

EVA, HARTMANN AND RAD ON KULLBACK-LEIBLER MINIMIZATION

ABSTRACT. We address problems with a recent attempt by Eva, Hartmann and Rad to justify the Kullback-Leibler divergence minimization solution to van Fraassen’s Judy Benjamin problem.

A preprint of Eva, Hartmann and Rad (2019) touts the virtues of van Fraassen’s (1981) solution of the Judy Benjamin Problem by Kullback-Leibler (KL) divergence minimization. The paper is forthcoming in *Mind* and deeply flawed. I’ll restrict my discussion to the latter point.

In fact, I’ll restrict my discussion to a section of the paper that is spoken of in that paper’s introduction as follows:¹

we present an analysis of van Fraassen’s Judy Benjamin example in terms of f -divergence minimization 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 Kullback-Leibler divergence) as the unique probabilistic distance measure that agents should utilize if they hope to achieve maximally accurate beliefs....

The portion of the argument outlined above runs as follows. Let W be a set of worlds and let $\mathcal{J}_S(Q, w)$ denote inaccuracy of credence function Q at world w according to some proper scoring rule, e.g. the logarithmic scoring rule, which is given by $\mathcal{J}_L(Q, w) = -\log(Q(w))$. Define

$$\exp_S(Q|P) = \sum_{w \in W} P(W) \cdot \mathcal{J}_S(Q, w),$$

i.e. the expectation of the inaccuracy of Q from the perspective of P . Eva, Hartmann and Rad now consider the following putative norm:

¹In particular, I won’t devote space in the main text to the fact that the authors follow a recent trend of treating sentences such as “If you are in Red Territory, then the odds are 3 : 1 that you are in Second Company area” as indicative conditionals. I don’t believe that they can be conditionals of any sort, because I don’t believe that any proposition can be said to serve as consequent. They’re more plausibly just instances of a common form of vernacular for reporting conditional probabilities.

Diachronic Accuracy Norm (Leitgeb and Pettigrew 2010): 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 \mathcal{J}_S is one's chosen proper scoring rule.

Eva, Hartmann and Rad go on to say:

In particular, the diachronic accuracy norm requires agents to adopt the posterior probability distribution Q that minimizes the quantity $\exp_S(Q|P) - \exp_S(P|P)$. (...) When $\mathcal{J}_S = \mathcal{J}_L$, $\exp_S(Q|P) - \exp_S(P|P)$ is equal to the KL-divergence between Q and P .

But the KL-divergence between Q and P is defined by

$$D_{KL}(Q|P) = - \sum_{w \in W} Q(w) \log \left(\frac{P(w)}{Q(w)} \right).$$

Eva, Hartmann and Rad are therefore mistaken. When $\mathcal{J}_S = \mathcal{J}_L$,

$$\exp_S(Q|P) - \exp_S(P|P) = - \sum_{w \in W} P(W) \left(\log \frac{Q(w)}{P(w)} \right),$$

which is not the KL-divergence between Q and P at all, but rather the KL-divergence between P and Q , i.e. the so called "IKL-divergence" (for *Inverse Kullback-Leibler*) between Q and P .

In other words, under logarithmic scoring the Diachronic Accuracy Norm recommends minimization of the IKL-divergence between the posterior Q and the prior P (a method somewhat cautiously put forth by Douven and Romeijn 2012 as a method of updating in Judy Benjamin type problems), not minimization of the KL-divergence between Q and P (the method employed by van Fraassen 1981, and the one that Eva, Hartmann and Rad are attempting to advocate for in their paper). A norm that would recommend minimization of the KL-divergence between Q and P would look more like this:

Alternate 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 , $\exp_S(P|Q) - \exp_S(Q|Q) \leq \exp_S(P|Q') - \exp_S(Q'|Q')$, where \mathcal{J}_S is one's chosen proper scoring rule.

In the terminology of Reinhard Selten, the Alternate Diachronic Accuracy Norm says that one should minimize the *expected score loss of the prior P at the posterior Q* . In other words, one should choose one's posterior in such a way that one's *prior* comes out as accurate as possible relative to (and from the perspective of) the posterior. If the Alternate Diachronic Accuracy Norm were in fact normative, it would provide the argument Eva, Hartmann and Rad desire. Of course, if these authors employed the Diachronic Accuracy Norm in the first place because they had independent reason for thinking it was a good norm, this option won't appeal to them. But if they were just using it because it was in the literature and they thought it gave them what they wanted, they shouldn't mind switching to the alternate norm. Especially now that *it's* in the literature, too. (You can find it near the top of this page.)

And, to be sure, the alternate norm is the more plausible of the two. The Diachronic Accuracy Norm actually makes no sense whatsoever; why would one treat one's prior P as actual in the process of replacing it by a presumably superior posterior Q ? In accepting constraint C , one has already acknowledged that P doesn't give the right probabilities. In a case where no probability distribution that's even close to P meets the constraint, P may in fact be known already to be a very bad estimate of the right probabilities. So no calculation of an "expected" inaccuracy that takes P as actual is going to give meaningful results.

But in any event the Alternate Diachronic Accuracy Norm can't be normative, either. For one thing, the constraint C plainly needs to be a convex set. Even if we fix that issue, however, it still can't be normative. Suppose for example that we toss a biased coin (it lands *heads* a third of the time). You go into another room and ask God some question providing you with indirect evidence as to how the coin landed. (What sort of question exactly I don't know...maybe something like "did either the coin land *heads* or the Mets win last night?" That sort of thing.) You emerge and I ask you "So...does your credence in *heads* now lie in the interval $[\cdot 5, \cdot 7]$? You say "yes". If I think you

are ideally rational and know more than I do, I should defer to your credence in *heads*. So what I have learned constrains my credence in *heads* to the convex set $[.5, .7]$. But on any reasonable distance measure (e.g. expected score loss), the credence function assigning *heads* probability .5 is “closest”, among those meeting the constraint, to my current function. Should I then update my credence in *heads* to .5? No...to do so would imply that I think your credence in *heads* is exactly .5 with probability 1. What I should rather do is update my credence to the expectation of yours, conditional on the constraint.

Of course what makes expected score loss (or any other plausible distance) minimization intuitively wrong here is that the nearest point is an extreme point of the constraint region. When the nearest point is safely in the interior and plausibly near the centroid of our distribution over the credences of the expert one defers to, updating to the nearest point might seem less “intuitively wrong”. But having the nearest point fall close to the centroid is really just an accident of geometry. There’s not really any philosophical reason to think that one should update to the nearest point—one should of course update to the centroid of the distribution. That is, one’s posterior credence should be the expectation that one assigns to the expert’s credence.²

References

- Eva, Benjamin, Stephen Hartmann and Soroush Rafiee Rad. 2019. Learning from Conditionals. Preprint. (Headed “Forthcoming in Mind”) <http://philsci-archive.pitt.edu/15781/> Accessed 3/15/2019.
- Grove, Adam J. and Joseph Y. Halpern. 1997. Probability Update: Conditioning vs. Cross-entropy, *Proceedings of the 13th Annual Conference on Uncertainty in Artificial Intelligence*, San Francisco. Morgan Kaufmann. 208-214.
- Leitgeb, H. and R. Pettigrew (2010). An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy. *Philosophy of Science* 77:236-272.
- Van Fraassen, Bas C. 1981. A Problem for Relative Information Minimizers in Probability Kinematics. *The British Journal for the Philosophy of Science* 32:375-379.

²Grove and Halpern (1997) first offered such an analysis of Judy Benjamin, in the (apparently futile) hope that it would teach others of “the dangers involved in indiscriminately applying supposedly ‘simple’ and ‘general’ rules” in updating.