Accuracy notes

PHIL 735 Week 11
Matt Kotzen and Jim Pryor

November 15, 2023

1 Accuracy and partial beliefs

- So far, we've seen arguments for various aspects of Bayesianism that (at least initially appear to) rest on practical premises
- Today, we're going to focus on arguments that are designed to rely only on considerations of accuracy, which (according to the proponents of these arguments) is properly understood in this context as a purely *epistemic* value
- It's reasonably clear how to think about accuracy for full beliefs: a full belief is (fully) accurate if its content is true, and a full belief is (fully) inaccurate if its content is false
- By contrast, for partial beliefs, there's an intuitive sense in which accuracy isn't an all-or-nothing affair—one partial belief can *be more accurate than* another partial belief, even if the first partial belief isn't maximally accurate
- So, the question then arises of how to measure the "gradational accuracy" of partial beliefs. Is there a unique correct way to measure how accurate some particular partial belief, or some credence function, is relative to a particular world? If not, are there plausible constraints that would at least rule out *some* ways of measuring gradational accuracy?

2 Joyce 1998

- Strategy: Articulate several constraints on any adequate function that endeavors to measure (in)accuracy. Then, show that, for any non-probabilistic credence function (defined over a given algebra), there is a probabilistic credence function (defined over that same algebra) that is guaranteed to be less inaccurate (as measured by that function), no matter which possible world is actual.

- "What I have in mind here is a kind of 'epistemic Dutch book argument' in which the relevant scoring rule assigns each credence function **b** and possible world ω a penalty $\mathbf{I}(\mathbf{b}, \omega)$ assessed in units of gradational inaccuracy. The rule **I** will gauge the extent to which the truth-value estimates sanctioned by **b** diverge from the truth-values that propositions would have were ω actual. My claim is going to be that, once we appreciate what **I** must look like, we will see that violations of the laws of probability always decrease the accuracy of partial beliefs."
  - And, for any probabilistic credence function, there does *not* exist an alternative credence function that is guaranteed to be less inaccurate, no matter which possible world is actual.
- Joyce's notation:
  - **B** is the family of all credence functions defined on a countable Boolean algebra of propositions Ω
  - **V** is the subset of **B** containing all consistent truth-value assignments to members of Ω (i.e., the possible worlds ω)
    - And the possible worlds ω are treated as functions that assign a value of 1 to propositions that are true at the world, and a value of 0 to propositions that are false at the world
  - **I** is defined over pairs in $\mathbf{B} \times \mathbf{V}$
  - The collection of all probability functions in **B** is **V**'s convex hull $\mathbf{V}^+$
    - The convex hull of a set is the set of all "weighted averages" of elements of the set where the weights are non-negative and sum to 1. So, the convex hull of the set of possible worlds is the set of all weighted averages of the possible worlds, using coefficients that sum to 1.
  - $\mathbf{B} \sim \mathbf{V}^+$ is everything that is in **B** but not in $\mathbf{V}^+$—i.e., the set of credence functions that violate the laws of probability
- Structure of **B**
  - Always contains a unique "line" $\mathbf{L} = \{\lambda\mathbf{b} + (1-\lambda)b^* : \lambda \in \mathbb{R}\}$ that passes through any two of its "points" **b** and $\mathbf{b}^*$
  - The line segment from **b** to $\mathbf{b}^*$, $\mathbf{bb}^*$, is the subset of **L** for which $0 \leq \lambda \leq 1$. A credence function $\lambda\mathbf{b} + (1-\lambda)\mathbf{b}^*$ that falls on $\mathbf{bb}^*$ is a "mixture" of **b** and $\mathbf{b}^*$. If $\lambda > 1/2$, then the mixture favors **b**; if $\lambda < 1/2$, then the mixture favors $\mathbf{b}^*$; if $\lambda = 1/2$, then the mixture is even, halfway in between **b** and $\mathbf{b}^*$.

- Constraints on **I**
  - Structure: For each $\omega \in \mathbf{V}, \mathbf{I}(\mathbf{b}, \omega)$ is a non-negative, continuous function of **b** that goes to infinity in the limit as $\mathbf{b}(X)$ goes to infinity, for any $X \in \Omega$
  - Extensionality: At each possible world ω, $\mathbf{I}(\mathbf{b}, \omega)$ is a function of nothing other than the truth-values that ω assigns to propositions in Ω and the degrees of confidence that **b** assigns these propositions.
  - Dominance: If $\mathbf{b}(Y) = \mathbf{b}^*(Y)$ for every $Y \in \Omega$ other than X, then $\mathbf{I}(\mathbf{b}, \omega) > \mathbf{I}(\mathbf{b}^*, \omega)$ if and only if $|\omega(X) - \mathbf{b}(X)| > |\omega(X) - \mathbf{b}^*(X)|$
  - Normality: If $|\omega(X) - \mathbf{b}(X)| = |\omega^*(X) - \mathbf{b}^*(X)|$ for all $X \in \Omega$, then $\mathbf{I}(\mathbf{b}, \omega) = \mathbf{I}(\mathbf{b}^*, \omega^*)$
  - Weak Convexity: Let $\mathbf{m} = (1/2\mathbf{b} + 1/2\mathbf{b}^*)$ be the midpoint of the line segment between **b** and $\mathbf{b}^*$. If $\mathbf{I}(\mathbf{b}, \omega) = \mathbf{I}(\mathbf{b}^*, \omega)$, then it will always be the case that $\mathbf{I}(\mathbf{b}, \omega) \geq \mathbf{I}(\mathbf{m}, \omega)$ with identity only if $\mathbf{b} = \mathbf{b}^*$
  - Symmetry: If $\mathbf{I}(\mathbf{b}, \omega) = \mathbf{I}(\mathbf{b}^*, \omega)$, then for any $\lambda \in [0, 1]$ one has $\mathbf{I}(\lambda\mathbf{b} + (1 - \lambda)\mathbf{b}^*, \omega) = \mathbf{I}((1 - \lambda)\mathbf{b} + \lambda\mathbf{b}^*, \omega)$

- Main Theorem: If gradational accuracy is measured by a function **I** that satisfies Structure, Extensionality, Dominance, Normality, Weak Convexity, and Symmetry, then for each $c \in \mathbf{B} \sim \mathbf{V}^+$ there is a $c^* \in \mathbf{V}^+$ such that $\mathbf{I}(\mathbf{c}, \omega) > \mathbf{I}(\mathbf{c}^*, \omega)$ for every $\omega \in \mathbf{V}$
  - Joyce notes that the Brier score satisfies all of his axioms
    - Brier score: where $\Omega = \{X_1, X_2, \ldots, X_n\}, \mathbf{I}(\mathbf{b}, \omega) = (\omega(X_1) - b(X_1))^2 + (\omega(X_2) - b(X_2))^2 + \cdots + (\omega(X_n) - b(X_n))^2$

- Titelbaum gives helpful illustrations of the proof strategy for all three Kolmogorov axioms, on pp. 348–50 (and endnote 15).

3 Proper scoring rules

- One putative advantage of the Brier score is that it is a "proper" scoring rule: an agent with a probabilistic credence distribution who uses the Brier score expects her own credences to be more accurate than any other distribution over the same set of propositions
- The dialectical role of properness here is a bit subtle
  - When we're trying to construct an argument for probabilism, it can seem circular to impose properness as an adequacy condition on scoring rules:

    - Titelbaum, pp. 359–60: "A proper scoring rule is one on which probabilistic distributions always expect themselves to be more accurate than the alternatives. But why focus on what *probabilistic* distributions expect? . . . [W]hen an inaccuracy measure is used to *argue* for probabilism—as in the Gradational Accuracy Argument—it seems question-begging to privilege probabilistic distributions in selecting that scoring rule."
    - Joyce 1998, p. 589: "Though I am happy to grant that [the principle that a rational agent should aim to minimize her expected inaccuracy and the principle that an agent's own credences are the ones that minimize expected inaccuracy] hold for partial beliefs that obey the axioms of probability, the problem is that they must also hold when the axioms are violated if they are to serve as premises in a justification for the fundamental dogma of probabilism."
  - Still, it would seem awkward if someone were to argue for some constraint on credences by appealing to a scoring rule (or class of scoring rules) which has the consequence that agents who obey that constraint are failing to maximize expected accuracy. And at least Joyce's argument for probabilism doesn't have that defect.
  - And, if we could generate an argument that probabilistic credences are rationally *permitted*, then perhaps that could be used as a premise in an argument for a proper scoring rule: the idea would be that an acceptable scoring rule shouldn't ever tell a rational agent to regard her own credences to be less accurate than an alternative. See Joyce 2009.
  - And, it's plausible that committed probabilists should be interested only in proper scoring rules. There are lots of proper scoring rules, and lots of improper ones, so considerations of propriety can at least help the probabilist to narrow down their search.

- Some possible scoring rules. Assume that the agent has a probabilistic credence function and that his actual $p(X) = a$. x is a variable that takes value $x = 1$ when X is true, and takes value $x = 0$ when X is false.
  - Linear rule: Accuracy of $p(X) = b$ is measured by $xb + (1 - x)(1 - b)$. An agent for whom $p(X) = b$ has accuracy b if X is true, and $1 - b$ if X is false.
    - Since the agent's actual $p(X) = a$, her expected accuracy of having credence b is: $ab + (1 - a)(1 - b)$
    - If $a > 1/2$, this is maximized when $b = 1$
    - If $a < 1/2$, this is maximized when $b = 0$
    - For example, suppose the agent's actual $p(X) = 3/4$. Their expected accuracy for $p(X) = 1$ is $(3/4)(1) + (1/4)(0) = 3/4$, whereas their expected accuracy for $p(X) = 3/4$ is only $(3/4)(3/4) + (1/4)(1/4) = 5/8$.

    - Improper!
- Another scoring rule: Accuracy of $p(X) = b$ is measured by $xb - b^2/2$. The agent's accuracy is $b - b^2/2$ if X is true, and $-b^2/2$ if X is false.
    - Since the agent's actual $p(X) = a$, her expected accuracy of having credence b is: $ab - b^2/2$
    - For any given a, this is at a maximum at $b = a$. Proper!
    - But, this scoring rule is weirdly asymmetric. For example, when X is false, accuracy of $p(X) = 1$ is $-1/2$, whereas when X is true, accuracy of $p(X) = 0$ is 0.
- Brier score: Inaccuracy of $p(X) = b$ is measured by $(x - b)^2$. So, accuracy of $p(X) = b$ is $1 - (x - b)^2$. An agent for whom $p(X) = b$ has accuracy $1 - (1 - b)^2$ if X is true, and $1 - b^2$ if X is false.
    - Since the agent's actual $p(X) = a$, her expected accuracy of having credence b is: $1 - (a - b)^2$
    - For any given a, this is at a maximum at $b = a$. Proper!
    - And it's nicely symmetric: For example, when X is false, accuracy of $p(X) = 1$ is 0, and when X is true, accuracy of $p(X) = 0$ is 0.