Learning From Conditionals

Benjamin Eva[*]   Stephan Hartmann[†]   Soroush Rafiee Rad[‡]

March 20, 2019

Abstract

In this article, we address a major outstanding question of probabilistic Bayesian epistemology: 'How should a rational Bayesian agent update their beliefs upon learning an indicative conditional?'. A number of authors have recently contended that this question is fundamentally underdetermined by Bayesian norms, and hence that there is no single update procedure that rational agents are obliged to follow upon learning an indicative conditional. Here, we resist this trend and argue that a core set of widely accepted Bayesian norms is sufficient to identify a normatively privileged updating procedure for this kind of learning. Along the way, we justify a privileged formalisation of the notion of 'epistemic conservativity', offer a new analysis of the Judy Benjamin problem and emphasise the distinction between interpreting the content of new evidence and updating one's beliefs on the basis of that content.

1   Introduction

Indicative conditionals have been the source of much trouble for logicians, epistemologists, cognitive scientists and philosophers of language. Although they appear to play a central role in logical and uncertain reasoning, there is still no general consensus regarding their semantics or even whether they should be thought of as propositions with determinate semantics at all (see e.g. Adams (1975), Edgington (1995), Bennett (2003)). It is unsurprising then that the problem of assigning probabilities to indicative conditionals is equally controversial (see e.g. Lewis (1976), Hajek (2012), Skovgaard Olsen *et al.* (2016)), and thus that the relationship between indicative conditionals and the norms of Bayesian epistemology remains largely opaque (see e.g. Douven (2015)). In this paper, we aim to clarify this relationship by addressing a specific question: How should an agent change their beliefs when they learn an indicative conditional from a perfectly reliable information source? This question is already the focus of a large literature (see e.g. Nayak *et al.* (1996), Uffink (1995), Grove & Halpern (1997), van Fraassen (1981), van Fraassen *et al.* (1986)). However, in a recent survey,

---

[*]Department of Philosophy, University of Konstanz, 78464 Konstanz (Germany) – http://be0367.wixsite.com/benevaphilosophy – benjamin.eva@uni-konstanz.de.

[†]Munich Center for Mathematical Philosophy, LMU Munich, 80539 Munich (Germany) – http://www.stephanhartmann.org – s.hartmann@lmu.de.

[‡]Department of Philosophy, University of Bayreuth, 95440 Bayreuth (Germany) – http://www.rafieerad.org – Soroush.RafieeRad@uni-bayreuth.de.

Douven (2012) concludes that all extant accounts are problematic (for a much more detailed discussion, see Douven & Romeijn (2012)), and hence that a proper general account of updating probabilistic belief states by indicative conditionals is yet to be formulated.

The most important existing approach to modeling the learning of indicative conditionals utilises the formal machinery of probabilistic distance-minimisation. The idea is this: An agent has beliefs about a set of propositions. These beliefs are represented by a prior probability distribution P. She then learns new information which imposes a constraint on the posterior distribution Q. If the learned information is the conditional A $\to$ B, then (it is commonly assumed) the new distribution has to satisfy a conditional probability constraint of the form $Q(\mathrm{B}|\mathrm{A}) = p' \leq 1$. The posterior probability distribution is then found by minimising some distance measure between Q and P. Intuitively, the idea is that the agent should stay as close as possible to their prior belief state whilst ensuring that the constraint implied by the new information is satisfied. This can be made formally precise by utilising probabilistic distance measures such as the Kullback leibler divergence (see e.g. van Fraassen (1981); Diaconis & Zabell (1982)). It turns out that minimising the Kullback leibler divergence between the prior and posterior distributions is equivalent to conditionalising on the material conditional A $\supset$ B when the learned conditional probability constraint is extreme ($Q(\mathrm{B}|\mathrm{A}) = 1$). Thus, the distance-minimisation approach reduces to simply conditionalising on the material conditional in the simplest case where the learned indicative conditional is interpreted as imposing an extreme conditional probability constraint.

However, this approach has been strongly criticised by Douven (2012), Douven & Romeijn (2012) and Douven & Dietz (2011), who forward a range of examples where conditionalising on the material conditional appears to yield strange and intuitively irrational updates. Furthermore, van Fraassen (1981) considers cases where the learned indicative conditional is naturally interpreted as imposing a non-extreme conditional probability constraint of the form $Q(\mathrm{B}|\mathrm{A}) = p' \in (0, 1)$. In these cases, minimising the distance between the prior and posterior distributions will generally diverge from simply Jeffrey conditionalizng on the material conditional, and it is unclear which probabilistic distance measure (if any) yields the intuitively rational update. Since these more general cases appear to involve learning experiences which are not easily describable in terms of conditionalisation or Jeffrey conditionalisation, they take us beyond the confines of standard Bayesian epistemology and seem to fall outside the purview of traditional Bayesian norms.

In this article, we defend (a particular variant of) the distance-minimisation approach against Douven *et al.*'s purported counter examples and argue that conditionalising on the material conditional produces reasonable and well motivated results in cases where the learned conditional is assumed to impose an extreme conditional probability constraint (Section 2). We then move on to consider the more general situation where

the learned conditional is allowed to impose a non-extreme constraint. We argue that in this case agents should be guided by a norm of epistemic conservativity, and compare ways in which this norm can be made precise. Specifically, we contend that epistemically conservative agents should update on new evidence by minimising a particular kind of probabilistic distance measure (an 'f-divergence') between their prior and posterior credences (Section 3). In Section 4 we present an analysis of van Fraassen's Judy Benjamin example in terms of f-divergence minimisation and note that different f-divergences give rise to different updates, none of which is obviously superior to its competitors. In order to resolve this impasse, we then appeal to resources from epistemic utility theory to identify a particular f-divergence (the Inverse Kullback leibler divergence) as normatively privileged, and argue that agents should utilise that divergence if they hope to achieve maximally accurate credences (Section 5). In Section 6 we utilise these considerations to develop a novel solution to the Judy Benjamin problem (and to learning conditionals in general), compare that solution to competing approaches from the recent literature, and draw an important and hitherto neglected distinction between the question of how rational agents should update their beliefs after learning indicative conditionals and the question of how rational agents should *interpret the content* of newly learned conditionals. We conclude by describing how contemporary work in the psychology of reasoning can be usefully brought to bear on the latter question, and raise some outstanding issues to be addressed in future work (Section 7).

2 The Material Conditional: A Defence

We are interested here in the question of how an epistemic agent should update their degrees of belief after learning an indicative conditional of the form 'If A, then B'. The existing literature on this question typically operates on the basis of a couple of idealizing assumptions. Firstly, it is standardly assumed that the information source from which the agent learns the conditional is completely trustworthy and reliable, so the agent doesn't have to worry about the possibility of being misled.[1]

Secondly, the agent has to translate the content of the learned conditional into a constraint on their posterior credences. In the kind of learning experience to which (Jeffrey) conditionalisation is standardly applied, the agent learns that a given truth-functional[2] proposition H is either more or less likely than they had previously thought, and this effectively imposes a constraint on the posterior probability of H (e.g. $P^*(\text{H}) = 1$).[3] In contrast, the dominant approach to modeling the learning of conditionals assumes that

---

[1]See Bovens & Hartmann (2003) for models of learning from partially reliable information sources. Note, however, that these models only consider the learning of truth functional propositions and cannot be straightforwardly applied to the learning of conditionals.

[2]A terminological clarification: by 'truth functional proposition', we mean either an atomic proposition or a proposition that can be generated from atomic propositions by the application of the standard truth-functional operations of classical logic (i.e. negation, disjunction and conjunction).

[3]We denote the posterior probability distribution which results from (Jeffrey) conditionalisation by P^*.

learning the conditional 'If A, then B' imposes a constraint, not on the probability of any single truth-functional proposition, but rather on the conditional probability of the consequent given the antecedent, $P^*(\text{B}|\text{A})$. This assumption is intuitively justified by the influential idea, commonly referred to as 'Adams' thesis', that the probability of the indicative conditional 'If A, then B' is given by the corresponding conditional probability $P(\text{B}|\text{A})$. Although Lewis' (1976) famous triviality results cast significant doubt on the general viability of this position, it still seems plausible to contend that learning the indicative conditional 'If A, then B' from a trustworthy source imposes a constraint on $P^*(\text{B}|\text{A})$. For the majority of the paper, we will accept this standard assumption as a given, and assume that epistemic agents can legitimately interpret the content of newly learned indicative conditionals as constraints on the posterior conditional probability of the consequent given the antecedent.[4]

Given these idealisations, the question of how one should update their credences after learning the conditional 'If A, then B' reduces to the problem of how one should update their credences upon learning a new value for the conditional probability $P^*(\text{B}|\text{A})$. For example, one might learn that the probability of it raining, given that there is a thunderstorm, is much lower than previously imagined, or that one's chances of developing heart disease given that one eats a lot of red meat is higher than expected. In the simplest case, one would learn that $P^*(\text{B}|\text{A}) = 1$, i.e. that A's truth provides an absolute guarantee of B's truth. But this is equivalent to the constraint that $P^*(\text{A} \wedge \neg\text{B}) = 0$, since both constraints simply state that there's no possibility of A being true and B being false. And of course, this means that the agent has really just learned the truth functional proposition $\neg\text{A} \vee \text{B} =_{df} \text{A} \supset \text{B}$ and hence that Bayesian norms demand that she should update her credences by conditionalising on $\text{A} \supset \text{B}$. So from a Bayesian perspective, there is no puzzle in the extreme case in which the learned conditional is interpreted as imposing the constraint $P^*(\text{B}|\text{A}) = 1$. The agent is duty bound to conditionalise on $\text{A} \supset \text{B}$, on pain of being susceptible to diachronic Dutch books (Lewis (1999)) and failing to minimise expected inaccuracy (Greaves & Wallace (2006)).

Although the situation looks clear cut from a Bayesian perspective, a number of authors have criticised the view that conditionalising on the material conditional provides a generally acceptable update procedure for agents who learn indicative conditionals. In the rest of this section, we provide a response to these criticisms and show that conditionalising on the material conditional performs remarkably well in a number of purportedly problematic cases.[5]

---

[4]As well as being standard in the philosophical literature (see e.g. van Fraassen (1981), van Fraassen *et al.* (1986), Douven & Dietz (2011), Douven & Romeijn (2012)), this assumption is also widely (although not universally) accepted in the psychology of reasoning, where it has played an important role in the empirical study of dynamic conditional inferences (see e.g. Oaksford and Chater (2007), Fernbach and Erb (2013)). But we stress again that we are concerned here only with how learning indicative conditionals causes a rational agent to alter their credences over time, and take no definite position regarding the controversial issue of how agents should attach probabilities to indicative conditionals at a given moment in time. We remain completely agnostic about Adams' thesis for the purposes of this paper.

[5]A similar conclusion is independently supported by Trpin (2018).

2.1 The Case of Two Propositions

We begin with some formal preliminaries. An agent has beliefs about the propositions A and B. We introduce binary propositional variables A and B (in italic script) which have the values A and ¬A, and B and ¬B (in roman script), respectively. These beliefs are represented by a prior probability distribution P over the possible values of A and B. In order to specify P, it is sufficient to fix the prior probability

$$P(A) = a, \tag{1}$$

and the conditional probabilities

$$P(B|A) = p \quad , \quad P(B|\neg A) = q. \tag{2}$$

With this, the prior probability distribution over the variables A and B is given by

$$P(A, B) = a\,p \quad , \quad P(A, \neg B) = a\,\overline{p}$$
$$P(\neg A, B) = \overline{a}\,q \quad , \quad P(\neg A, \neg B) = \overline{a}\,\overline{q}, \tag{3}$$

where we have used the shorthand notations $P(A, B)$ for $P(A \wedge B)$ and $\overline{x}$ for $1 - x$.

Next, the agent learns the natural language indicative conditional A → B. For now, we assume that the agent (legitimately) interprets the content of this conditional as imposing the extreme conditional probability constraint $P^*(B|A) = p' = 1$, which, as we noted above, is equivalent to learning

MC (Material Conditional): A ⊃ B, which is equivalent to ¬A ∨ B.

She then updates her probability distribution using Bayesian conditionalisation and arrives at the new probability distribution $P^*(\cdot) = P(\cdot|\neg A \vee B)$. This distribution is given by (all proofs are in Appendix A):

$$P^*(A, B) = \frac{a\,p}{a\,p + \overline{a}} \quad , \quad P^*(A, \neg B) = 0$$
$$P^*(\neg A, B) = \frac{\overline{a}\,q}{a\,p + \overline{a}} \quad , \quad P^*(\neg A, \neg B) = \frac{\overline{a}\,\overline{q}}{a\,p + \overline{a}} \tag{4}$$

From this, we can show that the following proposition holds.

Proposition 1 (Simple Conditional – MC) *An agent considers the propositions* A *and* B *and has a prior probability distribution* P *according to eqs. (3) (with* $0 < a, p, q < 1$*) defined over the corresponding propositional variables. Learning* **MC** *and updating by Bayesian conditionalisation then implies that the probability*

5

of A *decreases and that the probability of* B *increases, i.e.* $P^*(A) < P(A)$ *and* $P^*(B) > P(B)$.

Proposition 1 was first noticed by Popper & Miller (1983), and several authors have since cited the result as symptomatic of the representational deficiency of the material conditional. Specifically, it has been widely contended that it is implausible that the probability of the antecedent should decrease whenever we learn an indicative conditional, and hence that conditionalising on the material conditional cannot be the rational update procedure for this kind of learning. There are a couple of pertinent observations to be made here.

Recall that, following orthodoxy, we are working under the assumption that learning the indicative conditional 'If A, then B' imposes a constraint on the conditional probability $P^*(B|A)$. We have noted that, in the special case where this constraint is extreme, i.e. of the form $P^*(B|A) = 1$, standard Bayesian epistemology requires us to conditionalise on the proposition $A \supset B$. Now, there are at least two lessons that one could plausibly draw from Proposition 1 (which have not yet been distinguished in the literature).

Firstly, one could draw the conclusion that conditionalising on $A \supset B$ is not generally the rational response to learning the constraint $P^*(B|A) = 1$. Alternatively, one could draw the conclusion that it is not generally legitimate to interpret the content of the learned conditional as an extreme conditional probability constraint. The former conclusion concerns the norms of Bayesian epistemology, and the latter concerns the meaning of conditionals. In recent work, these two conclusions have been unduly conflated, which has unsurprisingly led to significant confusion.

We take the latter conclusion to be entirely uncontroversial and obvious from reflection on the actual use of conditionals in everyday discourse and reasoning. When somebody warns you that 'if you smoke, you'll get lung cancer', you typically interpret the content of that statement not as 'if you smoke, you will certainly get lung cancer', but rather as 'if you smoke, you will have a significantly increased risk of lung cancer', i.e. you interpret the conditional as imposing a non-extreme conditional probability constraint. In light of this observation, it is clear that conditionalising on the material conditional is not, in full generality, the rational update procedure for agents who learn indicative conditionals (since it will not generally respect the learned constraint). But, to the best of our knowledge, nobody is actually defending the position that conditionalising on the material conditional is always the rational response to learning an indicative conditional. Clearly, this update procedure is meant only for the special case in which the learned conditional *can* be legitimately interpreted as imposing an extreme conditional probability constraint. So Proposition 1 is only problematic if, *in those special cases*, it is sometimes intuitively irrational to become less confident in A after learning 'If A, then B'. If that were the case, then the import of Proposition 1 would be far greater than previously recognised, since it would seem to constitute an argument against Bayesian conditionalisation in general. Fortunately, it is, in full generality, intuitively rational to become less confident in A when we learn that A's

truth counts as a guarantee of B's truth.

To see this, consider the four maximal conjunctions or 'possible worlds' over which P is defined ($A \wedge B$, $A \wedge \neg B$, $\neg A \wedge B$ and $\neg A \wedge \neg B$). When we learn the indicative conditional 'If A, then B', interpreted as the constraint $P^*(B|A) = 1$, what we learn is that the actual world is not a world at which A is true and B is false. Thus, we rule out half of the logically possible worlds at which the antecedent is true. Any probability mass that was stored in those worlds will then have to be transferred to other worlds, some of which will be worlds at which the antecedent is false. So it is unsurprising that the antecedent becomes less probable when we learn the indicative conditional. For, the learned conditional explicitly rules out some of the ways in which the antecedent can be true. This is just an instance of the general phenomenon that the more informative a proposition is, the less likely it is to be true. When we learn the indicative conditional, the antecedent becomes more informative (and hence more easily falsifiable) and less probable. This is an intuitive and unproblematic observation. More generally, it is clear that when the agent interprets the learned evidence as imposing the constraint that $P^*(B|A) = 1$, they have really just learned the proposition $A \supset B$, and all the usual rationality arguments for conditionalisation apply. There is nothing special about this case that renders conditionalisation inapplicable.

However, Douven & Dietz (2011), Douven (2012) and Douven & Romeijn (2012) have presented a range of cases which purport to show that, even in the special case where the learned conditional can be legitimately interpreted by an extreme conditional probability constraint, conditionalising on the material conditional leads to problematic results. We turn now to providing a detailed analysis of these examples. In what follows, we call conditionals which can be interpreted as imposing extreme conditional probability constraints 'strict conditionals', and accept Douven *et* al.'s presumption that the conditionals learned in the prospective counterexamples can reasonably be construed as strict (even if this might reasonably be deemed an unrealistic idealisation).

2.2   The Case of More than Two Propositions

Douven and collaborators have presented several examples that are meant to show that conditionalising on the material conditional leads to counter-intuitive consequences, even when the learned conditional is strict. Specifically, the following example is supposed to show that learning a strict conditional can sometimes cause the probability of the antecedent to *increase*, which, in light of Proposition 1, is incompatible with conditionalising on the material conditional.

The Ski Trip Example. Harry sees his friend Sue buying a ski outfit. This surprises him a bit, because he did not know that she had any plans to go on a ski trip. He knows that she recently had an important

exam and thinks it unlikely that she passed it. Then he meets Tom, his best friend and also a friend of Sue's, who is just on his way to Sue to hear whether she passed the exam, and who tells him: "If Sue passed the exam, her father will take her on a ski vacation." Recalling his earlier observation, Harry now comes to find it more likely that Sue passed the exam. So in this example upon learning the conditional information Harry should intuitively increase the probability of the antecedent of the conditional.[6] (Douven & Dietz, 2011)

We begin by noting that in the ski trip example, the agent entertains three propositions, and not only two as in Proposition 1. What's more, the agent does not only learn an indicative conditional, but also another proposition, i.e. that Sue bought a new ski outfit. Clearly, a Bayesian agent also has to conditionalise on this information (a fact that was not taken into account by Douven & Dietz (2011), Douven (2012) and Douven & Romeijn (2012)). If one does so, then (modulo one further plausible assumption about Harry's prior beliefs) one obtains the intuitively correct result, as we will show now.

We introduce three binary propositional variables: The variable E has the values E: "Sue passed the exam", and ¬E: "Sue did not pass the exam". The variable S has the values S: "Sue's father invites her for a ski trip", and ¬S: "Sue's father does not invite her for a ski trip". The variable B has the values B: "Sue buys a new ski outfit", and ¬B: "Sue does not buy a new ski outfit". We represent Harry's prior and posterior credences by the probability distribution P and P^*, respectively. From the story, it is clear that Harry has low initial credence in B, E and S. Thus, we can safely assume that both $P(\neg E, S, B)$ and $P(\neg E, \neg S, B)$ are quite close to zero.[7] During the course of the example, Harry learns two new pieces of evidence:

Ski1: B

Ski2: E ⊃ S.

Armed with the assumption that both $P(\neg E, S, B)$ and $P(\neg E, \neg S, B)$ are small, it is possible to show (see Appendix B.1) that under a wide class of plausible formal reconstructions of the example, conditionalising on **Ski1** and **Ski2** will lead Harry to increase his credence in E, which is exactly the desired result. The probability of the antecedent will increase when we conditionalise on both the material conditional and the auxiliary proposition B. Thus, the ski trip example does not undermine the view that upon learning a (strict) indicative conditional, a Bayesian agent should simply conditionalise on the corresponding material conditional. Previous analyses of this example simply failed to take into account the fact that, in the example, the agent learns more than just the conditional.[8] It is unsurprising that when this extra information is neglected, the resulting

---

[6]Note that the interpretation of the conditional as strict in this example can easily be made more plausible if we imagine that the learned conditional is rather "If Sue passed the exam, her father will *definitely* take her on a ski vacation."

[7]Since it is specified that $P(S)$ and $P(B)$ are both low and the variables S and B are obviously (strongly) positively correlated.

[8]One might plausibly contend that B should be treated as part of Harry's background knowledge, rather than something he

update seems irrational. More generally, the ski trip example does nothing to cast doubt on the view that upon learning a strict indicative conditional, the antecedent will never increase in probability. Certainly, the net effect of a compound update in which a strict conditional is learned together with some other propositions might be to increase the probability of the antecedent (as shown above), but that increase is not due to the learning of the conditional itself. Rather, it occurs *in spite of* the learning of the conditional, and is due entirely to the acquisition of additional evidence (in this case, the auxiliary proposition B).

Let us now turn to the second prospective counterexample:[9]

The Sundowners Example. Sarah and her sister Marian have arranged to go for sundowners at the Westcliff hotel tomorrow. Sarah feels that there is some chance that it will rain, but thinks they can always enjoy the view from inside. To make sure, Marian consults the staff at the Westcliff hotel and finds out that in the event of rain, the inside area will be occupied by a wedding party. So she tells Sarah: "If it rains tomorrow, then we cannot have sundowners at the Westcliff." Upon learning this conditional, Sarah sets her probability for sundowners *and* rain to 0, but does not change her probability for rain. Thus, in this example, learning the conditional information has the effect of leaving the probability of the antecedent unchanged.[10] (Douven & Romeijn, 2012)

To begin, we introduce two binary propositional variables. The variable R has values R: "It will rain tomorrow", and ¬R: "It will not rain tomorrow". The variable S has the values S: "Sarah and Marian have sundowners", and ¬S: "Sarah and Marian do not have sundowners". According to the story, Sarah learns

Sun1: R → ¬S. We assume that Sarah represents this indicative conditional by the material conditional R ⊃ ¬S, which is logically equivalent to ¬R ∨ ¬S.

Assuming that $0 < P(\text{R}), P(\text{S}|\text{R}), P(\text{S}|\neg\text{R}) < 1$, we can then use Proposition 1 and conclude that $P^*(\text{R}) < P(\text{R})$, which conflicts with our intuitive judgment that the probability of rain should remain unchanged. At first glance, this example is more difficult to analyse than the ski trip example because it seems that Sarah only learns the conditional, and there is no additional evidence to explain why her update should conflict with Proposition 1.

---

explicitly updates on when he learns the conditional. But this poses no problems for our analysis. We can also assume that Harry starts with a credence function P, which he later updates after learning B (which will then form part of his background knowledge), and only subsequently comes to learn the conditional. Because conditionalisation is a commutative operation (the order in which propositions are learned makes no difference to the posterior function we end up with), this is formally equivalent to treating B as part of the learning experience in the example. Our argument still works in exactly the same way.

[9]Note that Douven & Romeijn (2012) also consider a third salient example (the 'driving test example') that is structurally analogous to the ski trip example, but in which it is intuitively rational for the agent to *decrease* their credence in the antecedent of the learned conditional. Since this does not yield any direct conflict with Proposition 1, we do not address it here.

[10]Again, we can make the analysis of the example in terms of strict conditionals more plausible by changing the learned conditional to "If it rains tomorrow, then we *definitely* cannot have sundowners at the Westcliff."

In order to see what's gone wrong here, it is helpful to consider what motivates the intuition that Sarah's credence in R should remain fixed across the learning experience. Importantly, we have stipulated that Sarah trusts Marian completely and so will assume that Marian has a good reason for asserting the conditional, for example that there is going to be a wedding party at the hotel. What's more, Sarah will certainly consider this reason (whatever it may be) to be probabilistically independent of R. Clearly, whether or not the sisters can have sundowners inside in the event of rain has nothing to do with whether it will in fact rain, and this is why we think Sarah's credence in R should remain unchanged. Whatever caused Marian to assert the conditional has nothing to do with the possibility of rain. But the example, as presented by Douven & Romeijn (2012) and formalised above, implicitly assumes this background structure in order to motivate the desired verdict (that Sarah's credence in R shouldn't change), whilst neglecting it entirely when it comes to formally modeling Sarah's update in terms of conditionalising on the material conditional. When this background structure is properly taken into account, conditionalising on the material conditional will straightforwardly yield the intuitive verdict, as we will shortly show.

Another way to illustrate the reasoning above is to employ the distinction between the semantic and pragmatic content of an utterance (see e.g. Szabo (2005)). Roughly, the semantic content is the content that is explicitly conveyed by the speaker's words whereas the pragmatic content encompasses all the extra content that would need to be conveyed in order for the utterance to be felicitous in the given context, or that would automatically be inferred by any competent listener upon hearing the utterance in that context. Thus, when Sarah says 'If it rains tomorrow, then we cannot have sundowners at the westcliff', the semantic content of her utterance can be interpreted by the material conditional $R \supset \neg S$. However, it is clear that upon hearing this utterance any competent listener would automatically infer that in the event of rain, something will prevent them from having Sundowner's, and this therefore forms part of the pragmatic content of the utterance. Alternatively, if we assume that what makes an utterance felicitous is its accordance with Grice's maxims (1975), then we can invoke the maxim of quantity – 'make your contribution as informative as is required' – to make the same point. Part of what Marian needs to convey in order for the conversation to be successful is that in the event of rain, something will stop them from having sundowners inside. So in order for her utterance to be felicitous, this must be part of its content. Specifically, it must be part of the utterance's pragmatic content. The plausibility of this contention is further reinforced when we notice that in ordinary discourse, it would seem strange to distinguish between the content of the utterances 'If it rains tomorrow, we cannot have sundowners' and 'It it rains tomorrow, something will prevent us from having sundowners'.

In order to model all the relevant aspects of Sarah's epistemic state, we introduce a third propositional variable D (for 'disabler') with the values D: "In the event of rain, something (e.g. a wedding party) will

prevent us from having sundowners inside", and ¬D: "In the event of rain, nothing will prevent us from having sundowners inside". Just like R, D is a disabler for S: R will prevent Sarah and Marian from having the sundowners outside, and D will prevent them having sundowners inside. If we make the simplifying assumption that Sarah is certain that the only thing that could prevent them from having sundowner inside is a wedding party, then D can be identified with the proposition 'there will be a wedding party at the hotel'.[11] It is part of Sarah's background knowledge that there will be no sundowners if R and D are instantiated, i.e.

$$R \wedge D \rightarrow \neg S. \tag{5}$$

In probabilistic terms, this means that the prior probability distribution satisfies the constraint $P(\neg S|R, D) = 1$ and hence

$$P(R, S, D) = 0. \tag{6}$$

Furthermore, it is clear (as we argued above) that Sarah regards R and D as being probabilistically independent. Whether or not it will be possible to have sundowners inside in the event of rain is entirely independent of the question whether or not it will rain in the first place. When Marian learns that a wedding party has reserved the inside area in the event of rain, she doesn't take this information to render the prospect of rain any more or less likely then before. Indeed, this is exactly the intuition that motivates the claim that Sarah should not change her credence in R when she learns the conditional. We can also employ the causal structure of the example to make the same point. Intuitively, both R and D causally influence S, but they do not causally influence one another (and there is no obvious non-trivial common cause). Applying the causal Markov condition to this intuitive picture of the causal structure licenses the conclusion that R and D are probabilistically independent (see e.g. Sprites *et al.* (2000)).

Finally, the example assumes that Sarah considers Marian to be a perfectly reliable information source and interprets the conditional R → ¬S as a strict conditional. Hence, she ensures that the posterior conditional probability of S given R is zero (which is ensured by conditionalising on R ⊃ ¬S). But, as we noted above, if Sarah truly regards Marian as perfectly reliable, then of course she should additionally conditionalise on D, i.e she should conditionalise on there being some reason that they won't be able to have sundowners inside in the event of rain. If she leaves the probability of D below 1, then in her posterior beliefs, she will simultaneously be certain that sundowners is impossible in the event of rain and be uncertain about whether there is any reason that they can't have sundowners inside in the event of rain. Clearly, when Marian tells Sarah that they can't have sundowners in the event of rain, part of the content of her assertion is that there is some

---

[11] Maybe Sarah does not think about a wedding party when she makes up her prior beliefs. But she will certainly take into account that there is *some* disabler that will have the result that they cannot have sundowners inside. Indeed, the story specifies that Sarah 'thinks they can always enjoy the view from inside', i.e. her prior for D is low.

reason that they can't have sundowners inside (what else could have caused her to assert the conditional?), and Sarah needs to take this into account when she changes her beliefs. Specifically, as well as conditionalising on **Sun 1**, Sarah should additionally conditionalise on

Sun2: D

Together, these observations are sufficient to obtain the following result.

Proposition 2 (Sundowners Example – MC) *An agent considers the propositions* R, S *and* D *and has a prior probability distribution* P *defined over the corresponding propositional variables.* R *and* D *are probabilistically independent and* P *satisfies eq. (6). Learning* **Sun1** *and* **Sun2** *and updating by Bayesian conditionalisation then implies that* $P^*(R) = P(R)$.

The crucial point here is that when Marian asserts the conditional 'If R then not-S', part of the conversational implicature of her assertion is that there is some reason that they won't be able to have sundowners in the event of rain, and it is part of the sisters' epistemic common ground that any such reason will be entirely independent of whether it will in fact rain. This is what motivates the intuition that Sarah should not change her credence in rain after learning the conditional. It is not surprising that when we omit these important structural features from our probabilistic representation of Sarah's epistemic state (and her subsequent update), we obtain peculiar results. Proposition 2 shows that when all of this extra structure is properly taken into account, Sarah's learning experience can be adequately modeled in a Bayesian framework. Overall, we have reached the following conclusions:[12]

1. Updating by conditionalising on the material conditional leads to perfectly plausible results if there are only two propositions involved.

2. If more than two relevant propositions are involved, then these have to be accounted for and modeled in a proper way. More specifically, the ski trip example turns out to be unproblematic if one conditionalises not only on the material conditional but also on the relevant additional evidence. To properly account for our intuitions, it is also important to take account of the probabilistic independencies that are implied by the intuitive causal structure of the examples.

3. Conditionalising on the material conditional also produces the desired result in the sundowners example. In this case the extra proposition D is not explicitly stated, but its presence is implicitly assumed in

---

[12]We should stress here that these conclusions do *not* commit us to the view that natural language indicative conditionals can in general be adequately modelled by the material conditional. Indeed, these conclusions are all perfectly compatible with Adams' thesis (according to which the probability of an indicative conditional is given by the relevant conditional probability, and *not* the probability of the material conditional), since conditionalising on the material conditional is equivalent to imposing an extreme conditional probability constraint on the posterior distribution. Furthermore, the next section makes clear that the material conditional is fundamentally incapable of adequately modeling the learning of non-strict conditionals.

motivating the 'intuitively correct' response to the learning experience.

3   Non-Strict Conditionals and Distance-Minimisation

In the previous section, we defended the position that conditionalising on the material conditional is generally the rational update procedure for learning strict indicative conditionals (although there may be situations in which the material conditional does not exhaust the conversational implicature of the strict conditional, in which case the implicit (pragmatic) content of the assertion should also be taken into account). For the remainder of the paper, we will focus on the more general case in which the learned conditional is non-strict, i.e. it cannot be legitimately interpreted by an extreme constraint of the form $Q(\text{B}|\text{A}) = 1$, but only by a non-extreme constraint of the form $Q(\text{B}|\text{A}) = p' \in (0,1)$. By way of illustration, consider van Fraassen's famous Judy Benjamin example:

The Judy Benjamin Example. A soldier, Judy Benjamin, is dropped with her platoon in a territory that is divided in two halves, Red Territory and Blue Territory, respectively, with each territory in turn being divided in equal parts, Second Company area and Headquarters Company area, thus forming four quadrants of roughly equal size. Because the platoon was dropped more or less at the center of the whole territory, Judy Benjamin deems it equally likely that they are in one quadrant as that they are in any of the others. They then receive the following radio message: "I can't be sure where you are. If you are in Red Territory, then the odds are 3 : 1 that you are in Second Company area." After this, the radio contact breaks down. Supposing that Judy accepts this message, how should she adjust her degrees of belief? (van Fraassen, 1981)

Here, Judy learns a conditional that should be interpreted as telling her to increase her estimate of the conditional probability of being in the Second Company area given that she is in Red Territory, but not all the way to one. Clearly, it would be a mistake to conditionalise on the material conditional 'Red Territory ⊃ Second Company area', since that would result in mistakenly ruling out the possibility of being in Red Territory and the Headquarter area. So we need a more general approach to learning conditionals that is able to guide agents through this kind of learning experience. In particular, we need to identify the rational updating procedure for agents that learn a constraint of the form $Q(\text{B}|\text{A}) = p' \in (0,1)$.[13]

The dominant approach to this problem in the literature is to argue that agents who learn non-extreme conditional probability constraints should be guided by a norm of conservativity. In particular, they should

---

[13]Huisman (2017) actually argues that the rational response to the Judy Benjamin example is to adopt indeterminate posterior credences. Here, we aim to stay within the standard Bayesian paradigm and identify a rational update procedure that yields precise and probabilistic posterior credences.

update their credences in a way which ensures that their new credences (i) satisfy the learned constraint, and (ii) are as close as possible to their original credences, i.e. they should not change their beliefs in ways which are not necessitated by the new evidence. Of course, this is just an instance of the general normative principle that agents should change their beliefs conservatively in the face of new evidence. This principle has played a fundamental role in both Bayesian epistemology and the theory of qualitative belief revision (see e.g. Gärdenfors & Rott (1995)). In what follows, we will take it for granted that rational agents should be guided by a norm of epistemic conservativity, and ignore recent controversies regarding whether that norm should ever be sacrificed in the interests of alternative epistemic goods (see e.g. Myrvold (2018)). Rather, we will focus on the question of how this norm should be made precise in a probabilistic setting and how it can be applied to guide Bayesian agents who learn non-strict conditionals, as in van Fraassen's Judy Benjamin example.

More specifically, we are interested in the task of identifying exactly which notion of 'closeness' is appropriate when we demand that epistemic agents should try to remain 'as close as possible' to their original credences when updating on new evidence. What's the right way to measure the distance between two different probabilistic belief states here? Without an answer to this question, the conservativity norm is not sufficiently precise to provide any concrete advice regarding how agents should update their beliefs in the face of new evidence.

3.1 Formalising Conservativity

It turns out that there are many ways to measure the distance between different probability distributions. For example, imagine that we want to measure the distance between the distributions P and Q, defined over some common set W of possible worlds. One natural (and widely employed) option is to use the Euclidean squared distance between P and Q, $D(P,Q) = \sum_{w \in W} (P(w) - Q(w))^2$. The problem with this approach is that if we apply the conservativity norm with this notion of closeness, we will end up advocating update procedures that conflict with existing Bayesian update rules in strange and problematic ways. In particular, suppose that an agent learns the constraint that the posterior probability of some truth functional proposition A should increase to a particular value $Q(A) = a' < 1$. Leitgeb & Pettigrew (2010) showed that minimising the Euclidean squared distance to this constraint yields a new update rule (now known as 'Leitgeb-Pettigrew updating' or 'LP-updating') according to which the agent should add some constant value to the probability of all possible worlds in which A is true. Levinstein (2012) showed that this update rule yields absurd and patently irrational updates in many cases. For example, if an agent initially believes two propositions A and B to be independent and then encounters some evidence that makes them far more confident in the truth of

A, LP-updating can lead them to radically revise their credence in B, even though they previously considered A and B to be independent and the learned evidence doesn't appear to bear on that consideration in any way.[14] The fundamental problem here is that LP-updating does not preserve the conditional probability ratios between the considered propositions, i.e. it is not a 'rigid' updating procedure.

It is well known that the only updating procedure that is guaranteed to preserve all relevant conditional probabilities in this kind of learning scenario is Jeffrey conditionalisation, which is widely accepted as the rational update procedure for Bayesian agents in cases of 'ineffable learning', i.e. cases where the agent gains evidence that makes them more confident of some proposition without becoming certain of any proposition. As well as having the important property of preserving conditional probabilities, Jeffrey conditionalisation has also been justified by diachronic dutch book arguments (Skyrms (1984)). In what follows, we take it for granted that Jeffrey conditionalisation is the rational update procedure for agents who undergo ineffable learning. Of course, we acknowledge that this might be a controversial step since some authors contend that Jeffrey conditionalisation is only a rational updating procedure in some circumstance (see e.g. Schwan & Stern (2017)), but given the widespread acceptance of Jeffrey conditionalisation in the Bayesian canon and the lack of any plausible alternative probabilistic updating rules, we take it that this assumption does not lessen the philosophical import of our conclusions.

Once we assume that rational agents should update by Jeffrey conditionalisation, we are able to rule out the Euclidean squared distance as a formalisation of epistemic conservativity. For, agents who try to be epistemically conservative in the sense captured by the Euclidean squared distance will employ LP-updating instead of Jeffrey conditionalisation in cases of ineffable learning, and will therefore end up committing the kinds of reasoning errors described by Levinstein (2012). The question now is whether there are alternative ways of formalising epistemic conservativity that do not conflict with Jeffrey conditionalisation and are therefore compatible with existing Bayesian update rules. Happily, there exists a whole class of probabilistic distance measures that satisfy exactly this property.

Definition 1 (f-divergence) *Let $S_1, \ldots, S_n$ be the possible values of a random variable S over which probability distributions P (= the prior distribution) and Q (= the posterior distribution) are defined. The f-divergence between Q and P is then given by*

$$D_f(Q||P) := \sum_{i=1}^{n} P(S_i) \, f\left(\frac{Q(S_i)}{P(S_i)}\right),$$

---

[14]Specifically, Levinstein presents a case in which the agent has very high prior high credences in the propositions R: "The colour of the car behind the door is red" and ¬G : "Ghosts don't exist". Of course, she considers R and ¬G to be probabilistically independent. She then catches a glimpse of the car and becomes confident that the car actually isn't red. According to LP-updating, this should cause her to become agnostic about the existence of ghosts (her credence in G will go from being close to zero to being close to 1/2).

where f is a convex function such that $f(1) = 0$.

f-divergences have a number of interesting properties (see Csiszár (1967, 1975, 1991, 2008) for a more detailed discussion): They are always non-negative and they vanish iff $Q = P$. However, unlike proper distance functions, they are not necessarily symmetric and they may violate the triangle inequality. Examples of f-divergences are the Kullback leibler divergence ($f(x) = x \log x$), the inverse Kullback leibler divergence ($f(x) = -\log x$), the Hellinger distance ($f(x) = (1 - \sqrt{x})^2$), and the χ^2-divergence ($f(x) = (x - 1)^2$). For us, the most important property of f-divergencies is that they imply Jeffrey conditionalisation. (See also Diaconis & Zabell (1982).)

Proposition 3 (Jeffrey Conditionalisation) *Let P be the prior probability distribution over the binary propositional variables H and E. The agent then learns that the probability of E shifts from $P(E)$ to $Q(E)$. Minimising the f-divergence between Q and P, the posterior probability of H is then given by $Q(H) = P(H|E) Q(E) + P(H|\neg E) Q(\neg E)$.*

Thus, in situations where the learned constraint on the posterior distribution is just the probability of a proposition (i.e. in the cases described by standard Bayesian updating), minimising the f-divergence between the prior and posterior distributions is equivalent to updating by (Jeffrey) conditionalisation. More pertinently, an epistemically conservative agent who formalises conservativity by means of an f-divergence will always reason in accordance with Jeffrey conditionalisation, and f-divergences can be seen as providing the formal resources needed to apply the norm of epistemic conservativity in a way that is continuous with the norms of standard Bayesian epistemology.

3.2   Learning Conditionals by Minimising f-Divergences

Returning to the central question of this paper, we now consider the implications of the preceding discussion for the problem of learning indicative conditionals. The first important observation is given by the following corollary of Proposition 3.

Corollary 1 (Simple Conditional – f-Divergence) *An agent considers the propositions A and B and has a prior probability distribution P according to eqs. (3) (with $0 < a, p, q < 1$) defined over the corresponding propositional variables. The agent then learns $A \rightarrow B$ which imposes the constraint $Q(B|A) = 1$ on the posterior distribution Q. Minimising the f-divergence between Q and P then implies that the posterior probability distribution Q is exactly the same as in eqs. (4).*

Conditionalising on the material conditional yields exactly the same result as minimising any f-divergence

(including the inverse KL divergence) with the relevant extreme conditional probability constraint. So the norm of epistemic conservativity (formalised in terms of f-divergence minimisation) yields the standard update procedure for learning strict conditionals.[15]

For non-strict conditionals, the situation is of course significantly more complicated. To illustrate, we consider the simple situation where an agent entertains the propositions A and B. She then learns the indicative conditional A $\rightarrow$ B. This, however, does not prompt her to increase the conditional probability of B given A to 1 (since the conditional is non-strict), but only to some value $p' > P(\text{B}|\text{A})$. She increases the conditional probability, but not to its maximal value. We assume that her prior beliefs P are represented by the probability distribution given in eqs. (3) and that the posterior probability distribution Q can be parameterized in the same way, i.e.

$$Q(\text{A}, \text{B}) = a'\, p' \quad , \quad Q(\text{A}, \neg\text{B}) = a'\, \overline{p'}$$
$$Q(\neg\text{A}, \text{B}) = \overline{a'}\, q' \quad , \quad Q(\neg\text{A}, \neg\text{B}) = \overline{a'}\, \overline{q'}. \tag{7}$$

The posterior probability distribution is then given by the following proposition:

Proposition 4 (Simple Conditional – Non-Extreme f-Divergence) *An agent considers the propositions* A *and* B *and has a prior probability distribution* P *according to eqs. (3) (with $0 < a, p, q < 1$) defined over the corresponding propositional variables. The agent then learns* A $\rightarrow$ B *which imposes the constraint* $Q(\text{B}|\text{A}) = p' < 1$ *on the posterior distribution* Q *(see eq. (7))). Minimising the f-divergence between* Q *and* P *then implies that* $q' = q$ *and* $a' = a/(a + \overline{a} \cdot l)$ *where the value of the likelihood ratio* l *depends on the particular f-divergence. More specifically, we obtain*

$$l_{IKL} = 1 \quad , \quad l_{KL} = \left(\frac{p'}{p}\right)^{p'} \left(\frac{\overline{p'}}{\overline{p}}\right)^{\overline{p'}}$$
$$l_H = \frac{1}{\left(\sqrt{p\,p'} + \sqrt{\overline{p}\,\overline{p'}}\right)^2} \quad , \quad l_{\chi^2} = \frac{p'^2}{p} + \frac{\overline{p'}^2}{\overline{p}}.$$

Furthermore, $P'(\text{B}) = P(\text{B}) + [(p' - p) - \overline{a}\,(l - 1)\,(p - q)] \cdot P'(\text{A})$ *can be smaller than, equal to, or larger than* $P(\text{B})$*, depending on the prior distribution.*

Proposition 4 gives a powerful general characterisation of the update obtained by minimising an f-divergence to a non-extreme conditional probability constraint. It also has important ramifications for the

---

[15]Note that this observation implies that in the sundowners and ski trip examples, minimising an f-divergence to the constraints **Ski1/Ski2** and **Sun1/Sun2** respectively will yield the same updates as those we obtained in the previous section using standard Bayesian conditionalisation. See also Appendix B.

analysis of van Fraassen's Judy Benjamin example, as we will now show.

4   Judy Benjamin Reloaded

Recall the Judy Benjamin example from the beginning of the previous section. In the literature, three main criteria have been forwarded as general constraints on Judy's response to the learned conditional. In order to state these criteria, we introduce two binary propositional variables. The variable R has the values R: "Judy lands in Red Territory", and ¬R: "Judy lands in Blue Territory". The variable S has the values S: "Judy lands in the Second Company Area", and ¬S: "Judy lands in Headquarters Area". P and Q represent Judy's prior and posterior beliefs, respectively (where it is assumed that P is a uniform distribution over the partition $\{R \wedge S, R \wedge \neg S, \neg R \wedge S, \neg R \wedge \neg S\}$).

JB1: $Q(S|R) = 3/4$ (and hence $Q(\neg S|R) = 1/4$).

JB2: $Q(A|X) = P(A|X)$ for any A and for any $X \in \{R \wedge S, R \wedge \neg S, \neg R\}$.

JB3: $Q(R) = P(R) = 1/2$.

JB1 simply requires that Judy updates her beliefs to take into account the content of the learned conditional. JB2 stipulates that Judy shouldn't change any conditional degrees of belief other than $P(S|R)$ and $P(\neg S|R)$ (since the learned information doesn't seem to tell her anything about the corresponding conditional probabilities). JB3 requires that Judy shouldn't change her degree of belief in R, i.e. she shouldn't become more or less confident that she is in Red Territory than she was before she learned the conditional (again, since the learned information doesn't seem to tell her anything about whether or not she is in Red Territory).[16]

Now, van Fraassen (1981) initially presented the Judy Benjamin example as a test case for the proposal that rational agents should always update their credences by minimising the Kullback leibler (KL) divergence between their prior and posterior degrees of belief (this is sometimes referred to as the 'infomin principle'). He noted that this proposal has two powerful motivations, namely:

(1) Minimising the KL divergence implies conditionalisation and Jeffrey conditionalisation when the learning experience concerns the probability of a truth-functional proposition.

(2) Minimising the KL divergence yields an intuitively conservative update.

---

[16]Note that the kind of learning described by the Judy Benjamin example clearly cannot be properly modeled by Jeffrey conditionalising to set the probability of the material conditional R ⊃ S to a new value. To see why, imagine that Judy rather learns the conditional "If you are in Red Territory, then the odds that you are in the Second Company Area are 50 : 50." Since her prior beliefs already satisfied the equality $P(S|R) = 1/2$, it seems that she has not really learned anything new and so has no reason to change her beliefs. However, if she Jeffrey conditionalises to set $Q(R ⊃ S) = 1/2$, her degrees of belief will change significantly (since $P(R ⊃ S) \neq 1/2$). This intuitively irrational response can be avoided if she rather minimises a distance measure to the learned conditional probability constraint.

However, he also noted that, despite these appealing properties, the proposal yields an unintuitive update in the Judy Benjamin case.[17] Specifically, if Judy updates by minimising the KL divergence to the constraint $Q(S|R) = 3/4$, her posterior credences will violate the JB3 criterion, i.e. her confidence that she is in Red Territory will decrease across the update. And this seems problematic since the learned conditional didn't seem to tell her anything about whether or not she was in Red Territory. Rather, the learned conditional only told her what she should expect *under the supposition that she in Red Territory*. Thus, it has been widely claimed that the Judy Benjamin example is a major problem for those that advocate minimising the KL divergence as a fully general updating procedure.

van Fraassen *et al.* (1986) provide an alternative analysis of the example. Specifically, they consider an amended version of the example in which Judy learns that the conditional probability of S given R is 1 (rather than 3/4), i.e. the learned conditional is strict.[18] They note that in this extreme version of the example, it seems natural for Judy to reduce her belief in R. For, the learned information is equivalent to the proposition $\neg(R \wedge \neg S) = R \supset S$, and conditionalising on that proposition will lead to a decrease in the probability of R (by Proposition 1). Now, there is nothing ostensibly 'special' about the learned odds in this case (1 : 0) as compared to those learned in the standard version of the example (3 : 1). So, the story goes, there is no reason to think that Judy can rationally decrease her belief in R in one case but not the other. Thus, we should drop the requirement that Judy's confidence in R stays fixed across the update.

But even if this response is correct and JB3 is not a genuine rationality constraint (we discuss this response in more detail in Section 6), van Fraassen *et al.* (1986) note that there is another more fundamental problem for advocates of KL minimisation. In particular, they present two alternative probabilistic distance measures such that (i) updating by minimising those measures yields Jeffrey conditionalisation, (ii) all three measures look like equally good formalisations of epistemic conservativity, and (iii) all three measures yield different updates in the Judy Benjamin example. The problem is that there doesn't seem to be any principled ground for choosing between the alternative distance measures, and hence that it's unclear why Judy should minimise the KL divergence instead of some other probability distance measure. Again, the norm of epistemic conservativity is not sufficiently precise to provide concrete guidance to epistemic agents.

It is immediately clear that the analysis provided by van Fraassen (1981) and van Fraassen *et al.* (1986) can be greatly generalised and refined in light of the results described in the previous section. Firstly, we know that all f-divergences (of which there are infinitely many) yield update procedures which are equivalent to Jeffrey

---

[17]Vasudevan (2018) presents an alternative rationale for updating by the KL divergence in the Judy Benjamin example. Specifically, he does not invoke any kind of conservativity norm, but rather argues that minimising the KL divergence is justified by the assumption that the informant's report omits no relevant information. For the purposes of this paper, we intend to focus on clarifying the relationship between the Judy Benjamin problem and norms of epistemic conservativity, although we think that additional epistemic principles such as that invoked by Vasudevan (2018) may also have an important role to play.

[18]This extreme version of the example is also considered by Bovens (2013), who similarly reaches the conclusion that Judy's credence in R should change across the update. However, this analysis is subsequently retracted in Bovens & Ferreira (2013).

conditionalisation when the learned information concerns the probability of a truth functional proposition, and we also know from Proposition 4 that the the resultant updates will generally all be different. Thus, the degree to which the rational update in the Judy Benjamin example is underdetermined is far greater than has previously been acknowledged. There are infinitely many update procedures that are both compatible with Jeffrey conditionalisation and intuitively conservative, and there's no obvious reason why Judy should prefer any one of them to any other. Or is there?

Recall Proposition 4, which shows that minimising the f-divergence to the constraint $Q(\mathrm{S}|\mathrm{R}) = 3/4$ will typically (excluding limit cases) lead to an update that violates JB3, with one prominent exception, namely the inverse KL divergence (IKL). Minimising the IKL divergence to a non-extreme conditional probability constraint will always leave the prior probability of the antecedent unchanged. This means that minimising the IKL divergence is the only way for Judy to satisfy all three criteria JB1–JB3 whilst being epistemically conservative and updating in a way that is consistent with Jeffrey conditionalisation. So if, contra van Fraassen *et al.* (1986), we accept the normative force of JB3, we are forced to conclude that minimising the IKL divergence is the rational solution to the Judy Benjamin problem. At first blush, this looks like an appealing solution to the problem. The IKL divergence appears to provide an update procedure that satisfies both of the main motivations that van Fraassen cites in favour of the KL divergence, and is also the only such update procedure that recommends what many consider to be the rational update in the Judy Benjamin example. The solution proposed by Douven & Romeijn (2012) is actually equivalent to minimising the IKL divergence (although they don't identify the important fact that the IKL divergence is an f-divergence that produces an update that is continuous with Jeffrey conditionalisation). However, Douven and Romeijn stop short of endorsing the minimisation of the IKL divergence as a generally applicable update strategy.[19] They eventually suggest that no single distance measure/update procedure is suitable for all learning scenarios, and tentatively advocate a kind of contextualism about update rules.

Here, we attempt to identify a single update procedure that provides a uniquely rational response to any given learning scenario. However, we think that the normative justification for such a procedure must be more general and philosophically principled than the simple observation that it yields what some take to be the rational update in the Judy Benjamin problem. The simple fact that the IKL divergence defines an update that satisfies JB3 does not constitute a serious argument in favour of using that distance measure in all cases of non-propositional learning.[20] For, we have already seen that some authors (e.g. van Fraassen *et al.* (1986)) question whether JB3 should even be accepted as a rationality constraint in the first place. Furthermore, as

---

[19]This is already clear from the fact that they don't think that rational agents should conditionalise on the material conditional in the extreme version of the example where the learned constraint is strict, i.e. $Q(\mathrm{S}|\mathrm{R}) = 1$.

[20]By non-propositional learning, we mean learning experiences such as Judy Benjamin's, in which the newly acquired evidence can't be adequately represented by a constraint on the probability of any one truth functional proposition.

we will see in the Section 6, it is possible to construct analogous examples in which it seems intuitive that the agent's credence in the antecedent of the learned non-strict conditional should change, and in these cases the IKL divergence is guaranteed to give what looks like the wrong answer. We need a principled and context independent way of deciding between alternative f-divergences that can subsequently be used to cast light on problematic learning scenarios such as that described in the Judy Benjamin example.

We turn now to giving a general normative argument for a single, universally applicable update procedure based on the minimisation of an f-divergence. In Section 6, we consider the implications of this argument for the analysis of the Judy Benjamin example.

5   Accuracy and Conservativity

One natural approach to trying to decide between alternative f-divergences is to employ considerations from epistemic utility theory. In recent years, a number of authors have attempted to utilise the notion of 'epistemic utility' or 'accuracy' to provide purely epistemic justifications for the basic norms of Bayesian epistemology, including probabilism (Leitgeb & Pettigrew (2010)) and conditionalisation (Greaves & Wallace (2006)). One might hope that a similar justification can be given for a particular f-divergence, in which case we would be furnished with a unique rational update procedure in Judy Benjamin style cases.[21]

Recall that the central idea behind epistemic utility theory is that epistemic agents should aim to approximate the truth with their degrees of belief or, in other words, that they should aim to minimise their inaccuracy. It is standardly assumed that the right way to quantify the inaccuracy of an agent's partial beliefs is to employ a strictly proper scoring rule, i.e. a measure of inaccuracy which ensures that agents always expect their own partial beliefs to uniquely minimise inaccuracy. Two of the most important and widely defended proper scoring rules are the Brier and logarithmic scoring rules, defined below (where W is the set of possible worlds).

- The logarithmic scoring rule is the rule which defines the inaccuracy of the credence function P at a world $w \in W$ by $\mathfrak{I}_L(P, w) = -\ln(P(w))$.

- The Brier score is the rule which defines the inaccuracy of the credence function P at a world $w \in W$ by $\mathfrak{I}_B(P, w) = (1 - P(w))^2 + \sum_{w' \neq w} P(w')^2$.

Now, it is well known that all strictly proper scoring rules, including $\mathfrak{I}_L$ and $\mathfrak{I}_B$, can be used to define a corresponding divergence measure, as we will soon show. Specifically, given a strictly proper scoring rule

---

[21]In what follows we assume a basic familiarity with epistemic utility theory. For a book-length introduction to the subject, see Pettigrew (2016).

$\mathfrak{I}_S$ and any two credence functions P and Q, we can define $\exp_S(Q|P) := \sum_{w \in W} P(W) \cdot \mathfrak{I}_S(Q, w)$, which encodes the extent to which an agent whose credences are given by P will expect the credences given by Q to be inaccurate. By strict propriety, $\exp_S(Q|P)$ is uniquely minimised when $Q = P$. Leitgeb & Pettigrew (2010) utilise the notion of 'expected inaccuracy' to articulate the following epistemic norm.

Diachronic Accuracy Norm: Suppose that at time t_1 an agent's prior credences are embodied by the probability distribution P, and that between times t_1 and t_2, she learns new information which imposes a constraint C on her posterior credences at time t_2. Then at time t_2 the agent should adopt credences encoded by a probability distribution Q such that (i) Q satisfies C, and (ii) for any Q' that satisfies C, $\exp_S(Q|P) \leq \exp_S(Q'|P)$, where $\mathfrak{I}_S$ is one's chosen proper scoring rule, i.e. Q should minimise expected inaccuracy (from the perspective of the agent's prior beliefs) amongst the class of possible probability functions that satisfy the learned constraint C.

The basic idea here is that when you update to take into account the new evidence, you should do so in the way that you expect to result in the most accurate posterior beliefs. In particular, the Diachronic Accuracy Norm requires agents to adopt the posterior probability distribution Q that minimises the quantity $D_S(P, Q) = \exp_S(Q|P) - \exp_S(P|P)$. When $\mathfrak{I}_S = \mathfrak{I}_B$, this quantity is just the square Euclidean distance between P and Q, as discussed in Section 3. When $\mathfrak{I}_S = \mathfrak{I}_L$, $\exp_S(Q|P) - \exp_S(P|P)$ is equal to the KL divergence between P and Q, or the IKL divergence between Q and P.[22] Now, we already know that the IKL divergence is an example of an f-divergence, so by trying to minimise expected inaccuracy as measured by the logarithmic scoring rule, we will end up minimising an f-divergence (the IKL divergence) to the learned constraint, which is exactly the learning procedure advocated in this paper. And this is where things get interesting. For, the following two facts have been established in the statistics literature (assuming that the sample space is finite):

1. For any strictly proper scoring rule $\mathfrak{I}_S$, the divergence measure defined by $D(P, Q) := \exp_S(Q|P) - \exp_S(P|P)$ is an example of a 'Bregmann divergence' (Gneiting & Raftery (2007)).

2. The KL divergence is the only element of the intersection of the set of Bregman divergences and the set of f-divergences (Amari (2009)).

Thus, if we assume that $D_S(P, Q)$ is equal to an f-divergence between P and Q, then that f-divergence must be the KL divergence. So the only way to minimise expected inaccuracy whilst minimising an f-divergence between P and Q is to minimise the KL divergence between P and Q, i.e. to minimise the IKL divergence between Q and P. In contrast, it is easily seen that minimising the KL divergence between Q and

---

[22]We are very grateful and indebted to Randall McCutcheon for correcting an error at this point in a previous version of this manuscript

P does not minimise expected inaccuracy for *any* strictly proper scoring rule. For, if $D_S(P,Q)$ were equal to $\text{KL}(Q,P)$, that would imply that $D_S(P,Q)$ was equal to $\text{IKL}(P,Q)$, which can't be the case since IKL is not a Bregman divergence. So the IKL divergence has the very special property of being an f-divergence such that minimising the divergence from Q to P minimises expected inaccuracy according to a strictly proper scoring rule (the log rule). We've seen that assuming conservativity and Jeffrey conditionalisation as basic epistemic norms motivates adopting an update procedure defined by the minimisation of an f-divergence. However, there are many alternative f-divergences, all of which define different updating strategies. We need extra normative principles if we hope to overcome this underdetermination. We've now seen that the Diachronic Accuracy Norm is enough to break the deadlock and identify the minimisation of the IKL divergence between Q and P as a normatively privileged rational updating procedure, since unlike i.e. the KL divergence, it is an f-divergence that minimises expected inaccuracy according to a strictly proper scoring rule. For those readers that prefer to think in terms of the KL divergence, we stress that this update procedure is equivalent to minimising the KL divergence between P and Q.[23]

Before moving on, it is worth pausing briefly to note that the preceding argument might be considered problematic by those that advocate quantifying epistemic utility by means of a scoring rule other than the log rule. For, the IKL divergence only minimises expected inaccuracy when the agent in question identifies inaccuracy with epistemic disutility as measured by $\mathfrak{I}_L$. And if one does not consider this conception of epistemic disutility to be plausible, then perhaps the argument will be less than entirely convincing. We have a couple of responses to these concerns. Firstly, we reiterate that the following two norms are widely accepted in the literature: (i) agents should employ update rules that are compatible with Jeffrey conditionalisation, and (ii) agents should always aim to minimise their expected accuracy as measured by a strictly proper scoring rule. And it is a mathematical theorem that these two norms can be satisfied in full generality if one adopts the updating rule that says 'minimise the Inverse Kullback Leibler divergence between your posterior and prior distributions (relative to the constraints given by the learning experience)'. The argument itself does not assert the superiority of the logarithmic scoring rule as a premise. Rather, it shows that if an agent wants to abide by two plausible norms relating to updating and epistemic utility, then they have good reason to conceive of epistemic utility in terms of the log rule. The argument does not rely on the admittedly controversial claim that the log rule encodes the most plausible conception of epistemic utility. Rather, it constitutes a new and compelling piece of evidence in support of that claim. Secondly, it should be noted that although the log score certainly has its critics, it is by no means universally rejected, and it is still very much a part of the very much open debate surrounding the proper formalisation of epistemic utility (see e.g. Fallis and Lewis (2016),

---

[23]Throughout the paper, we try to follow the common convention of minimising the divergence between Q and P, rather than between P and Q. For this reason, we talk in terms of minimising the IKL divergence between Q and P, rather than minimising the KL divergence between P and Q.

Levinstein (2017)). Thus, the fact that the conclusion of the argument seems to count in support of the log rule certainly shouldn't be considered a straightforward reductio against the premises.

6 Updating, Interpreting and Learning

Recall that van Fraassen *et al.* (1986) rejected JB3 as a constraint on prospective solutions to the Judy Benjamin example and viewed the upshot of that example to be that we need additional grounds for choosing between rival probabilistic distance measures. We've shown that the Diachronic Accuracy Norm allows us to solve this problem and identify minimising the IKL divergence as a normatively privileged update procedure for epistemically conservative Bayesian agents. In fact, this argument has the extra payoff that it vindicates the common intuition that it is irrational for Judy to change her credence in R after learning $Q(S|R) = 3/4$ (since minimising IKL ensures that her prior credence in R remains unchanged). It is interesting to recall that minimising any well known f-divergence other than IKL to the constraint $Q(S|R) = 3/4$ would cause Judy to change her credence in R. But there is an alternative way to ensure that JB3 is satisfied without employing the IKL divergence. Specifically, we could impose JB3 as an *additional constraint* on the update, so that Judy now interprets the content of the learned evidence not just by the constraint $Q(S|R) = 3/4$, but rather by the dual constraints that $Q(S|R) = 3/4$ *and* $Q(R) = 1/2 = P(R)$. If we do that, then we obtain the following result.

Proposition 5 (Simple Conditional − Non-Extreme f-Divergence^{+}) *An agent considers the propositions* A *and* B *and has a prior probability distribution* P *according to eqs. (3) (with* $0 < a, p, q < 1$*) defined over the corresponding propositional varables. The agent then learns* A $\rightarrow$ B *which imposes the constraints* $Q(B|A) = p' < 1$ *and* $Q(A) = P(A)$ *on the new distribution* Q *(see eq. (7)). Minimising any* f-*divergence between* Q *and* P *will result in the same posterior distribution as that which results from minimising the IKL divergence with the sole constraint that* $Q(B|A) = p' < 1$.

Proposition 5 tells us that if Judy updates by minimising an f-divergence to the dual constraints that $Q(S|R) = 3/4$ *and* $Q(R) = 1/2 = P(R)$, she will end up with the same credences that she would have had if she had rather updated by minimising the IKL divergence to the sole constraint that $Q(S|R) = 3/4$. Imposing JB3 as an additional constraint on the update completely eliminates the underdetermination of the rational update procedure. Importantly, this demonstrates that it is possible to obtain the solution to the Judy Benjamin problem forwarded by Douven & Romeijn (2012) (and independently justified by the Diachronic Accuracy Norm) by simply accepting JB3 as part of the content of Judy's evidence. And in the standard version of the example, it seems that there is a case for claiming that JB3 is indeed part of what Judy learns.

For, before informing Judy that $Q(S|R) = 3/4$, the radio operator explicitly states that 'I can't be sure where you are', which may be interpreted as telling Judy that she should remain agnostic about whether or not she's in Red Territory. However, one could remove this sentence from the example and it seems likely that many would have a strong intuition that Judy's credence in R should remain fixed, since the constraint $Q(S|R) = 3/4$ itself doesn't tell us anything definite about the probability of R being true. So why should Judy accept JB3 as part of the content of the learned conditional? To answer this question, it is useful to consider the following examples, which generalize the Judy Benjamin example in an illuminating way.

Lena the Scientist Example. Lena is considering a set of mutually exclusive (but not logically exhaustive) scientific hypotheses $H_1, \ldots, H_n$ and their relationship to a possible future experimental outcome E. Conditional on any of the H_i's, E has a low probability of being observed.

Case 1: Based on previous experience, Lena strongly believes that E will in fact be observed. In light of the fact that all of the considered theories conflict with this belief (since $P(E|H_i)$ is low for all i), she thinks that all of the considered theories are probably false, and a new theory is needed. However, one of her research assistants then informs her that the probability of E conditioned on H_1 was previously miscalculated, and is actually very high, i.e. Lena learns the conditional 'If H_1 then probably E'. Delighted to learn that one of her theories gives reasonable predictions of future experiments, Lena increases both the conditional probability of E given H_1 and the overall probability of H_1.

Case 2: Based on previous experience, Lena strongly believes that E will *not* be observed. She is satisfied that all of the candidate theories cohere with this belief. However, one of her research assistants then informs her that the probability of E conditioned on H_1 was previously miscalculated, and is actually very high, i.e. Lena learns the conditional 'If H_1 then probably E'. As well as increasing the conditional probability of E given H_1, Lena also dramatically reduces her belief in H_1, since it now conflicts with her belief that E won't be observed.

Case 3: Due to a lack of relevant evidence, Lena is agnostic about whether E will be observed. One of her research assistants then informs her that the probability of E conditioned on H_1 was previously miscalculated, and is actually very high, i.e. Lena learns the conditional 'If H_1 then probably E'. Lena increases the conditional probability of E given H_1, but keeps the probability of H_1 fixed, since she is indifferent about whether E will be observed, and so doesn't consider H_1's evidential relationship to E to bear on its plausibility.[24]

---

[24] One might wonder whether Lena's response in Cases 1 and 3 conflict with our defence of Proposition 1 in Section 2. There, we argued that it is natural that upon learning a strict conditional, an agent should typically become less confident in the

We take it that Lena's responses to the learned conditional are intuitively rational in all three of these cases. When she learns that H_1 coheres better with her other beliefs than she'd previously thought, she becomes more confident in its truth (Case 1). Conversely, when she learns that H_1 conflicts with her existing beliefs, she becomes more skeptical of its truth (Case 2). Finally, when the learned conditional has no implications for H_1's relation to her other beliefs (as in the Judy Benjamin example), she does not have any reason to change her belief in H_1. We've already seen that Lena's response in Case 3 can only be captured by minimising the IKL divergence or, equivalently, minimising an arbitrary f-divergence with the additional constraint that her credence in H_1 remains fixed. And this coheres with our observation that the IKL divergence is normatively privileged from the perspective of expected inaccuracy minimisation. But then cases 1 and 2 become problematic, since minimising the IKL divergence to ensure $Q(E|H_1) > P(E|H_1)$ will always result in the probability of H_1 remaining fixed (Proposition 4), and in cases 1 and 2 it is intuitively rational for Lena to *increase* and *decrease* her credence in H_1, respectively.[25] Thus it seems that epistemically conservative Bayesian agents are bound to deviate from the intuitively rational update in two of the three cases. This all suggests the following general updating strategy for learning non-strict conditionals.

Non-Strict Conditionals Updating Strategy: *An agent considers the propositions* A *and* B *and has a prior probability distribution* P *according to eqs. (3) (with* $0 < a, p, q < 1$*) defined over the corresponding propositional variables. The agent then learns* A $\rightarrow$ B *which imposes the constraint* $Q(B|A) = p' > P(B|A)$ *(*$p' < 1$*) on the posterior distribution* Q *(see eq. (7))). She will then update in one of three ways:*

(1) *If* $P(B) > 1/2$*, she will update by minimising the IKL divergence to the constraints (i)* $Q(B|A) = p'$ *and (ii)* $Q(A) > P(A)$*.*

(2) *If* $P(B) = 1/2$*, she will update by minimising the IKL divergence to the constraints (i)* $Q(B|A) = p'$ *and (ii)* $Q(A) = P(A)$*.*

(3) *If* $P(B) < 1/2$*, she will update by minimising the IKL divergence to the constraint (i)* $Q(B|A) = p'$ *and (ii)* $Q(A) < P(A)$*.*

According to this strategy, the probabilistic constraints that are used to interpret the learned conditional depend crucially on the prior beliefs of the agent. Whether the agent considers the learned conditional to

---

antecedent, since the learned conditional renders the antecedent more informative. However, the situation is different when the learned conditional is non-strict, i.e. when the learned conditional probability constraint is non-extreme. For, in the strict case, the extreme constraint actually rules out some of the worlds at which the antecedent is true. In the non-strict case, the constraint does not rule out any worlds but merely shifts the distribution of the probability mass across the worlds where the antecedent is true. There is no meaningful sense in which the antecedent becomes more informative in this case.

[25]The fact that some f-divergences will allow for an increase in the probability of the antecedent was pointed out to the authors on the *mathoverflow* thread at http://tinyurl.com/mathoverflow1. The f-divergences that satisfy this property are mathematically quite complex and have not been considered in the philosophical or psychological literature. Normatively, we view these divergences as being on a par with all other f-divergences except for the IKL divergence.

be evidence for, against or irrelevant to the antecedent will be determined by their prior credence in the consequent. This updating strategy has a couple of key advantages. Firstly, it yields the intuitively rational update in all of the Lena the scientist cases, and hence also in the Judy Benjamin example. Secondly, it allows us to maintain the position that minimising the IKL divergence provides the rational update in all learning scenarios. It is only the *interpretation* of the evidence that varies according to epistemic context. Once the evidence has been interpreted in terms of concrete constraints on the posterior distribution, the principles of Jeffrey conditionalisation, epistemic conservativity and the Diachronic Accuracy Norm jointly identify minimising the IKL divergence as a normatively privileged rational update procedure. The main philosophical upshot is that the evidential import of a learned conditional is partially determined by epistemic context. In some sense, the agent has to make a choice about how to interpret the evidence.

A similar conclusion was ultimately reached by Douven & Romeijn (2012), who also thought that the rational response to learning conditionals was partially determined by the agent's prior belief state, and the relative degrees to which their individual beliefs are 'entrenched'. However, the nature of their contextualism was of a fundamentally different sort to that advocated here. Specifically, they argued that the choice of distance measure to be minimised should be determined by context, and advocated both the IKL divergence and Hellinger distance as providing the rational responses to some learning scenarios. We contend that this approach is unprincipled and overly vague because, as Douven and Romeijn acknowledge, it is completely unclear exactly how one should go about choosing a particular distance measure for a given learning scenario and epistemic context. Here, we've argued that epistemically conservative Bayesian agents should always update by minimising the IKL divergence. There is no contextual relativity in terms of the rational probabilistic update rule to apply relative to a set of definite constraints on the posterior. Rather, the contextual relativity occurs only at the stage at which the agent interprets the learned conditional as a set of such constraints.

One possible objection to the proposed update strategy is that $1/2$ is an arbitrary and artificial threshold for determining the response to the learned conditional. An alternative approach would be to adopt a Lockean theory of belief (see e.g. chapter 4 of Foley (2000), Leitgeb (2017)), according to which an agent believes a proposition A if and only if $P(A)$ is above some pre-determined threshold value $t \in [1/2, 1]$, i.e. if and only if the agent has a sufficiently high confidence in A. This would suggest the following variation of the update strategy:

Lockean Non-Strict Conditionals Updating Strategy: *An agent considers the propositions* A *and* B *and has a prior probability distribution* P *according to eqs.* (3) *(with* $0 < a, p, q < 1$*) defined over the corresponding propositional variables. The agent then learns* A $\rightarrow$ B *which imposes the constraint* $Q(B|A) = p' > P(B|A)$ $(p' < 1)$ *on the posterior distribution* Q *(see eq.* (7)*)). She will then update in one of three ways:*

(1) *If $P(B) \geq t$, she will update by minimising the IKL divergence to the constraints (i) $Q(B|A) = p'$ and (ii) $Q(A) > P(A)$.*

(2) *If $P(B) \in (1-t, t)$, she will update by minimising the IKL divergence to the constraints (i) $Q(B|A) = p'$ and (ii) $Q(A) = P(A)$.*

(3) *If $P(B) \leq 1-t$, she will update by minimising the IKL divergence to the constraint (i) $Q(B|A) = p'$ and (ii) $Q(A) < P(A)$.*

The Lockean update strategy is equivalent to the original strategy in the special limiting case where $t = 1/2$. But it is widely acknowledged that a threshold of $1/2$ provides an overly permissive characterisation of the notion of belief (it seems unnatural to say that you *believe* that a coin with bias of 0.501 will land heads). When $t > 1/2$, the Lockean strategy instructs agents to increase (decrease) the probability of the antecedent if they have a prior belief in the consequent (the negation of the consequent). Again, the motivation is that when an agent learns a non-strict conditional $A \to B$, they should change their belief in A depending on whether the relevant conditional probability constraint makes A cohere better or worse with their prior beliefs. It is easy to see the Lockean strategy has all the same virtues as the standard strategy (of which it is a generalisation). The problem of choosing an appropriate threshold value t is of course a deep one that goes significantly beyond the scope of this paper (see e.g. Fitelson & Shear (forthcoming)).

7   Conclusion

According to the position defended here, an agent can never simply update on a non-extreme conditional probability constraint such as $Q(S|R) = 3/4$. Before updating on such a constraint, they need to consider how the new information relates to their prior beliefs, and use that consideration to determine whether what they've learned bears meaningfully on the plausibility of the conditioning event (and if so, in what way). Once they've gone through this pre-processing step, the epistemic norms described above will compel them to update by minimising the IKL divergence to the set of posterior distributions satisfying the resulting set of constraints (which will consist of the original conditional probability constraint and an additional constraint on the probability of the antecedent). The idea that agents need to go through this complex process of interpretation and updating is motivated primarily by the fact that no simple updating rule or distance-minimisation procedure is capable of giving the intuitively correct verdict in all of the relevant thought experiments. But of course, intuitions may differ from person to person, and one could simply reject what we take to be the intuitive solution to the Lena the Scientist cases. Whether this is an acceptable approach would seem to depend largely on the strength and uniformity of the intuitions it violates. One could turn

to the psychology of reasoning in order to gain more definite insights here. In recent years, there has been much empirical work analyzing the outcomes of dynamic conditional reasoning experiments in terms of the minimisation of the KL divergence (see e.g. Singmann *et* al. (2016), Ali *et al.* (2011)). If it were the case that the Lockean updating strategy described above provided a significantly better account of people's reasoning habits than simple KL divergence minimisation (which already performs quite well), that would suggest that the intuitions behind these thought experiments are quite deeply rooted, and hence that violating them would incur a high theoretical cost. Conversely, if it turned out that these intuitions weren't found to play a key role in people's reasoning habits, that would suggest that nothing substantial is lost if we rather advocate simply minimising a fixed f-divergence (such as the IKL divergence) to the given conditional probability constraint in every case. Whether or not one takes the Lockean updating strategy to be necessary will, we suspect, depend crucially on one's position on fundamental issues of philosophical methodology regarding i.e. how seriously to take common philosophical intuitions and the normative significance of people's actual reasoning habits. For now, we remain agnostic about these wider issues and are content to have presented an updating strategy that is consistent with foundational Bayesian norms and chimes well with pre-theoretic intuitions. If one considers these intuitions to be unimportant, then of course it is possible to disregard this more complex procedure in favour of simply minimising the IKL divergence to the relevant conditional probability constraint.

We have also highlighted the previously unrecognized fact that minimising the IKL divergence between Q and P seems to be the only way to generalize Bayesian updating methods to non-propositional learning scenarios in a manner that (1) is intuitively conservative, and (2) minimises expected inaccuracy as measured by a strictly proper scoring rule. This shows, contra van Fraassen *et al.* (1986), that it is possible to identify a normatively privileged formalisation of epistemic conservativity. In response to our arguments, one might also take it to be significant that minimising the KL divergence between Q and P has been shown to provide a good account of people's actual reasoning patterns in dynamic conditional inference tasks. However, the possibility of using alternative f-divergences to describe people's reasoning patterns is a new one that has not yet been seriously explored. If it turns out that other f-divergences provide a better fit to people's actual reasoning habits than the IKL divergence does, that could be taken to show that humans don't generally reason in a way that minimises expected inaccuracy. Again, whether or not one would interpret such a result as being normatively significant is likely to depend on one's stance regarding the prospects of naturalistic epistemology and the extent to which one views normative epistemology as having a duty to vindicate everyday reasoning practices (see e.g. Bishop and Trout (2005)). More generally, it is clear that the arguments presented here open a new frontier at the interface of the psychology of reasoning and normative Bayesian epistemology.

Acknowledgements

A Proofs

A.1 Proposition 1

We first compute the posterior distribution from eqs. (4). We begin with

$$
\begin{aligned}
P(\neg A \vee B) &= P(\neg A) + P(B) - P(\neg A, B) \\
&= \overline{a} + a\,p + \overline{a}\,q - \overline{a}\,q \\
&= a\,p + \overline{a}.
\end{aligned}
\tag{8}
$$

Next, we compute

$$
\begin{aligned}
P^*(A, B) &= P(A, B | \neg A \vee B) = \frac{P(A, B)}{P(\neg A \vee B)} = \frac{a\,p}{a\,p + \overline{a}} \\
P^*(A, \neg B) &= P(A, \neg B | \neg A \vee B) = 0 \\
P^*(\neg A, B) &= P(\neg A, B | \neg A \vee B) = \frac{P(\neg A, B)}{P(\neg A \vee B)} = \frac{\overline{a}\,q}{a\,p + \overline{a}} \\
P^*(\neg A, \neg B) &= P(\neg A, \neg B | \neg A \vee B) = \frac{P(\neg A, \neg B)}{P(\neg A \vee B)} = \frac{\overline{a}\,\overline{q}}{a\,p + \overline{a}}.
\end{aligned}
\tag{9}
$$

We can now calculate

$$
\begin{aligned}
P^*(\mathrm{A}) &= P^*(\mathrm{A},\mathrm{B}) + P^*(\mathrm{A},\neg\mathrm{B}) = \frac{a\,p}{a\,p+\overline{a}} = a - \frac{a\,\overline{a}\,\overline{p}}{a\,p+\overline{a}} < a \\
P^*(\mathrm{B}) &= P^*(\mathrm{A},\mathrm{B}) + P^*(\neg\mathrm{A},\mathrm{B}) = \frac{a\,p+\overline{a}\,q}{a\,p+\overline{a}} = \frac{P(\mathrm{B})}{a\,p+\overline{a}} > P(\mathrm{B}). \quad \blacksquare
\end{aligned}
$$

A.2   Proposition 2

$$
\begin{aligned}
P^*(\mathrm{R}) &= P(\mathrm{R}|\neg\mathrm{R}\vee\neg\mathrm{S},\mathrm{D}) \\
&= \frac{P(\mathrm{R},\neg\mathrm{R}\vee\neg\mathrm{S},\mathrm{D})}{P(\neg\mathrm{R}\vee\neg\mathrm{S},\mathrm{D})} \\
&= \frac{P(\mathrm{R},\neg\mathrm{S},\mathrm{D})}{P(\mathrm{R}\neg\mathrm{S},\mathrm{D}) + P(\neg\mathrm{R},\mathrm{S},\mathrm{D}) + P(\neg\mathrm{R}\neg\mathrm{S},\mathrm{D})} \\
&= \frac{P(\mathrm{R},\neg\mathrm{S},\mathrm{D})}{P(\mathrm{R}\neg\mathrm{S},\mathrm{D}) + P(\neg\mathrm{R},\mathrm{D})} \tag{10}
\end{aligned}
$$

Using eq. (6), we note that $P(\mathrm{R},\neg\mathrm{S},\mathrm{D}) = P(\neg\mathrm{S}|\mathrm{R},\mathrm{D})\,P(\mathrm{R},\mathrm{D}) = P(\mathrm{R},\mathrm{D})$. Hence,

$$
P^*(\mathrm{R}) = \frac{P(\mathrm{R},\mathrm{D})}{P(\mathrm{R},\mathrm{D}) + P(\neg\mathrm{R},\mathrm{D})}. \tag{11}
$$

Finally, we use the fact that R and D are probabilistically independent and obtain

$$
\begin{aligned}
P^*(\mathrm{R}) &= \frac{P(\mathrm{R})\,P(\mathrm{D})}{P(\mathrm{R})\,P(\mathrm{D}) + P(\neg\mathrm{R})\,P(\mathrm{D})} \\
&= P(\mathrm{R}). \quad \blacksquare
\end{aligned}
$$

A.3   Proposition 3

P and Q are given by Table 1 and we set $Q(\mathrm{E}) = e'$. Q satisfies the following two constraints:

H	E	P	Q
1	1	p_1	q_1
1	0	p_2	q_2
0	1	p_3	q_3
0	0	p_4	q_4

Table 1: The probability distributions P and Q

$$q_1 + q_3 = e' \tag{12}$$

$$q_1 + q_2 + q_3 + q_4 = 1 \tag{13}$$

Hence, we have to minimise

$$L = \sum_{i=1}^{4} p_i \, f(q_i/p_i) + \lambda(q_1 + q_3 - e') + \mu(q_1 + q_2 + q_3 + q_4 - 1). \tag{14}$$

To do so, we differentiate by q_k and obtain

$$\frac{\partial L}{\partial q_1} = f'(q_1/p_1) + \lambda + \mu = 0 \quad , \quad \frac{\partial L}{\partial q_3} = f'(q_3/p_3) + \lambda + \mu = 0 \tag{15}$$

$$\frac{\partial L}{\partial q_2} = f'(q_2/p_2) + \mu = 0 \quad , \quad \frac{\partial L}{\partial q_2} = f'(q_4/p_4) + \mu = 0. \tag{16}$$

Eqs. (15) imply that

$$q_1 = \alpha \, p_1 \quad , \quad q_3 = \alpha \, p_3. \tag{17}$$

Similarly, eqs. (16) imply that

$$q_2 = \beta \, p_2 \quad , \quad q_4 = \beta \, p_4. \tag{18}$$

We insert eqs. (17) into eq. (12) and obtain

$$\alpha = \frac{e'}{p_1 + p_3}. \tag{19}$$

Next, we insert eqs. (17), (18) and (19) into eq. (13) and obtain

$$\beta = \frac{\overline{e'}}{p_2 + p_4}. \tag{20}$$

Hence,

$$
\begin{aligned}
Q(\mathrm{H}) &= q_1 + q_2 \\
&= \frac{p_1}{p_1 + p_3} \, e' + \frac{p_2}{p_2 + p_4} \, \overline{e'} \\
&= P(\mathrm{H|E}) \, Q(\mathrm{E}) + P(\mathrm{H|\neg E}) \, Q(\neg \mathrm{E}). \quad \blacksquare
\end{aligned}
$$

A.4 Corollary 1

Following our discussion in Section 2, the agent learns the material conditional A ⊃ B with certainty. We can then use conditonalization (which is a special case of Jeffrey conditionalisation from Proposition 3) and obtain the same result as in Proposition 1. ▋

A.5 Proposition 4

To find a' and q', we compute the f-divergence between Q and P,

$$F = a\,p\,f\left(\frac{a'\,p'}{a\,p}\right) + a\,\overline{p}\,f\left(\frac{a'\,\overline{p'}}{a\,\overline{p}}\right) + \overline{a}\,q\,f\left(\frac{\overline{a'}\,q'}{\overline{a}\,q}\right) + \overline{a}\,\overline{q}\,f\left(\frac{\overline{a'}\,\overline{q'}}{\overline{a}\,\overline{q}}\right). \tag{21}$$

To find the minimum of F, we first differentiate F by q' and set the resulting expression equal to zero. Hence,

$$\overline{a'}\left[f'\left(\frac{\overline{a'}\,q'}{\overline{a}\,q}\right) - f'\left(\frac{\overline{a'}\,\overline{q'}}{\overline{a}\,\overline{q}}\right)\right] = 0. \tag{22}$$

The convexity of f then implies that $q' = q$. Inserting this into eq. (26) yields

$$F = a\,p\,f\left(\frac{a'\,p'}{a\,p}\right) + a\,\overline{p}\,f\left(\frac{a'\,\overline{p'}}{a\,\overline{p}}\right) + \overline{a}\,f\left(\frac{\overline{a'}}{\overline{a}}\right). \tag{23}$$

To find a', we differentiate this expression by a' and set the resulting expression equal to zero and obtain

$$p'\,f'\left(\frac{a'\,p'}{a\,p}\right) + \overline{p'}\,f'\left(\frac{a'\,\overline{p'}}{a\,\overline{p}}\right) = f'\left(\frac{\overline{a'}}{\overline{a}}\right). \tag{24}$$

This equation cannot be solved in full generality. The result will depend on the specific form of convex function f. In all cases considered here, eq. (24) simplifies to an expression of the form

$$\frac{a'\,\overline{a}}{\overline{a'}\,a} = \frac{1}{l}, \tag{25}$$

with the likelihood ratio l. Hence, $a' = a/(a + \overline{a}\,l)$. For the IKL divergence, we have $f(x) = -\log x$ and hence $f'(x) = -1/x$. Inserting this into eq. (24) then yields $l = 1$ and therefore $a' = a$. (Note that this result only holds for $p' < 1$. See Proposition 1.) The results for the other f-divergencies obtain accordingly.

Finally, we calculate

$$
\begin{aligned}
P'(\mathrm{B}) - P(\mathrm{B}) &= a'\,p' + \overline{a'}\,q - a\,p - \overline{a}\,q \\
&= \frac{1}{a + \overline{a}\,l} \cdot [a\,p + a\,(p'-p) + \overline{a}\,l\,q - (a+\overline{a}\,l)(a\,p + \overline{a}\,q)] \\
&= \frac{1}{a + \overline{a}\,l} \cdot [a\,\overline{a}\,p + a\,(p'-p) + a\,\overline{a}\,l\,q - a\,\overline{a}\,q - a\,\overline{a}\,l\,p] \\
&= \frac{a}{a + \overline{a}\,l} \cdot [(p'-p) - \overline{a}\,(l-1)\,(p-q)] \\
&= [(p'-p) - \overline{a}\,(l-1)\,(p-q)] \cdot P'(\mathrm{A}),
\end{aligned}
$$

from which the statement in the proposition follows. ∎

A.6   Proposition 5

We follow the same procedure as in the proof of Proposition 4. Noting the additional constraint $Q(\mathrm{A}) = P(\mathrm{A})$ (i.e. $a' = a$), the f-divergence between Q and P is given by

$$
F = a\,p\,f\left(\frac{p'}{p}\right) + a\,\overline{p}\,f\left(\frac{\overline{p'}}{\overline{p}}\right) + \overline{a}\,q\,f\left(\frac{q'}{q}\right) + \overline{a}\,\overline{q}\,f\left(\frac{\overline{q'}}{\overline{q}}\right). \tag{26}
$$

To find the minimum of F and to determine the value of the only free parameter (i.e. q'), we differentiate F by q' and set the resulting expression equal to zero. Hence,

$$
\overline{a}\left[f'\left(\frac{q'}{q}\right) - f'\left(\frac{\overline{q'}}{\overline{q}}\right)\right] = 0. \tag{27}
$$

The convexity of f then implies that $q' = q$.

Comparing this result with the results for the IKL reported in Proposition 4, we find that the corresponding posterior probability distributions are the same. ∎

B   The Ski Trip Example

We consider the three propositional variables E, S and B mentioned above. P denotes the prior probability distribution, P^* denotes the posterior probability distribution after learning **Ski1** and **Ski2** via conditionalisation, and Q denotes the posterior probability distribution that one obtains after minimising a f-divergence between Q and P taking the constraints implied by **Ski1** and **Ski2** into account. See Table 2.

E	S	B	P	P^*	Q
1	1	1	p_1	p_1^*	q_1
1	1	0	p_2	p_2^*	q_2
1	0	1	p_3	p_3^*	q_3
1	0	0	p_4	p_4^*	q_4
0	1	1	p_5	p_5^*	q_5
0	1	0	p_6	p_6^*	q_6
0	0	1	p_7	p_7^*	q_7
0	0	0	p_8	p_8^*	q_8

Table 2: The probability distributions P, P^* and Q for the ski trip example

B.1 Conditionalisation

To determine the posterior distribution P^*, we apply Bayesian conditionalisation on the two new pieces of information. As $P(\cdot, \cdot, \neg B | \cdot, B) = 0$, one sees immediately that $p_2^* = p_4^* = p_6^* = p_8^* = 0$. Furthermore, $P(E, \neg S, \cdot | \neg E \vee S, \cdot) = 0$. Hence, also $p_3^* = 0$. This leaves us with the computation of

$$
\begin{aligned}
p_1^* &:= P(\mathrm{E}, \mathrm{S}, \mathrm{B} | \neg \mathrm{E} \vee \mathrm{S}, \mathrm{B}) = \frac{P(\mathrm{E}, \mathrm{S}, \mathrm{B})}{P(\neg \mathrm{E} \vee \mathrm{S}, \mathrm{B})} = \frac{p_1}{p_1 + p_5 + p_7} \\
p_5^* &:= P(\neg \mathrm{E}, \mathrm{S}, \mathrm{B} | \neg \mathrm{E} \vee \mathrm{S}, \mathrm{B}) = \frac{P(\neg \mathrm{E}, \mathrm{S}, \mathrm{B})}{P(\neg \mathrm{E} \vee \mathrm{S}, \mathrm{B})} = \frac{p_5}{p_1 + p_5 + p_7} \\
p_7^* &:= P(\neg \mathrm{E}, \neg \mathrm{S}, \mathrm{B} | \neg \mathrm{E} \vee \mathrm{S}, \mathrm{B}) = \frac{P(\neg \mathrm{E}, \neg \mathrm{S}, \mathrm{B})}{P(\neg \mathrm{E} \vee \mathrm{S}, \mathrm{B})} = \frac{p_7}{p_1 + p_5 + p_7}.
\end{aligned}
\tag{28}
$$

Hence, $P^*(\mathrm{E}) = p_1^*$. Note further that $P(\mathrm{E}) = p_1 + p_2 + p_3 + p_4$. Hence, $P^*(\mathrm{E}) > P(\mathrm{E})$ iff $p_1 > (p_1 + p_2 + p_3 + p_4)(p_1 + p_5 + p_7)$. As $P(\mathrm{E})$ is small, this condition holds if p_5 and p_7 are small (as mentioned above).

To get a more intuitive understanding of the example, we make a plausible additional assumption about Harry's prior beliefs, namely that he considers the variables B and E to be probabilistically independent given S, i.e. $P(\mathrm{B}|\mathrm{E}, \mathrm{S}) = P(\mathrm{B}|\mathrm{S})$. Intuitively, this means that Harry only considers E to be evidentially relevant to B to the extent that it counts as evidence for S. Whether or not Sue passes the exam is only probabilistically relevant to whether or not she buys a new ski outfit when we do not know whether or not she is going on a ski trip. Sue's passing the exam is only relevant to the extent that it counts as evidence for her going on a ski trip. Another way to see this is to consider the intuitive causal relationships that hold between the variables. Whether Sue passes the exam causally influences whether her father takes her on a ski trip, which causally influences whether she buys a new ski outfit, and there are no other causal relationships between the

Figure 1: The Bayesian Network for the ski trip example

variables. But then the causal Markov condition, which tells us that 'causes screen of their effects' (see e.g. Sprites *et al.* (2000)), implies that B is probabilistically independent of E once we condition on any value of S. Thus, we can safely assume that B is probabilistically independent of E conditional on any value of S. This assumption is encoded in the Bayesian Network depicted in Figure 1.

To complete the Bayesian Network, we have to fix the prior probability of E, i.e.

$$P(\text{E}) = e, \tag{29}$$

and the conditional probabilities

$$P(\text{S}|\text{E}) = x_1 \quad , \quad P(\text{S}|\neg\text{E}) = y_1 \tag{30}$$

$$P(\text{B}|\text{S}) = x_2 \quad , \quad P(\text{B}|\neg\text{S}) = y_2.$$

With this, we obtain:

$$P^*(\text{E}) = \frac{e\,x_1\,x_2}{e\,x_1\,x_2 + \overline{e}\,(y_1\,x_2 + \overline{y_1}\,y_2)} \tag{31}$$

Hence, $P^*(\text{E}) > P(\text{E})$ iff

$$k_0 := \frac{x_1\,x_2}{y_1\,x_2 + \overline{y_1}\,y_2} > 1. \tag{32}$$

To proceed, we have to determine whether the condition $k_0 > 1$ holds. From the story it is clear that there is no reason for Harry to expect Sue to buy a skiing outfit if she is not invited on a skiing trip. Hence, $y_2 \approx 0$ and therefore

$$k_0 \approx \frac{x_1}{y_1}. \tag{33}$$

It is also clear from the story that it is more likely that Sue's father invites her for a ski trip if she passes the exam than if she does not pass the exam. Hence, $x_1 > y_1$ and thus $k_0 > 1$.

We have therefore shown that conditionalising on the material conditional E $\supset$ S and on the proposition B gives the intuitively correct result for the ski trip example.

B.2 Minimising the f-Divergence

Ski1 implies that $Q(\text{B}) = q_1 + q_3 + q_5 + q_7 = 0$. Hence, $q_1 = q_3 = q_5 = q_7 = 0$. **Ski1** implies that $Q(\text{S}|\text{E}) = 1$. Hence, $Q(\text{E}, \neg\text{S}) = q_3 + q_4 = 0$ and therefore $q_4 = 0$. With this, we can compute the f-divergence:

$$F = p_1\, f(q_1/p_1) + p_5\, f(q_5/p_5) + p_7\, f(q_7/p_7) \tag{34}$$

Differentiating F by q_1, q_5 and q_7, respectively, and setting the resulting expressions equal to zero yields (using the convexity of f) $q_1/p_1 = q_5/p_5 = q_7/p_7$. Hence, for $i = 1, 5, 7$,

$$q_i = \frac{p_i}{p_1 + p_5 + p_7}.$$

Comparing this result with eqs. (28), we confirm that $P^* = Q$ for the ski trip example.

References

Adams, E. (1975). *The Logic of Conditionals. Synthese Library*, Vol. 86. Boston: D. Reidel.

Ali, N., N. Chater and M. Oaksford (2011). The Mental Representation of Causal Conditional Reasoning. *Cognition* 119(3): 403–418.

Amari, S.I. (2009). α-Divergence Is Unique, Belonging to Both f-Divergence and Bregman Divergence Classes. *IEEE Transactions on Information Theory* 55(11). DOI 10.1109/TIT.2009.2030485.

Bennett, J. (2003). *A Philosophical Guide to Conditionals.* Oxford: Oxford University Press.

Bishop, M. and J.D. Trout (2005). The Pathologies of Standard Analytic Epistemology. *Nous* 39 (4): 694–714.

Bovens, L. (2010). Judy Benjamin is a Sleeping Beauty. *Analysis* 70 (1): 23–26.

Bovens, L. and J. L. Ferreira (2010). Monty Hall Drives a Wedge between Judy Benjamin and the Sleeping Beauty: A Reply to Bovens. *Analysis* 70 (3): 473–81.

Bovens, L. and S. Hartmann (2003). *Bayesian Epistemology.* Oxford: Oxford University Press.

Csiszár, I. (1967). Information-Type Measures of Difference of Probability Distributions and Indirect Observation. *Studia Scientiarum Mathematicarum Hungarica* 2: 299–318.

Csiszár, I. (1975). I-Divergence Geometry of Probability Distributions and Minimisation Problems. *Annals of Probability* 3(1):146–158.

Csiszár, I. (1991). Why Least Squares and Maximum Entropy? An Axiomatic Approach to Inference for Linear Inverse Problems. *Annals of Statistics* 19(4): 2032–2066.

Csiszár, I. (2008). Axiomatic Characterizations of Information Measures. *Entropy* 10(3): 261–273.

Diaconis, P. and S. Zabell (1982). Updating Subjective Probability. *Journal of the American Statistical Association* 77: 822–830.

Douven, I. (2012). Learning Conditional Information. *Mind & Language* 27 (3): 239–263.

Douven, I. (2015). *The Epistemology of Indicative Conditionals.* Cambridge: Cambridge University Press.

Douven, I. and R. Dietz (2011). A Puzzle about Stalnaker's Hypothesis. *Topoi* 30: 31–37.

Douven, I. and J. W. Romeijn (2012). A New Resolution of the Judy Benjamin Problem. *Mind* 120 (479): 637–670.

Edgington, D. (1995). On Conditionals. *Mind* 104: 235–329.

Fernbach, P. and C. Erb (2013). A Quantitative Causal Model Theory of Conditional Reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 39(5): 1327–1343.

Fallis, D. and Lewis, P. J. (2016). The brier rule is not a good measure of epistemic utility (and other useful facts about epistemic betterness). *Australasian Journal of Philosophy*, 94 (3): 576–590.

Fitelson, B. and T. Shear (forthcoming). Two Approaches to Belief Revision. To appear in *Erkenntnis.*

Foley, R. (2000). *Working Without a Net.* Oxford: Oxford University Press.

Fraassen, B. van (1981). A Problem for Relative Information Minimisers in Probability Kinematics. *British Journal for the Philosophy of Science* 32: 375–379.

Fraassen, B. van, R.I.G. Hughes and G. Harman (1986). A Problem for Relative Information Minimisers, Continued. *British Journal for the Philosophy of Science* 37: 453–463.

Gärdenfors, P. and H. Rott (1995). Belief Revision. In *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 4. Oxford: Oxford University Press, pp. 35–132.

Gneiting, T. and A. Raftery (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association* 102(477). DOI 10.1198/016214506000001437.

Greaves, H. and D. Wallace (2006). Justifying Conditionalisation: Conditionalisation Maximizes Expected Epistemic Utility. *Mind* 115(459): 607–632.

Grice, P. (1975). Logic and Conversation. In Cole, P. and Morgan, J. (eds), *Syntax and semantics. 3: Speech acts.* New York: Academic Press, pp. 41–58.

Grove, A. and J. Halpern (1997). Probability Update: Conditioning vs. Cross-Entropy. In *Proceedings of the 13th Annual Conference on Uncertainty in Artificial Intelligence.* San Francisco: Morgan Kaufmann, pp. 208–214.

Hajek, A. (2012). The Fall of 'Adam's Thesis'?. *Journal of Logic, Language and Information* 21(2): 145–161.

Huisman, L.M. (2017). Learning from Simple Indicative Conditionals. *Erkenntnis* 82(3): 583–601.

Kern-Isberner, G. (2001). *Conditionals in Nonmonotonic Reasoning and Belief Revision.* Berlin: Springer.

Leitgeb, H. (2017). *The Stability of Belief: How Rational Belief Coheres with Probability.* Oxford: Oxford University Press.

Leitgeb, H. and R. Pettigrew (2010). An Objective Justification of Bayesianism II: The Consequences of Minimising Inaccuracy. *Philosophy of Science* 77: 236–272.

Levinstein, B. (2012). Leitgeb and Pettigrew on Accuracy and Updating. *Philosophy of Science* 79(3): 413–424.

Levinstein, B. (2017). A Pragmatist's Guide to Epistemic Utility. *Philosophy of Science* 84(4): 613–638..

Lewis, D. (1976). Probabilities of Conditionals and Conditional Probabilities. *Philosophical Review* 85: 297–315.

Lewis, D. (1999). *Papers in Metaphysics and Epistemology.* Cambridge: Cambridge University Press.

Myrvold, W. (2018). Learning is a Risky Business. Preprint available at http://philsci-archive.pitt.edu/14254/.

Nayak, A., M. Pagnucco and P. Peppas (1996). Learning From Conditionals: Judy Benjamin's Other Problems. In W. Wahlster (ed.): *ECAI96: 12th European Conference on Artificial Intelligence.* New York: Wiley, pp. 75–79.

Oaksford, M. and N. Chater (2007). *Bayesian Rationality: The Probabilistic Approach to Human Reasoning.* Oxford: Oxford University Press.

Pettigrew, R. (2016). *Accuracy and the Laws of Credence.* Oxford: Oxford University Press.

Popper, K. and D. Miller (1983). A Proof of the Impossibility of Inductive Probability. *Nature* 302: 687–688.

Schwan, B. and R. Stern (2017). A Causal Understanding of When and When Not to Jeffrey Conditionalise. *Philosopher's Imprint* 17(8): 1–21.

Singmann, H., K.C. Klauer and S. Beller (2016). Probabilistic Conditional Reasoning: Disentangling Form and Content with the Dual-Source Model. *Cognitive Psychology* 88: 61–87.

Skovgaard Olsen, N., H. Singmann and K.C. Klauer (2016). The Relevance Effect and Conditionals. *Cognition* 150: 26–36.

Skyrms, B. (1984). *Pragmatics and Empiricism.* New Haven: Yale University Press.

Sprites, P., C. Glymour and R. Scheines. (2000). *Causation, Prediction and Search.* Springer: New York.

Szabo, Z (ed). (2005). *Semantics versus Pragmatics.* Oxford: Oxford University Press.

Trpin, B. (2018). *Influence of Conditionals on Belief Updating.* Unpublished PhD thesis. University of Ljubljana.

Uffink, J. (1995). Can the Maximum Entropy Principle be Explained as a Consistency Requirement? *Studies in the History and Philosophy of Modern Physics* 26: 223–261.

Vasudevan, A. (2018). Entropy and Insufficient Reason: a Note on the Judy Benjamin Problem. *British Journal for the Philosophy of Science* https://doi.org/10.1093/bjps/axy013.

Weisberg, J. (2011). Varieties of Bayesianism. In D. Gabbay, S. Hartmann and J. Woods (eds.): *Handbook of the History of Logic.* Amsterdam: Elsevier, pp. 477–551.