VASUDEVAN ON JUDY BENJAMIN

Abstract. Anubav Vasudevan characterized van Fraassen's "Infomin" solution to the Judy Benjamin Problem (i.e. the solution by way of minimizing the Kullback-Leibler divergence between the posterior and prior) as an implementation of a "brand of epistemic charity" taking "the form of an assumption on the part of Judy Benjamin that her informant's evidential report leaves out no relevant information". After an analysis of the example that led Vasudevan to this way of thinking about Infomin, as well as of a new one that supports the rival position of Douven and Romeijn in favor of minimizing the inverse Kullback-Leibler divergence between the posterior and the prior, we come to a different and more enlightening characterization of Infomin's implicit assumptions.

1. Introduction

Van Fraassen (1981) invites us to consider a probabilistic puzzle involving the fictional character Judy Benjamin (from the film "Private Benjamin"). The puzzle, in essence, runs as follows: Judy has prior credence function $P = (p_1, p_2, p_3, p_4) = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ over some measurable partition (E_1, E_2, E_3, E_4) of an event space. A duty officer, who is an expert relative to Judy (he knows more than she does), reports to Judy that, according to his credence function, E_3 is three times as likely as E_4. What should Judy's posterior credence function be? (Judy knows that the duty officer reports, always and only, this ratio.)

Recall that the Kullback-Leibler divergence (KL-divergence) between $Q = (q_1, q_2, q_3, q_4)$ and P (equivalently, the inverse IKL-divergence between P and Q) is given by

$$D_{KL}(Q|P) = \sum_{i=1}^{4} q_i \log\left(\frac{q_i}{p_i}\right) = \log(4) + \sum_{i=1}^{4} q_i \log(q_i) = \log(4) - E(Q).$$

Here E is the entropy of the partition Q. Van Fraassen suggests that Judy update to that credence function Q for which $D_{KL}(Q|P)$ attains its minimum value, subject to the constraint $q_3 = 3q_4$. (In this case, where Q attains maximum entropy, subject to the constraint.)

Vasudevan (2018) suggests a novel interpretation of Judy's actions in adopting this updating procedure:

> ...in applying the principle of maximum entropy, Judy Benjamin is not acting out of a concern to maximize uncertainty in the face of new evidence, but is rather exercising a certain brand of epistemic charity towards her informant. This epistemic charity takes the form of an assumption on the part of Judy Benjamin that her informant's evidential report leaves out no relevant information.

Two things bothers us about this line on Judy's actions. First, it's a central tenet of virtually all treatments of the Judy Benjamin Problem that the duty officer always and only reports the likelihood of E_3 relative to that of E_4. Indeed, it's fairly obvious that if one doesn't set the protocol in stone, there'd be no way to even begin to analyze the problem. So it isn't clear to us what Vasudevan means by "her informant's evidential report leaves out no relevant information", given that the report always contains information of precisely the same type. In fact it's a truism, so far as we can tell, that the informant *is* leaving out relevant information, for he could simply tell Judy what his credences in E_1 and E_2 are, and these values only fail the test for relevance if they are constant almost surely from Judy's perspective.

Second, we have discovered a phenomenon whereby Judy's posterior credence in $E_1 \vee E_2$ decreases as the duty officer's expected amount of additional information increases. And, whatever "her informant's evidential report leaves out no relevant information" in fact means, it doesn't plausibly imply that the informant's expected information exposure is vanishingly small–as it would at the least have to, to justify her use of Infomin.

In this note, we will review the Infomin-friendly example Vasudevan used to motivate his take on van Fraassen's solution. Then we will look at a second example in which it isn't minimization of the Kullback-Leibler divergence between Judy's posterior and her prior that gives the correct solution, but minimization of the inverse Kullback-Leibler divergence. We'll then make note of what is common to these examples (namely the low expected information exposure of the informant) and what is different, using these observations to renew discussion as to Infomin's implicit assumptions.

2. Two examples

In support of his rhetoric, Vasudevan imagines a scenario running more or less as follows. Let $n >> 1$ be large. Imagine that a four sided die has been subjected to n independent rolls. Judy and the duty officer both know that the die is fair, but only the duty officer knows the outcomes of the rolls. What happens now is that a roll is selected uniformly at random from the sample, and E_i is the event that the selected roll landed i. The duty officer does not know which roll was selected, so his credence function over the partition (E_1, E_2, E_3, E_4) will defer to the relative frequencies of the 4 outcomes in the sample. Judy's credence function over this partition will of course be $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. The duty officer now reports (always and only) the ratio of his credence in E_3 to his credence in E_4, that is the ratio of "3 rolls" to "4 rolls" in the sample. If the reported ratio is 3:1, what should be Judy's posterior credence function over the partition (E_1, E_2, E_3, E_4)?

This scenario is a "precisification" of the Judy Benjamin Problem; just enough extra detail has been provided to compute a solution directly. Since Judy will continue to regard the die as fair regardless of what happens in the sample, she will assign all words of length n on the alphabet $\{1, 2, 3, 4\}$ (considered as outcomes for the entire experiment) equal probability.

Denote by $R(n)$ the set of quadruples (r_1, r_2, r_3, r_4) with each r_i non-negative, $r_1+r_2+r_3+r_4 = n$ and $r_3 = 3r_4$. We introduce new variables a, b and θ. These represent (assuming that r_i represents the number of i rolls in the sample) the relative frequency of 1, 2 and 4 rolls in our sample, respectively. Also Q, which represents the duty office's probability function over (E_1, E_2, E_3, E_4), i.e.

$$Q = (a = \frac{r_1}{n}, b = \frac{r_2}{n}, 3\theta = \frac{r_3}{n}, \theta = \frac{r_4}{n}).$$

Next, some counting. For a fixed quadruple $(r_1, r_2, r_3, r_4) \in R(n)$ the number of words of length n having r_i occurrences of i, $i = 1, 2, 3, 4$ is:

$$\begin{aligned}
&N(r_1, r_2, r_3, r_4) = \frac{n!}{r_1!r_2!r_3!r_4!} \\
=&\frac{n!}{(an)!(bn)!(3\theta n)!(\theta n)!} \\
\approx&\frac{\sqrt{2\pi n}(n/e)^n}{\sqrt{2\pi an}(an/e)^{an}\sqrt{2\pi bn}(bn/e)^{bn}\sqrt{6\pi\theta n}(3\theta n/e)^{3\theta n}\sqrt{2\pi\theta n}(\theta n/e)^{\theta n}} \\
=&\frac{1}{\sqrt{3ab\theta^2}a^{an}b^{bn}(3\theta)^{3\theta n}\theta^{\theta n}}.
\end{aligned}$$

If we treat the triple (a, b, θ) as a continuous variable on the region

$$R = \{(x, y, z) : x, y, z \geq 0, x + y + 4z = 1\},$$

the final fraction in the above display is maximized where the log of its denominator is minimized, i.e. where

$$\begin{aligned}&\frac{1}{2}\log(3ab\theta^2)+n(a\log a+b\log b+3\theta\log(3\theta)+\theta\log\theta)\\=&\frac{1}{2}\log(3ab\theta^2)-nE(Q)\end{aligned}$$

is minimized. As $n\to\infty$, the $-nE(Q)$ term dominates and the point where the minimum occurs will approach the triple $(a^*,b^*,\theta^*)\approx(.26633,.26633,.11684)$ where $E(Q)$ is maximized.

Now let B be a small ball centered at (a^*,b^*,θ^*), and put

$$T=\inf_{(a,b,\theta)\in R\setminus B}-E(Q).$$

Let $B'\subset B$ be a neighborhood of (a^*,b^*,θ^*) having the property that for some $T'<T$, $-E(Q)<T'$ whenever $(a,b,\theta)\in B'$. Then whenever $(r_1,r_2,r_3,r_4),(r_1',r_2',r_3',r_4')\in R(n)$ with the corresponding triples satisfying $(a,b,\theta)\in R\setminus B$, $(a',b',\theta')\in B'$, one has

$$-E(Q')<T'<T\le -E(Q).$$

Pick T'' strictly between T' and T. For n large enough, one will have

$$\begin{aligned}&\frac{1}{2}\log(3a'b'\theta'^2)-nE(Q')\\<&nT'<nT''<\frac{1}{2}\log(3ab\theta^2)-nE(Q).\end{aligned}$$

Taking the exponential of both sides then inverting,

$$N(r_1',r_2',r_3',r_4')>e^{-nT'}>e^{-nT''}>N(r_1,r_2,r_3,r_4),$$

so that

$$\frac{N(r_1',r_2',r_3',r_4')}{N(r_1,r_2,r_3,r_4)}>e^{n(T''-T')}\to\infty \text{ as } n\to\infty.$$

Meanwhile it's obvious that there is some $M<\infty$ such that

$$\begin{aligned}&\Big|\{(r_1,r_2,r_3,r_4)\in R(n):(a,b,\theta)\in R\setminus B\}\Big|\\\le&M\Big|\{(r_1',r_2',r_3',r_4')\in R(n):(a',b',\theta')\in B'\}\Big|\end{aligned}$$

for all large enough n.

Denote by $W(n)$ the set of words of length n on the alphabet $\{1,2,3,4\}$ such that the letter 3 occurs exactly three times as often as the letter 4. For $w\in W(n)$, denote by $r_i(w)$ the number of occurences of the letter i in w. Then set $a(w) = \frac{r_1}{n}$, etc. What the above calculations show is that for any $\epsilon > 0$, there is a K so large that whenever $n \geq K$,

$$\Big|\{w \in W(n) : (a(w), b(w), \theta(w)) \in B\}\Big| > (1 - \epsilon)|W(n)|.$$

We may summarize as follows. Judy knows that the die is fair, so her prior credence function over the results of the "selected" roll is $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. An expert with additional evidence (statistics from an n roll sample including the selected roll) reports, in effect, that exactly three times as many rolls in this sample landed 3 as landed 4. Conditioning, Judy comes to have posterior credence function asymptotically (as $n \to \infty$) approaching $(a^*, b^*, 3\theta^*, \theta^*)$, in agreement with KL-divergence minimization and the maximum entropy principle.

This is a correct analysis of the precisification under consideration. The question, however, is whether setting its terms comes to nothing more than implementing "an assumption on the part of Judy Benjamin that her informants evidential report leaves out no relevant information". To shed light on that matter, we present an alternate scenario whose solution matches with a rival updating scheme (defended in Douven and Romeijn 2011 and Eva, Hartmann and Rad 2019) based on IKL-divergence minimization.

In the alternate scenario, there is again a four sided die and there is again a sample of size n. This time, however, Judy knows the statistics of the sample. Namely, the die landed i exactly $n/4$ times, $1 \leq i \leq 4$. Also, this time the die is not assumed to be fair; the duty officer knows the die's true chance function (a_0, b_0, c_0, d_0) but Judy begins with some continuous distribution, given by density function g, invariant over permutations of the four sides, over the possible values of (a, b, c, d). We will assume in particular that $g > 0$ a.e. (She hasn't completely ruled out any region.) Another change is that this time E_i denotes the event that the outcome of the *next* roll of the die (i.e. a roll that is independent of the sample) lands i. Finally, the duty officer reports (always and only) the ratio $c_0 : d_0$. On this occasion, the reported ratio is 3:1. What will be Judy's posterior over the partition (E_1, E_2, E_3, E_4)?

After observation of the sample, Judy's posterior density function is given by

$$h(a, b, c, d) = ka^{n/4}b^{n/4}c^{n/4}d^{n/4}g(a, b, c, d)$$

for an appropriate constant k. Since $n >> 1$, for reasonable g this distribution would in any event be tightly concentrated about $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$, but since g is assumed to be invariant under permutations of the variables a, b, c, d, we can assume more, namely that its expectation (that is, Judy's prior credence function over the partition (E_1, E_2, E_3, E_4)) is *exactly* $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. The report, meanwhile, teaches Judy that the true chances lie on the surface

$$R = \{(a, b, c, d) : a, b, c, d \geq 0,\ a + b + c + d = 1,\ c = 3d\}.$$

The mode of her distribution over the chances restricted to R is at that point $(a, b, c, d) \in R$ where $a^{n/4}b^{n/4}c^{n/4}d^{n/4}g(a, b, c, d)$ achieves its maximum. When n is large enough this will happen near the point $(a, b, c, d) \in R$ where $abcd = 3abd^2$ is maximized. By symmetry this maximum occurs at a point where $a = b$. The substitution $d = \frac{1}{4}(1 - a - b)$ followed by the substitution $a = b$ shows that the quantity to be maximized is $a^2(1 - 2a)^2$. This fourth degree polynomial has zeros of multiplicity two at 0 and $\frac{1}{2}$. The sought-for maximum therefore occurs at the midpoint of these, $a = \frac{1}{4}$. When n is large, meanwhile, the distribution given by the density function h is highly concentrated about its mode. So after hearing the duty officer's report when n is large, Judy's posterior over the partition (E_1, E_2, E_3, E_4) is $\approx (\frac{1}{4}, \frac{1}{4}, \frac{3}{8}, \frac{1}{8})$, consistent with IKL-divergence minimization.

There are notable differences in these two scenarios.[1] For example, in the first scenario the duty officer has deductive information that Judy lacks, but no inductive information. (Judy already knows the true chances, screening off whatever inductive information the duty officer's random sample might otherwise have provided.) In the second scenario the duty officer has inductive information that Judy lacks, but no deductive information. (The partition in question is causally independent of the evidence.) That's an interesting distinction, but we hardly sees how it justifies Vasudevan's rhetoric. We don't see, in particular, why compliance with "assume that your informant's evidential report leaves out no relevant information" ought to entail compliance with "assume that your informant's additional evidence is purely deductive". Indeed, we don't see any reason to think that the former injunction is more or less compatible with the latter than it is with "assume that your informant's additional evidence is purely inductive."

[1]Perhaps more interesting is what they have in common. As anyone following the Judy Benjamin literature at all closely will already know, both of the update methods being compared here (i.e. KL-minimization and IKL-minimization) expose Judy to Reflection violations. These worsen as the duty officer's expected amount of information exposure (information that is relevant to the partition (E_1, E_2, E_3, E_4), that is) increases, and vanish only as the duty officer's expected information exposure tends to zero. That is precisely what happens here as $n \to \infty$. (The larger n is, the more certain Judy is that the actual chances of the 4 possible outcomes lie close to their relative frequencies in the sample.) So both scenarios are at the same extreme in that sense, though they appear to lie at opposite extremes in another.

Elsewhere, Vasudevan writes:

> Such charity on the part of Judy Benjamin is analogous to that which a student affords to an examiner in assuming that a certain problem on an exam is well posed. Even if the problem statement does not include any explicit claim to this effect, the pragmatics of exam-sitting require the student to extend charity to his examiner by presupposing that the problem, as stated, does not leave out any information that is relevant for determining its solution.

So far as we know, it's only considerations of simplicity that can play this role. So if, on an exam, I am asked to give the next term in the sequence $1, 4, 9, \ldots$, I should choose $f(4)$ for the "simplest" function $f(x)$ such that $f(1) = 1$, $f(2) = 4$ and $f(3) = 9$; $f(x) = x^2$ perhaps, for which $f(4) = 16$. (Not, say, $f(x) = \frac{1}{6}x^3 + \frac{11}{6}x - 1$, for which $f(4) = 17$.) But even assuming that "simplicity" places Judy in a "low information" setting (see footnote 1), we don't see that the Vasudevan scenario is interestingly more simple than the alternate scenario we gave–though each does exhibit simplicity of a kind, in that each seems to represent an "extreme case". One possible take on these matters is that the two scenarios lie at opposite ends of a bipolar "low information spectrum"; in problems where there is a clear-cut "simplest" solution, by contrast, the spectrum of possible solutions is usually monopolar, with the simplest solution lying at the unique pole. At any rate, we don't find this second piece of rhetoric to be all that perspicuous, either.[2]

References

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.

Vasudevan, Anubav. 2018. Entropy and insufficient reason: a note on the Judy Benjamin problem. *The British Journal for the Philosophy of Science*. Forthcoming.

RMCCTCHN@MEMPHIS.EDU

[2]Low information Brownian motion scenarios (which favor minimization of Euclidean distance, resulting in posteriors tending to $(\frac{5}{18}, \frac{5}{18}, \frac{1}{3}, \frac{1}{9})$) may complicate the landscape of this spectrum. Such scenarios assume that the informant's posterior arises from an n-step random walk in which the variance of a single step is on the order of $\frac{\epsilon}{n}$ for small fixed $\epsilon > 0$. Vasudevan's scenario, by contrast, arises from an n-step random walk in which a single step is noticably asymmetric and has variance on the order of $\frac{1}{n^2}$. The extreme lowness of this variance makes the asymmetry relevant, arguably in a way that is antithetical to "simplicity" considerations.