(*) on ϕ produces precisely the same posterior credence function as Jeffrey Conditionalizing on ϕ'. For in that case, since it's possible that Holistically Conditionalizing(*) on ϕ produces a posterior credence function in which $P^+(e_i \& t_i) > P^+(e_i)$ – since Holistic Conditionalization(*) is able to respond to Weisberg's puzzle – and since Jeffrey Conditionalizing upon ϕ' produces that exact same posterior credence function, it is clear that Jeffrey Conditionalization too is able to respond to Weisberg's puzzle.

5 Problems for Holistic Conditionalization*

I began this essay by noting that Jeffrey Conditionalization is not a complete theory of perceptual learning. Posterior credences are produced by two distinct credence revisions: an exogenous revision on which an experience determines a weighted partition, and an endogenous revision on which a weighted partition together with a prior credence function determine a posterior credence function. Since only the endogenous revision is governed by Jeffrey Conditionalization, that rule is incomplete as a theory of perceptual learning. What's more, §2's response to Weisberg's puzzle requires that weighted partitions be composed of long conjunctions of evidence propositions and their potential underminers. Each of those elements must be identified and weighted exogenously, and hence the amount of work done outside of the formal model is greater than might have been expected.

Holistic Conditionalization and Holistic Conditionalization* are incomplete in roughly the same way, each requiring both exogenous and endogenous revisions to produce a posterior credence function. Like Jeffrey Conditionalization, the holistic update rules require very complex input partitions, here consisting of conjunctions of evidence propositions and the background theories. Hence in order to determine the input of Holistic Conditionalization one must first identify the evidence propositions, identify the the relevant background theories, and pair the evidence propositions with the theories that produced them. That determination is entirely exogenous, so here again the amount of work done outside of the formal model is greater than might have been expected.

Nonetheless, there's a case to be made that each holistic rule is *less* incomplete than Jeffrey Conditionalization because each requires *less* work to be done exogenously. The inputs to Jeffrey Conditionalization contain three components: (i) the prior credence function, (ii) the elements of the partition, and (iii) the weights of those elements. Importantly, none of the three elements is defined in terms of the others; each is specified independently. As we've seen, however, Holistic Conditionalization's partition elements are conjunctions of the form $e_i \& t_i$, each weighted according to $P(t_i)$. As a result, once (i) and (ii) are determined, (iii) is determined as well. Similarly, Holistic Conditionalization*'s partition elements are conjunctions of the form $e_i \& t_i$, this time weighted to $P(t_i) \cdot \Delta_i$. But since Δ_i is defined in terms of $P(\cdot)$ and $\{e_i \& t_i\}$, here again (i) and (ii) are sufficient to determine (iii). Hence both Holistic update rules have a prima facie explanatory advantage over Jeffrey Conditionalization.

In spite of this prima facie explanatory advantage, however, both holistic update rules prove to be problematic. As we saw in §3, the particular way that Holistic Conditionalization determines the weights of partition elements makes it impossible for experiences to affect confidence in background theories. That's an intolerable consequence, so Holistic Conditionalization must be rejected in spite of its prima facie explanatory advantage over Jeffrey Conditionalization. In this section I identify four problems for Holistic Conditionalization* that many will find intolerable, concluding that it too should be rejected.

The first problem is that, on its most natural interpretation,[17] Holistic Conditionalization* is committed to the theory dependence of perceptual learning, not just in updating on evidence propositions but also in the determination of evidence by experience.[18] Holistic Conditionalization*'s partition elements are conjunctions of background theories and evidence propositions, and the identity of those evidence propositions depends upon which background theories the agent accepts. For example, if Morgan's sole background theory is *the lighting is normal*, then her experience as of a red hat might produce evidence propositions *the hat looks red* and *the hat is red*, but if Scarlet's sole background

[17] This isn't the only possible interpretation of Holistic Conditionalization*. Learning e independent of background theories can be accommodated by Holistic Conditionalization* with input $\{<\top, e>\}$, and requiring all inputs to be of this form produces a rule quasi-identical to Strict Conditionalization (dox). So-interpreted, Holistic Conditionalization*, like Strict Conditionalization (dox), cannot accommodate inputs that are vulnerable to undermining defeat.

[18] All Bayesians accept the theory dependence of endogenous credence revision. That's just conditionalization, the results of which are partly determined by prior conditional probabilities that are (usually) defined in terms of prior unconditional probabilities. Put another way, the significance of an evidence proposition depends upon background beliefs. As a special case of Jeffrey Conditionalization, and hence as a version of Bayesianism, Holistic Conditionalization* shares this commitment.

theory is *the lighting is red*, then that same experience might produce only *the hat looks red*. The *identities* of evidence propositions are not all that depends on background theories: so too do the *credences* in those evidence propositions, which will be determined along with all other credences by $P(\cdot)$ together with the input partition.

For this reason Holistic Conditionalization* is likely to be rejected by those sympathetic to the *immediacy* of perceptual learning. Immediacy is a core commitment of Dogmatists (Pryor (2000)),[19] and is a natural fit for Phenomenal Conservatives (Huemer (2007)), Knowledge Firsters (Williamson (2000)) and Disjunctivists, some Process Reliabilists (Goldman (2008)) and others. Jeffrey Conditionalization is much more hospitable to immediacy of perceptual learning:[20] both the identities and the weights of evidence propositions are determined exogenously, and the rule is completely agnostic about the nature of exogenous credence revisions.

A second problem is that Holistic Conditionalization* is inconsistent with broadly Moorean treatments of perceptual learning and skepticism.[21] Mooreans hold that my red hat experience can dramatically increase my credence in *the hat is red* — e.g. from 1/5 to 9/10 — even if I don't start out confident that skeptical background theories are false. Suppose that I'm sure that either *the lighting is normal* ($=t_N$) or *the lighting is red* ($=t_R$) and that my credence in each is 1/2. If the lighting is normal, then my experience as of the red hat generates two evidence propositions: e_r = *the hat is red* and e_{Ar} = *the hat appears red*, but if the lighting is red, then my evidence is only e_{Ar}. By Holistic Conditionalization*:

$$\begin{aligned} P^{+}(e_r) =& [P(e_r \mid t_N \& (e_r \& e_{Ar})) \cdot P(t_N) \cdot \Delta_{t_N}] \\ &+ [P(e_r \mid t_R \& e_{Ar}) \cdot P(t_R) \cdot \Delta_{t_R}] \end{aligned}$$

We normally wouldn't expect a correlation between appearing red under a red light and actually being red, so plausibly $P(e_r \mid t_R \& e_{Ar}) = P(e_r)$. In that case, and assuming the values from the preceding paragraph, our equation simplifies

[19] Strictly speaking, Dogmatism is a theory of perceptual *justification* rather than rational credences, while Holistic Conditionalization governs updates to rational credences. If justification and rational credence are allowed to vary independently then there needn't be a conflict between Dogmatism and Holistic Conditionalization.

[20] At least in the initial determination of evidence propositions by experience; endogenous revisions are theory dependent as described in footnote 18.

[21] See Moore (1953).

to:

$$P^+(e_r) = 9/10 = [1 \cdot 1/2 \cdot \Delta_{t_N}] + [1/5 \cdot 1/2 \cdot \Delta_{t_R}]$$
$$= [1/2 \cdot \Delta_{t_N}] + [1/10 \cdot \Delta_{t_R}]$$

According to the Moorean, it should be possible that $P^+(e_r) = 9/10$, but in this case that's not possible. In order for t_i to be confirmed and hence for $\Delta_i > 1$, t_i must do a better job predicting the shared evidence than other background theories. Since background theories must form a partition, it follows that confirmation for one implies disconfirmation for another: $\Delta_{t_N} > 1$ iff $1 > \Delta_{t_R}$. As a result, $P^+(e_r) = 9/10$ iff the experience strongly confirms t_N and strongly disconfirms t_R. But this is the opposite of what Holistic Conditionalization* requires of the case. The only evidence proposition shared by t_N and t_R is e_{Ar}, that the hat appears red, and t_R actually does a better job of predicting e_{Ar} than t_N does; after all, I'm more likely to have red-hat experiences when the lighting is red than when the lighting is normal. That means that this episode of perceptual learning will confirm t_R and disconfirm t_N, precisely the opposite of what's needed. In other words, if Holistic Conditionalization* is correct, then my prior credences constrain my capacity to learn from experience in precisely the way that the Moorean rejects.[22] In contrast, on Jeffrey Conditionalization prior credences do not meaningfully constrain partition weighting, and hence it is consistent with Moorean accounts of perceptual learning.

A third problem for Holistic Conditionalization* is that it is committed to an implausible account of undermining defeat. In some cases, evidence supports a conclusion only when that evidence is combined with an auxiliary hypothesis. For example, *the gas tank is full* is plausibly confirmed by *the indicator points at 'F'* only in combination with the auxiliary hypothesis *the indicator is functioning properly*. Further, if I believe that the gas tank is full on the basis of the evidence together with the auxiliary hypothesis, and then I lose confidence in the auxiliary hypothesis, then my belief that the gas tank is full will suffer some sort of defeat. This case fits a general schema that Pryor (2013) labels 'quotidian undermining': conclusion h is supported by evidence together

[22]This result generalizes to any case in which (i) there's a non-skeptical hypothesis that regards e as evidence and a skeptical hypothesis SK that doesn't, (ii) all other perceptual evidence is shared, and (iii) the skeptical hypothesis does a better job predicting the shared evidence. White (2006, 531-7) employs a similar argument against the combination of Dogmatism and Bayesianism; see Miller (2016) for a response.

with auxiliary hypothesis AUX, and h suffers defeat when confidence in AUX decreases.

Cases relevant to Weisberg's puzzle differ from the gas tank example in that they involve an experience $\mathcal{E}$ supporting an evidence proposition e_i, rather than an evidence proposition supporting conclusion h. Nonetheless, if $\mathcal{E}$ supports e_i only relative to background theory t_i, and if decreased confidence in t_i means decreased confidence in e_i, then the schema is satisfied and e_i is vulnerable to quotidian undermining. According to Holistic Conditionalization*, *any* possible vulnerability to undermining is a product of theory dependence: $\mathcal{E}$'s support for e_i depends on t_i – the analogue of AUX – and support for e_i is undermined only when confidence in t_i decreases. Hence according to Holistic Conditionalization*, all perceptual undermining is quotidian.

But not all perceptual undermining is quotidian. That would imply that every potential undermining defeater for e_i is the negation of one of the t_i upon which $\mathcal{E}$'s support for e_i depends. The set of potential underminers for any proposition supported by an experience is very large. For example, when my perceptual experience as of the red hat supports high confidence in *the hat is red*, that proposition becomes vulnerable to the following undermining defeaters: *my color experience is generally reliable but not in this specific case; I was on color-distorting drugs X, Y, and Z and not on color-drug antidotes a, b, or c; I have a poor memory for color experiences*, and many many more. If each of these underminers is quotidian, as Holistic Conditionalization* requires, then each must be somehow included in the the background theories mediating the evidential significance of $\mathcal{E}$. But they can't themselves be the background theories, i.e. the t_i's: Holistic Conditionalization requires that the $\{t_i\}$, so its elements must be pairwise inconsistent. But many pairs of potential underminers are perfectly consistent with one another: it might be the case that that I'm on color-drugs *and* the hat is under a red light. That means that, when $\mathcal{E}$ supports e_i, the background theories mediating that support must include each of e_i's potential undermining defeaters (or its negation) as a conjunct in a long conjunction (see §2). Taking n as the number of potential underminers, the number of distinct background theories mediating $\mathcal{E}$'s support for e_i is at least 2^n.

There is nothing inconsistent about the resulting picture, and of course the defender of Holistic Conditionalization* is free to offer whatever account of background theories they prefer. The point is simply that the defender of Holistic Conditionalization* is forced to an account on which the background theories mediating perceptual learning are extremely fine-grained and extremely numerous. This is a far messier and less appealing picture than one might have expected.

As with Holistic Conditionalization*, the inputs to Jeffrey Conditionalization will include very fine-grained partitions. But because Jeffrey Conditionalization is agnostic about the origins of its inputs (and on the weights of its partition elements in particular), it needn't construe the partition elements as background theories mediating the episode of perceptual learning, and it needn't construe undermining defeat as resulting from a loss of confidence in the background theory. In other words, it needn't assimilate all cases of undermining defeat to the quotidian schema.

A fourth problem with Holistic Conditionalization* is that it requires an implausible account of uncertainty about evidence propositions. If that rule is correct, then a red hat experience might produce evidence proposition *the hat is red* relative to background theory *the lighting is normal* but not produce that evidence proposition relative to background theory *the lighting is red*. Assuming that $P(\text{hat red} \mid \text{lighting red}) < 1$ and $P^+(\text{lighting red}) > 0$, it follows that $P^+(\text{hat red}) < 1$. However, if *the hat is red* is evidence relative to *the lighting is normal*, then $P^+(\text{hat red} \mid \text{lighting normal}) = 1$.[23] In other words, although evidence propositions needn't be unconditionally certain, they are always certain conditional on the relevant background theories.

One consequence is that the only possible source of rational uncertainty about evidence propositions is uncertainty about background theories. If *the lighting is red* and *the lighting is normal* form a partition — as they must if they are the only background theories relevant to my red-hat experience — and then I definitively rule out *the lighting is red*, I must become certain that the lighting is normal. And since $P^+(\text{hat red} \mid \text{lighting normal}) = 1$, I must also become certain that the hat is red. Hence when I stop being uncertain about my background theories I stop being uncertain about my evidence propositions.

But uncertainty about background theories is not the only possible source of rational uncertainty about evidence propositions. Another possible source is the experience itself. For example, suppose that I inspect a cloth under dim

[23] Proof: suppose $\mathcal{E}$ produces e_1 as evidence relative to t_1. Then $P^+(t_1\&e_1) = \sum_i P(t_1\&e_1 \mid t_i\&e_i) \cdot P(t_i) \cdot \Delta_i$. One of the summands will be $P(t_1\&e_1 \mid t_1\&e_1) \cdot P(t_1) \cdot \Delta_1$, which is equal to $P(t_1) \cdot \Delta_1$. Every other summand will be of the form $P(t_1\&e_1 \mid t_n\&e_n) \cdot P(t_n) \cdot \Delta_n$, where $n \neq 1$. But each of those summands equals 0: t_1 and t_n are elements of a partition, so they are jointly inconsistent, so $t_1\&e_1$ and $t_n\&e_n$ are jointly inconsistent. Hence $P^+(t_1\&e_1) = P(t_1) \cdot \Delta_1$. By parallel reasoning, $P^+(t_1) = P(t_1) \cdot \Delta_1$. As a result, $P^+(t_1\&e_1) = P^+(t_1)$. By the definition of conditional probability it follows that $P^+(e_1 \mid t_1) = {}^{P^+(t_1\&e_1)}/_{P^+(t_1)} = 1$ (or undefined).

candlelight trying to discern its color, ultimately deciding that it's probably green, possibly blue, and only improbably violet. [24] My uncertainty in *the cloth is green* (etc.) is at least in part attributable to the nature of the experience itself.

Suppose that before I see the cloth I am certain about the lighting conditions and the condition of my own perceptual faculties and all the rest. Still, the rational response to my experience is to be uncertain about the color of the cloth.[25] According to Holistic Conditionalization* that's impossible. My experience has at least two effects: it increases my credence in *the cloth is green*, and it makes that increased credence vulnerable to new undermining defeaters, e.g. *I'm hallucinating.* On Holistic Conditionalization*, that's only possible if *the cloth is green* is evidence relative to some background theory t, but not evidence relative to some other background theory t'. (In that case t' is an undermining defeater for *the cloth is green.*) But if I were certain that each undermining background theory t' is false, and hence that t is true, then I must be certain that the cloth is green. But I shouldn't be certain that the cloth is green: the character of my experience makes that unreasonable.

If experience itself is a possible source of uncertainty about evidence propositions, then it must be possible to be certain about all background theories t_i while being uncertain about evidence proposition e_i. That's impossible on Holistic Conditionalization*, which requires that $P^+(e_i \mid t_i) = 1$ any time e_i is evidence relative to t_i. But if $P^+(t_i) = 1$, then $\{e_i \& t_i, \neg e_i \& t_i\}$ forms a partition equivalent to the simple $\{e_i, \neg e_i\}$ partition of Jeffrey's $n = 2$ case. (1983, 168-70) Jeffrey Conditionalization does not meaningfully constrain the weighting of partition elements, so there's no barrier to weighting e_i lower than 1.[26]

[24](Jeffrey, 1983, 165-6) uses this example to motivate his generalization of Strict Conditionalization.

[25]I'm not making any specific claim about the content of visual experience, e.g. that the content is vague. My claim is purely epistemic: at least sometimes, experiences affect evidence propositions without making them certain, and this uncertainty is not a product of uncertainty about background theories.

[26]As an anonymous referee points out, we could accommodate theory-dependent uncertain evidence with a generalization of Holistic Conditionalization*:

$$P^+(\cdot) = \sum_i P(t_i) \cdot \Delta_i \cdot \sum_j P(\cdot \mid t_i \& e_{ij}) \cdot \omega_{ij}$$

It's worth noting, however, that without an account of ω_{ij} the resulting rule (i) is a notational variant of Jeffrey Conditionalization, and (ii) loses Holistic Conditionalization*'s main advantage over Jeffrey Conditionalization: its capacity to determine $P^+(e_i \& t_i)$ (for each $e_i \& t_i \in \{e_i \& t_i\}$) from $P(\cdot)$ (see the beginning of this section). In the absence of further elaboration on the proposal, ω_{ij} values cannot be determined in terms of $P(\cdot)$, and hence neither can $P^+(e_i \& t_i)$. Further, if uncertainty about background theories (as reflected in $P(\cdot)$) is not the only source of rational uncertainty about theory-dependent evidence, then no further

6 Conclusion

Weisberg's puzzle illustrates that underminable perceptual learning combines awkwardly with rigid updating rules such as Jeffrey Conditionalization and Holistic Conditionalization*. If evidence is underminable, then the two must be probabilistically dependent. This dependence cannot be introduced endogenously by a rigid update rule, so it must be introduced exogenously into the weighted partition. Partition selection is mostly unconstrained by either update rule, and hence the lesson of Weisberg's puzzle is that rigid update rules face a previously unappreciated explanatory limitation: they cannot explain the probabilistic dependence relations between evidence and underminers.

Though subject to this limitation, both Holistic Conditionalization* and Jeffrey Conditionalization are consistent with underminable perceptual learning. The former rule is a special case of the latter, and it enjoys a prima facie explanatory advantage. But Holistic Conditionalization* faces a number of problems. First, it is inconsistent with immediate perceptual confirmation. Second, it is inconsistent with Moorean anti-skeptical approaches to perceptual learning. Third, it is committed to an implausible pan-quotidian account of perceptual undermining. And fourth, it identifies uncertainty about background theories as the only possible source of uncertainty in evidence propositions. Jeffrey Conditionalization avoids each of these problems, so in spite of its prima facie explanatory disadvantage Jeffrey Conditionalization is the better rule.

References

Christensen, D. (1992). Confirmational Holism and Bayesian Epistemology. *Philosophy of Science*, 59(4):540–557.

Field, H. (1978). A note on jeffrey conditionalization. *Philosophy of Science*, 45(3):361–367.

Gallow, J. D. (2014). How to Learn From Theory-Dependent Evidence; or Commutativity and Holism: A Solution for Conditionalizers. *British Journal for the Philosophy of Science*, 65(3):493–519.

Goldman, A. (2008). Immediate justification and process reliabilism. *Epistemology: new essays*, pages 63–82.

elaboration on the proposal can hope to regain this advantage.

Hartmann, S. and Sprenger, J. (2011). Bayesian epistemology. In Bernecker, S. and Pritchard, D., editors, *Routledge Companion to Epistemology*, Routledge Philosophy Companions, chapter 55, pages 609–620. Routledge.

Howson, C. and Urbach, P. (1993). *Scientific Reasoning: The Bayesian Approach*. Open Court.

Huber, F. (2014). For true conditionalizers Weisberg's paradox is a false alarm. *Symposion*, 1(1):111–119.

Huemer, M. (2007). Compassionate phenomenal conservatism. *Philosophy and Phenomenological Research*, 74(1):30–55.

Jeffrey, R. C. (1983). *The Logic of Decision*. University of Chicago Press.

Jeffrey, R. C. (2004). *Subjective Probability: The Real Thing*. Cambridge University Press.

Meacham, C. J. (2016). Understanding conditionalization. *Canadian Journal of Philosophy*, 45(5-6):767–797.

Miller, B. T. (2015). Updating, undermining, and perceptual learning. *Philosophical Studies*, pages 1–23.

Miller, B. T. (2016). How to be a bayesian dogmatist. *Australasian Journal of Philosophy*, 94(4):766–780.

Moore, G. E. (1953). Hume's theory examined. In *Some Main Problems of Philosophy*, pages 108–126. The Macmillan Company.

Pryor, J. (2000). The Skeptic and the Dogmatist. *Noûs*, 34(4):517–549.

Pryor, J. (2013). Problems for Credulism. In Tucker, C., editor, *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford University Press.

Talbott, W. (2016). Bayesian epistemology. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition.

Titelbaum, M. G. (ms.). *Fundamentals of Bayesian Epistemology*. Oxford University Press. Forthcoming.

van Fraassen, B. C. (1980). Rational belief and probability kinematics. *Philosophy of Science*, 47(2):165–187.

Weisberg, J. (2009). Commutativity or Holism? A Dilemma for Conditionalizers. *British Journal for the Philosophy of Science*, 60(4):793–812.

Weisberg, J. (2011). Varieties of bayesianism. In Gabbay, D., Hartmann, S., and Woods, J., editors, *Handbook of the History and Philosophy of Logic*, volume 10, pages 477–551. Elsevier.

Weisberg, J. (2014). Updating, Undermining, and Independence. *The British Journal for the Philosophy of Science*, 66:121–159.

White, R. (2006). Problems for Dogmatism. *Philosophical Studies*, 131(3):525–57.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.