

UPDATING, UNDERMINING, AND INDEPENDENCE

Jonathan Weisberg
University of Toronto

ABSTRACT. Sometimes appearances provide epistemic support that gets undercut later. In an earlier paper I argued that standard Bayesian update rules are at odds with this phenomenon because they are ‘rigid’. Here I generalize and bolster that argument. I first show that the update rules of Dempster-Shafer theory and ranking theory are rigid too, hence also at odds with the defeasibility of appearances. I then rebut three Bayesian attempts to solve the problem. I conclude that defeasible appearances pose a more difficult and pervasive challenge for formal epistemology than is currently thought.

SOMETIMES appearances provide initial support that gets undercut later. I might glimpse a red-looking sock but then learn the lighting is deceptive, for example. Call this phenomenon *perceptual undermining* since the undermined support comes from perception, at least in part. Here I argue that three leading models of belief-change cannot accommodate this phenomenon. The standard updating rules of Bayesianism, Dempster-Shafer theory, and ranking theory all mishandle perceptual undermining.

The worry that Bayesianism runs afoul of perceptual undermining originates with David Christensen (1992). Christensen argues that Bayesian update rules at best treat the interaction between perception and background belief as a black box. My response to a reddish glimpse of a sock should depend on what I think about the reliability of my vision, but Bayesianism does not model or regulate this interaction. If I suspect my vision is unreliable, the Bayesian can recommend that I Jeffrey Conditionalize (§1.1) on the proposition *The sock is red* with a middling probability instead of a high one. But the choice of a middling input instead of a high one is not something the Bayesian formalism explains or prescribes (Field, 1978).

In my (2009) I took Christensen’s argument a step further. Jeffrey Conditionalization doesn’t just fail to regulate perceptual undermining, it bungles it. Consider

I am grateful to David Christensen, Branden Fitelson, Alan Hájek, James Hawthorne, Franz Huber, Jim Pryor, Scott Sturgeon, Ralph Wedgwood, audience members of the 2009 Conditionals and Conditionalization Workshop at the University of Leuven, and several anonymous referees for their generous help.

the case where I first glimpse the sock, then learn the lighting is deceptive. Surprisingly, Jeffrey Conditionalization makes my discovery about the deceptive lighting ineffectual, failing to undermine my belief that the sock is red. The reason is that Jeffrey Conditionalization is ‘rigid’ (§1.1), and thus independence preserving. Before I glimpse the sock, the quality of the lighting and the actual colour of the sock are independent, meaning information about one factor has no bearing on my beliefs about the other. Because it is rigid, Jeffrey Conditionalization preserves this independence. Information about the lighting’s deceptiveness has no bearing on my beliefs about the sock’s actual colour, even after I come to believe it is red based on its looking that way.

My aim here is to generalize and strengthen this argument. I first prove a general “RIP theorem” encompassing Bayesianism, Dempster-Shafer theory, and ranking theory. The standard update rules of these frameworks are all Rigid, hence Independence Preserving (RIP). I then critique three solutions that have been proposed in response to the challenge for Bayesianism.

I’ll begin by rehearsing the challenge for Bayesianism in §1. I’ll then lay out the RIP theorem in §§2–3. Then, in §§4–6, I will address three Bayesian solutions to the problem. I conclude in §7 that perceptual undermining poses a more serious challenge than is currently thought.

1. THE CHALLENGE FOR BAYESIANISM

Bayesianism is traditionally anchored in the psychological assumption that we have degrees of belief measurable by real numbers, and the normative assumption that these degrees of belief should obey a synchronic constraint called *Probabilism*:¹

Probabilism: An agent’s degrees of belief ought to obey the probability axioms.

We will assume Probabilism in this section. Our focus is the diachronic question how probabilistic degrees of belief should change in response to new experiences.

1.1. Updating and Experience. Bayesians typically answer the diachronic question by looking to the agent’s prior conditional degrees of belief, those she had before the new experience. The traditional answer in this vein was that the agent should *conditionalize* on a proposition completely describing her observation:

¹Some Bayesians, like Levi (1974), allow imprecise degrees of belief represented by sets of probability functions instead of single probability functions. The problem of perceptual undermining applies to this more liberal view too, as fn. 5 demonstrates.

Conditionalization: If your degrees of belief are given by the probability function p and you make an observation completely described by B , your new degrees of belief should be given by $p'(A) = p(A|B)$ for all A .

The proposition B might be interpreted externally (*There is a red sock on the floor*) or internally (*There appears to be a red sock on the floor*). Either way, most Bayesians now think Conditionalization can't be right, at least not for fallible agents like us. It has the effect of making the agent certain of B , yet experience rarely (if ever) furnishes us with certainties, whether external or internal.

Many Bayesians therefore follow Richard Jeffrey (1965; 1968) in embracing a more liberal rule, one that doesn't require certainties but still revolves around the agent's prior conditional degrees of belief. Jeffrey conceives experience-based updating as having two parts: experience "directly affects" (1968, §2) some of the agent's credences, and these changes rationalize others which "propagate" (1965, p. 168) over her remaining beliefs. How should the indirect effects of experience propagate?

Jeffrey Conditionalization: If your degrees of belief are given by the probability function p , and (i) experience directly affects your credences over the partition $\{B_i\}$ changing them to the values $p'(B_i)$, but (ii) experience does not directly affect any other credences, then your new credences should be given by $p'(A) = \sum_i p(A|B_i)p'(B_i)$ for all A .

Intuitively the idea is this. When experience directly speaks only to the question of which B_i is true, and thus not to what each B_i indicates if true, one's new opinions should weight what each B_i would indicate if true by the new probability that it is true.

In later writings, Jeffrey (1983, pp. 136–7) abandons the two-part picture along with talk of "direct effects" and "propagation". He says instead that Jeffrey Conditionalization is merely a consequence of Probabilism in the special case where the agent responds to experience by changing her credences in such a way that:

$$p'(A|B_i) = p(A|B_i) \text{ for any } A \text{ and } B_i. \quad (1)$$

In this special case we can transform the law of total probability:

$$p'(A) = \sum_i p'(A|B_i)p'(B_i) \quad (2)$$

to obtain as a theorem:

$$p'(A) = \sum_i p(A|B_i)p'(B_i). \quad (3)$$

The idea seems to be that formula (3) is not a diachronic norm governing experience-based updates, but rather a useful tool for describing these updates in cases where it is merely a theorem of Probabilism.² (Cf. Bradley 2005; Wagner 2013.)

However Jeffrey may have understood his rule at various times, we will follow authors like Field and Christensen who understand Jeffrey Conditionalization as a substantive diachronic norm governing the indirect import of experience. In so doing, we adhere to a Bayesian tradition that includes authors like Carnap (Jeffrey, 1975, p. 44), Hacking (1967), and Lewis (1999), all of whom view Jeffrey Conditionalization's predecessor, Conditionalization, as a substantive diachronic norm. Viewed in this way, the substance of Jeffrey Conditionalization lies in its assertion that, when the direct effects of experience leave conditional probabilities on a partition untouched despite touching their conditions, those untouched conditional probabilities serve as the "arrows" along which experience's effects should "propagate".

On this view of things, (1) is an important effect of obeying Jeffrey Conditionalization. When an agent applies Jeffrey Conditionalization, not only does experience make no direct changes to her conditional probabilities given elements of $\{B_i\}$, it makes no *indirect* changes either. This property of Jeffrey Conditionalization is known as *rigidity*:³

Jeffrey Conditionalization is Rigid: If p is a probability function and p' comes from p by Jeffrey Conditionalization on the partition $\{B_i\}$, then $p'(A|B_i) = p(A|B_i)$ for every A and B_i .

Rigidity is a key ingredient in the problem of perceptual undermining.

1.2. **The Problem.** Rigidity looks like a desirable property at first. Suppose I start out thinking that, if the sock on the floor is red, then it is probably my roommate's. If I then observe that it is red, my conditional opinion should not change. I should continue to believe the sock is probably my roommate's if it is red. Other features of my opinion should change. I should come to believe it is probably my roommate's sock, for example. But my conditional beliefs about the owner of the sock given its colour should not change when I observe its colour.

²There are hints of this view in Jeffrey's earlier writings. His views may have been ambiguous all along, wavering between viewing equation (3) as a diachronic norm, as a handy tool in the art of judgment, and as a useful theorem for rationally reconstructing an agent's belief-changes.

³Jeffrey sometimes writes as if rigidity is not a consequence of applying Jeffrey Conditionalization, but rather a precondition for its applicability (Jeffrey, 2002). This understanding fits much better with the special-case-theorem interpretation than with our substantive-diachronic-norm interpretation. If rigidity were a pre-condition for following a diachronic norm, whether you ought to follow it would depend on your posterior credences, when it's those posterior credences the norm is supposed to govern.

Despite appearances though, complete rigidity is not desirable. There are some propositions whose probabilities conditional on the sock's redness should change. Propositions that should undermine my confidence in the sock's redness must change their conditional probabilities in order to function as underminers. To see why, let's work through the sock example from a Bayesian point of view.

Let E be the proposition *The sock is red* and F the undermining proposition *The lighting in the room is deceptive, so that apparent colour is no indication of actual colour*. Initially my credence in E is low; in response to my glance at the sock it becomes high; and if I learn F it should become low again. So we have the following constraints:

$$p(E) = \text{low}, \tag{4}$$

$$p'(E) = \text{high}, \tag{5}$$

$$p'(E|F) = \text{low}. \tag{6}$$

Notice that we also have the constraint:

$$p(E|F) = p(E), \tag{7}$$

because the trickiness of the lighting has no bearing on the actual colour of the sock at the outset. Initially, the lighting is only relevant to the sock's apparent colour. Only after I base my opinion about the sock's redness on its appearance does the quality of the lighting become relevant to its actual colour.

The problem is that (5), (6), and (7) are incompatible with p' coming from p by a rigid update on $\{E, \bar{E}\}$. The following is a theorem:

Rigidity is Independence Preserving: If the transition from p to p' is rigid on the partition $\{B_i\}$ and $p(B_i|A) = p(B_i)$ for every B_i , then $p'(B_i|A) = p'(B_i)$ for every B_i .⁴

The theorem tells us that, because E is independent of F at the outset as stipulated in (7), and the update from p to p' is rigid on the partition $\{E, \bar{E}\}$, E will remain independent of F under p' , contra (5) and (6). Intuitively of course, the trickiness of the lighting should be irrelevant to the sock's colour before the glance, and negatively relevant after. But, as long as the update in between is rigid with respect to the

⁴*Proof.* If the transition from p to p' is rigid with respect to $\{B_i\}$, then $p'(A|B_i) = p(A|B_i) = p(A)$ for every B_i . By the theorem of total probability, $p'(A)$ is a weighted sum of these $p'(A|B_i)$. So if they all have the same value, it must be that $p'(A) = p'(A|B_i)$, for any B_i .

sock's redness, this cannot happen. Rigidity prevents the introduction of a negative correlation between E and F .⁵

Stepping back from the sock example, we can see how rigidity clashes with perceptual undermining in general. Cases of perceptual undermining are ones where a perceptual state provides evidence for a proposition, but some undermining defeater for that support is then discovered. Because the defeater is an undermining one, as opposed to a rebutting one, it is not evidence against the proposition itself; it merely undercuts the evidential support offered by the perceptual state (Pollock, 1986, 2008). In probabilistic terms this means the underminer is irrelevant to the supported proposition at first, but negatively relevant after the perceptual state has lent its support. And this is precisely what Rigidity is Independence Preserving rules out. If the underminer is irrelevant before the perceptual state supports the proposition, it is irrelevant after as well. So rigidity prevents perceptual undermining when it obviously shouldn't.

1.3. **Objections.** The problem might seem easily answered at first, even misguided or marginal. Answering some preliminary objections shows it isn't so.

Objection. The problem results from updating on the wrong partition. If one's evidence is a perceptual state like a glimpse of a red sock, the appropriate evidential propositions are appearance propositions, not propositions about the sock's actual colour. So the right proposition to update on is not $E = \textit{The sock is red}$, but rather something like $E^* = \textit{The sock appears red}$.

Reply. For this objection to work there must be no underminers for appearance propositions. Otherwise the same problem just re-arises at the level of appearance propositions. Suppose, for example, that the proposition *I've just had a brain scan showing that I am an unreliable judge of my own colour experiences* is an underminer for the appearance proposition, *The sock appears red*. At the outset this proposition is probabilistically irrelevant to whether the sock will appear red to me in a moment. However, when the sock does appear red to me, and I become confident on that basis that it appears red to me, this underminer should become negatively relevant to the proposition that the sock appears red. But once again, because Rigidity is Independence Preserving, a negative correlation cannot be introduced, making it impossible for this underminer to act as it should. It cannot be irrelevant before the update and negatively relevant after.

⁵ The problem also affects imprecise credences represented by sets of probability functions. We may assume that the agent definitely regards E and F as independent before the glimpse, so that her representor (van Fraassen, 1990) contains only probability measures on which E and F are independent. Updating each member of her representor by Jeffrey Conditionalization on $\{E, \bar{E}\}$ then yields a posterior representor containing only probability functions on which E and F are independent.

One could maintain that there are no underminers for appearance propositions—that nothing could cast doubt on the proposition that there appears to be a red sock, not even a brain scan. I won't attempt a full refutation of this view here. Doing so would require a much longer discussion than would be appropriate here, and I expect most readers will not be sympathetic to this view anyway. Instead I will offer a few brief reasons for proceeding on the assumption that appearance propositions are not the solution to our problem.⁶

First, if the transition from perception to beliefs about perception is susceptible to error, then the discovery that one is prone to such errors is generally available as an underminer. It may be a difficult, empirical question whether that transition really is generally susceptible to error, but this is just grist for the mill. For then there are empirical discoveries that could cast doubt on our beliefs about our own sensory experiences. Suppose that after much study neurologists conclude sensory experience and belief-formation happen in distinct parts of the brain. They also find that the causal pathway connecting one part to the other can be interrupted by magnetic interference. If one then participates in a study that uses such interference to dupe subjects into believing they are having reddish experiences when they are actually having green ones, it would be entirely reasonable upon being debriefed to wonder what one actually experienced during the experiment.

Second, a number of classic arguments and examples have lead most contemporary philosophers, including many foundationalists, to be leary of the kind of Cartesian certitude we're entertaining. Ryle's famous example of the speckled hen is one widely-discussed reason: looking at a hen with 47 speckles on its facing side, couldn't you make a mistake about how many speckles are in your visual field? Some respond that one's experience of the hen has no determinate number of speckles (Ayer, 1940). One simply experiences it has having many speckles. But this is yet more grist for the mill. Some philosophers think the hen appears to have 47 speckles, others think the appearance has no determinate number. Somebody must be mistaken about what they are experiencing.

And there are other classic arguments for the fallibility of introspection into sensory experience. Reichenbach (1938, p. 176) and Ayer (1946, p. 89) argue that any attempt to classify one's sensory experience as (say) greenish presupposes that one is using this classification in the usual way, a presupposition which could be mistaken (cf. Christensen 1992, pp. 544–5). And Reichenbach (1952) offers a Bayesian

⁶Readers who remain unconvinced can view the arguments to follow as *modus tollens* instead of *modus ponens*: the difficulty handling perceptual undermining in our leading formal frameworks is an argument for the kind of Cartesian foundationalism they advocate.

argument that “phenomenal reports” are not certain. Experiences at an earlier time serve as predictors of what one is likely to experience later, so later observations can disconfirm that one had such-and-such experiences earlier.

Third and finally, if appearance propositions were certain this would eliminate much of the motivation for generalizing Conditionalization to Jeffrey Conditionalization. A major reason for adopting Jeffrey Conditionalization is the thought that, at least sometimes, experience does not provide us with certain information, not even about appearances. There are other reasons one might be interested in Jeffrey Conditionalization, but this one is central for many authors. For these authors, solving our undermining problem by rejecting the defeasibility of appearance propositions is not an appealing option. (Cf. also Christensen 1992, §§3–4.)

Objection. Not all cases of undermining defeat are as pure as in the sock example. Sometimes a proposition is both an undermining defeater and a rebutting defeater. For example, suppose the housekeeper testifies that the butler murdered the master of the house, but we learn later that the butler reported the housekeeper to the master for stealing. This later discovery undermines the housekeeper’s credibility (at least partially), since she has motive to smear the butler. But it is also evidence that the butler did not murder the master, just in virtue of being a display of loyalty. So it undermines *and* rebuts the proposition that the butler committed the murder. When it comes to perceptual undermining, if an undermining defeater is also a rebutting defeater, we cannot assume that the underminer is probabilistically irrelevant to the supported proposition at the outset. So an assumption analogous to (7) does not hold in every case of perceptual undermining.

Reply. The fact that Rigidity is Independence Preserving may not rule out all perceptual underminers, but it still rules out “pure” perceptual underminers. Unless we are prepared to say there is no such thing as a pure underminer of perceptual support, the problem remains. And there do seem to be cases of pure perceptual underminers, as in the sock example. More generally, there is always the possibility of learning that one is prone to misperception in the present environment. Since at the outset this possibility (or something more specific) will usually be irrelevant to what the environment contains, we can expect pure perceptual underminers to be ubiquitous.

Objection. It was a mistake to apply Jeffrey Conditionalization to the sock example in the first place, since clause (ii) of Jeffrey Conditionalization is not satisfied there. Experience doesn’t just directly affect my credences in E and \bar{E} , it also directly affects my conditional credences *given* E and \bar{E} , especially my credence in F given E . Indeed,

this is how the desired negative correlation between E and F gets introduced into p' . (Cf. Wagner 2013.)

Reply. If Jeffrey Conditionalization doesn't apply here, then it never applies (or hardly ever does). Actually discovering a perceptual underminer may be a rare occurrence, but the potential for perceptual underminers is ubiquitous, as the replies to the previous two objections showed. I should always be prepared to give up my perceptually-based beliefs in response to underminers. But then I can't obey Jeffrey Conditionalization, since it gets the conditional probabilities of these potential underminers wrong. Whether or not I actually learn F , $p'(E|F)$ should be low, yet Jeffrey Conditionalization makes it high. So this objection saps Jeffrey Conditionalization's substance by making it generally inapplicable.

Objection. The perceptual undermining problem only shows that Jeffrey Conditionalization is not a global update rule, in the sense that it is not the way to update one's credence in every proposition. We may need some other rule to tell us how to update our credences in perceptual underminers, but Jeffrey Conditionalization is still adequate for less exotic beliefs. For example, I can still count on Jeffrey Conditionalization to dictate my new credence that the sock belongs to my roommate.

Reply. Jeffrey Conditionalization is popular because its recommendations are plausible for the most part. So whatever the correct diachronic norm is, it will probably agree with many of Jeffrey Conditionalization's recommendations. But the fact remains that perceptual undermining exposes Jeffrey Conditionalization to be only partly correct at best. And as long as we do not have a complete account, we should worry about why our account is incomplete. Maybe pursuing rules that revolve around prior conditional credences was wrongheaded. Maybe the correct, general diachronic norm will explain why Jeffrey Conditionalization works when it does and doesn't when it doesn't. In general, philosophers do not rest content with a partial account when counterexamples emerge. We look for a more general account, so that the counterexamples will illuminate what our earlier account missed.

There are more ways Bayesians might respond to the problem of perceptual undermining. But before we consider them, I want to explore another avenue. There are other models of rational belief change, and we might hope that they avoid the problem altogether. In the next two sections I show that two prominent alternatives face the same problem.

2. THE CHALLENGE FOR DEMPSTER-SHAFER THEORY

Dempster-Shafer theory (DST) represents doxastic states differently than Bayesianism, and it uses a different rule to update those states. So we might hope DST can avoid the challenge perceptual undermining poses for Bayesianism. Here I argue that it can't. The RIP theorem applies to DST too: updating in DST is Rigid, hence Independence Preserving. I'll first lay out the elements of DST, then formulate the RIP theorem and perceptual undermining problem for DST.

2.1. Background on Dempster-Shafer Theory. DST represents doxastic states using belief functions instead of probability functions:

Belief Function: A function $bel : \wp(\Omega) \rightarrow [0, 1]$ ^{7,8} is a *belief function* just in case:

- (B1) $bel(\emptyset) = 0$,
- (B2) $bel(\Omega) = 1$,
- (B3) For any subsets A_1, \dots, A_n of Ω ,

$$bel(A_1 \cup \dots \cup A_n) \geq \sum_{I \subseteq \{1, \dots, n\}, I \neq \emptyset} (-1)^{|I|+1} bel(\cap_{i \in I} A_i).$$

While the first two axioms are familiar from probability theory, (B3) is less familiar and certainly more difficult to read. It helps to notice that it is the same as the generalized addition rule from probability theory, except that '=' is replaced by '≥'. So, where probability theory says that:

$$p(A \cup B) = p(A) + p(B) - p(A \cap B),$$

for any propositions A and B , DST only requires that:

$$bel(A \cup B) \geq bel(A) + bel(B) - bel(A \cap B).$$

For this reason, belief functions are sometimes called *super-additive measures*, as opposed to the *additive measures* of probability theory.

To understand how belief functions are updated in DST, we first need to introduce a second way of representing doxastic states, namely via mass functions:

Mass Function: A function $m : \wp(\Omega) \rightarrow [0, 1]$ is a *mass function* just in case:

- (M1) $m(\emptyset) = 0$,
- (M2) $\sum_{A \subseteq \Omega} m(A) = 1$.

⁷For technical reasons (fn. 9) we'll restrict ourselves to finite Ω . The infinite case is discussed in Appendix A.II.

⁸ $\wp(\Omega)$ is the powerset of Ω .

Mass functions are an alternative way of representing doxastic states in DST because every belief function corresponds to a unique mass function and vice versa:⁹

Proposition. If m is a mass function and we define:

$$bel(A) := \sum_{B:B \subseteq A} m(B),$$

then bel is a belief function. And if bel is a belief function, there is a unique mass function m such that:

$$bel(A) = \sum_{B:B \subseteq A} m(B).$$

A proposition's mass $m(A)$ should not be confused with its degree of belief, which is given instead by $bel(A)$. What $m(A)$ represents is more like the amount of evidence pointing specifically to A . To illustrate, suppose a somewhat reliable source at NASA predicts that a meteor will strike Canada tomorrow, while another, slightly less reliable source believes it will strike the United States. Then I might assign mass 0.3 to its landing in Canada and mass 0.2 to its landing in the U.S., with the remaining 0.5 mass "unused", i.e. assigned to Ω . My degree of belief that the meteor will strike somewhere in North America is then the sum of the masses for all logically stronger propositions, in this case $0.5 = 0.3 + 0.2$. And this is so despite the fact that 0 mass is assigned to North America. Neither piece of evidence points directly to the proposition that the asteroid will land in North America. Each speaks instead to a more specific possibility: Canada or the U.S.

Bayesians typically represent evidence with a proposition, a probability distribution over a partition, or more generally as an expected value. While DST can also represent evidence as a single proposition, or even as a belief function over a partition, its most general representation allows evidence to be *any* kind of belief function over $\wp(\Omega)$. In general, evidence is updated on by combining the belief function representing your prior credal state with a belief function representing the evidence.

The rule for performing these combinations is defined using mass functions. We first define a combination operator, \oplus , as follows:

Dempster Combination: Let the non-zero points of the mass function m_1 be the B_i 's, and the non-zero points of mass function m_2 the C_j 's. The *Dempster*

⁹This fact does not hold when Ω is infinite because some belief functions will not have corresponding mass functions (Halpern, 2003, p. 36). For more on the infinite case, see Appendix A.II.

combination of m_1 and m_2 , written $m_1 \oplus m_2$, is defined:

$$\begin{aligned} (m_1 \oplus m_2)(\emptyset) &= 0, \text{ and} \\ (m_1 \oplus m_2)(A) &= \frac{1}{c} \sum_{B_i \cap C_j = A} m_1(B_i)m_2(C_j) \text{ for } A \neq \emptyset, \end{aligned}$$

where c is the normalization constant:

$$c = 1 - \sum_{B_i \cap C_j = \emptyset} m_1(B_i)m_2(C_j).$$

The Dempster combination of the belief functions corresponding to m_1 and m_2 , call them bel_1 and bel_2 , is similarly written $bel_1 \oplus bel_2$.

One can prove that the combination of two mass functions is always a mass function, and thus the combination of two belief functions is always a belief function. The dynamics of DST can thus be given by:

Dempster's Rule: If your credences are represented by the belief function bel and your new evidence by the belief function bel_B , then your new credences should be represented by the belief function $bel' = bel \oplus bel_B$.

Just as Bayesians must be careful to separate the definition of conditional probability from the dynamic rule of Conditionalization, in DST we must separate the definition of \oplus from this dynamic rule for updating.

Like in probability theory, we can look to the special case where a single proposition is learned with certainty for a notion of conditional belief, and write $bel(A|B)$. There is even an analogue of the ratio rule from probability theory.¹⁰ However, this notion of conditional belief plays a very minor role in DST, unlike the notion of conditional probability in probability theory. For example, there is no analogue of Bayes' rule, relating $bel(A|B)$ and $bel(B|A)$ in a simple, helpful way. There is also no analogue of the rule of total probability, relating $bel(A)$ to $bel(A|B)$ and $bel(A|\bar{B})$. Indeed, this latter fact will prove a minor stumbling block for our discussion of rigidity, though one we can work around (see the discussion of conglomerability in the next subsection).

2.2. The Problem for Dempster-Shafer Theory. Now that we have DST on the table, we can ask whether it correctly handles perceptual undermining. My argument that it does not parallels the one for Bayesianism: when updating on E by Dempster's Rule, independence is preserved between E and its underminer, F . Thus, having

¹⁰In probability theory, $p(A|B) = p(A \cap B)/p(B)$. In DST, $bel(A|B) = [bel(A \cup \bar{B}) - bel(\bar{B})]/[1 - bel(\bar{B})]$.

updated on E in response to the appearance of the red sock, subsequently learning F will not lower one's degree of belief in E .

The argument requires three things. First we need to formalize “updating on E ” in DST. When I glimpse the red-looking sock, what belief function do I combine with my prior belief function? Second, we need to formalize the notion of independence in DST. When does a belief function treat two propositions as independent? And third, we need theorems. We need to show that Dempster's Rule is Rigid, and that Rigidity is Independence Preserving. Let's tackle each of these tasks in turn.

How should we understand “updating on E ” in the context of DST? For Bayesians the answer was straightforward: Jeffrey Conditionalize using the partition $\{E, \bar{E}\}$, with a high value for E and a low value for \bar{E} . But DST is more liberal. It allows that one can have evidence for E without having evidence against or about \bar{E} . So there are a number of candidates for representing the evidence when we “update on E ”. We could use a mass function that:

- assigns all mass to E ,
- assigns some mass to E , the rest to \bar{E} ,
- assigns some mass to E , the rest to Ω , or
- assigns some mass to E , some to \bar{E} , and the rest to Ω .

Fortunately, it won't matter which representation we choose. In all four cases, the corresponding belief function is said to be “focused on” the partition $\{E, \bar{E}\}$:

Focus: A belief function bel is *focused on* $\{E, \bar{E}\}$ just in case $m(A) = 0$ whenever $A \notin \{E, \bar{E}, \Omega\}$, where m is the mass function corresponding to bel .

Since our results will apply whenever the evidence is focused on $\{E, \bar{E}\}$, we needn't worry about which representation is correct.

How should we understand independence in the context of DST? Importing the usual definitions from probability theory would be unwise. For example, we might say that A and B are independent just in case $bel(A \cap B) = bel(A)bel(B)$. But this definition turns out to be very weak: it would allow propositions to be independent when information about one still affects one's beliefs about the other, and it also allows that A and B are independent even though \bar{A} and B are not. (Ben Yaghlane et al. (2002) go so far as to call this definition “useless”.)

Better definitions are available though. The salient definition for our purposes¹¹ comes from Shafer (1976):

Independence (DST): A and B are *independent* under bel if and only if every belief function bel_B focused on $\{B, \bar{B}\}$ is such that:

$$\begin{aligned} (bel \oplus bel_B)(A) &= bel(A), \\ (bel \oplus bel_B)(\bar{A}) &= bel(\bar{A}). \end{aligned}$$

Intuitively, A and B are independent just in case no information about $\{B, \bar{B}\}$ will affect one's beliefs about $\{A, \bar{A}\}$. As one might expect, independence is symmetric in DST: if information about $\{B, \bar{B}\}$ does not inform one's beliefs about $\{A, \bar{A}\}$, the reverse is also true.¹²

At this point we might appear ready to put the perceptual undermining problem to DST. At the outset of the sock case, E and F should be independent, so that $bel(F|E) = bel(F|\bar{E}) = bel(F)$. And we can show that updating on E —i.e. combining bel with a belief function focused on $\{E, \bar{E}\}$ —preserves these conditional beliefs, so that $bel'(F|E) = bel'(F|\bar{E})$.¹³ But it would be fallacious to then infer that E and F are independent under bel' . For even though $bel'(F|E) = bel'(F|\bar{E})$, it does not follow that $bel'(F|E) = bel'(F)$. That is, belief functions are not “conglomerable” the way probabilities are: sometimes one's conditional beliefs in A given B and given \bar{B} are the same, yet one's unconditional belief in A is something else.¹⁴

One might just stop there, concluding that DST is fatally flawed for not having such an elementary property. But we will press on, for two reasons. First, there may be cases that favour abandoning conglomerability. Cases like Derek Parfit's miner puzzle might be used to argue that conglomerability should not always hold (Kolodny & MacFarlane, 2010). Second, DST does have a closely related property: if

¹¹This definition of ‘independence’ is usually called “cognitive independence”. Shafer also defines a notion he calls “evidential independence.” Dempster's rule preserves evidential independence too, as I prove in an extended version of this paper available online. The proof is not included here because it is quite long, and because (despite its name) evidential independence is not the appropriate conception of independence for our purposes. Rather, evidential independence captures the idea that one's beliefs about $A \cap B$ can be “factored” into beliefs just about A and beliefs just about B . In probability theory the two kinds of independence coincide, since $p(B|A) = p(B)$ just in case $p(A \cap B) = p(A)p(B)$. Factorizability is a useful property for simplifying representation and computation. But only the first kind of independence, “cognitive independence”, is relevant to our epistemological discussion. (The labels “cognitive” and “evidential” can be very misleading here, especially to a philosophical audience. Unfortunately, this terminology has stuck in the literature.)

¹²This follows from Shafer's (1976) Theorem 7.9, which states that independence is equivalent to the relation $1 - bel(A_i \cap B_j) = (1 - bel(A_i))(1 - bel(B_j))$ holding for any $A_i \in \{A, \bar{A}\}$, $B_j \in \{B, \bar{B}\}$.

¹³See Lemma 2 in (Shafer, 1981).

¹⁴See Example 1 of Appendix A for a proof.

one's credence in A would be x no matter what evidence one got about $\{B, \bar{B}\}$, then one's credence in A is x right now.¹⁵ And this property may be what really draws us to conglomerability. Perhaps we find conglomerability plausible because we think that, when one's credence in A would be the same no matter what information one received about $\{B, \bar{B}\}$, one should have that credence in A now.¹⁶ In a Bayesian context, that's equivalent to having $p(A) = x$ when $p(A|B) = p(A|\bar{B}) = x$. But the dynamics work differently in DST. In DST, it can be that one would have credence x in A if one learned B with certainty, and likewise for \bar{B} , yet less-than-certain information about $\{B, \bar{B}\}$ would yield some other credence in A . Thus DST's proponents might argue that the failure of synchronic, probability-style conglomerability is appropriate given their diachronic commitments. Rather than try to settle the matter here, we will spot DST correctness on this point, and show that it nonetheless runs afoul of the perceptual undermining problem.

Because DST's dynamics differ in this way, we must broaden our conception of rigidity accordingly. One's credences in A given certainty in B and given certainty in \bar{B} no longer capture all of one's B -conditional attitudes, since they do not fix one's credences given less-than-certain information about B and \bar{B} . So let us conceive of rigidity as something stronger: preserving all B -conditional attitudes, even conditional on less-than-certain information about $\{B, \bar{B}\}$. We can then show:

Dempster Combination is Rigid: If bel is a belief function, bel_1, bel_2, bel_3 are belief functions focused on $\{B, \bar{B}\}$, and

$$\begin{aligned} (bel \oplus bel_1)(B) &= ((bel \oplus bel_2) \oplus bel_3)(B), \\ (bel \oplus bel_1)(\bar{B}) &= ((bel \oplus bel_2) \oplus bel_3)(\bar{B}), \end{aligned}$$

then $bel \oplus bel_1 = (bel \oplus bel_2) \oplus bel_3$.¹⁷

Intuitively, this says that updating on bel_2 doesn't change one's B -conditional attitudes, where B -conditional attitudes are individuated according to the posterior credences in B and \bar{B} . We can also show:

Rigidity is Independence Preserving (DST): Suppose the transition from bel to bel' is rigid on the partition $\{B, \bar{B}\}$. Then if A and B are independent under bel , they are independent under bel' too.¹⁸

¹⁵See the proof of "Total Conglomeration" in Appendix A.I

¹⁶This thought has been a recurring theme in the literature on the Reflection Principle (van Fraassen, 1984, 1995). See (Elga, 2007) and (Weisberg, 2007), for example.

¹⁷See Appendix A.II for a proof.

¹⁸See Appendix A.II for a proof.

Thus we have the RIP theorem for DST: Dempster’s Rule is Rigid, hence Independence Preserving.

Now we can put the perceptual undermining problem to DST. Before glimpsing the sock, my credence function bel should regard E and F as independent. Information about $\{E, \bar{E}\}$ should not affect my credences about $\{F, \bar{F}\}$, and vice versa. When I then glimpse the sock and “update on E ”, I combine bel with a belief function focused on $\{E, \bar{E}\}$ using Dempster’s Rule.¹⁹ But because Dempster’s Rule is Rigid, hence Independence Preserving, the resulting bel' will also regard E and F as independent. So subsequently learning F cannot lower my credence in E . Thus DST faces the same challenge Bayesianism does. Representing credences by belief functions and updating by Dempster’s rule is not a solution.

There are variants of DST we have not considered, and which may yet prove better at handling perceptual undermining. Smets’ (1990) Transferable Belief Model uses “conjunctive combination”, an unnormalized version of Dempster combination. Although we won’t prove it here, inspecting the conjunctive combination rule and our proofs in Appendix A.II makes it clear that the RIP theorem readily extends to conjunctive combination. But Dubois and Prade (1986) also offer a disjunctive combination rule, and Fagin & Halpern (1991) offer an alternative definition of $bel(B|A)$ inspired by viewing belief functions as lower envelopes on sets of probability distributions. And there are still more options. Whether any of these is better suited to handle perceptual undermining is an open question.

3. THE CHALLENGE FOR RANKING THEORY

Ranking theory (Spohn, 1988, 2012) departs from Bayesianism and from DST, both in the way it represents doxastic states and in the rules it uses to update them.²⁰ So we might hope it can come through where Bayesianism and DST have so far failed. We will see, however, that the RIP theorem extends to ranking theory too.

In ranking theory, doxastic states are represented by ranking functions:

Ranking Functon: A function $\kappa : \wp(\Omega) \rightarrow \mathbb{N} \cup \{\infty\}$ is a *ranking function* just in case:

$$(R1) \quad \kappa(\emptyset) = \infty,$$

¹⁹*Objection.* The right way to respond to the appearance of the sock is not to combine with an evidential belief function focused on $\{E, \bar{E}\}$, but to combine with a more complicated belief function that will introduce the necessary negative correlation between E and F . *Reply.* This move is parallel to the Bayesian’s reply that we must use a richer partition, and is subject to the same objections (see §5).

²⁰It also differs in the doxastic attitudes it is chiefly designed to represent, viz. full beliefs rather than degrees of belief. We will consider the significance of this difference once we have the RIP theorem on the table.

$$(R_2) \quad \kappa(\Omega) = 0,$$

$$(R_3) \quad \kappa(A \cup B) = \min[\kappa(A), \kappa(B)] \text{ for any } A \text{ and } B.$$

A ranking function represents degrees of *disbelief* on an integer scale. Propositions not disbelieved at all are ranked 0, and disbelieved propositions are then ranked with greater and greater degrees up through ∞ . Notice though, propositions ranked 0 are not necessarily believed: if A and \bar{A} are both ranked 0, the subject suspends judgment about A vs. \bar{A} . Rather, a proposition is believed just in case its negation is ranked greater than 0 (A and \bar{A} can't both be ranked above 0, on pain of violating (R₂) or (R₃)).

Ranking theory employs a notion of conditional rank similar to the notion of conditional probability:

Conditional Rank: If κ is a ranking function, the *conditional rank* of A given B , written $\kappa(A|B)$, is defined:

$$\kappa(A|B) = \kappa(A \cap B) - \kappa(B).$$

Using conditional ranks, we have two rules for updating that parallel Conditionalization and Jeffrey Conditionalization respectively:²¹

Rank Conditionalization: If your ranking function is κ and you then make an observation completely described by B , then for any A your new rank should be: $\kappa'(A) = \kappa(A|B)$.

Spohn Conditionalization: If your ranking function is κ and (i) experience directly affects your ranks on the partition $\{B_i\}$ changing them to the values $\kappa'(B_i)$, but (ii) experience does not directly affect any other ranks, then for any A your new rank should be:

$$\kappa'(A) = \min[\kappa(A|B_1) + r_1, \dots, \kappa(A|B_n) + r_n].$$

Rank Conditionalization is the special case of Spohn Conditionalization where the evidential partition contains B and \bar{B} and the new ranks are 0 and ∞ , respectively. So we will focus on Spohn Conditionalization, for which we have the following result:

²¹I've adapted the usual statements of these update rules to make them explicitly rules for updating in response to *experience*. Whether Spohn and other ranking theorists intend this interpretation, it is appropriate in the present context. Our question is whether the resources of ranking theory do better with perceptual undermining than those of Bayesianism or DST. See (Huber, 2013, §3.3) or (Huber, 2009, §4) for a more standard statement.

Spohn Conditionalization is Rigid: If κ' comes from κ by Spohn Conditionalization on the partition $\{B_i\}$, then $\kappa'(A|B_i) = \kappa(A|B_i)$ for all A, B_i .²²

So we face the question whether rigidity is independence preserving in ranking theory.

Independence for ranking functions is defined much like in probability theory (Spohn, 1999):

Rank Independence: If κ is a ranking function, then A is *independent* of B under κ just in case both of the following conditions hold:

$$\begin{aligned}\kappa(A|B) &= \kappa(A|\bar{B}), \\ \kappa(\bar{A}|B) &= \kappa(\bar{A}|\bar{B}).\end{aligned}$$

And parallel to probability theory, we have the result:

Rigidity is Independence Preserving (Ranking Theory): Suppose the transition from κ to κ' is rigid on the partition $\{B, \bar{B}\}$. Then, if A and B are independent under κ , they are independent under κ' too.²³

This completes our explication of the general RIP theorem. All three frameworks under consideration—Bayesianism, Dempster-Shafer theory, and ranking theory—use update rules that are Rigid, hence Independence Preserving.

Before we move on though, a few words about interpreting ranking functions are in order. While Bayesianism and DST are commonly interpreted as theories governing partial belief, the standard interpretation of ranking theory treats it as a theory governing full belief. On the standard interpretation, an agent fully believes A just in case she disbelieves its negation to some degree greater than 0: $\kappa(\neg A) > 0$. The theory does deal in degrees of *disbelief*, but whether/how these degrees of disbelief relate to Bayesian degrees of belief is not straightforward. Ranking theory's degrees of disbelief are operationalized in terms of the agent's conditional full beliefs, or in terms of the number of independent sources needed to warrant full belief. This contrasts sharply with Bayesianism's operationalization of degrees of belief in terms of preferences. Since Bayesianism does not include a standardized theory of full belief, and ranking theory does not include a decision theory, it's not clear how to relate degrees of belief to degrees of disbelief. Thus ranking theory may have a different subject matter than Bayesianism and DST, and may not be a competitor or alternative to those theories.

²²See Appendix B for a proof, though the result is already implicit in (Spohn, 1988, §5).

²³See Appendix B for a proof, though this result is also implicit in (Spohn, 1988, §6).

One can still interpret ranking theory's formalism as a theory governing degrees of belief (just upside down). If one does, then the results of this section show that the resulting theory of partial belief faces the same perceptual undermining problem as Bayesianism. But if one sticks to the standard interpretation, the moral of these results is importantly different. They then show that perceptual undermining isn't just a problem for degrees of belief, but for the theory of full belief too. An agent who fully believes the sock is red based on its appearance cannot have that belief undermined by the subsequent discovery that the lighting is deceptive. This makes the problem of perceptual undermining that much more pressing.

It's natural to wonder whether perceptual undermining is also a problem for other theories of full belief, especially the well-known AGM theory of belief revision. The prospects are not good.²⁴

AGM represents an agent's beliefs with a set of sentences, on which it imposes various synchronic and diachronic constraints. For example, the set should be logically closed, and adding a sentence should have the same effect as adding any logically equivalent sentence. (Alchourrón et al., 1985) However, the standard set of constraints massively underdetermines how an agent should update her beliefs. So a second layer of representation is often added, an "entrenchment ranking", which describes how easy various beliefs are for the agent to give up. Some natural assumptions then yield a unique update rule, mapping each entrenchment ranking and input sentence to a new set of beliefs. (Gärdenfors & Makinson, 1988; Gärdenfors, 1988; Grove, 1988)

Infamously though, this framework only handles one revision. When the agent learns a second proposition that requires her to revise her beliefs again, the way forward is once again massively underdetermined. The reason is that the output of the first update is just a set of beliefs, sans entrenchment ranking. Without a new entrenchment ranking to work with, AGM's weak diachronic constraints are all we have to fall back on. (Boutilier, 1993; Darwiche & Pearl, 1994)

This limitation is crucial for us because the problem of perceptual undermining involves two revisions, a response to perception followed by a response to an underminer. Because AGM and entrenchment rankings are effectively silent on the second update, they are silent about the problem of perceptual undermining.

Notably though, ranking theory can be viewed as an implementation of AGM, one that solves the problem of iterated revisions. Ranking Theory's states of full belief, and its rules for updating them, obey the standard AGM constraints. Indeed, we can view ranking functions as refined entrenchment rankings, providing cardinal rather than merely ordinal information about degrees of entrenchment. (Spohn,

²⁴I am indebted to Franz Huber for his generous help with the following discussion.

1999; Huber, 2009, 2013) So, because ranking theory is an implementation of AGM and faces the problem of perceptual undermining, the problem is a threat to AGM.

Must any reasonable implementation of AGM run afoul of perceptual undermining? That we can't say. For all I've said here, it remains possible that some other way of strengthening AGM yields a theory different enough from ranking theory to avoid the problem. We can note, however, that one other prominent strengthening of AGM, Darwiche & Pearl's (1994), is not enough to prevent the problem. Darwiche & Pearl propose four constraints to be added to the standard AGM ones. But they also show that ranking theory respects their postulates just as it respects the weaker AGM postulates. Whether some other strengthening or implementation of AGM might do the job we must leave as an open question.

4. THE APPEAL TO METACOGNITION

Let's return to the challenge for Bayesianism and see what solutions might be available. In my presentation of the problem, I assumed that the apparent redness of the sock should increase my credence in E (*The sock is red*), and the subsequent discovery that F (*The lighting is deceptive*) should decrease my credence in E . But, one might object, the discovery that the lighting is deceptive should not decrease my credence that the sock is red unless I think my credence in the sock's redness is based on its appearance. After all, if I think my credence in the sock's redness is based on testimony, finding out that the lighting is deceptive should have no effect on my credence that the sock is red. Only if I think my credence in the sock's redness is appearance-based should facts about misleading appearances affect my credence.

This observation, continues the objection, exposes the point where a negative correlation between F and E gets introduced into p' . After the first update in response to the sock appearance, there is an update in response to that update, on the proposition $E' = \text{My credence at } t' \text{ that the sock is red is based on its having appeared red at } t$. That is, I make a metacognitive observation that prompts an update in between the updates on E and on F . And that intermediate update is where the negative correlation between E and F is introduced. Thus, when I finally do learn F , it will reduce my credence in E . Unless, of course, I do not make the metacognitive observation that my credence in E is appearance-based. Then there is no intermediate update and learning F does not affect my credence in E , as is appropriate.²⁵

This solution has prima facie appeal because it acknowledges what seems plausible: that perceptual beliefs are often accompanied by metacognitive information about

²⁵This proposal is based on a suggestion of Scott Sturgeon's from his talk, "Undercutting Defeat & Edgington's Burglar," given at the *Conditionals and Paradox: Celebrating the Work of Dorothy Edgington* conference.

their sources, and that the rational response to subsequent underminers should depend on what meta-data we have about a belief's sources. Nevertheless, it fails for a fairly simple reason. The metacognitive observation that my credence in E is appearance-based cannot introduce the desired negative correlation. To see this, notice that at the outset, before I look at the sock, the conjunction *The lighting is deceptive* \wedge *My credence at t' that the sock is red will be based on its appearing red at t* is probabilistically independent of the proposition *The sock is red*. The fact that in a moment my credence in E will be high based on misleading appearances has no bearing on what the actual colour of the sock is. In other words, at the outset E is independent of $F \wedge E'$, i.e. $p(E|F \wedge E') = p(E)$. Because Rigidity will Preserve this Independence, learning E' in addition to F cannot lower the probability of E . Thus the intermediate, metacognitive update does not provide an opportunity to introduce the negative correlation between E and F that we are after.

There is a second way of pursuing this strategy. Instead of accommodating the metacognitive observation by conditionalizing on E' , I might apply an update rule like Information Minimization (Shore & Johnson, 1980; Williams, 1980), feeding it the constraint that $p'(E|F)$ should be low. On this approach, the transition from p to p' happens in two steps: $p \rightarrow p^* \rightarrow p'$, with the first $p \rightarrow p^*$ transition using Jeffrey Conditionalization on $\{E, \bar{E}\}$, and the second $p^* \rightarrow p'$ transition happening by Information Minimization with the constraint that $p'(E|F)$ be low.

This approach gives away the game though. One way to see this is to note that the first $p \rightarrow p^*$ step is now superfluous. The entire transition from p to p' can be achieved in a single step by applying Information Minimization once with the constraints that $p'(E)$ shall be high and $p'(E|F)$ low. Another way of seeing the point: we are effectively conceding that Information Minimization must be used in all (or nearly all) cases that were traditionally handled using Jeffrey Conditionalization. For, as we saw in §1.3, perceptual underminers are ubiquitous. Any judgment we make in response to experience could, potentially, be undermined by the later discovery that our faculties are unreliable when it comes to observations of the kind in question. If we solve the problem this way, then all (or nearly all) cases of observation must be handled by Information Minimization.

Fans of Information Minimization might be okay with this result, but there are reasons not to be sanguine. One reason is that we must now always feed constraints like “make $p'(E)$ high but keep $p'(E|F)$ low” into our update rule. The more such constraints we put in by hand, the less epistemological work the rule does for us, and the more we should wonder about the depth of our approach. And as we'll see in the next section, every update will require many such constraints, since there will

generally be many possible underminers. We'll see some additional reasons to be wary of replacing Jeffrey Conditionalization with Information Minimization in §7.

5. THE APPEAL TO RICHER INPUTS

Let's look elsewhere for an opportunity to introduce a negative correlation between E and F . What about the first update, the one that happens in response to the sock's red appearance? If we apply Jeffrey Conditionalization to the partition $\{E, \bar{E}\}$, the fact that Rigidity is Independence Preserving dashes any hopes of introducing the desired negative correlation. But what if we use a different partition, one that already captures the defeat relations between underminer and underminee? For example, instead of applying Jeffrey Conditionalization to a distribution like:

E	\bar{E}
.99	.01

we could instead use:

	E	\bar{E}
F	.001	.009
\bar{F}	.989	.001

Then we get the desired results: $p'(E) = .99$ but $p'(E|F) = 0.1$, so F will function as an underminer after the first update.²⁶

My main concern about this proposal is that it gets the right results but in the wrong way. To see why, we must first consider a lesser problem, namely that it shifts a tremendous amount of work to an as-yet unarticulated part of Bayesian theory.

It has long been appreciated that Jeffrey Conditionalization is incomplete as a rule for updating probabilities in response to experience, since it does not specify what distribution one should update on in response to a given experience (Field, 1978;

²⁶This proposal has been suggested by several philosophers in personal communication, and is discussed by Pryor (forthcoming, §10). Relatedly, Wagner (2013) endorses Jeffrey Conditionalizing on $\neg F \supset E$. But as Gallow (forthcoming) points out, this lowers $p(F)$. Surely a glimpse of a red-looking sock shouldn't affect your credence that the lighting is deceptive.

Christensen, 1992).²⁷ Jeffrey Conditionalization presupposes that the distribution of probabilities over some partition is “directly affected” (Jeffrey, 1968, §§2–3), but it does not say when these direct effects should²⁸ happen in response to experience, or what they should be like. In order for our theory of updating to be complete, we would need a rule supplementing Jeffrey Conditionalization, a rule that specifies what distribution to plug into Jeffrey Conditionalization in each case.

Despite being incomplete in this way, Jeffrey Conditionalization is interesting because, for many problems of interest, we can make plausible, innocuous assumptions about what input distribution to use. That is, we can apply the rule in obvious ways that allow us to proceed to address orthogonal problems of interest. The current context is not like that, however. The solution on offer trades on precisely this incompleteness of Jeffrey Conditionalization, suggesting that we fill out the gap in Bayesian theory one way rather than another in order to solve a problem *for* Jeffrey Conditionalization. And if we pursue this solution seriously, we will find that filling out the gap is a surprisingly complex and non-obvious task. To see what a complex and difficult job it is, consider in detail what we would need to do in the sock example to get all the desired results.

First, notice that the input partition must be much finer than the simple, 4-cell partition suggested above, $\{EF, E\bar{F}, \bar{E}F, \bar{E}\bar{F}\}$. Many different propositions besides F can undermine the support the appearance of the sock lends to E . The fact that the lights are deceptive is one underminer, but so are propositions about the health of my eyes, the functioning of my visual system, and even the quality of the air in the room. In short, there is a long list of potential underminers for E , and we will need an input partition fine enough to capture them all. Second, notice that the distribution over this rich partition must be quite complex, since different propositions will undermine E to different degrees. For example, discovering that the lighting renders colour vision entirely unreliable would seem to defeat E completely, causing me to return my credence that the sock is red to its original value. But what if, instead, I learn that

²⁷Something very similar is true of Conditionalization of course. But the problem seems to have been felt more severely by those interested in Jeffrey Conditionalization, presumably because they are concerned to reject an epistemology based on protocol-sentences, thereby ruling out the most obvious solution to the parallel problem for Conditionalization.

²⁸Some authors would object to this “should” talk, on the grounds that there is no fact of the matter how one’s credences over the partition *ought* to change in response to experience, only a fact about how they do change. On this view, normative questions only arise after an experience has had its direct effects on one’s credences. Jeffrey himself takes this view explicitly in his (1968, §3), though Carnap clearly disagreed (Jeffrey, 1975, p. 44), and (Christensen, 1992) follows Carnap, joining many contemporary epistemologists in rejecting that view. I will follow Carnap and Christensen, as Jeffrey’s view here denies our starting assumption that experience provides *prima facie* evidential support. My arguments are addressed to those who think experience plays a justificatory, hence normative, role.

the lighting is deceptive but only certain times of day? Then the defeat shouldn't be as extreme. And if I learn that the air in the room is unusual in ways that can cause optical illusions, I might want to reduce my credence in E by a bit, but perhaps not by very much. Third, notice that the distribution chosen must be sensitive to the agent's background beliefs. Which propositions are underminers for which others is largely sensitive to background beliefs. On one theory of optics, the quality of the air might be irrelevant to the trustworthiness of visual perception, while on another it might be highly relevant.

So we might be able to handle the problem that perceptual undermining poses for Jeffrey Conditionalization by shifting the work to the as-yet unspecified, supplementary input rule. But we would be shifting a tremendous amount of complex, precise, and subtle work to that rule.²⁹ That is the lesser problem.

The greater problem is that we would not just be shifting a lot of work, we would also be shifting the *wrong kind* of work. An update rule is supposed to determine our new credences as a function of our old beliefs and the new evidence. But on the current proposal, "the new evidence" is not really the new evidence. The complex distribution we would be plugging into Jeffrey Conditionalization would be produced by considering how an experience as of a red-looking sock and our background beliefs about optics combine to warrant new beliefs about the quality of the air and the colour of the sock. And this is precisely the kind of work our update rule was supposed to do.³⁰ The current proposal tries to save our update rule, Jeffrey Conditionalization, by turning its supplementary input rule into an update rule too.

This is bad for a number of reasons. First, we drain substance from Jeffrey Conditionalization by taking away work it was supposed to do for us and moving it to another part of our theory. Second, we move the work to a part of the theory we don't actually have yet. Finally, in shifting the work this way we undermine the paradigm underlying rules like Conditionalization and Jeffrey Conditionalization. Updating is no longer a matter of aligning new unconditional probabilities with old conditional probabilities on the evidence. Rather, it is a matter of doing whatever the as-yet

²⁹Cf. Pryor (forthcoming), who worries that we would be moving epistemologically important work "offstage".

³⁰Worse yet, those new beliefs might not be determined by any formal rule. Conditionalization and Jeffrey Conditionalization use purely logical and probabilistic features of our priors and "the new evidence" to determine new credences. But there might be no logico-probabilistic relationship between a perceptual state, our prior credences, and the probabilities one should input to Jeffrey Conditionalization.

unspecified input rule does to determine new credences based on prior credences and sensory experience.³¹

6. THE APPEAL TO A GENERIC UNDERMINER

Maybe instead of enriching the input distribution in such drastic ways, we can get what we need with just one enrichment. Let G be the “generic” underminer, *My colour experience is not a reliable guide to the sock’s colour*, and let’s use as our input distribution:

	E	\bar{E}
G	.001	.009
\bar{G}	.989	.001

Then, when I later learn some fact that supports G , such as that the lighting is deceptive, the probability of E will be reduced to a low value (0.1 in this toy example). In general, learning any particular underminer—about the quality of the lighting, the health of my vision, the makeup of the air in the room—will support G and consequently undermine E . So we get the desired undermining without having to craft a fine-grained and carefully tuned input distribution. And, unlike the distribution proposed in the previous section, this one might plausibly be taken to represent the new evidence alone (as opposed to representing the joint effects of the new evidence and our background beliefs).³²

My objection here is that no proposition can legitimately play the role of the generic defeater, G . G must be a proposition that is made probable by all underminers, and which reduces the probability of E back to its original value, $p(E)$. More specifically, because underminers come in various strengths, G must be made probable to different degrees by different underminers. What proposition could have these features? Three proposals come to mind.

First we might try $G_1 = \textit{My colour experience is not a reliable indicator of actual colours}$. But this proposition is too general. The fact that my colour experience is not reliable in general does not entail that it is probabilistically irrelevant in this

³¹This objection is very much in the spirit of Christensen’s (1992) original concerns about (Jeffrey) Conditionalization, perhaps bringing the discussion full circle.

³²Thanks to Ralph Wedgwood for his helpful correspondence here. This proposal is based on his post at http://el-prod.baylor.edu/certain_doubts/?p=843.

one instance. And some underminers will not have such broad scope. For example, the proposition that the air in the room induces optical illusions is at least a partial underminer of E , but it does not mean that my colour vision is unreliable in general. Broadly speaking, many propositions will undermine E without sufficiently probabilifying G , and G itself is not strong enough to act as a perfect underminer of E . Thus G cannot stand as a proxy between E and its various potential underminers.

Second we might try $G_2 = \textit{The objective chance of } E \textit{ is low}$. G_2 has the advantage over G_1 that, assuming compliance with the Principal Principle (Lewis, 1980), it will reduce the probability of E to the desired level. And it's plausible that many underminers will support G_2 to the extent that they are underminers. Still, it's contentious whether all underminers will support G_2 . For example, inadmissible propositions might be capable of undermining E but would not do so by influencing my credences about the chance of E . In general, how underminers will interact with my beliefs about the chances will depend on one's views about chances and the Principal Principle, and such matters are contentious to say the least (see (Meacham, 2010) for some recent discussion of the relevant issues). And this worry points to the really fundamental concern about G_2 : it is too theoretically loaded. There is disagreement about whether there really are such things as chances, but even those who think there are chances must concede that some others do not. These others will, presumably, be sensitive to underminers in the same way as those who believe in chances. But for them, G_2 cannot play the role of G , so the true story about perceptual undermining cannot be that it works by updating on $\{EG_2, \overline{EG_2}, \overline{EG_2}, \overline{EG_2}\}$.

Third and finally we might try $G_3 = \textit{My colour experience is not a reliable indicator in this one instance}$. G_3 has all the virtues of G_2 but without the complications that come with appealing to the controversial notion of objective chance. Still G_3 is dangerously trivializing. How do we understand the notion that something is “an unreliable indicator in just one instance”, except maybe as meaning either “the objective chance is low” or “should not be trusted”? The first option takes us back to G_2 and its troubles, while the second brings us back to the worry that Jeffrey Conditionalization is being sapped of any substance. It is well-known that any transition from one probability function to another can be viewed as the result of Conditionalization, if we just enrich the domain of the probability function appropriately (see (van Fraassen, 1989, p. 322), for example). And interpreting G_3 as saying something like “I should reduce my credence in E ” looks like it amounts to doing just that. We enrich the domain of the credence function with proxy propositions that do nothing more than allow us to model the desired transitions as instances of Conditionalization. These propositions have no content beyond saying that we should have a certain credence in a certain

proposition. Conditionalizing on them may give us the desired credences, but then rules like Jeffrey Conditionalization become trivial, saying merely that you should adopt the credences you think you should adopt.

7. CONCLUSION

I have argued for two claims. First, the clash between perceptual undermining and rigid updating is not confined to Bayesianism. The same problem arises in Dempster-Shafer theory and in ranking theory. Second, all the proposed ways of avoiding the clash in the context of Bayesianism were inadequate. I conclude that perceptual undermining presents formal epistemology with a broader and more resilient challenge than one might have thought.

How should we respond to this challenge? Maybe we just have yet to see how to correctly apply rules like Jeffrey Conditionalization. Or, maybe, we need different update rules to handle perceptual undermining.

If we take the second route, more abstract rules like Information Minimization (Shore & Johnson, 1980; Williams, 1980) are a natural place to look. But even with such proposals to fall back on, the significance of perceptual undermining should not be underestimated. First, appealing to a rule like Information Minimization raises worries similar to those that emerged for the appeal to richer inputs in §5. Relying on Information Minimization to solve our problem means feeding it constraints like “ $p'(E)$ should be high but $p'(E|F)$ low”. The more such constraints we put in by hand, the less epistemological work the rule does for us, and the more we should question the depth of our theory. The supposed centrality of prior conditional credences to updating is also called into question.

Second, rules like Information Minimization are often viewed as generalizations of (Jeffrey) Conditionalization, generalizations that we only need for uncommon or *recherché* cases like the Judy Benjamin problem (van Fraassen, 1981). But given the ubiquity of perceptual undermining noted in §1.3, we would need to change our view here. Even simple cases of observation would now have to be handled by a highly abstract rule like Information Minimization.

Third and relatedly, the controversies around these more abstract rules become more pressing. Information Minimization has been criticized for delivering counter-intuitive results in the Judy Benjamin problem. It is also defended on the grounds that it makes the “minimal revision” to the agent’s beliefs that respect the imposed constraints. But the measure of distance between belief-states used in this appeal is not symmetric, hence not a proper metric (Howson & Franklin, 1994). There are other ways of measuring distance, and competing update rules corresponding

to them (see (Douven & Romeijn, 2011), for example). The ubiquity of perceptual underminers makes the need to settle this competition more pressing. It also raises questions about whether having Jeffrey Conditionalization as a special case provides evidence for the correctness of these rules, as is sometimes suggested.

Gallow (forthcoming) offers a different sort of Bayesian proposal. He proposes an update-rule specifically designed to handle perceptual underminers, while still preserving the idea that prior conditional credences should be our guide to new credences. Indeed, Gallow shows that his rule has Conditionalization as a special case when there is no chance of undermining.

Whichever approach Bayesians prefer, Dempster-Shafer theory and ranking theory may need to follow suit. In the case of Dempster-Shafer theory, a successful alternative may even be on the books already. As mentioned at the end of §2.2, existing alternatives to Dempster's rule, like those of Dubois & Prade (1986) and Fagin & Halpern (1991), have yet to be explored as potential answers to the challenge posed by perceptual undermining.

APPENDIX A: RESULTS FOR DEMPSTER-SHAFER THEORY

Here we prove the formal results used in §2.2. The central results are that Dempster's Rule is Rigid, and that Rigidity is Independence Preserving in DST. We will also see a short, direct proof that Dempster's Rule Preserves Independence. The first route to independence-preservation, via rigidity, verifies that DST is of a piece with Bayesianism and ranking theory: all three are rigid, hence independence-preserving. The second, shorter route illuminates the case where Ω is infinite, which we bracketed back in fn. 9. But first, we verify minor claims promised in fns. 14 and 15.

Appendix A.I Minor Results. As promised in footnote 14:

Example 1: Non-conglomerability in DST: Let $\Omega = \{AB, A\bar{B}, \bar{A}B, \bar{A}\bar{B}\}$ and let *bel* be given by the mass function:

$$m(S) = \begin{cases} 1/10 & \text{if } S = B \text{ or } S = \bar{B}, \\ 2/10 & \text{if } S = (A_i \cup \bar{B}_j) \text{ for some } A_i \in \{A, \bar{A}\}, B_j \in \{B, \bar{B}\}, \\ 0 & \text{otherwise.} \end{cases}$$

Then for any $A_i \in \{A, \bar{A}\}$ and $B_j \in \{B, \bar{B}\}$:

$$\begin{aligned} \text{bel}(A_i|B_j) &= 2/9, \\ \text{bel}(A_i) &= 0. \end{aligned}$$

And as promised in footnote 15:

Total Conglomeration: If $(bel \oplus bel_B)(A) = x$ for every belief function bel_B focused on $\{B, \bar{B}\}$, then $bel(A) = x$.

Proof. Suppose $(bel \oplus bel_B)(A) = x$ for every belief function bel_B focused on $\{B, \bar{B}\}$. By the definition of \oplus , we have for any m_B and $S \subseteq A$:

$$(m \oplus m_B)(S) = \frac{\sum_{X: X \cap B = S} m(X)m_B(B)}{c} + \frac{\sum_{Y: Y \cap \bar{B} = S} m(Y)m_B(\bar{B})}{c} + \frac{m(S)m_B(\Omega)}{c}.$$

Now consider what happens as $m_B(B) + m_B(\bar{B})$ approaches zero: c approaches 1, so each of the left two summands approaches 0 while the right summand approaches $m(S)$. Now recall that:

$$(bel \oplus bel_B)(A) = \sum_{S: S \subseteq A} (m \oplus m_B)(S).$$

We just saw that, as $m_B(B) + m_B(\bar{B})$ approaches 0, each term in this sum approaches $m(S)$, so the sum approaches $bel(A)$. At the same time, by hypothesis the sum remains constant at x , so $bel(A) = x$ too. \square

Appendix A.II Rigidity and Independence. Here we show that Dempster's Rule is Rigid, and that Rigidity is Independence Preserving in DST.

Dempster Combination is Rigid: If bel is a belief function, bel_1, bel_2, bel_3 are belief functions focused on $\{B, \bar{B}\}$, and

$$\begin{aligned} (bel \oplus bel_1)(B) &= ((bel \oplus bel_2) \oplus bel_3)(B), \\ (bel \oplus bel_1)(\bar{B}) &= ((bel \oplus bel_2) \oplus bel_3)(\bar{B}), \end{aligned}$$

then $bel \oplus bel_1 = (bel \oplus bel_2) \oplus bel_3$.

Proof. Suppose bel, bel_1, bel_2, bel_3 are belief functions satisfying the theorem's hypotheses. Because \oplus is associative (Shafer, 1976, Theorem 3.3),

$$\begin{aligned} (bel \oplus bel_1)(B) &= (bel \oplus (bel_2 \oplus bel_3))(B), \\ (bel \oplus bel_1)(\bar{B}) &= (bel \oplus (bel_2 \oplus bel_3))(\bar{B}). \end{aligned}$$

Moreover, because bel_2 and bel_3 are both focused on $\{B, \bar{B}\}$, their combination is too. Let bel_4 be that combination, i.e. $bel_4 := bel_2 \oplus bel_3$. We then have:

$$\begin{aligned} (bel \oplus bel_1)(B) &= (bel \oplus bel_4)(B), \\ (bel \oplus bel_1)(\bar{B}) &= (bel \oplus bel_4)(\bar{B}), \end{aligned}$$

with bel_1 and bel_4 both focused on $\{B, \bar{B}\}$.

We will show that, except in a special case, $bel_4 = bel_1$, and thus $bel \oplus bel_1 = bel \oplus bel_4 = bel \oplus (bel_2 \oplus bel_3) = (bel \oplus bel_2) \oplus bel_3$, as desired. The special case is when bel 's mass function is only positive on subsets of B and \bar{B} . In this case we can show directly that $bel \oplus bel_1 = bel \oplus bel_4$. We begin there.

Suppose bel 's mass function, m , is positive only on subsets of B and \bar{B} (not necessarily strict). Let the B_i be the subsets of B , and the \bar{B}_j those of \bar{B} . By the definition of \oplus :

$$\begin{aligned} (bel \oplus bel_1)(B) &= \frac{\sum_i m(B_i)m_1(B) + \sum_i m(B_i)m_1(\Omega)}{c_1} \\ &= \sum_i m(B_i) \frac{m_1(B) + m_1(\Omega)}{c_1}, \end{aligned}$$

where c_1 is our normalization constant. Similarly:

$$(bel \oplus bel_4)(B) = \sum_i m(B_i) \frac{m_4(B) + m_4(\Omega)}{c_4},$$

Since by hypothesis $(bel \oplus bel_4)(B) = (bel \oplus bel_1)(B)$,

$$\sum_i m(B_i) \frac{m_1(B) + m_1(\Omega)}{c_1} = \sum_i m(B_i) \frac{m_4(B) + m_4(\Omega)}{c_4}.$$

So:

$$\frac{m_1(B) + m_1(\Omega)}{c_1} = \frac{m_4(B) + m_4(\Omega)}{c_4}.$$

Thus $(m \oplus m_1)(B_i) = (m \oplus m_4)(B_i)$ for each $B_i \subseteq B$:

$$\begin{aligned} (m \oplus m_1)(B_i) &= \frac{m_1(B)m(B_i) + m_1(\Omega)m(B_i)}{c_1} \\ &= m(B_i) \frac{m_1(B) + m_1(\Omega)}{c_1} \\ &= m(B_i) \frac{m_4(B) + m_4(\Omega)}{c_4} \\ &= (m \oplus m_4)(B_i). \end{aligned}$$

Parallel reasoning involving \bar{B} and the \bar{B}_j 's shows that the same holds there: $(m \oplus m_1)(\bar{B}_j) = (m \oplus m_4)(\bar{B}_j)$ for every $\bar{B}_j \subseteq \bar{B}$. Thus $m \oplus m_1 = m \oplus m_4$, hence $bel \oplus bel_1 = bel \oplus bel_4$ as desired.

For the remainder of the proof then we bracket this special case. Let the N_k be the sets that are neither subsets of B nor of \bar{B} (' N ' for 'Neither'). We now show that $bel_1 = bel_4$ under the assumption that $m(N_k) > 0$ for some k .

By the definition of \oplus :

$$(bel \oplus bel_4)(B) = \frac{1}{c_4} \left[m_4(B) \sum_i m(B_i) + m_4(\Omega) \sum_i m(B_i) + m_4(B) \sum_k m(N_k) \right].$$

Adopting the following shorthand will prove useful:

$$a := \sum_i m(B_i), \quad b := \sum_k m(N_k), \quad c := \sum_j m(\bar{B}_j),$$

allowing us to tidy things up considerably:

$$(bel \oplus bel_4)(B) = \frac{m_4(B)(a + b) + m_4(\Omega)a}{c_4}.$$

We can then afford to expand c_4 :

$$(bel \oplus bel_4)(B) = \frac{m_4(B)(a + b) + m_4(\Omega)a}{1 - m_4(B) \sum_j m(\bar{B}_j) - m_4(\bar{B}) \sum_i m(B_i)},$$

which our shorthand renders:

$$(bel \oplus bel_4)(B) = \frac{m_4(B)(a + b) + m_4(\Omega)a}{1 - m_4(B)c - m_4(\bar{B})a}.$$

Parallel reasoning about \bar{B} yields a similar equation:

$$(bel \oplus bel_4)(\bar{B}) = \frac{m_4(\bar{B})(c + b) + m_4(\Omega)c}{1 - m_4(B)c - m_4(\bar{B})a}.$$

Now by hypothesis, $(bel \oplus bel_4)(B) = (bel \oplus bel_1)(B)$, and likewise for \bar{B} . Let's introduce constants for these values:

$$k := (bel \oplus bel_1)(B), \quad \bar{k} := (bel \oplus bel_1)(\bar{B}).$$

And we'll use the variables x , y , and z for m_4 's assignments:

$$x := m_4(B), \quad y := m_4(\bar{B}), \quad z := m_4(\Omega).$$

We then have the following two constraints on m_4 :

$$k = \frac{x(a + b) + za}{1 - xc - ya}, \quad (8)$$

$$\bar{k} = \frac{y(c + b) + zc}{1 - xc - ya}. \quad (9)$$

A third constraint falls out of the mass $m \oplus m_4$ assigns to the N_k 's, which must total to $1 - k - \bar{k}$, and comes from products of the form $m_4(\Omega)m(N_k)$:

$$1 - k - \bar{k} = \frac{zb}{1 - xc - ya}. \quad (10)$$

And finally, because m_4 is a mass function, we have a fourth constraint:

$$x + y + z = 1. \quad (11)$$

These four constraints form a set of linear equations in the unknowns x , y , and z . Except in the special case where $b = 0$, we will see that there can be only one solution. But we already dealt with the special case where $b := \sum_k m(N_k) = 0$ at the beginning of the proof, and are now operating under the assumption that $b \neq 0$.

To show that (8)–(9) have a unique solution in x , y , and z when $b \neq 0$, we first consider the case where $z \neq 0$. In this case $1 - k - \bar{k} \neq 0$ too, for otherwise (10) would entail $b = 0$. Thus we can divide each of equations (8) and (9) by equation (10):

$$\frac{x(a+b) + za}{zb} = \frac{k}{1 - k - \bar{k}}, \quad (12)$$

$$\frac{y(c+b) + zc}{zb} = \frac{\bar{k}}{1 - k - \bar{k}}. \quad (13)$$

Solving for x and for y :

$$x = z \frac{\left(\frac{bk}{1-k-\bar{k}} - a\right)}{a+b}, \quad (14)$$

$$y = z \frac{\left(\frac{b\bar{k}}{1-k-\bar{k}} - c\right)}{c+b}. \quad (15)$$

Substituting (14) and (15) for x and y in (11) and solving for z , we find a unique solution:

$$z = \frac{1}{\left(\frac{\frac{bk}{1-k-\bar{k}} - a}{a+b}\right) + \left(\frac{\frac{b\bar{k}}{1-k-\bar{k}} - c}{c+b}\right) + 1}. \quad (16)$$

Substituting for z back into equations (12) and (13) yields unique values for x and y as well. Because there is exactly one solution, and by hypothesis $x = m_1(B)$, $y = m_1(\bar{B})$, $z = m_1(\Omega)$ is a solution, $m_4 = m_1$ as promised.

What if $z = 0$? Then equations (8), (9), and (11) can be rewritten:

$$k = \frac{x(a+b)}{x(a+b) + y(c+b)}, \quad (17)$$

$$\bar{k} = \frac{y(c+b)}{x(a+b) + y(c+b)}, \quad (18)$$

$$1 = x + y \quad (19)$$

Since any increase in x must be accompanied by a decrease in y , the right hand side of (17) is strictly increasing in x . So only one value of x can satisfy (17), and thus only one value of y can satisfy (18). So the solution is still unique. \square

Having shown that Dempster Combination is Rigid, we now verify that Rigidity is Independence Preserving in DST:

Rigidity is Independence Preserving (DST): Suppose the transition from bel to bel' is rigid on the partition $\{B, \bar{B}\}$. Then if A and B are independent under bel , they are independent under bel' too.

Proof. Suppose the transition from bel to bel' is rigid on the partition $\{B, \bar{B}\}$, and A and B are independent under bel . Let bel_3 be a belief function focused on $\{B, \bar{B}\}$. By rigidity:

$$\begin{aligned} (bel \oplus bel_1)(A) &= (bel' \oplus bel_3)(A), \\ (bel \oplus bel_1)(\bar{A}) &= (bel' \oplus bel_3)(\bar{A}), \end{aligned}$$

for any bel_1 such that $(bel \oplus bel_1)(B) = (bel' \oplus bel_3)(B)$ and $(bel' \oplus bel_1)(\bar{B}) = (bel \oplus bel_3)(\bar{B})$. If such a bel_1 exists for every bel_3 , our independence assumption entails:

$$\begin{aligned} bel(A) &= (bel \oplus bel_1)(A), \\ bel(\bar{A}) &= (bel \oplus bel_1)(\bar{A}). \end{aligned}$$

Thus:

$$\begin{aligned} bel(A) &= (bel' \oplus bel_3)(A), \\ bel(\bar{A}) &= (bel' \oplus bel_3)(\bar{A}), \end{aligned}$$

for every bel_3 focused on $\{B, \bar{B}\}$, i.e. independence is preserved as desired.

Moreover, such a bel_1 is guaranteed to exist for every such bel_3 . For any bel, bel', m_3 we can always find an m_1 such that:

$$\begin{aligned} (bel' \oplus bel_3)(B) &= \sum_i m(B_i) \frac{m_1(B) + m_1(\Omega)}{c_1}, \\ (bel' \oplus bel_3)(\bar{B}) &= \sum_j m(\bar{B}_j) \frac{m_1(\bar{B}) + m_1(\Omega)}{c_1}. \end{aligned}$$

If we just substitute $1 - m_1(B)$ for $m_1(\bar{B}) + m_1(\Omega)$, this boils down to solving two linear equations of the following form for x and z :

$$\begin{aligned} k &= \frac{a}{c_1}(x + z) \\ \bar{k} &= \frac{b}{c_1}(1 - x) \end{aligned}$$

□

This establishes the RIP theorem for DST. Notice though, if we just wanted to show that Dempster’s rule preserves independence, we could have taken a much shorter route:

Dempster’s Rule Preserves Independence: Let A and B be independent under bel , and let bel_1 , bel_2 , and bel_3 be focused on $\{B, \bar{B}\}$. Then:

$$\begin{aligned} (bel \oplus bel_1)(A) &= ((bel \oplus bel_2) \oplus bel_3)(A), \\ (bel \oplus bel_1)(\bar{A}) &= ((bel \oplus bel_2) \oplus bel_3)(\bar{A}). \end{aligned}$$

Proof. Let A and B be independent under bel and let bel_1 , bel_2 , bel_3 be focused on $\{B, \bar{B}\}$. First, by the associativity of \oplus :

$$((bel \oplus bel_2) \oplus bel_3)(A) = (bel \oplus (bel_2 \oplus bel_3))(A). \quad (20)$$

Second, because the combination of two belief functions focused on the same partition is also focused on that partition, $bel_2 \oplus bel_3$ is focused on $\{B, \bar{B}\}$. So since A and B are independent relative to bel :

$$(bel \oplus (bel_2 \oplus bel_3))(A) = bel(A). \quad (21)$$

Third, since A and B are independent under bel , and bel_1 is focused on $\{B, \bar{B}\}$:

$$bel(A) = (bel \oplus bel_1)(A). \quad (22)$$

Chaining together (20)–(22) we have:

$$((bel \oplus bel_2) \oplus bel_3)(A) = (bel \oplus bel_1)(A),$$

as desired. And parallel reasoning shows the same for \bar{A} . \square

This proof is not only mercifully short, it also enables us to say something about the case where Ω is infinite, which we bracketed back in fn. 9,

As mentioned in fn. 9, when Ω is infinite some belief functions have no corresponding mass function. Since Dempster’s rule is usually defined using mass functions, how these “massless” belief functions combine is left open by DST. Different generalizations are possible, and one might conceivably devise a generalization that, at least for massless belief functions, escapes the theorem we just proved for finite Ω .

But any such generalization of Dempster’s rule would have to be very strange. The short proof relies only on very elementary properties of \oplus , properties that ought not depend on Ω being finite. Specifically, we assumed that (i) \oplus is associative, and (ii) the combination of two belief functions focused on the same partition is also focused

on that partition. It would be very surprising, and puzzling, if these properties held when Ω is finite, yet failed for infinite Ω .

Also noteworthy: even if we did devise such a generalization of Dempster's rule, the RIP theorem would still cover those cases where the agent's belief function does have a corresponding mass function. For these agents, the two properties necessary to prove the theorem will hold. So only agents whose belief functions are "massless" would escape the problem. Unless there is some reason "massive" agents are rational to ignore perceptual undermining, there is no shelter in the massless edge-case.

APPENDIX B: RESULTS FOR RANKING THEORY

Here we show that Spohn Conditionalization is Rigid, and that Rigidity is Independence Preserving in ranking theory.

Spohn Conditionalization is Rigid: If κ' comes from κ by Spohn Conditionalization on the partition $\{B_i\}$, then $\kappa'(A|B_i) = \kappa(A|B_i)$ for all A, B_i .

Proof. Suppose κ' comes from κ by Spohn Conditionalization on the partition $\{B_i\}$ with the input values $\{r_i\}$.

First, by the definition of conditional ranks, we have for each i :

$$\kappa'(A|B_i) = \kappa'(A \cap B_i) - \kappa'(B_i). \quad (23)$$

Second, since κ' comes from κ by Spohn Conditionalization on $\{B_i\}$ with the values $\{r_i\}$, we know that:

$$\kappa'(A \cap B_i) - \kappa'(B_i) = \min_j \{\kappa(A \cap B_i|B_j) + r_j\} - \min_j \{\kappa(B_i|B_j) + r_j\}. \quad (24)$$

When are the terms on the right-hand side minimized? Notice that $\kappa(B_i|B_j) = 0$ when $j = i$, since $\kappa(B_i|B_i) = 0$ by the definition of conditional rank. Whereas $\kappa(B_i|B_j) = \infty$ when $i \neq j$, since then $B_i \cap B_j = \emptyset$. Thus both terms are minimized when $i = j$, giving us our third equation:

$$\begin{aligned} \min_j \{\kappa(A \cap B_i|B_j) + r_j\} - \min_j \{\kappa(B_i|B_j) + r_j\} &= \kappa(A \cap B_i|B_i) + r_i - r_i \\ &= \kappa(A \cap B_i|B_i). \end{aligned} \quad (25)$$

Fourth, by the definition of conditional rank:

$$\begin{aligned} \kappa(A \cap B_i|B_i) &= \kappa(A \cap B_i \cap B_i) - \kappa(B_i) \\ &= \kappa(A \cap B_i) - \kappa(B_i) \\ &= \kappa(A|B_i). \end{aligned} \quad (26)$$

Chaining together equations (23)–(26), we get:

$$\kappa'(A|B_i) = \kappa(A|B_i),$$

as desired. \square

Rigidity is Independence Preserving (Ranking Theory): Suppose the transition from κ to κ' is rigid on the partition $\{B, \bar{B}\}$. Then, if A and B are independent under κ , they are independent under κ' too.

Proof. Suppose the transition from κ to κ' is rigid on $\{B, \bar{B}\}$ and that A and B are independent under κ . Let $A_i \in \{A, \bar{A}\}$. We will derive a series of eight equations which we'll then chain together.

First, by elementary set theory:

$$\kappa'(A_i) = \kappa'([A_i \cap B] \cup [A_i \cap \bar{B}]). \quad (27)$$

Then by the third axiom of ranking theory, (R3):

$$\kappa'([A_i \cap B] \cup [A_i \cap \bar{B}]) = \min[\kappa'(A_i \cap B), \kappa'(A_i \cap \bar{B})]. \quad (28)$$

And by the definition of conditional ranks:

$$\begin{aligned} \min[\kappa'(A_i \cap B), \kappa'(A_i \cap \bar{B})] &= \min[\kappa'(A_i|B) + \kappa'(B), \\ &\quad \kappa'(A_i|\bar{B}) + \kappa'(\bar{B})]. \end{aligned} \quad (29)$$

Since the transition from κ to κ' is rigid:

$$\begin{aligned} \min[\kappa'(A_i|B) + \kappa'(B), \kappa'(A_i|\bar{B}) + \kappa'(\bar{B})] &= \min[\kappa(A_i|B) + \kappa'(B), \\ &\quad \kappa(A_i|\bar{B}) + \kappa'(\bar{B})]. \end{aligned} \quad (30)$$

Then, because A and B are independent under κ :

$$\begin{aligned} \min[\kappa(A_i|B) + \kappa'(B), \kappa(A_i|\bar{B}) + \kappa'(\bar{B})] &= \min[\kappa(A_i) + \kappa'(B), \\ &\quad \kappa(A_i) + \kappa'(\bar{B})]. \end{aligned} \quad (31)$$

But it's an elementary theorem of ranking theory that $\kappa'(B)$ or $\kappa'(\bar{B})$ must be 0, so:

$$\min[\kappa(A_i) + \kappa'(B), \kappa(A_i) + \kappa'(\bar{B})] = \kappa(A_i). \quad (32)$$

By the independence of A and B under κ again:

$$\kappa(A_i) = \kappa(A|B). \quad (33)$$

And since the update from κ to κ' is rigid:

$$\kappa(A_i|B) = \kappa'(A_i|B). \quad (34)$$

Chaining together equations (27)–(34), we get:

$$\kappa'(A_i) = \kappa'(A_i|B).$$

And parallel reasoning with \bar{B} in place of B in the last two equations yields:

$$\kappa'(A_i) = \kappa'(A_i|\bar{B}).$$

Thus:

$$\kappa'(A|B) = \kappa'(A) = \kappa'(A|\bar{B}),$$

$$\kappa'(\bar{A}|B) = \kappa'(\bar{A}) = \kappa'(\bar{A}|\bar{B}).$$

That is, A and B are independent under κ' , as desired. □

REFERENCES

- Alchourrón, Carlos E., Peter Gärdenfors & David Makinson. 1985. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic* 50(2): 510–30.
- Ayer, A. J. 1940. *The foundations of empirical knowledge*. London: Macmillan.
- Ayer, A. J. 1946. *Language, truth, and logic*. New York: Dover.
- Ben Yaghlane, Boutheina, Philippe Smets & Khaled Mellouli. 2002. Belief function independence: I. the marginal case. *International Journal of Approximate Reasoning* 29(1): 47–70.
- Boutilier, Craig. 1993. Revision sequences and nested conditionals. *Proceedings of the 13th International Joint Conference on Artificial intelligence* 1: 519–25.
- Bradley, Richard. 2005. Radical probabilism and bayesian conditioning. *Philosophy of Science* 72(2): 342–64.
- Christensen, David. 1992. Confirmational holism and bayesian epistemology. *Philosophy of Science* 59: 540–57.
- Darwiche, Adnan & Judea Pearl. 1994. On the logic of iterated belief revision. *Proceedings of the 5th Conference on Theoretical Aspects of Reasoning About Knowledge* 5–23.
- Douven, Igor & Jan-Willem Romeijn. 2011. A new resolution of the judy benjamin problem. *Mind* 120(479): 637–70.
- Dubois, Didier & Henri Prade. 1986. A set-theoretic view of belief functions. *International Journal of General Systems* 12(3): 193–226.
- Elga, Adam. 2007. Reflection and disagreement. *Noûs* 41(3): 478–502.

- Fagin, Ronald & Joseph Y. Halpern. 1991. A new approach to updating beliefs. In P. P. Bonissone, M. Henrion, L. N. Kanal & J. F. Lemmer (eds.), *Uncertainty in artificial intelligence*, vol. 6, 347–74.
- Field, Hartry. 1978. A note on jeffrey conditionalization. *Philosophy of Science* 45: 361–7.
- Gallow, J. Dmitri. forthcoming. How to learn from theory-dependent evidence; or commutativity and holism: A solution for conditionalizers. *British Journal for the Philosophy of Science*.
- Gärdenfors, Peter. 1988. *Knowledge in flux: Modelling the dynamics of epistemic states*. Cambridge, MA: The MIT Press.
- Gärdenfors, Peter & David Makinson. 1988. Revisions of knowledge systems using epistemic entrenchment. In Moshe Vardi (ed.), *Theoretical aspects of reasoning about knowledge*, 83–95. Los Altos: Morgan Kaufmann.
- Grove, Adam. 1988. Two modellings for theory change. *Journal of Philosophical Logic* 17(2): 157–70.
- Hacking, Ian. 1967. Slightly more realistic personal probability. *Philosophy of Science* 34(4): 311–25.
- Halpern, Joseph Y. 2003. *Reasoning about uncertainty*. Cambridge, MA: The MIT Press.
- Howson, Colin & Allan Franklin. 1994. Bayesian conditionalization and probability kinematics. *British Journal for the Philosophy of Science* 45: 451–66.
- Huber, Franz. 2009. Belief and degrees of belief. In Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of belief*, 1–33. Springer.
- Huber, Franz. 2013. Formal representations of belief. Stanford Encyclopedia of Philosophy. plato.stanford.edu/entries/formal-belief/.
- Jeffrey, Richard C. 1965. *The logic of decision*. Chicago: University of Chicago Press.
- Jeffrey, Richard C. 1968. Probable knowledge. In Imre Lakatos (ed.), *The problem of inductive logic*, 166–80. Amsterdam: North Holland.
- Jeffrey, Richard C. 1975. Carnap's empiricism. *Minnesota Studies in the Philosophy of Science* 6: 37–49.
- Jeffrey, Richard C. 1983. Bayesianism with a human face. In John Earman (ed.), *Testing scientific theories*, 133–56. University of Minnesota Press.
- Jeffrey, Richard C. 2002. *Subjective probability: The real thing*. www.princeton.edu/~bayesway/Book*.pdf.
- Kolodny, Niko & John MacFarlane. 2010. Ifs and oughts. *Journal of Philosophy* 107(3): 115–43.
- Levi, Isaac. 1974. On indeterminate probabilities. *Journal of Philosophy* 71: 391–418.

- Lewis, David. 1980. A subjectivist's guide to objective chance. In Richard C. Jeffrey (ed.), *Studies in inductive logic and probability*, vol. II, 263–93. Berkeley: University of California Press.
- Lewis, David. 1999. Why conditionalize? In *Papers in metaphysics and epistemology*, 403–7. Cambridge University Press.
- Meacham, Christopher J. G. 2010. Two mistakes regarding the principal principle. *British Journal for the Philosophy of Science* 61(2): 407–31.
- Pollock, John L. 1986. *Contemporary theories of knowledge*. Lanham, MD: Rowman and Littlefield.
- Pollock, John L. 2008. Defeasible reasoning. In Jonathan E. Adler & Lance J. Rips (eds.), *Reasoning: Studies of human inference and its foundations*, 451–70. Cambridge; New York: Cambridge University Press.
- Pryor, James. forthcoming. Problems for credulism. In Chris Tucker (ed.), *Seemings and justification: New essays on dogmatism and phenomenal conservatism*, Oxford: Oxford University Press.
- Reichenbach, Hans. 1938. *Experience and prediction: An analysis of the foundations and the structure knowledge*. Chicago: University of Chicago Press.
- Reichenbach, Hans. 1952. Are phenomenal reports absolutely certain? *The Philosophical Review* 61(2): 147–59.
- Shafer, Glenn. 1976. *A mathematical theory of evidence*. Princeton: Princeton University Press.
- Shafer, Glenn. 1981. Jeffrey's rule of conditioning. *Philosophy of Science* 48(3): 337–62.
- Shore, J. E. & R. W. Johnson. 1980. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory* IT-26: 26–37.
- Smets, Philipe. 1990. The combination of evidence in the transferable belief model. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12(5): 447–58.
- Spohn, Wolfgang. 1988. Ordinal conditional functions: A dynamic theory of epistemic states. In William L. Harper & Brian Skyrms (eds.), *Causation in decision, belief change, and statistics*, 105–34. Dordrecht: Kluwer.
- Spohn, Wolfgang. 1999. Ranking functions, agm style. In *Internet festschrift for peter gärdenfors*, 1–20. www.lucs.lu.se/spinning/categories/dynamics/Spohn/Spohn.pdf.
- Spohn, Wolfgang. 2012. *The laws of belief: Ranking theory and its philosophical applications*. Oxford: Oxford University Press.
- van Fraassen, Bas. 1981. A problem for relative information minimizers in probability kinematics. *British Journal for the Philosophy of Science* 32(4): 375–9.
- van Fraassen, Bas. 1984. Belief and the will. *The Journal of Philosophy* 81(5): 235–56.

- van Fraassen, Bas. 1989. *Laws and symmetry*. Oxford: Clarendon Press.
- van Fraassen, Bas. 1990. Figures in a probability landscape. In J. Dunn & A. Gupta (eds.), *Truth or consequences*, 345–56. Dordrecht: Kluwer.
- van Fraassen, Bas. 1995. Belief and the problem of ulysses and the sirens. *Philosophical Studies* 77: 7–37.
- Wagner, Carl. 2013. Is conditioning really incompatible with holism? *Journal of Philosophical Logic* 42(2): 409–14.
- Weisberg, Jonathan. 2007. Conditionalization, reflection, and self-knowledge. *Philosophical Studies* 135(2): 179–97.
- Weisberg, Jonathan. 2009. Commutativity or holism: A dilemma for conditionalizers. *British Journal for the Philosophy of Science* 60(4): 793–812.
- Williams, P. M. 1980. Bayesian conditionalisation and the principle of minimum information. *British Journal for the Philosophy of Science* 32(2): 131–44.