



The value of cost-free uncertain evidence

Patryk Dziurosz-Serafinowicz¹ · Dominika Dziurosz-Serafinowicz¹

Received: 14 February 2021 / Accepted: 20 August 2021 / Published online: 7 September 2021
© The Author(s) 2021

Abstract

We explore the question of whether cost-free *uncertain* evidence is worth waiting for in advance of making a decision. A classical result in Bayesian decision theory, known as the value of evidence theorem, says that, under certain conditions, when you update your credences by conditionalizing on some cost-free and *certain* evidence, the subjective expected utility of obtaining this evidence is never less than the subjective expected utility of not obtaining it. We extend this result to a type of update method, a variant of Judea Pearl's virtual conditionalization, where uncertain evidence is represented as a set of likelihood ratios. Moreover, we argue that focusing on this method rather than on the widely accepted Jeffrey conditionalization enables us to show that, under a fairly plausible assumption, gathering uncertain evidence not only maximizes expected pragmatic utility, but also minimizes expected epistemic disutility (inaccuracy).

Keywords The value of evidence theorem · Expected utility · Accuracy · Jeffrey conditionalization · Virtual conditionalization

1 Introduction

Can it be that a rational agent would postpone a decision in order to acquire cost-free but *uncertain* evidence? Suppose that there are two types of urns containing marbles. Urn *X* contains 8 blue and 2 violet marbles, while urn *Y* contains 2 blue and 8 violet marbles. A given urn is selected at random by the toss of a fair coin. Ann is about to guess the type of the selected urn, and her possible guesses have different practical consequences that depend on whether the selected urn is of type *X* or *Y*. But before guessing, she has the opportunity to draw, cost free, a marble from the

✉ Patryk Dziurosz-Serafinowicz
patryk.dziurosz-serafinowicz@ug.edu.pl

Dominika Dziurosz-Serafinowicz
dominika.dziurosz-serafinowicz@phdstud.ug.edu.pl

¹ Institute of Philosophy, University of Gdansk, Jana Bażyńskiego 4 Street, 80-309 Gdańsk, Poland

selected urn. Yet Ann knows in advance that the lighting is so dim that it would be difficult to discern what her experience says: given her background information about the lighting conditions, she expects that learning experience would make her uncertain about whether the drawn marble is blue, B , or violet, V .

A classical result in Bayesian decision theory (Savage 1954, ch. 4; Raiffa and Schlaifer 1961, ch. 4.5; Good 1967; Ramsey 1990), known as the value of evidence theorem (VET), says that, under certain conditions, when an agent updates her credences upon the receipt of cost-free evidence, the subjective expected pragmatic utility of obtaining this evidence is never less than the subjective expected pragmatic utility of not obtaining it. That is, expecting to obtain cost-free evidence cannot lead you to expecting to make worse practical decisions.¹ The original VET, however, is limited to cases where the agent learns proposition E for certain from a set \mathcal{E} of mutually exclusive and jointly exhaustive propositions, and hence may update her credences by dint of Bayesian conditionalization (BCondi, for short). Crucially, in such cases the agent is certain *ex ante* that exactly one proposition E from \mathcal{E} will be true. But often, like in Ann's case, we undergo learning experiences where it is hard to discern what that proposition is, and so we become uncertain about which element of \mathcal{E} is true. Could cost-free uncertain evidence so understood be worth waiting for in advance of making a decision?

This question has not gone unnoticed in the literature. Graves (1989) showed that we can extend VET so that it holds for cases where we become uncertain about what the logically strongest proposition we learn is, and update our credences by using a rule called Jeffrey conditionalization (JCondi, for short). This rule requires uncertain evidence to be specified as a redistribution of the agent's credences over the propositions in some partition \mathcal{E} of a set of possibilities, without assigning absolute certainty to any particular proposition (hereafter, a Jeffrey shift). For example, Ann's learning experience can be understood as a Jeffrey shift over the partition $\{B, V\}$. To accommodate this type of uncertain evidence, Graves's argument, as we will argue, is mobilized by two conceptual moves. The first one is that any Jeffrey shift can be specified as a sort of propositional certainty, i.e. as a proposition that receives posterior credence 1 in an enriched subjective probability space. This enrichment is achieved by adding to the original smaller space propositions about one's posterior credences attached to the members of a partition \mathcal{E} . The second key move is to show that, under certain conditions, BCondi on the proposition specifying the posterior credences over \mathcal{E} in the enriched space is equivalent to JCondi in the original small space.

After challenging Graves's argument, this paper offers an alternative extension of VET to the case of learning from uncertain evidence. To preview, instead of recasting uncertain evidence as certain in an enriched subjective probability space, the proposed view retains the uncertainty of one's evidence in the original smaller space, and provides a specification of this uncertainty by utilizing the method of virtual evidence proposed by Pearl (1988, 1990) and developed in Chan and Darwiche (2005). According to this method, uncertain evidence can be specified as a set of likelihood ratios, where each likelihood ratio tells you how well some virtual evidence fits with some

¹ Note that VET does not say that the agent will actually do better by learning a proposition. It says that she will do better in expectation, from the perspective of her *ex ante* credence function.

proposition in partition \mathcal{E} as compared to how well it fits with another proposition in that partition. The virtual evidence is meant to be an auxiliary proposition that bears on the truth of propositions in \mathcal{E} . We supply this method with a specific understanding of what the auxiliary proposition could be in the context of learning from uncertain evidence. Our proposal is that it can be understood as the proposition $U_{c_\lambda}^E$ which says that you take your evidence to be E and adopt the posterior credence function c_λ in response. For short, we will refer to this proposition as saying that *you update on E* .² Importantly, when you take E as your evidence and adopt a posterior credence function in response, you foresee the possibility that $U_{c_\lambda}^E$ could be true, even if E is in fact false. This is because, in the context of learning from uncertain evidence, you may mistake the true evidence E' for some other $E \in \mathcal{E}$. So for example, in Ann's case, when looking at the drawn marble in a dim light, she might take her evidence to be B , and adopt the posterior c_λ in response, when in fact V is true.

There is, we suggest, a way to incorporate this possibility of mistake in one's learning into the update mechanism. The key idea is that we can express it by determining the extent to which the proposition $U_{c_\lambda}^E$ is more likely under E than E' , or the extent to which $U_{c_\lambda}^E$ favours E over E' . Intuitively, if Ann has updated on B , she can express the extent to which $U_{c_\lambda}^B$ is more likely under B than V . One natural way to determine this extent is to settle on a likelihood ratio which tells Ann to what extent the proposition $U_{c_\lambda}^B$ is more expected under B than V . This likelihood ratio can be any non-negative number she considers reasonable in the light of her background knowledge. As will be explained in more detail later on, there is a reasonable way to plug these likelihood ratios into an update rule, without the need of determining the absolute likelihoods for $U_{c_\lambda}^E$. As will be shown, the proposed update method gives the expected result in Ann's case in the sense that it leaves B and V uncertain in Ann's posteriors.

Armed with the method of virtual evidence so understood, we show how VET can be extended to the context of uncertain evidence. Three basic ideas underpin this extension. First, updating on a set of likelihood ratios for the proposition $U_{c_\lambda}^E$ can be modelled as a variant of what Pearl called 'virtual conditionalization' (VCondi, for short). Second, under a plausible assumption, updating by this variant of VCondi is equivalent to a version of Bas van Fraassen's (1984) reflection principle. This principle says that one's posterior credence function should be equal to one's prior conditional on the proposition $U_{c_\lambda}^E$. Third, once we assume that the propositions $U_{c_\lambda}^E$ form a partition, we can show that the expected worth of accomodating cost-free uncertain evidence by the proposed variant of VCondi cannot be negative. And this expectation is calculated relative to the agent's prior credences over the propositions $U_{c_\lambda}^E$.

We proceed as follows. In Sect. 2, we discuss three different ways of specifying uncertain evidence that play a crucial role in both Graves's and our alternative extension of VET. In Sect. 3, we present three updating rules, each attuned to a different way of specifying uncertain evidence. In Sect. 4, we spell out in detail Graves's extension of VET. In Sect. 5, we show that Graves's argument is problematic when we care not only about the practical rationality of our decisions, but also about the epistemic rationality of our beliefs. In Sect. 6, we lay down our alternative extension of VET to

² A very similar proposition has been employed by Gallow (2019) in his account of externalist updating which covers situations where it is rational for an agent to be uncertain about what her evidence is.

the case of learning from uncertain evidence and show that it dovetails with a purely epistemic approach that aims to vindicate updating on uncertain evidence. Section 7 concludes.

2 Specifying uncertain evidence: three ways

Let $(\mathcal{W}, \mathcal{F}, c)$ be an agent's subjective probability space (her *credal space*), where \mathcal{W} represents the possibilities that the agent can distinguish between, \mathcal{F} is an algebra of subsets of \mathcal{W} that can be understood as the propositions the agent can express, and c is the agent's credence function which assigns numbers from $[0, 1]$, called credences, to propositions in \mathcal{F} . We will assume throughout that, at any given time, the credence function c is a probability function over \mathcal{F} . Since we will be mostly interested in the dynamics of credences in a decision context, we require that \mathcal{F} includes a finite partition \mathcal{S} of \mathcal{W} , $\mathcal{S} \subseteq \mathcal{F}$, which contains propositions S representing the states of the world upon which the consequences of the agent's actions depend. As will be apparent later on, we also need to require that the algebra \mathcal{F} is sufficiently rich so that it includes a finite number of propositions $U_{c_\lambda}^E$. Recall that a proposition of this sort says that you update on E .

Now, learning experience can provide an agent with various types of evidential input λ that may prompt a revision of their credence in any $X \in \mathcal{F}$, $c(X)$, resulting in a posterior credence in X , $c_\lambda(X)$. But how can we characterize an evidential input more precisely? There are at least two ways of answering this question that are well entrenched in Bayesian epistemology. Both assume that learning experiences do not provide evidential inputs all by themselves. Rather they provide evidential inputs because they impact on our credences about evidence propositions, which may in turn provide support for other propositions. However, these two approaches differ on how this impact should be cashed out.

According to the first, somewhat more popular view, any evidential input takes the form of a direct change in one's credences over some set of propositions. Thus, any evidential input acts like a *constraint* on the set of possible credence functions \mathcal{C} and restricts the candidates for the agent's posterior credence function. More formally, an evidential input of this sort can be understood as a set of posterior credence functions over \mathcal{F} , \mathcal{C}_λ , which contains those possible credence functions of the agent that are consistent with this input, that is, $\mathcal{C}_\lambda \subseteq \mathcal{C}$ (see, e.g. van Fraassen 1989, ch. 13; Uffink 1996; Joyce 2010; Dietrich et al. 2016; van Fraassen and Halpern 2017). Importantly, this view allows us to treat as evidential input information or data that cannot be expressed as propositional evidence.

According to the second view, we may think of evidential input as a sort of *update factor*, which can then bring about changes in one's credences (see, e.g. Wagner 2002; Hawthorne 2004). Some update factors may be essentially relative, i.e. they tell us how well an outcome of one's learning experience fits with the proposition X as compared to how well it fits with the proposition Y . A classical example of such a factor is given by the likelihood ratio of propositions X and Y given the evidence proposition E ,

$\frac{c(E|X)}{c(E|Y)}$ (see, e.g. Good 1950). A more general representation of this factor is given by the odds ratio of X and Y , also called the Bayes factor:³

Bayes factor: If c_λ is a posterior credence function and $X, Y \in \mathcal{F}$, then the Bayes factor of X and Y is given by:

$$\mathbb{B}_{c_\lambda, c}(X, Y) = \frac{c_\lambda(X)}{c_\lambda(Y)} / \frac{c(X)}{c(Y)}. \quad (1)$$

That is, the Bayes factor of X and Y is a ratio of new-to-old odds for X against Y , where the new odds are $\frac{c_\lambda(X)}{c_\lambda(Y)}$, and the old odds are $\frac{c(X)}{c(Y)}$. This ratio is meant to capture the factor by which the old odds for X against Y can be multiplied to get the new odds. Observe that the agent's Bayes factors do not uniquely specify her posterior credences, but only impose overall constraints on these posteriors.

In order to show that the likelihood ratio—which isolates the full import of learning experience, with prior credences factored out—is an instance of the Bayes factor, assume that one's posterior credence for any $X \in \mathcal{F}$, $c_\lambda(X)$, comes from one's prior credence in X by conditionalizing on the proposition $E \subseteq \mathcal{W}$, i.e. $c_\lambda(X) = c(X|E)$. Then,

$$\mathbb{B}_{c_\lambda, c}(X, Y) = \frac{\frac{c(X|E)}{c(Y|E)}}{\frac{c(X)}{c(Y)}} = \frac{\frac{\frac{c(X)c(E|X)}{c(E)}}{\frac{c(Y)c(E|Y)}{c(E)}}}{\frac{c(X)}{c(Y)}} = \frac{\frac{c(X)c(E|X)}{c(Y)c(E|Y)}}{\frac{c(X)}{c(Y)}} = \frac{c(E|X)}{c(E|Y)}. \quad (2)$$

Given the above preliminaries, how can we represent uncertain evidence? In what follows, we will characterize three ways of specifying uncertain evidence: the first one is due to Jeffrey (1983), the second one is due to Skyrms (1980), and was used in Graves's extension of VET, and the third one is an amended method proposed in Pearl (1988, 1990) and developed in Chan and Darwiche (2005).

Jeffrey (1983) argued that in many cases learning experience does not constrain an agent's credences in a way that is tailor-made for the orthodox Bayesian conditioning, that is, when the evidential input can be modelled as the logically strongest evidence proposition E in \mathcal{W} that receives the posterior credence of 1, or simply as the Bayes input:

Bayes input: $\mathcal{C}_\lambda^E = \mathcal{C}_\lambda = \{c_\lambda : c_\lambda(E) = 1\}$ for some $E \subseteq \mathcal{W}$ such that $E \neq \emptyset$.

Jeffrey believed that, in cases like Ann's, although learning experience does not single out an evidence proposition E that receives posterior credence 1, $c_\lambda(E) = 1$, it nevertheless directly affects the agent's prior credences over the propositions in some set \mathcal{E} , which is a partition of \mathcal{W} , shifting them to posterior credences $c_\lambda(E)$, for all $E \in \mathcal{E}$. Thus, uncertain evidence so understood can be modelled as the following evidential input:

³ Some philosophers deny that a Bayes factor represents an evidential input. Most notably, Bradley (2005) claims that the Bayes factors must be regarded as features of the agent's response to experience, not as what experience delivers to her. And this is so because they lack the kind of objectivity which obligates others using them as constraints on their posterior credences. But, as Jeffrey (1992, pp. 7–9) pointed out convincingly, so long as others trust your response to your experience, they would use your Bayes factors to move to a posterior credence function, whether or not they share your prior credence function.

Jeffrey shift: $\mathcal{C}_\lambda^{\mathbb{E}} = \mathcal{C}_\lambda = \{c_\lambda : c_\lambda(E) \geq 0 \text{ for all } E \in \mathcal{E}\}$ for some set \mathbb{E} of ordered pairs $\{\langle E, c_\lambda(E) \rangle\}$ such that $E \in \mathcal{E}$ and $\sum_{E \in \mathcal{E}} c_\lambda(E) = 1$.

That is, uncertain evidence so understood is a redistribution of the agent's credences over the propositions in some partition \mathcal{E} of the set of worlds she considers possible.

To provide Skyrms's characterization of uncertain evidence, we need to consider a particular extension of $(\mathcal{W}, \mathcal{F}, c)$. Given two algebras \mathcal{F} and \mathcal{F}^* , let the injective map $*$: $\mathcal{F} \rightarrow \mathcal{F}^*$ be an algebra embedding, that is, a function that preserves all Boolean operations. Then, let $(\mathcal{W}^*, \mathcal{F}^*, c^*)$ be an extension of $(\mathcal{W}, \mathcal{F}, c)$ such that:

E_1 For every $X \in \mathcal{F}$, $c^*(X^*) = c(X)$.

E_2 For some $\mathcal{E} \subseteq \mathcal{F}$ and a finite set of posterior credence functions \mathcal{C}_λ , \mathcal{F}^* contains a finite number of propositions $R_{c_\lambda}^{\mathcal{E}}$.

That is, the algebra \mathcal{F}^* in the extended credal space $(\mathcal{W}^*, \mathcal{F}^*, c^*)$ contains the copies X^* of all the propositions X in the original smaller algebra \mathcal{F} , and it also contains a finite number of propositions $R_{c_\lambda}^{\mathcal{E}}$, each saying that the posterior credences over \mathcal{E} are given by the credence function c_λ .

Now, we can also consider learning experiences in the extended credal space that prompt revisions of c^* resulting in a new credence function c_λ^* . In particular, given the set of posterior credence distributions over \mathcal{F}^* , $\mathcal{C}_\lambda^* \subseteq \mathcal{C}^*$, that are consistent with the evidential input λ in that extended credal space, uncertain evidence may be presented as the following evidential input:

Skyrms input: $\mathcal{C}_\lambda^{*R} = \mathcal{C}_\lambda^* = \{c_\lambda^* : c_\lambda^*(R_{c_\lambda}^{\mathcal{E}}) = 1\}$ for some $R_{c_\lambda}^{\mathcal{E}}$ such that $R_{c_\lambda}^{\mathcal{E}} \neq \emptyset$.

Thus, uncertain evidence modelled as Skyrms input is the agent's assignment of posterior credence 1 in the extended algebra to the proposition $R_{c_\lambda}^{\mathcal{E}}$ that specifies a Jeffrey shift over some partition \mathcal{E} .

Unlike Jeffrey's and Skyrms's methods, Pearl's method of virtual evidence specifies uncertain evidence as an evidential input that directly constrains quantities different from the absolute values of posterior credences. The core idea of Pearl's method is that we can interpret the uncertainty of every proposition E in some partition \mathcal{E} of \mathcal{W} as the uncertainty of E 's relevance to some auxiliary proposition in \mathcal{W} . And to specify how uncertain this relevance is, Pearl proposes to use a set of likelihood ratios. Here we will amend Pearl's method by assuming that the auxiliary proposition is the proposition $U_{c_\lambda}^E$, which says that you update on E . More precisely,

Pearl-style input: $\mathcal{L}^E = \left\{ \alpha_E : \alpha_E = \frac{c(U_{c_\lambda}^E|E)}{c(U_{c_\lambda}^E|E')}, \alpha_E \in [0, \infty) \text{ and } \alpha_{E'} = 1 \right\}$,

for some $U_{c_\lambda}^E \subseteq \mathcal{W}$ and some partition $\{E : E \in \mathcal{E}\}$ of \mathcal{W} .

That is, the set \mathcal{L}^E contains the likelihood ratios α_E of the proposition $U_{c_\lambda}^E$ that are relative to a pair of propositions E and E' , with the likelihood ratio $\frac{c(U_{c_\lambda}^E|E')}{c(U_{c_\lambda}^E|E')}$ set equal to 1. And each likelihood ratio in this set is given by the prior credence of $U_{c_\lambda}^E$ conditional on E divided by the prior credence of $U_{c_\lambda}^E$ conditional on E' .

Importantly, to determine a Pearl-style input, we only need to specify a set of likelihood ratios, not the absolute likelihoods for $U_{c_\lambda}^E$ given E , for every $E \in \mathcal{E}$. But since every likelihood ratio α_E in the set \mathcal{L}^E is proportional to the absolute likelihood of $U_{c_\lambda}^E$ given E , i.e. $\alpha_E \propto c(U_{c_\lambda}^E|E)$, there exists a positive constant l such that, for all $E \in \mathcal{E}$, $c(U_{c_\lambda}^E|E) = l \cdot \alpha_E$. Thus, the absolute likelihoods can be determined indirectly from the set of likelihood ratios, albeit not uniquely.

Before moving on, let us elaborate on the proposition $U_{c_\lambda}^E$. Recall that this proposition says that you take your evidence to be E and adopt the posterior credence function c_λ in response. Firstly, following Gallow (2019), we assume that ‘taking your evidence to be E ’ does not involve any belief that E is true, and hence does not mean that you assign to E the posterior credence 1. Secondly, we want to emphasize that even if the agent already knows that she has updated on E , it does not entail that $c(U_{c_\lambda}^E|E) = 1$, for all E in \mathcal{E} . That is, though $U_{c_\lambda}^E$ can be regarded as ‘old evidence’ after the agent has updated on E , we can still reasonably inquire about its evidential impact on the propositions in \mathcal{E} .⁴ If so, we should not interpret the conditional credence $c(U_{c_\lambda}^E|E)$ as the agent’s actual credence in $U_{c_\lambda}^E$, supposing E to be true, for if you already know that $U_{c_\lambda}^E$ is true, your prior actual credence in $U_{c_\lambda}^E$ conditional on any proposition is 1. Instead, we should understand it as a kind of counterfactual credence. That is, in determining these credences, we should answer the question: How probable would the actual evidence $U_{c_\lambda}^E$ be if E were true? Naturally, we might expect that assigning precise counterfactual credences to propositions involves a lot of conceptual and formal intricacies.⁵ Fortunately, in our approach, there is no need to determine the precise values of these counterfactual credences. It suffices that the agent will express only their ratios, leaving the precise values unspecified.

Let us now show how uncertain evidence recast as Pearl-style input can be applied to Ann’s case. Recall that Ann knows beforehand that because she would observe the drawn marble in a dim light, she would neither become certain that B is true, nor that V is true. Nevertheless, foreseeing the possibility of error, she can take B or V as her evidence and adopt some posterior credences over \mathcal{F} in response. Suppose that, after looking at the drawn marble, she updates on B . But, due to her knowledge about the dim lighting, she foresees the possibility that $U_{c_\lambda}^B$ could be true, even if B is false. Still, she can interpret $U_{c_\lambda}^B$ as providing evidence for B against V whose strength can be given by a set of likelihood ratios. Let us suppose that Ann thinks that it is twice as likely that $U_{c_\lambda}^B$ would be true if B were true as if V were true. If this is so, then we may say that learning experience provides Ann with the following likelihood ratios: $\alpha_B = \frac{c(U_{c_\lambda}^B|B)}{c(U_{c_\lambda}^B|V)} = 2$ and $\alpha_V = \frac{c(U_{c_\lambda}^B|V)}{c(U_{c_\lambda}^B|V)} = 1$, and so $\mathcal{L}^B = \{1, 2\}$. This set of likelihood ratios enables her in turn to specify, though not uniquely, the absolute

⁴ Consider the following analogy. You are about to toss a coin, but you don’t know whether it is fair (H_1) or double-headed (H_2). Suppose that you have observed that the coin lands heads 100 times in a row (E). Even though you already know that E is true, your observation provides evidence in favour of H_2 . For the more times you see heads, the more evidence you have for the coin being double-headed. Hence, you are far from saying that $c(E|H_1) = c(E|H_2) = 1$.

⁵ See, e.g. Sprenger (2015).

Table 1 Ann's likelihoods expressed in terms of her error credences $c(U_{c_\lambda}^B|\neg B)$ and $c(U_{c_\lambda}^{\neg B}|B)$

	B	$\neg B$
$U_{c_\lambda}^B$	$1 - c(U_{c_\lambda}^{\neg B} B)$	$c(U_{c_\lambda}^B \neg B)$
$U_{c_\lambda}^{\neg B}$	$c(U_{c_\lambda}^{\neg B} B)$	$1 - c(U_{c_\lambda}^B \neg B)$

likelihoods $c(U_{c_\lambda}^B|B)$ and $c(U_{c_\lambda}^B|V)$. For example, she might assign $c(U_{c_\lambda}^B|B) = 0.4$ and $c(U_{c_\lambda}^B|V) = 0.2$, or $c(U_{c_\lambda}^B|B) = 0.8$ and $c(U_{c_\lambda}^B|V) = 0.4$.

The above understanding of Pearl-style input can also be explained in terms of Ann's error credences, as shown in Table 1. Think of $U_{c_\lambda}^B$ and $U_{c_\lambda}^{\neg B}$, where $\neg B = V$, as the possible, mutually exclusive noisy signals of Ann's learning experience. That is, before looking at the drawn marble, Ann thinks that she could mistake B for $\neg B$, and likewise she could mistake $\neg B$ for B . Thus, Ann thinks that there is some non-zero probability—Ann's false positive credence—that $U_{c_\lambda}^B$ would be true, even if B were not, and some non-zero probability—Ann's false negative credence—that $U_{c_\lambda}^{\neg B}$ would be true even if B were true. Then, when Ann takes B as her evidence and adopts the posterior c_λ in response, the differential support that $U_{c_\lambda}^B$ provides can be expressed as a likelihood ratio α_B of false negative and false positive credences, $\frac{1 - c(U_{c_\lambda}^{\neg B}|B)}{c(U_{c_\lambda}^B|\neg B)}$.

Why should we think that the problem of specifying uncertain evidence is philosophically important? Firstly, as it turns out, the way we choose to specify uncertain evidence helps us to resolve the problem of non-commutativity of JCondi, which says that, once you update sequentially by dint of JCondi, switching the order in which a pair of Jeffrey shifts over partitions \mathcal{E} and \mathcal{E}' is learned can yield different posterior credences in the end. This feature of JCondi is often regarded as its flaw. But, as shown by Field (1978) and developed by Wagner (2002), JCondi is commutative when identical learning is interpreted as identical Bayes factors. Note also that if uncertain evidence is represented by Skyrms input, then sequential updating on such uncertain evidence by a variant of BCondi, as given in chapter 3, would also be commutative, since BCondi is essentially commutative. Secondly, as argued in Wagner (2009), a certain parametrization of JCondi which uses the agent's Bayes factors rather than her posterior credences enables us to show that JCondi and the so-called opinion pooling—a method of aggregating probabilistic credences—commute. That is, the result of pooling and then updating by JCondi is the same as first updating by JCondi and then pooling.

As we will show in what follows, the way we specify uncertain evidence bears also on whether or not updating on uncertain evidence both maximizes expected pragmatic utility and minimizes expected epistemic disutility (inaccuracy).

3 Updating on uncertain evidence

A widespread position in Bayesian epistemology is that JCondi is an appropriate update rule when the agent undergoes a learning experience which does not rationalize absolute certainty in any proposition. This rule may be presented as follows:

JCondi: Given a Jeffrey shift $C_\lambda^{\mathbb{E}}$, the agent's posterior credence function c_λ should be such that, for every $X \in \mathcal{F}$,

$$c_\lambda(X) = \sum_{E \in \mathcal{E}} c(X|E) \cdot c_\lambda(E)$$

Although any Jeffrey shift $C_\lambda^{\mathbb{E}}$ which does not rationalize absolute certainty in any proposition in \mathcal{E} cannot be mediated by way of Bayes input in the credal state $(\mathcal{W}, \mathcal{F}, c)$, it might be tempting to think that it can be mediated by some other proposition in $(\mathcal{W}, \mathcal{F}, c)$ that we learn for certain. But if this were possible, there would be no need for JCondi in the first place. After all, this rule was motivated by the thought that there is no proposition in \mathcal{F} to conditionalize upon. Consider Jeffrey's candlelight case in which the agent inspects a piece of cloth by candlelight and gets the impression that it is green, although she concedes that it might be blue, or even violet. As Jeffrey (1983, p. 165) argued convincingly, there is no proposition that can convey the precise quality of this learning experience. And even if we allow for the possibility that the agent learns with certainty the proposition that the cloth *appears* green, this proposition would be too vague, for various learning experiences that fit this proposition would justify different credences in the proposition that the cloth is green (Christensen 1992). But we might well think that a Jeffrey shift can be mediated by the proposition $R_{c_\lambda}^{\mathcal{E}}$ that the agent learns with certainty in $(\mathcal{W}, \mathcal{F}, c)$. After all, it seems plausible to think that, in the context of JCondi, the agent becomes certain that her credences over some partition \mathcal{E} shifted in a certain way. But it seems that, in cases like Jeffrey's candlelight example, this proposition does not describe the precise content of the agent's learning experience, but merely summarizes the effect of this experience on her posterior credences over \mathcal{E} . And, again, if $(\mathcal{W}, \mathcal{F}, c)$ included this kind of proposition, there would be no need for JCondi in the first place. After all, one could just then use only BCondi and conditionalize on the proposition $R_{c_\lambda}^{\mathcal{E}}$, so that, for every $X \in \mathcal{F}$, $c_\lambda(X) = c(X|R_{c_\lambda}^{\mathcal{E}})$.

Schwan and Stern (2017) have recently argued that what the agent learns with certainty in the context of JCondi can be represented by a dummy proposition, D , which says what the agent would have learned with certainty were she capable of expressing it. They claim that this in turn allows us to represent updating on uncertain evidence as conditionalization on a dummy proposition in \mathcal{F} . For example, when Ann looks at the drawn marble in a dim light and undergoes a Jeffrey shift over $\{B, V\}$, we can represent her as if she becomes certain of some ineffable-colour-proposition D . But though this view has considerable merit when combined with Schwan and Stern's causal understanding of when the condition called rigidity⁶ is satisfied, it does not

⁶ Rigidity says that for all $X \in \mathcal{F}$ and all $E \in \mathcal{E}$, $c_\lambda(X|E) = c(X|E)$. As shown in Jeffrey (1983), rigidity is equivalent to JCondi.

provide much help in establishing VET in the context of uncertain evidence. For VET requires an agent to determine prior credences over various evidence propositions that she might learn in the future, but the ineffability of the dummy proposition prevents it from entering into the calculation of her prior credences. In particular, we cannot determine Ann's prior credence in the ineffable-colour-proposition D because we lack the ability to describe what Ann sees in a dim light.

Still, however, it seems reasonable to think that a Jeffrey shift can be mediated by way of a Skyrms input, C_{λ}^{*R} , i.e. by the proposition $R_{c_{\lambda}}^{\mathcal{E}}$ that one learns with certainty in the extended credal state $(\mathcal{W}^*, \mathcal{F}^*, c^*)$.⁷ If so, then we may think that when the agent updates by JCondi, she ought to update as if she were (i) expanding \mathcal{F} to \mathcal{F}^* and extending c to c^* , (ii) conditionalizing c^* upon some $R_{c_{\lambda}}^{\mathcal{E}}$ in \mathcal{F}^* to get the posterior c_{λ}^* , and then (iii) recovering her posterior c_{λ} over \mathcal{F} by restricting c_{λ}^* to \mathcal{F} . This idea is captured by the extended Bayesian conditionalization:

EBCondi: Given a Skyrms input C_{λ}^{*R} , the agent's posterior credence function c_{λ} should be such that, for every $X \in \mathcal{F}$:

$$c_{\lambda}(X) = c^*(X^* | R_{c_{\lambda}}^{\mathcal{E}}).$$

Skyrms (1980) has argued that, under some auxiliary assumptions, updating by EBCondi in $(\mathcal{W}^*, \mathcal{F}^*, c^*)$ is equivalent to JCondi in $(\mathcal{W}, \mathcal{F}, c)$. This result can be stated more formally as follows:

Proposition 1 Suppose that \mathcal{E} is a finite partition of \mathcal{W} and $R_{c_{\lambda}}^{\mathcal{E}} \in \mathcal{F}^*$. And suppose that c^* satisfies the following two conditions:

- C_1 For every $E \in \mathcal{E}$, $c^*(E^* | R_{c_{\lambda}}^{\mathcal{E}}) = c_{\lambda}(E)$, provided that $c^*(R_{c_{\lambda}}^{\mathcal{E}}) > 0$.
 C_2 For every $X \in \mathcal{F}$ and every $E \in \mathcal{E}$, $c^*(X^* | E^* \wedge R_{c_{\lambda}}^{\mathcal{E}}) = c^*(X^* | E^*)$, provided that $c^*(E^* \wedge R_{c_{\lambda}}^{\mathcal{E}}) > 0$.

Then, $\text{EBCondi} \Leftrightarrow \text{JCondi}$.

Proof Given the conditions (C1) and (C2), the proof of Proposition 1 is straightforward:

$$\begin{aligned} c_{\lambda}(X) &= c^*(X^* | R_{c_{\lambda}}^{\mathcal{E}}) \\ &= \frac{c^*(X^* \wedge R_{c_{\lambda}}^{\mathcal{E}})}{c^*(R_{c_{\lambda}}^{\mathcal{E}})} \\ &= \frac{c^*(\bigvee_{E \in \mathcal{E}} (X^* \wedge R_{c_{\lambda}}^{\mathcal{E}} \wedge E^*))}{c^*(R_{c_{\lambda}}^{\mathcal{E}})} \\ &= \frac{\sum_{E \in \mathcal{E}} c^*(X^* \wedge R_{c_{\lambda}}^{\mathcal{E}} \wedge E^*)}{c^*(R_{c_{\lambda}}^{\mathcal{E}})} \\ &= \frac{\sum_{E \in \mathcal{E}} c^*(X^* | R_{c_{\lambda}}^{\mathcal{E}} \wedge E^*) c^*(R_{c_{\lambda}}^{\mathcal{E}} \wedge E^*)}{c^*(R_{c_{\lambda}}^{\mathcal{E}})} \end{aligned}$$

⁷ For this view, see Skyrms (1980).

$$\begin{aligned}
&= \frac{\sum_{E \in \mathcal{E}} c^*(X^* | R_{c_\lambda}^\mathcal{E} \wedge E^*) c^*(E^* | R_{c_\lambda}^\mathcal{E}) c^*(R_{c_\lambda}^\mathcal{E})}{c^*(R_{c_\lambda}^\mathcal{E})} \\
&= \sum_{E \in \mathcal{E}} c^*(X^* | R_{c_\lambda}^\mathcal{E} \wedge E^*) \cdot c^*(E^* | R_{c_\lambda}^\mathcal{E}) \\
&= \sum_{E \in \mathcal{E}} c^*(X^* | E^*) \cdot c_\lambda(E) \\
&\quad (\text{by (C1) and (C2)}) \\
&= \sum_{E \in \mathcal{E}} c(X|E) \cdot c_\lambda(E) \\
&\quad (\text{since } c^*(X^* | E^*) = c(X|E))
\end{aligned}$$

□

More informally, Proposition 1 tells us that, given a fixed extended credal state $(\mathcal{W}^*, \mathcal{F}^*, c^*)$ which satisfies both (C1) and (C2), updating by EBCondi in that space gives the same result as if the agent were updated her credences by JCondi in the smaller credal state $(\mathcal{W}, \mathcal{F}, c)$.

Note that condition (C1) may be understood as an instance of what Bas van Fraassen (1984) called reflection principle, i.e. a principle which requires one's current credences to defer to one's future credences. After all, (C1) says that the agent's prior conditional credence in proposition $E^* \leftrightarrow E \in \mathcal{F}$ given that her posterior credences over \mathcal{E} are determined by c_λ should be equal to the posterior credence in E , $c_\lambda(E)$. The second condition (C2) tells us that once you learn that some E in the partition \mathcal{E} is true, the information about your posterior credences over \mathcal{E} should have no bearing on the truth of X .

Importantly, Proposition 1 is not the only way to establish how JCondi can be represented as Bayesian conditioning in some extended credal state. Another influential approach has been given by (Diaconis and Zabell 1982, Theorem 2.1). They have identified a necessary and sufficient condition, sometimes called *superconditioning*, for one's posterior credence to be the result of conditioning one's prior credence in some larger credal space. It says that c_λ comes from c by conditioning in the extended credal space just in case there exists a number $b \geq 0$ such that, for every $X \in \mathcal{F}$, $c_\lambda(X) \leq b \cdot c(X)$. However, contrary to Skyrms, Diaconis and Zabell's approach places no constraints on how the extended credal state should look like. It only says that we can construct an extended credal space by adding two propositions a and b to the original credal space, where a says that the agent had the learning experience she had, and b indicates its absence. For this reason, Skyrms's approach appears to be better suited for the task of providing VET in the case of uncertain evidence. For VET requires an agent to assign prior credences to propositions describing the possible outcomes of her learning experience. And in order to do this, the agent should be in a position to grasp the content of these propositions, prior to her learning experience. But a is a proposition the agent is in a position to grasp once she has already had the learning experience. Skyrms's approach appears to deliver the required element by telling us that these propositions, represented by the $R_{c_\lambda}^\mathcal{E}$'s in the extended credal state,

describe the possible Jeffrey shifts the agent might undergo, and hence their content can be grasped before the agent's learning experience.

As we have argued in the previous section, we can also think of uncertain evidence in terms of Pearl style-input. But how should the agent update her credences in response to this input? Here we propose a variant of what Pearl calls virtual conditionalization:⁸

VCondi: Given a Pearl-style input \mathcal{L}^E , an agent's posterior credence in every $X \in \mathcal{F}$ should be:

$$c_\lambda(X) = \frac{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(X \wedge E)}{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(E)}.$$

Let us look at how VCondi works in Ann's case. Suppose that Ann assigns equal credences to the proposition that the selected urn is of type X (S) and to the proposition that the selected urn is of type Y ($\neg S$). Given the experimental set-up she faces, her prior credences for the propositions B , V , $S \wedge B$, and $S \wedge V$ can be determined as follows:

- $c(B) = c(S) \cdot c(B|S) + c(\neg S) \cdot c(B|\neg S) = \frac{1}{2} \cdot \frac{8}{10} + \frac{1}{2} \cdot \frac{2}{10} = \frac{1}{2}$
- $c(V) = c(S) \cdot c(V|S) + c(\neg S) \cdot c(V|\neg S) = \frac{1}{2} \cdot \frac{2}{10} + \frac{1}{2} \cdot \frac{8}{10} = \frac{1}{2}$
- $c(S \wedge B) = c(S) \cdot c(B|S) = \frac{1}{2} \cdot \frac{8}{10} = \frac{2}{5}$
- $c(S \wedge V) = c(S) \cdot c(V|S) = \frac{1}{2} \cdot \frac{2}{10} = \frac{1}{10}$

Ann then looks at the drawn marble in dim light, takes B as her evidence and adopts the posterior credence function c_λ in response. This in turn provides her with evidence

for B and V whose differential support is settled by Ann as $\alpha_B = \frac{c(U_{c_\lambda}^B|B)}{c(U_{c_\lambda}^B|V)} = 2$ and

$\alpha_V = \frac{c(U_{c_\lambda}^B|V)}{c(U_{c_\lambda}^B|B)} = 1$. Then, VCondi prescribes the following posterior credence for S :

$$\begin{aligned} c_\lambda(S) &= \frac{\alpha_B \cdot c(S \wedge B) + \alpha_V \cdot c(S \wedge V)}{\alpha_B \cdot c(B) + \alpha_V \cdot c(V)} \\ &= \frac{2 \cdot \frac{2}{5} + 1 \cdot \frac{1}{10}}{2 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2}} \\ &= \frac{3}{5}. \end{aligned}$$

⁸ Since the likelihood ratio α_E is a special kind of the Bayes factor, our variant of VCondi can be understood as an instance of Wagner's (2002) updating which says that upon learning Bayes factors of the form $b_E = \mathbb{B}_{c_\lambda, c}(E, E') = \frac{c_\lambda(E)}{c_\lambda(E')} / \frac{c(E)}{c(E')}$, an agent's posterior credence in every $X \in \mathcal{F}$ should be:

$$c_\lambda(X) = \frac{\sum_{E \in \mathcal{E}} b_E \cdot c(X \wedge E)}{\sum_{E \in \mathcal{E}} b_E \cdot c(E)}.$$

Note also that Wagner's updating rule is a reformulation of Field's (1978) updating—yet another updating rule in the context of learning from uncertain evidence.

Similarly, VCondi shows that, when after looking at the drawn marble in a dim light, Ann updates on B , she should become more confident in B , though still less than certain. That is:

$$\begin{aligned} c_\lambda(B) &= \frac{\alpha_B \cdot c(B \wedge B) + \alpha_V \cdot c(B \wedge V)}{\alpha_B \cdot c(B) + \alpha_V \cdot c(V)} \\ &= \frac{\alpha_B \cdot c(B)}{\alpha_B \cdot c(B) + \alpha_V \cdot c(V)} \\ &\quad (\text{since } B \text{ and } V \text{ are mutually exclusive}) \\ &= \frac{2 \cdot \frac{1}{2}}{2 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2}} \\ &= \frac{2}{3}. \end{aligned}$$

Although VCondi seems to deliver intuitively correct answers in the context of learning from uncertain evidence, one might still ask if it is a rational update method to follow. Specifically, VCondi hinges on the assumption that, after learning, the agent takes her evidence to be some E from \mathcal{E} and adopts the posterior credence function c_λ in response. But why shouldn't we think that the agent adopts this posterior also for reasons other than taking E as her evidence? For example, suppose that you take E as your evidence and you believe it provides a strong support for your scientific hypothesis. But you are a scientist in a small town, and you suspect that, like most small-town scientists, you will come very soon to justifiably doubt your hypothesis. So your posterior credence in X is not only a result of taking E as your evidence, but also a result of expecting evidence against X in the near future. But this is hardly acceptable, for expecting evidence against X is not the same as possessing or taking E as evidence against X .⁹ Yet there is nothing in the update mechanism of VCondi that precludes this possibility.

But there seems to be a reasonable way to prevent the above possibility. It rests on the following condition:

CIndi: For every $E \in \mathcal{E}$ and $X \in \mathcal{F}$,

$$c(U_{c_\lambda}^E | E \wedge X) = c(U_{c_\lambda}^E | E).$$

That is, CIndi, when imposed on one's prior credences, expresses a relation of conditional independence between $U_{c_\lambda}^E$ and X given E . It says that once we suppose E , the propositions $U_{c_\lambda}^E$ and X cease to bear any information about one another. Recall the small-town scientist's example. Suppose you think that it is more likely to have a high posterior credence in X when you don't expect evidence against X than when you do. What CIndi says is that, given E , your posterior credence function which assigns a high credence to X should be insensitive to whether or not you expect evidence against X . This seems intuitively right, for you adopt your posterior credence in X in

⁹ This is a slightly modified example due to Bovens (1995).

response to taking E as your evidence which provides a strong support for X , and not in response to expecting evidence against X .

As it turns out, when we assume CIndi, an agent who updates by VCondi would satisfy an instance of reflection principle which goes as follows:

Reflection: For every $X \in \mathcal{F}$,

$$c(X|U_{c_\lambda}^E) = c_\lambda(X),$$

provided that $c(U_{c_\lambda}^E) > 0$.

That is, this principle says that your prior credence in X given your anticipated posterior credence function over \mathcal{F} which results from taking E as your evidence should be equal to your anticipated posterior credence in X .¹⁰ Note that the proposition $U_{c_\lambda}^E$ does not specify the mechanism by which you arrive at the posterior credence function in response to taking E as your evidence. It thus describes your anticipated posterior credence function after the black-box learning event, where the content of your learning experience is left opaque (see Huttegger 2017, ch. 5.4). For your learning experience does not tell you which proposition in \mathcal{E} is true.

More precisely, we can state the following proposition:

Proposition 2 *Suppose that an agent's prior credences satisfy CIndi. Then, VCondi \Leftrightarrow Reflection.*

Proof (VCondi \Rightarrow Reflection)

Suppose c satisfies CIndi and c_λ satisfies VCondi. Then:

$$\begin{aligned} c_\lambda(X) &= \frac{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(X \wedge E)}{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(E)} \\ &\quad \text{(by VCondi)} \\ &= \frac{\sum_{E \in \mathcal{E}} (l \cdot \alpha_E) \cdot c(X \wedge E)}{\sum_{E \in \mathcal{E}} (l \cdot \alpha_E) \cdot c(E)} \\ &\quad \text{(multiplying } \alpha_E \text{ by the constant } l) \\ &= \frac{\sum_{E \in \mathcal{E}} c(U_{c_\lambda}^E|E) \cdot c(X \wedge E)}{\sum_{E \in \mathcal{E}} c(U_{c_\lambda}^E|E) \cdot c(E)} \\ &\quad \text{(since } c(U_{c_\lambda}^E|E) = l \cdot \alpha_E) \\ &= \frac{\sum_{E \in \mathcal{E}} c(U_{c_\lambda}^E|E \wedge X) \cdot c(X \wedge E)}{\sum_{E \in \mathcal{E}} c(U_{c_\lambda}^E|E) \cdot c(E)} \\ &\quad \text{(by CIndi)} \\ &= \frac{\sum_{E \in \mathcal{E}} c(X \wedge U_{c_\lambda}^E \wedge E)}{\sum_{E \in \mathcal{E}} c(U_{c_\lambda}^E \wedge E)} \end{aligned}$$

¹⁰ A similar version of Reflection has been introduced by Gallow (2019).

$$\begin{aligned}
&= \frac{c(X \wedge U_{c_\lambda}^E)}{c(U_{c_\lambda}^E)} \\
&= c(X|U_{c_\lambda}^E).
\end{aligned}$$

(Reflection \Rightarrow VCondi)

Suppose c satisfies CIndi and c_λ satisfies Reflection. Then:

$$\begin{aligned}
c_\lambda(X) &= c(X|U_{c_\lambda}^E) \\
&\text{(by Reflection)} \\
&= \frac{c(X \wedge U_{c_\lambda}^E)}{c(U_{c_\lambda}^E)} \\
&= \frac{\sum_{E \in \mathcal{E}} c(X \wedge U_{c_\lambda}^E \wedge E)}{\sum_{E \in \mathcal{E}} c(U_{c_\lambda}^E \wedge E)} \\
&= \frac{\sum_{E \in \mathcal{E}} c(U_{c_\lambda}^E|E \wedge X) \cdot c(X \wedge E)}{\sum_{E \in \mathcal{E}} c(U_{c_\lambda}^E|E) \cdot c(E)} \\
&= \frac{\sum_{E \in \mathcal{E}} c(U_{c_\lambda}^E|E) \cdot c(X \wedge E)}{\sum_{E \in \mathcal{E}} c(U_{c_\lambda}^E|E) \cdot c(E)} \\
&\text{(by CIndi)} \\
&= \frac{\sum_{E \in \mathcal{E}} (l \cdot \alpha_E) \cdot c(X \wedge E)}{\sum_{E \in \mathcal{E}} (l \cdot \alpha_E) \cdot c(E)} \\
&\text{(since } c(U_{c_\lambda}^E|E) = l \cdot \alpha_E) \\
&= \frac{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(X \wedge E)}{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(E)}.
\end{aligned}$$

It is important to emphasize that the above result may be understood as providing a guidance to rational updating in the context of uncertain evidence. For it is often argued that reflection principles regulate rational learning in the sense that an agent cannot violate some instance of reflection principle and at the same time think that she will form her posterior credences in a rational way (see, e.g. Skyrms 1990, 1997; Huttigger 2013, 2017). For example, when you update your credences in cases involving forgetting or memory loss, you often violate reflection principles and hence you do not form your posterior credences by rational learning.¹¹ If so, then the above result tells us how learning uncertain evidence can be deemed rational. After all, it implies that when VCondi is combined with CIndi, the agent's credences satisfy an instance of reflection principle, and hence are mandated by a rational learning process.

¹¹ For an excellent survey of these cases, see Briggs (2009).

4 Graves's extension of VET

In this section, we present a variant of Graves's extension of VET. Before doing so, however, let us highlight a few assumptions of this result.

First, we assume that an agent faces a decision problem $\langle \mathcal{A}, \mathcal{S}, c, u \rangle$, in which \mathcal{A} is a finite partition of propositions A_1, \dots, A_n representing the actions available to her, \mathcal{S} is a finite set of propositions S_1, \dots, S_n representing the possible states of the world upon which the consequences of the actions depend, c is the agent's credence function in her credal state $(\mathcal{W}, \mathcal{F}, c)$, and u is the agent's utility function which assigns to propositions of the form $A \wedge S$ a number that measures the pragmatic utility that would result for the agent were the act A to be performed in state S .¹²

Second, following Graves, we assume that before undergoing a learning experience, the agent contemplates a finite set \mathcal{R} of possible Jeffrey shifts. And each of these possible Jeffrey shifts can be represented as proposition $R_{c_\lambda}^\mathcal{E}$ that specifies the agent's assignment of posterior credences to the members of \mathcal{E} . Thus, what is crucial to Graves's argument is that there exists an extension $(\mathcal{W}^*, \mathcal{F}^*, c^*)$ of $(\mathcal{W}, \mathcal{F}, c)$ such that:

(G₁) For every $S \in \mathcal{F}$, $c^*(S^*) = c(S)$.

(G₂) For some $\mathcal{E} \subseteq \mathcal{F}$ and a finite set of posterior credence functions \mathcal{C}_λ , $\mathcal{R} \subseteq \mathcal{W}^*$.

(G₃) c^* satisfies conditions (C1) and (C2).

With the above in mind, observe first that, prior to undergoing a Jeffrey shift, the agent would choose the prior Bayes act, i.e. the act A that maximizes

$$\mathbb{E}_c[u(A)] = \sum_{S \in \mathcal{S}} c(S) \cdot u(A \wedge S). \quad (3)$$

By (G1), (3) is equivalent to¹³

$$\mathbb{E}_{c^*}[u(A)] = \sum_{S^* \in \mathcal{S}^*} c^*(S^*) \cdot u(A \wedge S^*). \quad (4)$$

After undergoing a learning experience which results in a Jeffrey shift $\mathcal{C}_\lambda^\mathbb{E}$, the agent would choose an act A to maximize

$$\mathbb{E}_{c_\lambda}[u(A)] = \sum_{S \in \mathcal{S}} \sum_{E \in \mathcal{E}} c(S|E) \cdot c_\lambda(E) \cdot u(A \wedge S). \quad (5)$$

¹² We also assume that the decision problem is framed within a Savage-style decision theory where states and acts are probabilistically independent. Moreover, we assume that (i) the states, acts and utilities are the same before and after the learning experience, (ii) the agent is an expected utility maximizer before and after the learning experience, and (iii) the states are probabilistically independent of the learning experience, though not of the particular outcomes of this learning experience. For the importance of these assumptions, see (Skyrms 1990, pp. 90–91).

¹³ Since $S^* \leftrightarrow S \in \mathcal{F}$, it is natural to assume that the utility function for propositions $A \wedge S^*$ in the extended credal space is the same as the utility function for propositions $A \wedge S$ in the original smaller credal space. Thanks to an anonymous referee for urging us to clarify this point.

That is, she would choose the act A that maximizes expected pragmatic utility calculated relative to the agent's posterior credence function, c_λ , recommended by JCondi. By Proposition 1, (5) is equivalent to

$$\mathbb{E}_{c_\lambda^*} [u(A) | R_{c_\lambda}^\mathcal{E}] = \sum_{S^* \in \mathcal{S}^*} c^*(S^* | R_{c_\lambda}^\mathcal{E}) \cdot u(A \wedge S^*). \quad (6)$$

Since we want to know *ex ante* whether updating by JCondi is always helpful in making practical decisions, we need to determine an expectation of $\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda^*} [u(A) | R_{c_\lambda}^\mathcal{E}]$, which is the maximal value of making a choice after undergoing a Jeffrey shift (or the posterior Bayes act). Because it is assumed that the agent considers *ex ante* a finite set of possible Jeffrey shifts, and these shifts can be represented by propositions $R_{c_\lambda}^\mathcal{E}$ from the partition \mathcal{R} , this can be achieved by weighing the posterior value of Bayes act by the prior credence of $R_{c_\lambda}^\mathcal{E}$. Hence, the expectation of posterior Bayes act is:

$$\mathbb{E}_{c^*} \left[\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda^*} [u(A) | R_{c_\lambda}^\mathcal{E}] \right] = \sum_{R_{c_\lambda}^\mathcal{E} \in \mathcal{R}} c^*(R_{c_\lambda}^\mathcal{E}) \cdot \max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda^*} [u(A) | R_{c_\lambda}^\mathcal{E}]. \quad (7)$$

Note that, by the equivalence of (5) and (6), $\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda^*} [u(A) | R_{c_\lambda}^\mathcal{E}] = \max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A)]$, and hence (7) is equivalent to

$$\mathbb{E}_{c^*} \left[\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A)] \right] = \sum_{R_{c_\lambda}^\mathcal{E} \in \mathcal{R}} c^*(R_{c_\lambda}^\mathcal{E}) \cdot \max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A)]. \quad (8)$$

Now let us introduce the quantity $\Delta_{c_\lambda^*}^u(A_R, A_{\max})$, which is the difference between the maximizer of (6), A_R , and the maximizer of (4), A_{\max} , as assessed by the agent's posterior credence $c_\lambda^*(S^*) = c^*(S^* | R_{c_\lambda}^\mathcal{E})$ for every $S^* \in \mathcal{F}^*$:

$$\Delta_{c_\lambda^*}^u(A_R, A_{\max}) = \sum_{S^* \in \mathcal{S}^*} c^*(S^* | R_{c_\lambda}^\mathcal{E}) \cdot u(A_R \wedge S^*) - \sum_{S^* \in \mathcal{S}^*} c^*(S^* | R_{c_\lambda}^\mathcal{E}) \cdot u(A_{\max} \wedge S^*). \quad (9)$$

Importantly, notice that $\Delta_{c_\lambda^*}^u(A_R, A_{\max}) \geq 0$, for if $A_{\max} = A_R$, then $\Delta_{c_\lambda^*}^u(A_R, A_{\max}) = 0$, and if $A_{\max} \neq A_R$, then $\Delta_{c_\lambda^*}^u(A_R, A_{\max}) > 0$.

We can also determine the expectation of $\Delta_{c_\lambda^*}^u(A_R, A_{\max})$ relative to the agent's extended prior function c^* over \mathcal{R} :

$$\mathbb{E}_{c^*} [\Delta_{c_\lambda^*}^u(A_R, A_{\max})] = \sum_{R_{c_\lambda}^\mathcal{E} \in \mathcal{R}} c^*(R_{c_\lambda}^\mathcal{E}) \cdot \Delta_{c_\lambda^*}^u(A_R, A_{\max}). \quad (10)$$

Given the above notions, it is now not difficult to show that the maximal value of choosing now cannot be greater than the expected value of choosing after updating by JCondi. To begin with, observe that since $\Delta_{c_\lambda^*}^u(A_R, A_{\max}) \geq 0$, its expected value is

also non-negative, and thus:

$$0 \leq \mathbb{E}_{c^*} \left[\Delta_{c_\lambda}^u (A_R, A_{\max}) \right] = \sum_{R_{c_\lambda}^\mathcal{E} \in \mathcal{R}} c^* (R_{c_\lambda}^\mathcal{E}) \cdot \Delta_{c_\lambda}^u (A_R, A_{\max}). \quad (11)$$

Then, by the definition of (9), we get

$$0 \leq \sum_{R_{c_\lambda}^\mathcal{E} \in \mathcal{R}} c^* (R_{c_\lambda}^\mathcal{E}) \left(\sum_{S^* \in \mathcal{S}^*} c^* (S^* | R_{c_\lambda}^\mathcal{E}) \cdot u(A_R \wedge S^*) - \sum_{S^* \in \mathcal{S}^*} c^* (S^* | R_{c_\lambda}^\mathcal{E}) \cdot u(A_{\max} \wedge S^*) \right). \quad (12)$$

By using (6), (12) reduces to

$$0 \leq \sum_{R_{c_\lambda}^\mathcal{E} \in \mathcal{R}} c^* (R_{c_\lambda}^\mathcal{E}) \left(\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda^*} [u(A) | R_{c_\lambda}^\mathcal{E}] - \sum_{S^* \in \mathcal{S}^*} c^* (S^* | R_{c_\lambda}^\mathcal{E}) \cdot u(A_{\max} \wedge S^*) \right). \quad (13)$$

Then, we get

$$0 \leq \sum_{R_{c_\lambda}^\mathcal{E} \in \mathcal{R}} c^* (R_{c_\lambda}^\mathcal{E}) \max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda^*} [u(A) | R_{c_\lambda}^\mathcal{E}] - \sum_{S^* \in \mathcal{S}^*} \sum_{R_{c_\lambda}^\mathcal{E} \in \mathcal{R}} c^* (R_{c_\lambda}^\mathcal{E}) \cdot c^* (S^* | R_{c_\lambda}^\mathcal{E}) \cdot u(A_{\max} \wedge S^*). \quad (14)$$

By the law of total probability we have $c^*(S^*) = \sum_{R_{c_\lambda}^\mathcal{E} \in \mathcal{R}} c^*(R_{c_\lambda}^\mathcal{E}) \cdot c^*(S^* | R_{c_\lambda}^\mathcal{E})$, and hence (14) reduces to

$$0 \leq \sum_{R_{c_\lambda}^\mathcal{E} \in \mathcal{R}} c^* (R_{c_\lambda}^\mathcal{E}) \max_{A \in \mathcal{A}} \mathbb{E} [u(A) | R_{c_\lambda}^\mathcal{E}] - \sum_{S^* \in \mathcal{S}^*} c^* (S^*) \cdot u(A_{\max} \wedge S^*). \quad (15)$$

And by (4) and (7), we get

$$0 \leq \mathbb{E}_{c^*} \left[\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda^*} [u(A) | R_{c_\lambda}^\mathcal{E}] \right] - \max_{A \in \mathcal{A}} \mathbb{E}_{c^*} [u(A)]. \quad (16)$$

Hence, we have

$$\max_{A \in \mathcal{A}} \mathbb{E}_{c^*} [u(A)] \leq \mathbb{E}_{c^*} \left[\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda^*} [u(A) | R_{c_\lambda}^\mathcal{E}] \right]. \quad (17)$$

And since (3) is equivalent to (4) and (5) is equivalent to (6), we get

$$\max_{A \in \mathcal{A}} \mathbb{E}_c [u(A)] \leq \mathbb{E}_{c^*} \left[\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A)] \right], \quad (18)$$

as required. Thus, we have shown that if we utilize an extended credal state in which c^* satisfies conditions (C1) and (C2), then the agent's update in response to uncertain evidence recommended by JCondi won't result in making foreseeably harmful practical decisions.

However, Graves's extension of VET seems problematic for at least two reasons. First, one might claim that it imposes unrealistic cognitive demands on decision-makers. That is, it requires that they must have *ex ante* prior credences over all the possible Jeffrey shifts they could undergo in the future. But it is doubtful that, in Ann's case, before looking at the drawn marble, she is aware of all the possible ways by which her visual experience can prompt changes in her initial credences. For a Jeffrey shift is understood as the agent's probabilistic *judgement* over \mathcal{E} or what she takes away from experience rather than what experience delivers to her. And it is perfectly possible that, when undergoing her learning experience, Ann would actually judge probabilistically in a way that she has not considered prior to her learning experience. But how fatal is this objection?

In response, one may argue that this is just an instance of a more general problem, which says that it is implausible that each evidence you might learn is one about which you had a prior credence. After all, it is tempting to think that you might learn a proposition that you do not already grasp. But if this is so, then this problem also affects the original VET which shows that updating by BCondi maximizes expected pragmatic utility. After all, it assumes that each proposition that you might conditionalize upon is already contained in the partition \mathcal{E} over which you have prior credences. Hence, the cognitive demands imposed by Graves's extension of VET are no more unrealistic than those imposed by the original VET.

Still, however, we think that this problem faces Graves's approach more severely. For, as it is often claimed, JCondi does not even tell us what partition of propositions *should be* affected by a given learning experience (Christensen 1992; Weisberg 2009). So, prior to her learning experience, the agent might not be able to even identify correctly the partition of propositions over which she would actually undergone a Jeffrey shift. That is, the agent might contemplate *ex ante* various Jeffrey shifts over partition \mathcal{E} when in fact her learning experience would directly affect an entirely different partition \mathcal{E}' . For example, although in Ann's case, she may stipulate *ex ante* that her learning experience would directly affect the partition $\{B, V\}$, there is no normative guidance as to how she could prevent the possibility that it would actually affect the partition $\{B \wedge L, V \wedge L\}$, where L is the proposition that the lighting is dim. Therefore, by assuming that, prior to a learning experience, the agent is always in a position to grasp correctly the set \mathcal{E} that would be directly affected by her anticipated learning experience, Graves's approach imposes far more unrealistic cognitive demands than the original VET.

The second problem involves accuracy considerations, and, as we will argue, is more threatening to Graves's approach than the first one. If we grant that the only goal of

updating our credences is to improve our practical decisions, i.e. to maximize expected pragmatic utility, then Graves's approach seems successful. But many philosophers claim that, alongside practical goals, we should update our credences in a way that could also be regarded as epistemically rational. In particular, some argue that our updated credences should minimize expected inaccuracy, i.e. they should be as close as possible to the truth, from the standpoint of our prior credences (e.g. Greaves and Wallace 2006; Leitgeb and Pettigrew 2010). If this is the case, then we will argue that Graves's extension of VET cannot guarantee that updating on uncertain evidence would be both practically and epistemically rational. Since this argument requires more scrutiny, we devote the next section to it.

5 Graves's extension of VET and accuracy

In the previous section, we have seen that if the agent obeys JCondi, then updating on uncertain evidence would never lead, in expectation, to worse *practical* decisions. In this section, we argue that this approach is in tension with a purely *epistemic* approach which seeks to justify the agent's dynamic norms for credences as a consequence of the rational pursuit of accuracy. That is, if we apply the machinery of accuracy-first epistemology, then we must conclude that there appears to be no acceptable way to show that JCondi minimizes expected inaccuracy. Therefore, for an expected inaccuracy minimizer, Graves's extension of VET cannot be satisfactory, since it establishes VET for JCondi and this update rule is hardly justifiable, if ever, in terms of the rational pursuit of accuracy.

To make this point more precise, we assume, following accuracy-firsters, that we have a local epistemic disutility function (or local inaccuracy measure) for each proposition X , \mathfrak{s}_X , which takes X 's truth-value at w , $w(X)$, and the credence $c(X)$ and returns the local epistemic disutility (or inaccuracy) of having that credence in X at a world in which X 's truth-value is $w(X)$, $\mathfrak{s}_X(w(X), c(X))$. There are some desirable properties that this function should have, and they single out the class of strictly proper scoring rules. These properties are the following:

Extensionality: \mathfrak{s}_X is extensional if it can be thought of as two functions: $\mathfrak{s}_X(1, x)$ which gives the local inaccuracy of having credence x in X when X is true, and $\mathfrak{s}_X(0, x)$ which gives the local inaccuracy of having credence x when X is false.

Strict Propriety: \mathfrak{s}_X is strictly proper if for all $0 \leq p, x \leq 1$,

$$p \cdot \mathfrak{s}_X(1, p) + (1 - p) \cdot \mathfrak{s}_X(0, p) \leq p \cdot \mathfrak{s}_X(1, x) + (1 - p) \cdot \mathfrak{s}_X(0, x),$$

with equality iff $x = p$.

Continuity: \mathfrak{s}_X is a continuous function of x on $[0, 1]$.

More informally, Extensionality says that the inaccuracy of having credence x in proposition X depends only on whether X is true or false. Strict Propriety tells us that

an agent with probabilistic credence p in X proposition expects only that credence to have the lowest inaccuracy. And Continuity says that the inaccuracy of having credence x in X varies continuously with that credence.

A standard expected inaccuracy-minimization argument for an update rule says that, before the evidence is in, an agent should expect to have less inaccurate posterior credences recommended by that rule than by any other update rule. To establish such an argument, accuracy firsters say that we should evaluate our posterior credences by looking at their expected inaccuracies, where the expectation is taken relative to a prior credence function. Can we establish such an argument in the case where your posterior credence in X is recommended by JCondi?

As shown by Leitgeb and Pettigrew (2010), the answer is negative. Their point is that, under the quadratic scoring rule, one's posterior credence in any $X \in \mathcal{F}$ that results from JCondi does not minimize expected local inaccuracy given by

$$\mathbb{E}_c [\mathfrak{s}_X (w(X), c_\lambda(X))] = \sum_{w \in \mathcal{W}} c(\{w\}) \cdot \mathfrak{s}_X (w(X), c_\lambda(X)). \quad (19)$$

That is, if our goal is to choose a posterior credence function c_λ that assigns credence $c_\lambda(E)$ for all $E \in \mathcal{E}$ (i.e. satisfies the constraints imposed by a Jeffrey shift) and is minimal with respect to the expected local inaccuracy of the credence it assigns to each $X \in \mathcal{F}$ by the lights of one's prior credence function c over the set of possible worlds \mathcal{W} , then this cannot be achieved by selecting the posterior credence function that results from JCondi. A crucial assumption of this result is that the local inaccuracy is measured by the quadratic scoring rule, $q_X = (i - x)^2$, where $i = 1$ or 0 . That is, this rule gives (i) the squared difference $(1 - x)^2$ between the credence x in X and the value $w(X) = 1$ of the indicator function of w ($w(X) = 1$ if X is true at w), and (ii) the squared difference $(0 - x)^2$ between the credence x in X and the value $w(X) = 0$ of the indicator function of w ($w(X) = 0$ if X is false at w). Importantly, q_X is a strictly proper scoring rule.

To make Leitgeb and Pettigrew's point more concrete, let us first introduce their alternative update rule:

LPCondi: Given a Jeffrey shift $\mathcal{C}_\lambda^\mathbb{E}$, let d_E be the unique real number such that

$$\sum_{\{w \in E : c(\{w\}) + d_E > 0\}} c(\{w\}) + d_E = c_\lambda(E).$$

Then, the agent's posterior credence function c_λ should be such that, for $w \in E$,

$$c_\lambda(\{w\}) = \begin{cases} c(\{w\}) + d_E, & \text{for } c(\{w\}) + d_E > 0 \\ 0, & \text{for } c(\{w\}) + d_E \leq 0. \end{cases}$$

Now, suppose that initially Ann does not rule out any of the following possible worlds:

- w_1 , in which the selected urn is X and the drawn marble is blue.
- w_2 , in which the selected urn is X and the drawn marble is violet.

- w_3 , in which the selected urn is Y and the drawn marble is blue.
- w_4 , in which the selected urn is Y and the drawn marble is violet.

More precisely, she assigns the following prior credences to these worlds:

$$c(\{w_1\}) = \frac{1}{4}, c(\{w_2\}) = \frac{1}{2}, c(\{w_3\}) = \frac{1}{8}, c(\{w_4\}) = \frac{1}{8}.$$

Ann, then, looks at the drawn marble in dim lighting which results in the following Jeffrey shift over the partition $\{B, V\}$:

- $c_\lambda(B) = c_\lambda(\{w_1, w_3\}) = \frac{1}{2}$.
- $c_\lambda(V) = c_\lambda(\{w_2, w_4\}) = \frac{1}{2}$.

In response to this Jeffrey shift, Ann applies JCondi and gets the following posterior credences:

- $c_\lambda(\{w_1\}) = c_\lambda(\{w_1, w_3\}) \cdot \frac{c(\{w_1\} \cap \{w_1, w_3\})}{c(\{w_1, w_3\})} + c_\lambda(\{w_2, w_4\}) \cdot \frac{c(\{w_1\} \cap \{w_2, w_4\})}{c(\{w_2, w_4\})} = \frac{1}{3}$.
- $c_\lambda(\{w_2\}) = c_\lambda(\{w_1, w_3\}) \cdot \frac{c(\{w_2\} \cap \{w_1, w_3\})}{c(\{w_1, w_3\})} + c_\lambda(\{w_2, w_4\}) \cdot \frac{c(\{w_2\} \cap \{w_2, w_4\})}{c(\{w_2, w_4\})} = \frac{2}{5}$.
- $c_\lambda(\{w_3\}) = c_\lambda(\{w_1, w_3\}) \cdot \frac{c(\{w_3\} \cap \{w_1, w_3\})}{c(\{w_1, w_3\})} + c_\lambda(\{w_2, w_4\}) \cdot \frac{c(\{w_3\} \cap \{w_2, w_4\})}{c(\{w_2, w_4\})} = \frac{1}{6}$.
- $c_\lambda(\{w_4\}) = c_\lambda(\{w_1, w_3\}) \cdot \frac{c(\{w_4\} \cap \{w_1, w_3\})}{c(\{w_1, w_3\})} + c_\lambda(\{w_2, w_4\}) \cdot \frac{c(\{w_4\} \cap \{w_2, w_4\})}{c(\{w_2, w_4\})} = \frac{1}{10}$.

Let us then apply the quadratic scoring rule, $q_{\{w_1\}}$, as our local inaccuracy measure in order to determine the local expected inaccuracy of Ann's posterior credence in proposition $\{w_1\}$ that results from JCondi, $\mathbb{E}_c[q_{\{w_1\}}(w(\{w_1\}), c_\lambda(\{w_1\}))]$. If we let $x = c_\lambda(\{w_1\})$, then we get:

$$\begin{aligned} \mathbb{E}_c[q_{\{w_1\}}(w(\{w_1\}), c_\lambda(\{w_1\}))] &= \sum_{w \in \mathcal{W}} c(\{w\}) \cdot q_{\{w_1\}}(w(\{w_1\}), c_\lambda(\{w_1\})) \\ &= \frac{1}{4} \left(1 - \frac{1}{3}\right)^2 + \frac{1}{2} \left(-\frac{1}{3}\right)^2 + \frac{1}{8} \left(-\frac{1}{3}\right)^2 + \frac{1}{8} \left(-\frac{1}{3}\right)^2 \\ &= \frac{7}{36} = 0,194. \end{aligned}$$

Now, suppose that, instead of JCondi, Ann uses LPCondi in response to the above Jeffrey shift. If we let $d_E = \frac{(c_\lambda(E) - c(E))}{|E|}$, then Ann's posterior credences recommended by LPCondi are the following:

- $c'_\lambda(\{w_1\}) = c(\{w_1\}) + \frac{(c_\lambda(B) - c(B))}{|B|} = \frac{1}{4} + \frac{(\frac{1}{2} - \frac{3}{8})}{2} = \frac{5}{16}$.
- $c'_\lambda(\{w_2\}) = c(\{w_2\}) + \frac{(c_\lambda(V) - c(V))}{|V|} = \frac{1}{2} + \frac{(\frac{1}{2} - \frac{5}{8})}{2} = \frac{7}{16}$.
- $c'_\lambda(\{w_3\}) = c(\{w_3\}) + \frac{(c_\lambda(B) - c(B))}{|B|} = \frac{1}{8} + \frac{(\frac{1}{2} - \frac{3}{8})}{2} = \frac{3}{16}$.
- $c'_\lambda(\{w_4\}) = c(\{w_4\}) + \frac{(c_\lambda(V) - c(V))}{|V|} = \frac{1}{8} + \frac{(\frac{1}{2} - \frac{5}{8})}{2} = \frac{1}{16}$.

Then, if we apply the quadratic scoring rule, $q_{\{w_1\}}$, and let $x = c'_\lambda(\{w_1\})$, then we get the following local expected inaccuracy of Ann's posterior credence in $\{w_1\}$ that

results from LPCondi:

$$\begin{aligned}\mathbb{E}_c [\mathfrak{q}_{\{w_1\}} (w(\{w_1\}), c'_\lambda(\{w_1\}))] &= \sum_{w \in \mathcal{W}} c(\{w\}) \cdot \mathfrak{q}_{\{w_1\}} (w(\{w_1\}), c'_\lambda(\{w_1\})) \\ &= \frac{1}{4} \left(1 - \frac{5}{16}\right)^2 + \frac{1}{2} \left(-\frac{5}{16}\right)^2 + \frac{1}{8} \left(-\frac{5}{16}\right)^2 + \frac{1}{8} \left(-\frac{5}{16}\right)^2 \\ &= \frac{49}{256} \approx 0,191.\end{aligned}$$

But since $\mathbb{E}_c [\mathfrak{q}_{\{w_1\}} (w(\{w_1\}), c'_\lambda(\{w_1\}))] < \mathbb{E}_c [\mathfrak{q}_{\{w_1\}} (w(\{w_1\}), c_\lambda(\{w_1\}))]$, it follows that Ann's posterior credence in $\{w_1\}$ determined by JCondi does not minimize the local expected inaccuracy.

Moreover, as shown by Leitgeb and Pettigrew (2010), the situation is not better when instead of focusing on the expected local inaccuracy, we focus on the expected global inaccuracy. Firstly, given a strictly proper scoring rule (or local inaccuracy measure) for proposition X , \mathfrak{s}_X , we may define a global inaccuracy measure for the credence function c as follows:

$$\mathfrak{J}_\mathfrak{s}(c, w) = \sum_{X \in \mathcal{F}} \mathfrak{s}_X(w(X), c(X)). \quad (20)$$

If we define $\mathfrak{J}_\mathfrak{s}$ in this way, we may say that the global inaccuracy measure is generated from \mathfrak{s} . As it is easy to see, $\mathfrak{J}_\mathfrak{s}$ is an additive inaccuracy measure, for the inaccuracy of c is the sum of the inaccuracies of individual credences that c assigns to the propositions in \mathcal{F} . Given an additive and strictly proper¹⁴ $\mathfrak{J}_\mathfrak{s}$, we can define the expected global inaccuracy of c_λ from the standpoint of c as follows:

$$\mathbb{E}_c [\mathfrak{J}_\mathfrak{s}(c_\lambda, w)] = \sum_{w \in \mathcal{W}} c(\{w\}) \cdot \mathfrak{J}_\mathfrak{s}(c_\lambda, w). \quad (21)$$

Specifically, Leitgeb and Pettigrew adopt the Brier score, i.e. a global inaccuracy measure generated from the quadratic scoring rule \mathfrak{q}_X :

$$\mathfrak{B}(c, w) = \sum_{X \in \mathcal{F}} \mathfrak{q}_X(w(X), c(X)). \quad (22)$$

Secondly, suppose that you are governed by the following norm: you should adopt the posterior credence function that satisfies constraints given by the $c_\lambda(E)$'s for all $E \in \mathcal{E}$, and is minimal amongst the posterior credence functions thus constrained with respect to expected global inaccuracy given by

$$\mathbb{E}_c [\mathfrak{B}(c_\lambda, w)] = \sum_{w \in \mathcal{W}} c(\{w\}) \cdot \mathfrak{B}(c_\lambda, w). \quad (23)$$

¹⁴ $\mathfrak{J}_\mathfrak{s}$ is strictly proper in the sense that, for any probabilistic credence function p and every credence function $c \neq p$, $\mathbb{E}_p [\mathfrak{J}_\mathfrak{s}(p, w)] = \sum_{w \in \mathcal{W}} p(\{w\}) \cdot \mathfrak{J}_\mathfrak{s}(p, w) < \mathbb{E}_p [\mathfrak{J}_\mathfrak{s}(c, w)] = \sum_{w \in \mathcal{W}} p(\{w\}) \cdot \mathfrak{J}_\mathfrak{s}(c, w)$.

Then, Leitgeb and Pettigrew show that the posterior credence that satisfies the above norm does not result from JCondi, but is mandated by LPCondi. So, again, if your goal is to minimize expected global inaccuracy, you should not update by JCondi.

Is there any way to salvage the idea of inaccuracy minimization for the case of JCondi? Levinstein (2012) has suggested that this could be achieved if we replace the Brier score with the following logarithmic global inaccuracy measure:

$$\mathfrak{L}(c, w) = -\ln c(\{w\}). \quad (24)$$

More precisely, \mathfrak{L} takes the inaccuracy of a credence function c at world w to be the negative of the natural logarithm, \ln , of the credence it assigns to w . Armed with the inaccuracy measure \mathfrak{L} , Levinstein has shown that one's probabilistic posterior credence function recommended by JCondi satisfies constraints given by the $c_\lambda(E)$'s for all $E \in \mathcal{E}$ and minimizes expected global inaccuracy given by

$$\mathbb{E}_c[\mathfrak{L}(c_\lambda, w)] = \sum_{w \in \mathcal{W}} c(\{w\}) \cdot \mathfrak{L}(c_\lambda, w). \quad (25)$$

However, this accuracy-based vindication of JCondi has a number of complications. First, \mathfrak{L} is not generated from a strictly proper scoring rule, and hence is not itself strictly proper. But without a strictly proper scoring rule, neither an expected-accuracy argument for BCondi (Greaves and Wallace 2006) nor an accuracy-dominance argument for BCondi (Briggs and Pettigrew 2020) can be established. Thus, if we adopt \mathfrak{L} , we rule out at least two well-trodden accuracy-based justifications of the most popular updating rule, BCondi. Second, \mathfrak{L} is not additive, for it only considers credences assigned to singleton propositions $\{w\}$, and says nothing about credences in the more coarse-grained propositions in \mathcal{F} . Consequently, this measure cannot distinguish between probabilistic and non-probabilistic credence functions, and hence cannot be used to establish the accuracy-dominance argument for probabilism, i.e. the norm which says that one's credences ought to be probabilities.¹⁵ Thus, it appears that \mathfrak{L} is hardly defensible.

The above considerations show that an expected inaccuracy minimizer cannot accept JCondi as an updating rule, since it does not minimize expected inaccuracy under the class of strictly proper scoring rules, and even if it does so under the inaccuracy measure \mathfrak{L} , this accuracy-based justification of JCondi is hardly acceptable. Hence, although Graves's argument establishes VET in the context of uncertain evidence, it does so for JCondi, and so cannot be accepted by the expected inaccuracy minimizer. If we wish to show that gathering uncertain evidence is both practically and epistemically rational in expectation, we should consider some alternatives to Graves's approach.

¹⁵ The accuracy-dominance argument for probabilism says that every non-probabilistic credence function is strongly accuracy-dominated by a probabilistic credence function, while no probabilistic credence function is even weakly accuracy-dominated by any other credence function. It thus shows that probabilism is necessary and sufficient for avoiding accuracy-dominance. For a detailed presentation of this argument, see Predd et al. (2009).

6 An alternative approach

In this section, we present an alternative account of the idea that learning from uncertain evidence cannot lead an agent to expect to make worse practical decisions. Our approach utilizes the following assumptions. First, we assume that the agent assigns prior credences to propositions $U_{c_\lambda}^E$ from the finite set \mathcal{U} , which is a partition of \mathcal{W} . And since \mathcal{U} is a partition, exactly one proposition $U_{c_\lambda}^E$ will be true after the agent's learning experience. Importantly, the finite set \mathcal{U} is already included in \mathcal{W} in the small space $(\mathcal{W}, \mathcal{F}, c)$. Hence, contrary to Graves's approach, we do not need to invoke an extended credal state $(\mathcal{W}^*, \mathcal{F}^*, c^*)$ in order to account for the practical value of updating on uncertain evidence. Recall that, in Graves's approach, if the propositions $R_{c_\lambda}^E$ were already included in the small space $(\mathcal{W}, \mathcal{F}, c)$, then, as we have argued, there would be no need to use JCondi in the first place, and so Graves's approach would in fact establish VET in the case where one updates by conditionalizing on the proposition $R_{c_\lambda}^E$. But since we do not assume that the agent updates by JCondi, and hence that no proposition in $(\mathcal{W}, \mathcal{F}, c)$ that describes your experience can be learned for certain, we take it to be plausible that \mathcal{U} is already included in \mathcal{W} and the agent learns exactly one proposition $U_{c_\lambda}^E$ for certain.

Second, our approach assumes that when experience results in the proposition $U_{c_\lambda}^E$ being true, the agent settles on the likelihood ratios for $U_{c_\lambda}^E$, and updates her credences by dint of VCondi. Moreover, in order to ensure that the agent adopts the posterior credence c_λ in response to evidence E alone, we assume that her prior credences obey CIndi. Thus, we are assured that updating by dint of VCondi is a rational learning process.

With the above assumptions in mind, let us now show how VET can be established for the agent who faces a decision problem $\langle \mathcal{A}, \mathcal{S}, c, u \rangle$ and updates, in response to uncertain evidence, by VCondi. Before receiving uncertain evidence, the agent would choose the act A which maximizes:

$$\mathbb{E}_c [u(A)] = \sum_{S \in \mathcal{S}} c(S) \cdot u(A \wedge S). \quad (26)$$

After undergoing a learning experience which provides a set of likelihood ratios for some $U_{c_\lambda}^E$ in \mathcal{U} , the agent would choose the posterior Bayes act, i.e. an act A which maximizes

$$\mathbb{E}_{c_\lambda} [u(A) | U_{c_\lambda}^E] = \sum_{S \in \mathcal{S}} \frac{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(S \wedge E)}{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(E)} \cdot u(A \wedge S) \quad (27)$$

Note that the posterior Bayes act is an act that maximizes expected pragmatic utility relative to the agent's posterior credence function, c_λ , mandated by VCondi. Since \mathcal{U} is a partition, we can determine the expectation of posterior Bayes act as follows:

$$\mathbb{E}_c \left[\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A) | U_{c_\lambda}^E] \right] = \sum_{U_{c_\lambda}^E \in \mathcal{U}} c(U_{c_\lambda}^E) \cdot \max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A) | U_{c_\lambda}^E]. \quad (28)$$

Now the difference, $\Delta_{c_\lambda}^u(A_U, A_{\max})$, between the maximizer of (27), A_U , and the maximizer of (26), A_{\max} , as assessed by the agent's posterior credence $c_\lambda(S) = \frac{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(S \wedge E)}{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(E)}$, can be given as follows:

$$\begin{aligned} \Delta_{c_\lambda}^u(A_U, A_{\max}) &= \sum_{S \in \mathcal{S}} \frac{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(S \wedge E)}{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(E)} \cdot u(A_U \wedge S) \\ &\quad - \sum_{S \in \mathcal{S}} \frac{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(S \wedge E)}{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(E)} \cdot u(A_{\max} \wedge S). \end{aligned} \quad (29)$$

Note that $\Delta_{c_\lambda}^u(A_U, A_{\max}) \geq 0$, for if $A_{\max} = A_U$, then $\Delta_{c_\lambda}^u(A_U, A_{\max}) = 0$, and if $A_{\max} \neq A_U$, then $\Delta_{c_\lambda}^u(A_U, A_{\max}) > 0$. We can also calculate the expectation of $\Delta_{c_\lambda}^u(A_U, A_{\max})$ as follows:

$$\mathbb{E}_c[\Delta_{c_\lambda}^u(A_U, A_{\max})] = \sum_{U_{c_\lambda}^E \in \mathcal{U}} c(U_{c_\lambda}^E) \cdot \Delta_{c_\lambda}^u(A_U, A_{\max}). \quad (30)$$

Now in order to establish that, in expectation, updating on uncertain evidence by VCondi cannot lead you to make worse practical decisions, we will show that

$$\max_{A \in \mathcal{A}} \mathbb{E}_c[u(A)] \leq \mathbb{E}_c \left[\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A) | U_{c_\lambda}^E] \right]. \quad (31)$$

To begin with, observe first that since $\Delta_{c_\lambda}^u(A_U, A_{\max}) \geq 0$, the expectation of its value must also be non-negative, and so:

$$0 \leq \mathbb{E}_c[\Delta_{c_\lambda}^u(A_U, A_{\max})] = \sum_{U_{c_\lambda}^E \in \mathcal{U}} c(U_{c_\lambda}^E) \cdot \Delta_{c_\lambda}^u(A_U, A_{\max}). \quad (32)$$

By (29), we have

$$\begin{aligned} 0 &\leq \sum_{U_{c_\lambda}^E \in \mathcal{U}} c(U_{c_\lambda}^E) \cdot \\ &\quad \left(\sum_{S \in \mathcal{S}} \frac{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(S \wedge E)}{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(E)} \cdot u(A_U \wedge S) - \sum_{S \in \mathcal{S}} \frac{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(S \wedge E)}{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(E)} \cdot u(A_{\max} \wedge S) \right). \end{aligned} \quad (33)$$

By using (27), we get

$$0 \leq \sum_{U_{c_\lambda}^E \in \mathcal{U}} c(U_{c_\lambda}^E) \left(\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A) | U_{c_\lambda}^E] - \sum_{S \in \mathcal{S}} \frac{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(S \wedge E)}{\sum_{E \in \mathcal{E}} \alpha_E \cdot c(E)} \cdot u(A_{\max} \wedge S) \right). \quad (34)$$

Now, assuming CIndi, we get, by Proposition 2,

$$0 \leq \sum_{U_{c_\lambda}^E \in \mathcal{U}} c(U_{c_\lambda}^E) \left(\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A) | U_{c_\lambda}^E] - \sum_{S \in \mathcal{S}} c(S | U_{c_\lambda}^E) \cdot u(A_{\max} \wedge S) \right). \quad (35)$$

Then, we have

$$0 \leq \sum_{U_{c_\lambda}^E \in \mathcal{U}} c(U_{c_\lambda}^E) \max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A) | U_{c_\lambda}^E] - \sum_{S \in \mathcal{S}} c(U_{c_\lambda}^E) \cdot c(S | U_{c_\lambda}^E) \cdot u(A_{\max} \wedge S). \quad (36)$$

And, by the law of total probability, we have $c(S) = \sum_{U_{c_\lambda}^E \in \mathcal{U}} c(U_{c_\lambda}^E) \cdot c(S | U_{c_\lambda}^E)$, and hence

$$0 \leq \sum_{U_{c_\lambda}^E \in \mathcal{U}} c(U_{c_\lambda}^E) \max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A) | U_{c_\lambda}^E] - \sum_{S \in \mathcal{S}} c(S) \cdot u(A_{\max} \wedge S). \quad (37)$$

Finally, by (26) and (28), we get

$$0 \leq \mathbb{E}_c \left[\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A) | U_{c_\lambda}^E] \right] - \max_{A \in \mathcal{A}} \mathbb{E}_c [u(A)]. \quad (38)$$

Hence, we have

$$\max_{A \in \mathcal{A}} \mathbb{E}_c [u(A)] \leq \mathbb{E}_c \left[\max_{A \in \mathcal{A}} \mathbb{E}_{c_\lambda} [u(A) | U_{c_\lambda}^E] \right], \quad (39)$$

as required.

So far we have seen that, under the assumption of CIndi, expecting to update by VCondi cannot lead you to expect to make worse practical decisions. Since VCondi, when coupled with CIndi, can be regarded as a rational update mechanism in the context of uncertain evidence, our result establishes VET in that context. But could our argument be accepted by an expected inaccuracy minimizer? It turns out that so long as the inaccuracy of your credences is measured by a strictly proper scoring rule and your credences obey CIndi, your posterior credences recommended by VCondi minimize a kind of expected inaccuracy. More precisely:

Proposition 3 *Suppose that c_λ and c are probability functions and c obeys CIndi. If \mathfrak{s}_X strictly proper, extensional and continuous, then, for any $U_{c_\lambda}^E \in \mathcal{U}$ with positive probability, and any $X \in \mathcal{F}$, the agent's posterior credence recommended by VCondi minimizes the expected local inaccuracy given by*

$$\mathbb{E}_c [\mathfrak{s}_X(w(X), c_\lambda(X))] = \sum_{w \in U_{c_\lambda}^E} c(\{w\}) \cdot \mathfrak{s}_X(w(X), c_\lambda(X)). \quad (40)$$

Proof Choose an arbitrary $X \in \mathcal{F}$ and $U_{c_\lambda}^E \in \mathcal{U}$ with positive probability. In order to establish Proposition 3, we need to find $c_\lambda(X)$ which minimizes (40). Let $x = c_\lambda(X)$. Since \mathfrak{s}_X is extensional, the choice of $c_\lambda(X)$ must minimize

$$\begin{aligned} & \sum_{w \in U_{c_\lambda}^E \cap X} c(\{w\}) \cdot \mathfrak{s}_E(1, c_\lambda(X)) + \sum_{w \in U_{c_\lambda}^E \cap \neg X} c(\{w\}) \cdot \mathfrak{s}_X(0, c_\lambda(X)) \\ &= c(U_{c_\lambda}^E \cap X) \cdot \mathfrak{s}_X(1, c_\lambda(X)) + c(U_{c_\lambda}^E \cap \neg X) \cdot \mathfrak{s}_X(0, c_\lambda(X)) \end{aligned} \quad (41)$$

If a choice of $c_\lambda(X)$ minimizes (41), it will continue to minimize it if we divide it by $c(U_{c_\lambda}^E)$ (which is positive by assumption). That is, the choice of $c_\lambda(X)$ will minimize:

$$\begin{aligned} & \frac{c(U_{c_\lambda}^E \cap X)}{c(U_{c_\lambda}^E)} \cdot \mathfrak{s}_X(1, c_\lambda(X)) + \frac{c(U_{c_\lambda}^E \cap \neg X)}{c(U_{c_\lambda}^E)} \cdot \mathfrak{s}_X(0, c_\lambda(X)) \\ &= c(X|U_{c_\lambda}^E) \cdot \mathfrak{s}_X(1, c_\lambda(X)) + c(\neg X|U_{c_\lambda}^E) \cdot \mathfrak{s}_X(0, c_\lambda(X)) \end{aligned} \quad (42)$$

Since $c(\cdot|U_{c_\lambda}^E)$ is a probability function, we may write (42) as:

$$c(X|U_{c_\lambda}^E) \cdot \mathfrak{s}_X(1, c_\lambda(X)) + (1 - c(X|U_{c_\lambda}^E)) \cdot \mathfrak{s}_X(0, c_\lambda(X)) \quad (43)$$

Now, since \mathfrak{s}_X is strictly proper, $c_\lambda(X)$ minimizes (43) iff $c_\lambda(X) = c(X|U_{c_\lambda}^E)$. Thus, for any $U_{c_\lambda}^E \in \mathcal{U}$ with positive probability, $c_\lambda(X)$ minimizes (40) iff $c_\lambda(X)$ obeys Reflection. Hence, given CIndi, $c_\lambda(X)$ minimizes (40) iff $c_\lambda(X)$ satisfies VCondi.

Thus, unlike Graves's argument, our alternative extension of VET dovetails with an accuracy-centered approach to vindicating updating on uncertain evidence. Under the current proposal, gathering uncertain evidence is not only pragmatically rational in expectation. It is also epistemically rational to do so.

Another advantage of our approach is that it appears to be less cognitively demanding than Graves's extension of VET. The agent is no longer required to assign *ex ante* prior credences over propositions describing the possible Jeffrey shifts she might undergo in the future. Hence, she is also no longer required to grasp correctly *ex ante* the partition of propositions that her learning experience would affect directly. Recall that this is especially problematic for Graves's approach, since neither EBCondi nor JCondi provides a normative guidance as to how to do this. However, one might object here that our approach assumes that the agent is in a position to specify *ex ante* the anticipated posterior credence functions, which itself appears to be highly demanding. In response, we would like to note that our approach works well even if the agent anticipates one and the same posterior credence function that she would adopt in response to taking E as well as to taking E' as her evidence. That is, the propositions $U_{c_\lambda}^E$ and $U_{c_\lambda}^{E'}$ would be still mutually exclusive, even if the agent adopted the same posterior credence function c_λ , irrespective of whether she took E or E' as her evidence. This is

so because taking E as one's evidence is essentially different from taking E' as one's evidence.¹⁶

One may also worry that the use of VCondi in our approach is redundant, for if the agent learns a proposition $U_{c_\lambda}^E$ for certain, she might simply conditionalize on $U_{c_\lambda}^E$ and set her posterior credence function equal to $c(\cdot|U_{c_\lambda}^E)$. In particular, she might determine her posterior credences in every $E \in \mathcal{E}$ as follows: $c_\lambda(E) = c(E|U_{c_\lambda}^E)$. To answer this objection, we would like to stress out that VCondi appears to be a more user-friendly update rule than Bayes conditioning on $U_{c_\lambda}^E$. After all, Bayes theorem tells

us that $c(E|U_{c_\lambda}^E) = \frac{c(E) \cdot c(U_{c_\lambda}^E|E)}{c(U_{c_\lambda}^E)}$, and so requires the agent to determine the absolute likelihood for $U_{c_\lambda}^E$, $c(U_{c_\lambda}^E|E)$. And this is a hard task if one wants this quantity to be determined in a reasonable way. For example, it is hard to say what the absolute likelihood $c(U_{c_\lambda}^B|B)$ should be in Ann's case. VCondi enables us to mitigate this problem, for it allows the agent to express only the likelihood ratios of the form $\frac{c(U_{c_\lambda}^E|E)}{c(U_{c_\lambda}^E|E')}$, without the need of determining the absolute likelihoods.

7 Conclusion

The question of how to capture learning without certainties is one of the pressing problems within Bayesian epistemology. A prevalent view among Bayesian epistemologists is that JCondi is an appropriate updating mechanism in the context of learning from uncertain evidence. This rule requires uncertain evidence to be specified as a redistribution of the agent's credences over the propositions in some partition of a set of possibilities. Within a decision-theoretic context, Graves has shown that gathering uncertain evidence so understood is pragmatically rational, i.e. updating by dint of JCondi is worth waiting for in advance of making a practical decision. We have argued that Graves's approach is problematic, for it imposes highly unrealistic demands on the agent, and, more importantly, is in tension with a purely epistemic vindication of updating on uncertain evidence. In its stead, we have suggested replacing JCondi with a different updating rule in the context of uncertain evidence, VCondi. This rule requires uncertain evidence to be specified as a particular set of likelihood ratios. As we have tried to show, VCondi gets cases like Ann's right, and when combined with a fairly plausible assumption, CIndi, is equivalent to a particular sort of reflection principle. Armed with this rule, we have shown how VET can be established when one acquires uncertain evidence, and how updating on uncertain evidence minimizes expected epistemic disutility.

We want to emphasize that it was not our goal to show that VCondi is *the* correct updating rule in the context of learning from uncertain evidence. Rather, our aim was to show that the way we specify uncertain evidence matters to whether updating on uncertain evidence is both pragmatically and epistemically rational. When we specify uncertain evidence as a set of likelihood ratios and employ VCondi, we can show that

¹⁶ For a similar view, see Gallow (2019).

gathering uncertain evidence maximizes expected pragmatic utility and minimizes expected epistemic disutility. But this goal is hard to achieve, if ever, when uncertain evidence is understood as a Jeffrey shift in response to which we use JCondi.

Acknowledgements We would like to thank two anonymous reviewers for this journal for helpful comments. Thanks also to the participants at *Bayesian Epistemology: Perspectives and Challenges* in Munich and at *ECAP 10* in Utrecht for helpful discussion. Funding for this research was provided by the National Science Centre, Poland (grant no. 2017/26/D/HS1/00068).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bovens, L. (1995). 'P and I Will Believe that Not-P': Diachronic constraints on rational belief. *Mind*, 104, 737–760.
- Bradley, R. (2005). Radical probabilism and Bayesian conditioning. *Philosophy of Science*, 72(2), 342–364.
- Briggs, R. (2009). Distorted reflection. *Philosophical Review*, 118(1), 59–85.
- Briggs, R., & Pettigrew, R. (2020). An accuracy-dominance argument for conditionalization. *Nous*, 54(1), 162–181.
- Chan, H., & Darwiche, A. (2005). On the revision of probabilistic beliefs using uncertain evidence. *Artificial Intelligence*, 163(1), 67–90.
- Christensen, D. (1992). Confirmation holism and Bayesian epistemology. *Philosophy of Science*, 59(4), 540–557.
- Diaconis, P., & Zabell, S. L. (1982). Updating subjective probability. *Journal of the American Statistical Association*, 77(380), 822–830.
- Dietrich, F., List, C., & Bradley, R. (2016). Belief revision generalized: A joint characterization of Bayes's and Jeffrey's rules. *Journal of Economic Theory*, 162, 352–371.
- Field, H. (1978). A note on Jeffrey conditionalization. *Philosophy of Science*, 45(3), 361–367.
- Gallow, J. D. (2019). Updating for externalists. *Nous*. <https://doi.org/10.1111/nous.12307>.
- Good, I. (1950). *Probability and the weighing of evidence*. Charles Griffin.
- Good, I. J. (1967). On the principle of total evidence. *British Journal for the Philosophy of Science*, 17(4), 319–321.
- Graves, P. R. (1989). The total evidence theorem for probability kinematics. *Philosophy of Science*, 56(2), 317–324.
- Greaves, H., & Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(495), 607–632.
- Hawthorne, J. (2004). Three models of sequential belief updating on uncertain evidence. *Journal of Philosophical Logic*, 33(1), 89–123.
- Huttegger, S. M. (2013). In defence of reflection. *Philosophy of Science*, 80(3), 413–433.
- Huttegger, S. M. (2017). *The probabilistic foundations of rational learning*. Cambridge University Press.
- Jeffrey, R. (1983). *The logic of decision*. Chicago: University of Chicago Press.
- Jeffrey, R. (1992). *Probability and the art of judgement*. Cambridge University Press.
- Joyce, J. M. (2010). The development of subjective Bayesianism. In S. H. Dov, M. Gabbay & J. Woods (Eds.), *Handbook of the history of logic* (Vol. 10, pp. 415–476). Elsevier.
- Leitgeb, H., & Pettigrew, R. (2010). An objective justification of Bayesianism II: The consequences of minimizing inaccuracy. *Philosophy of Science*, 77(2), 236–272.
- Levinstein, B. A. (2012). Leitgeb and Pettigrew on accuracy and updating. *Philosophy of Science*, 79(3), 413–424.

- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. Morgan Kaufmann.
- Pearl, J. (1990). Jeffrey's rule, passage of experience, and neo-bayesianism. In H. E. Kyburg, R. P. Loui, & G. N. Carlson (Eds.), *Knowledge representation and defeasible reasoning* (pp. 245–265). Springer.
- Predd, J., Seiringer, R., Lieb, E., Osherson, D., Poor, H. V., & Kulkarni, S. (2009). Probabilistic coherence and proper scoring rules. *IEEE Transactions on Information Theory*, 55(10), 4786–4792.
- Raiffa, H., & Schlaifer, R. (1961). *Applied statistical decision theory*. Harvard University Press.
- Ramsey, F. P. (1990). Weight or the value of knowledge. *British Journal for the Philosophy of Science*, 41, 1–4.
- Savage, L. J. (1954). *The foundations of statistics*. Dover Publications.
- Schwan, B., & Stern, R. (2017). A causal understanding of when and when not to Jeffrey conditionalize. *Philosophers' Imprint*, 17(8), 1–21.
- Skyrms, B. (1980). Higher-order degrees of belief. In D. H. Mellor (Ed.), *Prospects for pragmatism* (pp. 109–137). Cambridge University Press.
- Skyrms, B. (1990). *The dynamics of rational deliberation*. Harvard University Press.
- Skyrms, B. (1997). The structure of radical probabilism. *Erkenntnis*, 45(2–3), 285–297.
- Sprenger, J. (2015). A novel solution to the problem of old evidence. *Philosophy of Science*, 82(3), 383–401.
- Uffink, J. (1996). The constraint rule of the maximum entropy principle. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 27(1), 47–79.
- van Fraassen, B. (1984). Belief and the will. *Journal of Philosophy*, 81(5), 235–256.
- van Fraassen, B. C. (1989). *Laws and symmetry*. Oxford University Press.
- van Fraassen, B. C., & Halpern, J. Y. (2017). Updating probability: Tracking statistics as criterion. *The British Journal for the Philosophy of Science*, 68(3), 725–743.
- Wagner, C. (2002). Probability kinematics and commutativity. *Philosophy of Science*, 69(2), 266–278.
- Wagner, C. (2009). Jeffrey conditioning and external Bayesianity. *Logic Journal of IGPL*, 18(2), 336–345.
- Weisberg, J. (2009). Commutativity or holism? A dilemma for conditionalizers. *British Journal for the Philosophy of Science*, 60, 793–812.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.