Joyce 1998: 'A Nonpragmatic Vindication of Probabilism'

KEVIN DORST 9.14.15

I. Probabilism and Pragmatism

Probabilism: recognize degrees of belief; require that they conform to probability axioms.

Won't worry about conditionalization.

Joyce suspects that credences haven't (hadn't) taken center-stage because they haven't been given an *epistemological* foundation:

Cashed out in pragmatic terms

Defended on pragmatic grounds: Dutch Book Argument.

Betting ratios; willingness to act.

Joyce wants to give them an epistemological foundation:

(i) An account of how credences (in)accurately represent the world.

(ii) Show how this norm of accuracy vindicates probabilism.

A **"criterion of epistemic success."**

Full beliefs: **Norm of Truth:**

> An epistemically rational agent must strive to hold a system of full beliefs that strikes the best attainable overall balance between the epistemic good of fully believing truths and the epistemic evil of fully believing falsehoods. (577)

Q: He says this is the "well-known and uncontroversial" criterion of epistemic success for full beliefs. Is that so? What about knowledge?

Proposal: what distinguishes credences from full beliefs is the *standard of accuracy*. Beliefs are assessed as guesses, whereas credences are assessed as estimates (578).

Guesses only get credit if they are exactly right; estimates get more credit the closer they are to correct. Examples: Batting predictions; lottery selection.

The **Norm of Gradational Accuracy:**

> An epistemically rational agent must evaluate partial beliefs on the basis of their gradational accuracy, and she must strive to hold a system of partial beliefs that, in her best judgment, is likely to have an overall level of gradational accuracy at least as high as that of any alternative system she might adopt. (579)

Crucial: "best judgment" means "best estimate", not "best guess." Extremal credences and hedging epistemic bets.

Strategy: (1) impose/justify formal constraints on a measure of gradational accuracy, and (2) show that probabilism is required to promote this measure.

Note: "scoring rule" is explicit, but "decision procedure" is not.

Q: What is the story behind these constraints? Is it *conceptual* – constitutive of our notion of accuracy that it follows these rules? Is it *epistemological* – promoting accuracy would license irrational behavior if we didn't have these constraints. (Cart-before-horse feeling? Joyce 2009 and reflective equilibrium method.)

II. Problems with Dutch Books

Imagine a *prevision game* played by a *miser* (values only money, linearly): give her \$1 and a list of propositions $X_1, X_2, ..., X_n$ and she chooses prevision, which is a sequence of numbers $p = p_1, p_2, ..., p_n$; pay portion back according to a *scoring rule* $S(p, w)$ (581).

De Finetti uses a *quadratic loss function*:

$$S(p, w) = \sum_i (\lambda_i (w(X_i) - p_i)^2)$$

.

De Finetti proves that any prevision p that is not probabilistic is dominated by some probabilistic p^*, *so misers will choose probabilistic previsions.* **Problem:** What does this have to do with their credences? (583)

Define credences in terms of previsions? No: (1) failures of behaviorism; (2) Why use quadratic loss functions to define credences, instead of others?; (3) Nothing stops them from reporting different (probabilistic) previsions for different functions.

Analogy: 'temperature' and 'the quantity measured by thermometers.'

Prudential rationality requires maximizing expected utility? This guarantees that she will report her (probabilistic) credences. But: (1) Relies on expected utility maximization; (2) *prudential* justification.

Q: He defines $Exp(p) = b(X)S(p, 1) + (1 - (b(X))S(p, 0)$, but this only captures her *estimate* of the value of p if $(1 - b(X)) = b(\neg X)$. Doesn't that presuppose probabilism?

De-pragmatized dutch books? — The argument reveals that they make inconsistent value judgments. (586)

Joyce: but why is this an *epistemic* defect if the inconsistency is in preferences?

What about claiming inconsistent beliefs about values? Joyce (fn. 8): (1) Locate epistemic flaw in wrong subject matter; (2) Serious difficulties with view that desire is a kind of belief.
Q: But can't we just assert a rational connection, rather than a reduction? You rationally desire p more than q only if you can rationally believe that p is better than q.

III. A New Way

Following Jeffrey (1986), take credences to be estimates of truth-values. (587)

But we *don't* assume that such estimates conform to laws of mathematical expectation, since when a random variable only takes the values 0 and 1 (like truth-values), the laws of expectation are just the laws of probability.

Following Rosenkrantz (1981), construct an *epistemic Dutch book*, using an inaccuracy measure I instead of a scoring rule S.

But be careful here. De Finetti was free to assume that a miser would maximize relative to any scoring rule he chose, but he had problems connecting their chosen previsions to their credences.

Joyce has no trouble with the latter — it is the credences themselves that are being scored. But he can't impose whatever scoring rule he likes — it has to be the *actual* measure of accuracy the subject is trying to promote. So he has to argue that the axioms he imposes on I are satisfied by any reasonable metric of accuracy (589).

Note that we basically assume that the agent is an *accuracy*-miser.

This is where Rosenkrantz goes wrong.

(Notion of "accuracy" imprecise? Supervaluate — any precisified version should satisfy the proposed axioms.)

IV. The Axioms

I is an inaccuracy measure; a real-valued function of a credence-function $b \in B$ and a possible world $w \in V$; V^+ is the set of probability functions, which are weighted averages of members of V.

Importantly, B has enough structure to talk about the "line segment" between any two "points" b and b^*: $bb^* = \{\lambda b + (1-\lambda)b^* | \lambda \in [0,1]\}$. Any member of this line segment is a mixture of b and b^*.

[Diagram?]

1) Structure: For any w, $I(b,w)$ is a non-negative, continuous function of b that goes to infinity in the limit as $b(X)$ goes to infinity for any X.

Joyce: this should be uncontroversial. You can get as inaccurate as you like; and small changes in credences do not make for large changes in accuracy.

Q: Continuity is a bit substantive, right? Even if you're being scored for your estimates, maybe we put a premium on being within .5 of the right answer (you "leaned the right way"), with a discontinuous drop-off past that point.

2) Extensionality: At each w, $I(b,w)$ is a function of the truth-values of at w and the credences assigned by b to the $X \in \Omega$.

Joyce: most objections conflate finding a measure of *accuracy* with that of finding a measure of *epistemic utility*. But since we're after a notion of *closeness to the truth*, this is kosher.

"Accuracy is only one virtue among many" (592).

Q: Is that so? How can he admit the existence of other epistemic values yet still move from the fact that probabilism promotes *accuracy* to the conclusion that it is rationally required? How is showing that non-probabilistic credences are accuracy-dominated different than showing that your career is monetarily-dominated?

This is the source of the Easwaran & Fitelson worry. Joyce's (ms) response is something like: whatever other epistemic values there are, they have to play well with accuracy. It makes no sense for the evidence to support a state that is accuracy-dominated, for instance — evidence is after the truth too!

Objection: verisimilitude. Copernicus vs. Kepler.
Response: leads to other credences that are accurate.

3) Dominance: If b and b^* are identical except on X, then
$I(b,w) > I(b^*,w)$ iff $|w(X) - b(X)| > |w(X) - b^*(X)|$. (593)

Says that moving a single credence closer to the actual value always increases accuracy, no matter what one's other credences are.

Example of this failing: the calibration index used by van Fraasen (1983) and Shimony (1988).

Basically, you are calibrated when: of the set of propositions you assign credence t to, t of them are true. Then we measure "distances" from this ideal.

Two problems: (1) Ill-motivated structural assumptions on propositions; (2) bad predictions about how you should go about trying to be accurate [look to table on 595].

Q: Dominance does more than rule this out, though. Say I'm (incorrectly) certain in X and am $\frac{1}{2}$ in $\neg X$. If I lower my credence in $\neg X$, there's a sense in which I'm being epistemically better — my beliefs are more coherent; I believe the consequences of what I believe. (Maybe this is just the *epistemic utility* point again?)

4) Normality: If $|w(X) - b(X)| = |w^*(X) - b^*(X)|$ for all X, then $I(b,w) = I(b^*,w^*)$.

Joyce: gradational accuracy cares only about distance to truth, not (e.g.) which truths are true.

Q: Am I right? Sometimes "Extensionality" is used nowadays to include Normality and Joyce's Extensionality.

5) Weak Convexity: If $m = (\frac{1}{2}b + \frac{1}{2}b^*)$ is the midpoint of the line segment between distinct b and b^* and $I(b,w) = I(b^*,w)$, then $I(b,w) > I(m,w)$.

Note that the move from b to m is in the same direction — but only half as large — as the move from b to b^*.

> Weak Convexity is motivated by the intuition that extremism in the pursuit of accuracy is no virtue. It says that if a certain change in a person's degrees of belief does not improve accuracy then a more radical change in the same direction and of the same magnitude should not improve accuracy either (596)

So if move from b to m doesn't raise accuracy, then the move from m to b^* shouldn't raise it either (or *bump it back up*).

Q: It says more, doesn't it? Why can't b, b^*, and m all have *the same* accuracy? (I get it that accuracy shouldn't go *down* in move to m, and up in move to b^*; but why need it change at all?)

6) Symmetry: If $I(b,w) = I(b,w^*)$, then for any $\lambda \in [0,1]$ we have $I(\lambda b + (1-\lambda)b^*, w) = I((1-\lambda)b + \lambda b^*, w)$.

Given Weak Convexity, there will be a point between b and b^* that minimizes inaccuracy. If this weren't the midpoint, then we'd have an "unmotivated bias in favor of one set of beliefs or the other" (597). We ensure this by requiring that:

> The change in belief that moves an agent a proportion λ along the line segment form b toward b^* has the same overall effect on her accuracy as the 'mirror image' change that moves her the same proportion λ along the line segment from b^* to b. (597)

Q: Even if we buy that the midpoint should minimize inaccuracy, Symmetry is much stronger: the inaccuracy scores from b and b^* have to *follow the same path* to that minimum at the midpoint.

Upshots

> Structure, Extensionality, Normality, Weak Convexity, and Symmetry are the only constraints on measures of gradational accuracy we need to vindicate the fundamental dogma of probabilism. (597)
>
> It is thus established that degrees of belief that violate the laws of probability are invariably less accurate than they could be. [Strictly: for each such b, there's a probabilistic b^* that accuracy-dominates it.] Given that an epistemically rational agent will always strive to hold partial beliefs that are as accurate as possible, this vindicates the fundamental dogma of probabilism. (600)

Note the implicit appeal to a dominance decision-procedure here.

Question: Can we re-interpret this whole argument as applying to the *evidential support relation* b? As far as I can tell, all this takes is reformulating the norm of gradational accuracy.

À la "...the evidence must provide a system of support relations that, by it's own lights, is likely to have an overall level of gradational accuracy...[etc.]." This requires understanding the accuracy of a particular support relation for a particular proposition, but presumably that's entirely parallel to that of a credence.

Worry: Even if accuracy *does* need to be measured by a metric that obeys these axioms, might that be a fact open to rational doubt? <u>If not</u>, then it looks like the argument is committed to some sort of Fixed-Point-esque thesis in the background. <u>If so</u>, then the decision procedure relied upon becomes more problematic: even though improbabilistic credences are *in fact* accuracy-dominated, an agent could (permissibly) fail to know this. Or, put another way: they are accuracy-dominated relative to metaphysically possible worlds, but not relative to epistemically possible ones.