



# Bayesians Still Don't Learn from Conditionals

Mario Günther<sup>1,2</sup>  · Borut Trpin<sup>1,3</sup>

Received: 13 December 2021 / Accepted: 29 August 2022 / Published online: 19 September 2022  
© The Author(s) 2022

## Abstract

One of the open questions in Bayesian epistemology is how to rationally learn from indicative conditionals (Douven, 2016). Eva et al. (*Mind* 129(514):461–508, 2020) propose a strategy to resolve this question. They claim that their strategy provides a “uniquely rational response to any given learning scenario”. We show that their updating strategy is neither very general nor always rational. Even worse, we generalize their strategy and show that it still fails. Bad news for the Bayesians.

## 1 Introduction

How should a rational Bayesian agent update her beliefs when she learns an indicative conditional? Philosophers and cognitive scientists alike have tried hard to provide an answer.<sup>1</sup> Douven (2016, Ch. 6) surveys the extant answers to this question and comes to the conclusion that none is satisfactory. Should Bayesians then abandon all hope and admit defeat when it comes to updating on conditionals? Not according to Eva et al. (2020) who put forth “a normatively privileged updating procedure for this kind of learning” (p. 461). Their procedure can be roughly glossed as follows. A Bayesian agent learns “If  $A$  then  $C$ ” by minimizing the Inverse Kullback-Leibler (IKL) divergence subject to two constraints, one on the conditional probability  $Q(C|A)$ , the other on the probability of the antecedent  $Q(A)$ .

Eva et al. argue that learning by minimizing the IKL divergence is normatively privileged. Their reasons are that this minimization method generalizes Jeffrey conditionalizing, an update rule that is widely embraced as rational, and that it minimizes expected epistemic inaccuracy. However, to learn “If  $A$  then  $C$ ” by minimizing

---

<sup>1</sup> See, e.g., the efforts of Evans and Over (2004) and Oaksford and Chater (2007) in the psychology of reasoning, and one of its critical assessments by Skovgaard-Olsen et al. (2016).

✉ Mario Günther  
Mario.Guenther@lmu.de

<sup>1</sup> LMU Munich, München, Germany

<sup>2</sup> Australian National University, Canberra, Australia

<sup>3</sup> University of Maribor, Maribor, Slovenia

the IKL divergence only subject to the constraint on  $Q(C|A)$  keeps the probability of the antecedent fixed. In general, as Eva et al. (2020, p. 492) note, this is implausible. In many examples, the probability of the antecedent should intuitively change. Hence, they impose an additional constraint on  $Q(A)$ .

Here we will show that Eva et al. (2020) general updating strategy for learning conditionals fails for two reasons.<sup>2</sup> Firstly, their updating strategy does not account for cases where propositions besides the antecedent and consequent are relevant. This is surprising, to say the least, because they write in the first half of their paper “If more than two relevant propositions are involved, then these have to be accounted for and modelled in a proper way” (p. 475). Secondly, the constraint on  $Q(A)$  leads in many scenarios to intuitively false results. We generalize their updating strategy to overcome the restriction to two variables. And yet, as we will see, the constraint on  $Q(A)$  is still inappropriate for non-evidential conditionals.

We will present Eva et al.’s 2020 updating strategy before we investigate its claimed generality. Then we will show that their updating strategy fails for a number of examples, including the Ski Trip Example. We generalize their updating strategy to give some hope to the Bayesians. Yet we conclude that Bayesians still don’t learn from conditionals.

## 2 Eva et al.’s (2020) Updating Strategy

Here is the framework of Eva et al. (2020). An agent’s initial degrees of beliefs, or simply beliefs, are represented by a joint probability distribution  $P$  over binary propositional variables  $A$  and  $C$ . Her beliefs after learning a conditional are represented by the posterior probability distribution  $Q$ . The prior probability distribution  $P$  is assumed to be given by the values for  $P(A)$ ,  $P(C|A)$ , and  $P(C|\neg A)$ .

The idea behind Eva et al.’s 2020 updating strategy goes as follows: an agent learns an indicative conditional “If  $A$  then  $C$ ” by imposing the constraint  $Q(C|A) > P(C|A)$ . If her prior belief in  $C$  is high (or low), this constraint indicates that  $A$  is rather coherent (or incoherent) with her prior beliefs. As a result, the probability of  $A$  should increase (or decrease). The question upon learning a conditional is thus how well the antecedent coheres with the agent’s prior belief in the consequent.

This is their updating strategy (see pp. 492–493). An agent has a prior probability distribution  $P$  specified by values in the open interval  $(0,1)$  for  $P(A)$ ,  $P(C|A)$ , and  $P(C|\neg A)$ . The learning of the indicative conditional “If  $A$ , then  $C$ ” imposes the

<sup>2</sup> We focus our criticism on Eva et al.’s proposal because we take it to be the most elaborate Bayesian account of learning conditionals. As Collins et al. (2020, p. 14) observe, Eva et al.’s procedure belongs to the distance-minimization approach which they call “the standard method to model updating on conditionals”. We call an account of learning conditionals Bayesian if and only if it is compatible with the core Bayesian principles: probabilistic coherence and (Jeffrey) conditionalization (Hájek 2019, Sec. 3.3.4). The accounts of learning conditionals due to Douven (2016) and Günther (2017, 2018), for example, are not Bayesian on this understanding. Their rules for updating propositional beliefs deviate from (Jeffrey) conditionalization, and so may lead to posterior distributions that are different from and thus incompatible with (Jeffrey) conditionalization.

constraint  $1 > Q(C|A) > P(C|A)$  on the posterior distribution. Upon learning the conditional, she will update her beliefs by minimizing the IKL divergence between  $Q$  and  $P$ . This happens in one of three ways:

- (1) If  $P(C) > 1/2$ , she will update by minimizing the IKL divergence to the constraints
  - (i) the value of  $Q(C|A)$ , and
  - (ii)  $Q(A) > P(A)$ .
- (2) If  $P(C) = 1/2$ , she will update by minimizing the IKL divergence to the constraints
  - (i) the value of  $Q(C|A)$ , and
  - (ii)  $Q(A) = P(A)$ .
- (3) If  $P(C) < 1/2$ , she will update by minimizing the IKL divergence to the constraints
  - (i) the value of  $Q(C|A)$ , and
  - (ii)  $Q(A) < P(A)$ .

Notice that the posterior probability value of the antecedent increases, remains the same, or decreases depending on the agent's prior belief in the consequent  $C$  (constraint (ii) in each clause). If an agent's prior belief in the consequent is low, for instance, then to learn that the probability of the consequent given the antecedent is higher than previously thought should decrease the belief in the antecedent.<sup>3</sup> Before we show why the updating strategy fails, let us have a look at its generality.

### 3 Generality

Eva et al. (2020) claim that their updating strategy is “a single update procedure that provides a uniquely rational response to any given learning scenario” (p. 485). However, their updating strategy is not as general as they purport it to be. Here are some preliminary notes to better understand the restrictions of generality.

In their “general updating strategy”, an agent considers only two binary propositional variables  $A$  and  $C$  — let us call them basic propositions — and comes

<sup>3</sup> As Eva et al. (2020, p. 494) acknowledge, their threshold of  $1/2$  seems to be somewhat arbitrary. In response, they generalize their updating strategy by replacing the threshold value  $1/2$  with a Lockean threshold  $t \in [1/2, 1]$ . (For details about Lockean thresholds and Lockean theories of belief, see Foley (1993) and Leitgeb (2014, 2017). The outcome of their generalized strategy is still predetermined by the value of  $P(C)$  relative to the picked threshold. It will become clear that the Lockean updating strategy is thus as well susceptible — *mutatis mutandis* — to the arguments that follow. We therefore confine ourselves to discussing the non-Lockean version of their updating strategy.

equipped with a prior probability distribution  $P$  fully specified by the values for  $P(A)$ ,  $P(C|A)$ , and  $P(C|\neg A)$  (see pp. 492 and 466). Any probability distribution  $P$  is defined relative to a probability space  $\langle \Omega, \mathcal{F}, P \rangle$ , where  $\Omega$  is a set of “elementary events”, “possible worlds”, or “maximally specific propositions”, and  $\mathcal{F}$  is a  $\sigma$ -algebra of subsets of  $\Omega$  that is closed under disjunction, conjunction, and negation. Importantly,  $P$  is defined on all members of  $\mathcal{F}$ , but on *no other* propositions (see, e.g., Hájek (2003, p. 278)).

The basic propositions  $A$  and  $C$  determine the following set of maximally specific propositions:

$$\Omega = \{A \wedge C, A \wedge \neg C, \neg A \wedge C, \neg A \wedge \neg C\}.$$

It is easy to see that the  $\sigma$ -algebra  $\mathcal{F}$  of subsets of  $\Omega$  does not contain any basic proposition other than  $A$  and  $C$ , and also no Boolean combination including any other basic proposition.  $\mathcal{F}$  contains, for example, no third basic proposition  $D$ . And while  $P(A)$  and  $P(A \vee C)$ , for instance, are defined,  $P(D)$  and  $P(D \wedge \neg A)$  are not.

We have just seen that the  $\sigma$ -algebra underlying their “general” updating strategy is limited to two basic propositions, the antecedent  $A$  and the consequent  $C$ . The crux is that, when learning “If  $A$  then  $C$ ”, their agent comes equipped with a joint distribution *only* defined over the  $\sigma$ -algebra that has been generated by the binary variables  $A$  and  $C$ . As a consequence, their updating strategy cannot consider that there might be another basic proposition — apart from the antecedent and the consequent — relevant to the learning of a conditional. This is surprising because they claim in the first half of the paper that there are learning scenarios where additional basic propositions “have to be accounted for and modelled in a proper way”.<sup>4</sup> It remains unclear, however, how their updating strategy can account for additional basic propositions. The underlying algebra, metaphorically speaking, has simply no space for those additional propositions.

One may object to our point that it is obvious how the underlying  $\sigma$ -algebra is to be extended when an additional basic proposition is or becomes relevant. But we have learned from Williamson (2003) that extending the underlying language is a serious challenge to Bayesianism. Many authors think the extension should not affect the probability ratios of those propositions which have already been members of the  $\sigma$ -algebra (see, e.g., Wenmackers and Romeijn (2016) and Bradley (2017, Ch. 12)). However, Steele and Stefánsson (2020) convincingly argue against this obvious strategy: there are cases where algebra extensions should affect the probability ratios of the “old” propositions. So far, it is an unresolved challenge for Bayesian accounts how the underlying algebra should be extended in general. Eva et al. are silent on how the  $\sigma$ -algebra should be extended by additional basic propositions, or how the probabilities should be distributed when a new basic proposition comes into play. And so they fail to specify in their final account how additional basic propositions can influence what is learned. We consider a potential solution in Section 5.

Here is another limitation. Their updating strategy is restricted to the open interval, or what they call “non-strict” conditionals. In the first half of their paper,

<sup>4</sup> A case in point is the Ski Trip Example, which we will discuss in the next section.

they extensively deal with “strict” conditionals — conditionals that impose the strict constraint  $Q(C|A) = 1$ . They defend the view that strict conditionals can be learned by minimizing the IKL divergence, or any other  $f$ -divergence for that matter (see p. 480). Crucially, the strategy for strict conditionals does not invoke a constraint like clause (ii) in their “general” updating strategy. And so their strategy for strict conditionals is not guided by the rationale of their “general” updating strategy: the coherence with the prior belief in the conditional’s consequent plays no special role.

We do not see why strict conditionals require their own Bayesian updating strategy based on a different rationale. To save the generality of their updating strategy, we thus proceed upon the assumption that it applies to strict conditionals as well. We believe this is reasonable — consider a case where an agent learns a conditional with probability  $1 - \epsilon$  with  $\epsilon$  very close to 0 instead of learning a strict conditional. There is no obvious reason that rational learning ought to come with additional learning constraints regarding the antecedent just because one learns a conditional with  $1 - \epsilon$  probability instead of 1. In any case, we run into counterexamples even if we bite the bullet and restrict the generality of their updating strategy to non-strict conditionals as we show in the next section.<sup>5</sup>

Finally, learning a conditional “If  $A$  then  $C$ ” imposes the constraint that the posterior value of  $Q(C|A)$  is higher than the prior value of  $P(C|A)$ . This constraint seems to preclude any Bayesian agent from learning that “If  $A$  then  $C$ ” is less likely than previously thought. One “might learn that the probability of it raining, given that there is a thunderstorm, is much lower than previously imagined”, as Eva et al. (2020, p. 465) observe themselves. They show that learning from conditionals should proceed by minimizing the IKL divergence with a constraint on  $Q(C|A) = k$  regardless of whether  $k$  is greater than, equal to or less than  $P(C|A)$  (see Corollary 1, p. 480 for strict conditionals and Proposition 5 (p. 490) for non-strict conditionals).<sup>6</sup> But their learning strategy goes beyond this corollary and the proposition because it also includes a constraint on the antecedent and it is not clear what the constraint on the antecedent ought to be in such cases. However, this limitation might be quite innocent. If  $Q(C|A) < P(C|A)$ , the rationale of their updating strategy would presumably prescribe that the constraint (ii) is reversed in the clauses (1)-(3). Alternatively, learning that “If  $A$  then  $C$ ” is less likely than previously thought could be understood as learning the constraint that  $Q(\neg C|A)$  is more likely than previously thought.<sup>7</sup> While this limitation is relatively easy to fix this way or another, Eva et al. make us guess what their preferred fix is.

We have seen that their final updating strategy has no space in the underlying algebra for basic propositions in addition to the antecedent and consequent. Furthermore, they restrict their final updating strategy to non-strict conditionals without any

<sup>5</sup> Eva et al.’s updating method is anyways not as general as the non-Bayesian account put forth by Günther (2018). His method of learning conditionals applies to strict and non-strict conditionals, and to conditionals whose antecedents have zero probability.

<sup>6</sup> Thanks to an anonymous reviewer for pointing this out.

<sup>7</sup> We would like to thank an anonymous referee for this alternative idea.

need to do so. We will see more clearly in the next section why their updating strategy is not as general as they claim it to be.

#### 4 The Constraint on the Antecedent

Eva et al. (2020) make the following claim. Relative to the two constraints on the posterior, their updating strategy applies without contextual relativity. Specifically, “the contextual relativity occurs only at the stage at which the agent interprets the learned conditional as a set of such constraints” (p. 494). This claim clashes with their claim that relevant propositions other than the antecedent and the consequent are to be considered. Let us illustrate this clash by revisiting 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. (Douven and Dietz 2011, p. 33)

In this example, the prior beliefs  $P(B)$ ,  $P(E)$ ,  $P(S)$  of “Sue buying a ski outfit”, of “Sue passing the exam”, and of “Sue’s father taking her on a ski vacation”, respectively, are rather low. In the first half of their paper, Eva et al. (2020) observe that the agent must learn both in the Ski Trip Example, the conditional “If  $E$ , then  $S$ ” and the proposition  $B$  (see pp. 469–471). In their Appendix B.2, they show that the minimization of the IKL divergence between  $Q$  and  $P$  subject to the constraints  $Q(B) = 1$  and  $Q(S|E) = 1$  leads to the desired result: the posterior probability of Sue passing the exam is higher than its prior.<sup>8</sup>

Their “general” updating strategy, however, does either not apply to the Ski Trip Example, leads to the intuitively false result, or shows that the constraint (ii) in clauses (1)–(3) should be given up, at least in general. Since the prior probability of Sue’s father taking her to a ski vacation is low, clause (3) of their updating strategy applies. By constraint (ii), the posterior probability of Sue passing the exam must be lower than its prior. But this is the intuitively false result contradicting (Eva et al., 2020) solution of the Ski Trip Example in the first half of the paper. And notice, the constraint (ii) says  $Q(E) < P(E)$  independent of the value of  $Q(B)$ . In fact, the proposition  $B$  cannot be considered because their updating strategy is set up such that the agent considers only a  $\sigma$ -algebra over the antecedent  $E$  and the consequent  $S$ . Hence, their final updating strategy fails for the Ski Trip Example, an example they investigate in the same paper.

<sup>8</sup> In fact, they show the more general result that the minimization of any  $f$ -divergence leads to the desired result. Beware the typos at the beginning of their Appendix B.2, e.g.,  $Q(B) = 0$  should read  $Q(B) = 1$  and the second occurrence of “Ski 1” should read “Ski 2”.

You might think that their updating strategy does not apply to the Ski Trip Example because a strict conditional is learned. Perhaps, the approach presented by Eva et al. (2020) may be charitably interpreted as twofold: the “general” strategy we discuss here is only general with respect to non-strict conditionals, while learning from strict conditionals is governed by different rules described in the first half of their paper (pp. 461–475). After all, their strategy for strict conditionals handles the case with ease (pp. 469–471). However, we can easily imagine a very minor adaptation of the Ski Trip Example to show that the issue does not go away. To see this, just modify the example by replacing the strict constraint  $Q(S|E) = 1$  by the non-strict constraint  $Q(S|E) = .99$ . This is not an unreasonable possibility either: be it for testimonial or perceptual reasons. The source who tells the conditional might not be fully reliable or the statement is delivered in less than ideal circumstances (e.g., on a somewhat noisy street or over the phone). In any case, the situation is almost the same. Yet their “general” updating strategy still requires constraint (ii),  $Q(E) < P(E)$ , which directly contradicts the desired result. Ignoring the restriction to the antecedent and consequent for a moment, the constraint (ii) excludes that the relevant proposition  $B$  can have the desired impact.

To give up constraint (ii) would collapse their updating strategy to the minimization of the IKL divergence subject to the constraint  $Q(C|A)$ .<sup>9</sup> Yet this is no option either. Minimizing the IKL divergence keeps the probability of the antecedent fixed when there are no explicit constraints on the probability of the antecedent. This is bad news in light of examples like this, where the probability of the antecedent should intuitively change upon learning a conditional.

Consider a variation of the Ski Trip Example. Suppose Harry initially thinks it is just as likely as not that Sue’s father will take her on a ski vacation, that is  $P(S) = 1/2$ . He might simply be totally uncertain in this respect, or he might know the father to be very spontaneous. In either case, after learning that “If Sue passed the exam, her father will take her on a ski vacation” and seeing her buying a ski outfit, Harry should become more confident that she passed the exam. Yet, because of  $P(S) = 1/2$ , their updating strategy imposes the constraint  $Q(E) = P(E)$ , which is trivially satisfied when minimizing the IKL divergence subject to the conditional probability constraint. Harry must apparently remain just as doubtful about his friend’s success as he initially was. Hence, simply giving up on the constraint (ii) and returning to the minimization of the IKL divergence subject to the constraint  $Q(C|A)$  only also does not solve the problem.

Finally, we can easily construct a counterexample to show that the constraint (ii) of clause (1) also has unwanted implications:

Bob has a friend Laura who regularly performs a rain dance but he does not know why. Their common friend Sarah explains to him that Laura believes the dance may bring rain and says ‘If it is sunny tomorrow, Laura will do a rain dance’. Upon learning this conditional, Bob’s degree of belief that it will be sunny tomorrow remains unchanged.

<sup>9</sup> Indeed, Eva et al. (2020, p. 36) flirt with the idea that one could simply minimize the IKL divergence to the relevant conditional probability constraint.

The relevant prior beliefs here are  $P(S)$  and  $P(D)$  of “It will be sunny tomorrow” and “Laura will do a rain dance”, respectively. Because  $P(D) > 1/2$  — Bob expects Laura to perform a rain dance — clause (1) applies and with it the constraint (ii):  $Q(S) > P(S)$ . In other words, Bayesians who follow Eva et al. (2020) learning strategy shall become more confident that it will be sunny because of (the high expectation of) the rain dance. Whether or not the dance exerts influence on the weather, this change in view seems rather absurd. We have thus demonstrated that all three clauses of their learning strategy are prone to counterexamples: clause (1) is threatened by our Raindancer Example, (2) by our variant of the Ski Trip Example, and (3) by the original Ski Trip Example.

Another example Eva et al. investigate is Van Fraassen’s (1981) Judy Benjamin Problem. Judy learns the conditional “If you are in Red Territory ( $R$ ), the odds are 3:1 that you are in Second Company Area ( $S$ )”. The constraint on the posterior is thus  $Q(S|R) = 3/4$ . They remark that “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 is in Red Territory*” (p. 483, emphasis in original). From this, they infer that the constraint  $Q(R) = P(R)$  is “part of the content of the learned conditional” (p. 490). Their updating strategy is designed to impose this additional constraint since  $P(S) = 1/2$  in the Judy Benjamin Problem. But it is widely recognized in the literature that the “additional constraint” is rather a desideratum (see Van Fraassen 1981; Douven and Romeijn 2011). They explicitly build in their learning procedure what should follow from it.<sup>10</sup>

Eva et al. (2020) back up the claim that the learning of conditionals imposes the respective constraint (ii) on the antecedent by a generalized Judy Benjamin Problem, the Lena the Scientist Example. It is obvious that their updating strategy is motivated by the three cases of the Lena Example. Unsurprisingly, their updating strategy “yields the intuitively rational update in all of the Lena the scientist cases, and hence also in the Judy Benjamin example” (p. 493). But this argument seems to involve reverse engineering: the updating strategy is designed to capture the examples, and the examples then support the strategy. The argument would only convince if the updating strategy would apply successfully to other examples. But we have just shown that, as it stands, it faces a variety of counterexamples.

The previous paragraph suggests that their updating strategy is tailor-made for a few hand-picked examples. And indeed, in the Lena the Scientist Example with its three cases, Lena’s prior belief in the consequent determines whether the learned conditional is “evidence for, against, or irrelevant to the antecedent” (p. 493). But, of course, this evidential idea behind their updating strategy fails in general. As long as a conditional “does not tell anything” about the probability of the antecedent, but says only what should be expected under the supposition of its antecedent, their updating strategy leads generally to the intuitively false results. The conditional “If

<sup>10</sup> Consider a scenario that is just like the Judy Benjamin Problem except that  $P(S) = .6$ . For this scenario, their updating strategy implies that the probability of being in Red Territory increases. But why should  $Q(R) > P(R)$  be the case if “the learned conditional didn’t seem to tell her anything about whether or not she is in Red Territory”?



it rains, the match will be cancelled”, for instance, arguably says not much about the probability of rain. If the prior belief in the cancellation of the match is high or low, their updating strategy prescribes that the posterior of the antecedent is high or low, respectively. In particular, suppose your initial belief that the match will be cancelled is low. Their updating strategy then commands that, upon learning the conditional, you should decrease your belief that it rains. But, on a causal reading of the conditional, the probability for rain should not depend on whether the match will be cancelled.<sup>11</sup>

## 5 Generalizing Eva et al.'s, 2020 Updating Strategy

Is there a way to save Eva et al.'s updating strategy? We argue that there is no obvious way. Yet we may generalize their updating strategy in response to our criticisms. Let  $*$  be a variable ranging over  $\{<, =, >\}$ . There is a function  $*^{-1}$  which gives  $>$  if  $*$  takes the value  $<$ ,  $=$  if  $*$  takes the value  $=$ , and  $<$  if  $*$  takes the value  $>$ . The generalization can be stated as follows.

An agent is represented by a prior probability distribution  $P_0$  that is defined on all members of a  $\sigma$ -algebra generated by all relevant basic propositions. Let  $P$  be the result from Jeffrey conditionalizing on any proposition  $F$  the agent comes to believe (to a certain degree). Learning an indicative conditional “If  $A$  then  $C$ ” imposes a constraint  $Q(C|A) * P(C|A)$  on the posterior distribution  $Q$ , if  $P(A) > 0$ . The agent then updates her beliefs by minimizing the IKL divergence between  $Q$  and  $P$  subject to the value of  $Q(C|A)$ , and

- (a) in case  $P(C) > 1/2$ ,  $Q(A) * P(A)$  if  $Q(C|A) * P(C|A)$ ,
- (b) in case  $P(C) = 1/2$ ,  $Q(A) = P(A)$ , and
- (c) in case  $P(C) < 1/2$ ,  $Q(A) *^{-1} P(A)$  if  $Q(C|A) * P(C|A)$ .<sup>12</sup>

Let us revisit the generality issues. Our generalization of Eva et al.'s updating strategy makes explicit how an agent learns that a conditional is less likely than previously thought. When you learn  $Q(C|A) < P(C|A)$  and  $P(C) > 1/2$ , clause (a) implies the constraint  $Q(A) < P(A)$ . Provided  $C$  is more likely than not, you lower your belief in  $A$  upon learning that  $A$  and  $C$  do not cohere as good as you thought.

Similarly, when you learn  $Q(C|A) < P(C|A)$  and  $P(C) < 1/2$ , clause (c) implies the constraint  $Q(A) > P(A)$ . Provided  $C$  is less likely than not, you increase your belief in  $A$  upon learning that  $A$  and  $C$  do not cohere as good as you thought. By means of an example, suppose that some mother is quite confident that if her child

<sup>11</sup> We acknowledge that the reading of the conditional, which their updating strategy presupposes, does make sense. If you are (nearly) certain that the match will not be cancelled, this is good evidence that it will not rain. However, we would maintain that a causal reading makes sense as well, and so no general learning strategy for conditionals should simply exclude it. Günther (2022) proposes a distinction between causal and evidential conditionals. A similar distinction figures in the debate on evidential vs. causal decision theory. See, for instance, Gibbard and Harper (1978) and Joyce and Gibbard (1998).

<sup>12</sup> Like Eva et al.'s updating strategy, the generalization is amenable to replace the threshold  $1/2$  by a Lockean threshold.

is vaccinated, the child will be autistic. She also believes that autism is relatively uncommon. After learning that the conditional “if you vaccinate your child, the child will become autistic” is a myth, she lowers her belief in the corresponding conditional probability. By clause (c) of our generalization, she becomes more certain that she should vaccinate her child. The generalized strategy thus provides the intuitively correct responses in the spirit of Eva et al.: the generalization keeps the evidential idea that your belief in  $A$  should change depending on your prior belief in  $C$ . At the same time, clauses (a) and (c) allow to reverse Eva et al.’s constraint (ii) when you learn that  $A$  and  $C$  cohere less than previously thought.

Our generalization represents an agent by a probability distribution that is defined on all members of a  $\sigma$ -algebra generated by all relevant basic propositions. This allows the agent to Jeffrey conditionalize on any relevant proposition the agent learns before learning a conditional. As compared to Eva et al.’s updating strategy, the prior probability distribution is replaced by the probability distribution after having conditionalized on all relevant non-conditional propositions. We can illustrate this difference by revisiting the Ski Trip Example. On our generalization, the prior probability  $P_0(S)$  that Sue’s father takes her on a ski vacation is conditionalized on the observation  $B$  that Sue buys a skiing outfit. So  $P(S) = P_0(S|B)$ . Intuitively,  $P_0(S|B)$  seems to be more likely than  $P_0(S)$ . But does  $P_0(S|B)$  surpass the threshold of  $1/2$ ?<sup>13</sup> If so, how could one show this? This is but one remaining challenge.

Apart from this remaining challenge, the original strategy and the generalization still fail for non-evidential conditionals. Causal readings of conditionals — recall the above “If it rains, the match will be cancelled” — are one example of non-evidential conditionals. The rain causes the cancellation of the match. But the cancellation of the match does not cause it to rain. When you read the conditional causally, you do not believe that the weather depends on the cancellation. Believed causal relations may trump the evidential coherence of beliefs.

Finally, the generalization is not restricted to conditional probabilities in the open interval. The only restriction is that the probability of the antecedent must be greater than zero,  $P(A) > 0$ . Edgington (1995, pp. 264–265) thinks that this restriction is a defining feature of indicatives rather than a bug. And yet we can understand and learn conditionals like “If the spinner landed exactly at noon, it was spinning slightly before noon”. We can learn this conditional although the probability of the antecedent is zero.<sup>14</sup> As to the zero probability intolerance, we see no easy way out for the Bayesian.<sup>15</sup>

The generalization is still in the spirit of Eva et al.’s updating strategy. Both strategies neglect by how much the probability of the antecedent changes if it does. And still, we admit that they might be approximately appropriate for many coherence or evidential conditionals. However, both cannot handle “if it rains” cases like the example of match cancellation and the Sundowners Example. Some conditionals are simply not evidential.

<sup>13</sup> Our intuitions are not so clear: many people are not invited to a ski vacation by their father even though they buy a skiing outfit.

<sup>14</sup> There are uncountably many points in time. So you cannot assign non-zero and equal probability to each point. Otherwise your probability measure would exceed 1 by far.

<sup>15</sup> Perhaps, the Bayesian can retreat to non-standard probability measures, such as Popper measures.

One might object that Eva et al. can handle causal conditionals by introducing suitable enabling and disabling variables that have not been mentioned in the respective examples. To illustrate this objection, consider the Sundowners Example.

Sarah and her sister Marian have arranged to go for sundowners at the Westcliff hotel tomorrow. Sarah feels 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 ( $R$ ), we cannot have sundowners at the Westcliff ( $\neg S$ )’. Upon learning this conditional, Sarah sets her probability for sundowners and rain to 0, but she does not adapt her probability for rain. Douven & Dietz (2011, pp. 645–646)

In the first half of their paper, Eva et al. treat “If  $R$ , then  $\neg S$ ” as a strict conditional. Upon learning the conditional and nothing else, Sarah’s belief in rain,  $P(R)$ , would decrease. And note that Sarah does not know of any wedding party, she “thinks they can always enjoy the view from inside”. However, Eva et al. argue that Sarah tacitly learns about a disabler  $D$ : “In the event of rain, something (for example a wedding party) will prevent us from having sundowners inside”. (p. 473) They furthermore argue that Sarah thinks  $R$  and  $D$  are probabilistically independent. This gives them the desired result: if Sarah conditionalizes her beliefs on the conditional *and* the disabler,  $P(R)$  remains unchanged.<sup>16</sup>

Eva et al. take quite some liberty in postulating presumed facts that are not mentioned in the Sundowners Example, like the assumption that Sarah learns about a disabler and that this disabler is probabilistically independent of rain. Even if we were to grant this resolution of the Sundowners Example, it does not transfer to their updating procedure for non-strict conditionals, or our generalization thereof. Sarah starts out with a prior probability distribution  $P_0$ , which she then (Jeffrey) conditionalizes on  $D$  in order to obtain  $P$ . Her belief in rain remains unchanged because  $P(R) = P_0(R|D) = P_0(R)$  — the disabling condition  $D$  is by Eva et al.’s own lights independent of rain.

By learning the conditional Sarah imposes the constraint  $Q(\neg S|R) > P(\neg S|R)$  on her posterior distribution  $Q$ . She then updates her beliefs by minimizing the IKL divergence between  $Q$  and  $P$  subject to the value of  $Q(\neg S|R)$  and the constraint  $Q(R) < P(R)$  of case (c). Her belief that they cannot have sundowners,  $P(\neg S)$ , is rather low: she thinks “there is some chance” that it will rain, and the disabler only says that they cannot have sundowners *if* it rains. Indeed, Sarah learns that rain and having sundowners do not cohere as good as she previously thought. But this shouldn’t lower her probability of rain on a causal understanding of the conditional. Our

<sup>16</sup> Two observations may be of interest. First, the disabler  $D$  itself has an internal conditional structure that looks pretty much like the conditional Sarah is supposed to learn. Second, why do Eva et al. treat the Sundowners Example in a different way than the Judy Benjamin Problem? There, they don’t introduce a disabler, but say that the constraint  $Q(R) = P(R)$  is “part of the content of the learned conditional” because “the learned conditional didn’t seem to tell her anything about whether or not she was in Red Territory”. But in the Sundowners Example, the conditional Sarah learned likewise “didn’t seem to tell her anything about whether or not” it rains.

consideration of the Sundowners Example has also illustrated that Eva et al.'s two updating strategies for strict and non-strict conditionals differ significantly.

## 6 Parting Thoughts

Eva et al. (2020) claim that they offer a rational and general updating strategy for the learning of conditionals. We have shown that their updating strategy is neither very general nor always rational: in many cases it is not applicable, and where it is, it sometimes delivers the intuitively false results.

We have extended their updating strategy to cover the cases, where more than two propositions need to be accounted for. Otherwise their updating strategy has not even the chance to deliver the intuitively correct results for the Ski Trip Example. We admit that their strategy, and more so our generalization in their spirit, might be appropriate for many coherence or evidential conditionals. The point is that both cannot handle weather conditionals like “if it is sunny” or “if it rains”, as in the Raindancer Example and the example of match cancellation. The culprit is the respective constraint on the posterior probability of the antecedent required by their updating strategy. Hence, even if their updating strategy is “normatively privileged”, it is inappropriate for the learning of indicative conditionals — at least in general. Some conditionals are simply not evidential.

Orthodox and heterodox Bayesians have put great effort to model the learning of indicative conditionals. However, the diverse and technically advanced tools of Bayesianism, such as conditional probabilities, probabilistic independences or Bayesian networks, and distance minimization do not seem to solve the problem of learning conditionals in its most general form.<sup>17</sup> In a similar vein, Eva et al. (2020) write in their Conclusion “that no simple updating rule or distance-minimization procedure is capable of giving the intuitively correct verdict in all of the relevant thought experiments” (p. 495). While there might be a Bayesian solution a long way down the line, it may be time to look for other tools to model the learning of conditionals. Great effort does not replace a hammer which fits the problem.<sup>18</sup>

**Funding** Open Access funding enabled and organized by Projekt DEAL.

### Declarations

**Conflict of Interest** The authors declare no competing interests.

<sup>17</sup> Collins et al. (2020) show that no present Bayesian approach accounts satisfactorily for the behavioral data, including a precursor of Eva et al.'s updating strategy based on distance minimization to be found in Eva and Hartmann (2018). Vandeburgh (2021, p. 2428) puts forth a promising account, where the first step of learning a conditional is to learn a causal model, or a set of structural equations, that expresses “the right dependence” between the conditional's antecedent and consequent. However, his proposal is not yet general in the sense that he does not tell us *which* structural equations are to be learned upon learning a conditional (p. 2429). Instead he takes some liberty in selecting the structural equations for the examples he discusses.

<sup>18</sup> For a different hammer, see Günther (2017, 2018).

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

## References

- Bradley, R. (2017). *Decision theory with a human face*. Cambridge University Press.
- Collins, P.J., Krzyżanowska, K., Hartmann, S., Wheeler, G., & Hahn, U. (2020). Conditionals and testimony. *Cognitive Psychology*, *122*, 101329.
- Douven, I. (2016). *The epistemology of indicative conditionals: Formal and empirical approaches*. Cambridge University Press.
- Douven, I., & Dietz, R. (2011). A puzzle about Stalnaker's hypothesis. *Topoi*, *30*(1), 31–37.
- Douven, I., & Romeijn, J.-W. (2011). A new resolution of the Judy Benjamin problem. *Mind*, *120*(479), 637–670.
- Edgington, D. (1995). On conditionals. *Mind*, *104*(414), 235–329.
- Eva, B., & Hartmann, S. (2018). Bayesian argumentation and the value of logical validity. *Psychological Review*, *125*(5), 806–821.
- Eva, B., Hartmann, S., & Rad, S.R. (2020). Learning from conditionals. *Mind*, *129*(514), 461–508.
- Evans, J.S.B.T., & Over, D.E. (2004). *If*. Oxford University Press.
- Foley, R. (1993). *Working without a net*. Oxford University Press.
- Gibbard, A., & Harper, W.L. (1978). Counterfactuals and two kinds of expected utility. In *Ifs* (pp. 153–190). Springer.
- Günther, M. (2017). Learning conditional and causal information by Jeffrey imaging on Stalnaker conditionals. *Organon F*, *24*(4), 456–486.
- Günther, M. (2018). Learning conditional information by Jeffrey imaging on Stalnaker conditionals. *Journal of Philosophical Logic*, *47*(5), 851–876.
- Günther, M. (2022). Causal and evidential conditionals. *Minds & Machines*. <https://doi.org/10.1007/s11023-022-09606-w>.
- Hájek, A. (2003). What conditional probability could not be. *Synthese*, *137*(3), 273–323.
- Hájek, A. (2019). Interpretations of probability. In E.N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. Fall 2019 edn.
- Joyce, J., & Gibbard, A. (1998). Causal decision theory. In S. Barbera, P. Hammond, & C. Seidl (Eds.) *Handbook of Utility Theory*, (Vol. 1: Principles pp. 627–666). Kluwer Academic Publishers.
- Leitgeb, H. (2014). The stability theory of belief. *The Philosophical Review*, *123*(2), 131–171.
- Leitgeb, H. (2017). *The stability of belief: how rational belief coheres with probability*. Oxford University Press.
- Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford cognitive science series oxford. Oxford University Press.
- Skovgaard-Olsen, N., Singmann, H., & Klauer, K.C. (2016). The relevance effect and conditionals. *Cognition*, *150*, 26–36.
- Steele, K., & Stefánsson, H.O. (2020). Belief revision for growing awareness. *Mind*, *130*(520), 1207–1232.
- Van Fraassen, B.C. (1981). A problem for relative information minimizers in probability kinematics. *The British Journal for the Philosophy of Science*, *32*(4), 375–379.
- Vandenburgh, J. (2021). Conditional learning through causal models. *Synthese*, *199*(1), 2415–2437.
- Wenmackers, S., & Romeijn, J.-W. (2016). New theory about old evidence. *Synthese*, *193*(4), 1225–1250.
- Williamson, J. (2003). Bayesianism and language change. *Journal of Logic Language and Information*, *12*(1), 53–97.

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