Accuracy & Coherence II

Branden Fitelson[1]

Department of Philosophy
&
Center for Cognitive Science (RuCCS)
Rutgers University

branden@fitelson.org
http://fitelson.org/

[1]This talk includes joint work with Rachael Briggs (Sydney/ANU), Fabrizio Cariani (Northwestern), Kenny Easwaran (USC) & David McCarthy (Edinburgh).

- Recall our definition of "inaccuracy"/"accuracy dominance" for a (complete) set of comparative judgments $\mathfrak{C}$.
  - **Comparative**. Let $\mathfrak{C}$ be the full set of S's comparative judgments over $\mathcal{B} \times \mathcal{B}$. The *innaccuracy* of $\mathfrak{C}$ at a world w is given by the number of incorrect judgments in $\mathfrak{C}$ at w.
    - $p \sim_S q$ is (in)correct at w iff $p \equiv q$ is true (false) at w.
    - $p \succ_S q$ is (in)correct at w iff $p \,\&\, {\sim}q$ is true (false) at w.
  - $\mathfrak{C}'$ *accuracy-dominates* $\mathfrak{C}$ iff $\mathfrak{C}'$ contains *strictly fewer* incorrect judgments than $\mathfrak{C}$ at *some* w's, and $\mathfrak{C}'$ contains *at most as many* incorrect judgments as $\mathfrak{C}$ at *every* w.
- This simple, 2-valed scoring scheme may seem overly simplistic. It is based on the following underlying norm:

  (†) S should be more confident in truths than falsehoods.
- So, if p is **T** and q is **F**, then the judgments $q \succ_S p$ and $p \sim_S q$ are in violation of this basic underlying norm (†).
- But, (†) *alone* does not justify our choice of 2-valued scheme. Indeed, other scoring schemes seem plausible.

- Let's use "+1" to denote *best* epistemic status, "−1" to denote *worst* epistemic status, and "0" to denote "middling" epistemic status. Our simplest, 2-valued scheme is:

	P	Q	$P \succ_S Q$	$Q \succ_S P$	$Q \sim_S P$
w_1	**T**	**T**	−1	−1	+1
w_2	**T**	**F**	+1	−1	−1
w_3	**F**	**T**	−1	+1	−1
w_4	**F**	**F**	−1	−1	+1

- If we're going to use only 2-values ("correct/incorrect"), then it seems to me that this scheme is *forced* on us, by (†).
- But, one might think that a 3-valued scheme makes more sense. David Christensen makes the following observation.

  > Suppose I'm going to flip a coin. Can I rationally be indifferent between heads (H) and tails (T)? It seems that $H \sim_S T$ would be dominated by $H \succ_S T$ (or $T \succ_S H$), since $H \sim_S T$ is guaranteed to be "incorrect" and the latter aren't.
- Christensen is right. And, he suggests a 3-valued scheme.

	P	Q	$P \succ_S Q$	$Q \succ_S P$	$Q \sim_S P$
w_1	**T**	**T**	0	0	+1
w_2	**T**	**F**	+1	−1	0
w_3	**F**	**T**	−1	+1	0
w_4	**F**	**F**	0	0	+1

- I agree that D.C.'s scheme does seem superior (intuitively) to our simplest 2-valued scoring scheme (in various ways).
- If we use this (or some other) 3-valued scheme, the obvious way to calculate the score of $\mathfrak{C}$ (at w) is to take the *sum* of these 3-valued scores for all the propositions in $\mathfrak{C}$ (at w).
- Then, we would define accuracy-dominance as follows:
  - $\mathfrak{C}'$ *accuracy-dominates* $\mathfrak{C}$ iff $\mathfrak{C}'$ has a *higher* score than $\mathfrak{C}$ at *some* w, and $\mathfrak{C}'$ doesn't have a lower score than $\mathfrak{C}$ at *any* w.
- In any event, moving to a 3-valued scheme can *not* fill the gap in de Finetti's justification/grounding of subjective probability theory. Indeed, we have an *impossibility result*.

Theorem. *No* 2 *or* 3-valued scoring scheme is such that:

(0) S entails (at least *some* instances of) *both* transitivity *and* additivity as (weak) dominance norms.

and, the the following eight (8) scoring *desiderata* are met:

(1) Having a subset of judgments $\{p \succ_S q, p \succ_S r, q \sim_S r\}$ should not — *in and of itself* — ensure "incoherence".

(2) *Ditto* for subsets of the form $\{p \succ_S q, p \succ_S r, q \succ_S r\}$.

(3) $p \succ_S q$ should get a "worst" score when p is **F** and q is **T**.

(4) $p \succ_S q$ should get the same score when p and q are both **T** as it does when p and q are both **F**.

(5) $p \sim_S q$ should get the same score when p and q are both **T** as it does when p and q are both **F**.

(6) $p \sim_S q$ should get the same score when p is **T** and q is **F** as it does when p is **F** and q is **T**.

(7) The score of $p \succ_S q$ when p is **T** and q is **F** should not be strictly worse than the score of $p \succ_S q$ when p, q are both **T**.

(8) The score of $p \succ_S q$ when p is **T** and q is **F** should be strictly better than the score of $p \succ_S q$ when p is **F** and q is **T**.

- These eight *desiderata* seem to be sacrosanct (Christensen and everyone else I've talked to seems to accept all of them).
- The upshot of our Theorem is that — *it doesn't matter which scoring scheme you use. No* scoring scheme can ground *all* of de Finetti's axioms for comparative probability.
  - In fact, our simplest 2-valued scheme *gets as close as any 2 or 3-valued scheme* to grounding all of de Finetti's axioms. [This is why I introduced it first. It is *simple*, and *maximally charitable* to de Finetti (with respect to his project).]
- So, it seems there is no accuracy-dominance justification of all of de Finetti's intuitive axioms (much less the *unintuitive* Scott Axiom — see Extras slides). This *re*-raises a question:

☞ *Why* should we think $\succ_S$ has a numerical Pr-representation?

- There seems to be no compelling reason to suppose that our comparative confidence orderings are (numerically) probabilistically representable. This is an important *lacuna*.
- Next: Quantitative judgments (*viz.*, numerical credences).

- Let's suppose (*arguendo*) that S has a numerical credence function $b : \mathcal{B} \mapsto \mathbb{R}$ (these b's are *opinionated*, of course, and so we're ignoring suspension of judgment here, once again).
- As usual, we need to settle on a way of *scoring* b's for inaccuracy at each possible world w — call this $I(b, w)$.
- For simplicity, I'll assume $I(b, w)$ is an *additive* function, which *sums-up* the inaccuracies of b, for each $p \in \mathcal{B}$ at w.
- If we associate the number 1 with **T** and the number 0 with **F** (at each world w), then the inaccuracy of $b(p)$ at world w will be b's "distance (d) from the 0/1-truth-value of p" at w.
- **Example.** Suppose S has just two (contingent) propositions $\{P, \sim P\}$ in their doxastic space. Then, there are two salient possible worlds (w_1 in which P is **T**, and w_2 in which P is **F**).
  - $I(b, w_1) = d(b(P), 1) + d(b(\sim P), 0)$.
  - $I(b, w_2) = d(b(P), 0) + d(b(\sim P), 1)$.

- Various measures (d) of "distance from 0/1-truth-value" have been proposed/defended in the historical literature.
- The most popular choice (for giving an accuracy-dominance justification of probabilism) has been the *squared-difference* measure of "distance from 0/1-truth-value", which is:
  - $\mathfrak{s}(x, y) = (x - y)^2$.
- The distance measure $\mathfrak{s}$ gives rise to a measure of inaccuracy ($I_{\mathfrak{s}}$), which is known as the Brier Score. In our toy example, the Brier Scores of b in worlds w_1 and w_2 are:
  - $I_{\mathfrak{s}}(b, w_1) = \mathfrak{s}(b(P), 1) + \mathfrak{s}(b(\sim P), 0) = (b(P) - 1)^2 + b(\sim P)^2$.
  - $I_{\mathfrak{s}}(b, w_2) = \mathfrak{s}(b(P), 0) + \mathfrak{s}(b(\sim P), 1) = b(P)^2 + (b(\sim P) - 1)^2$.
- If one adopts the Brier Score as one's measure of b's inaccuracy, then one can give an accuracy-dominance argument for the axioms of the probability calculus.
- de Finetti [1] was the first to prove such a *Brier*-dominance theorem. Joyce [6, 5] interprets this as *accuracy*-dominance.

- **Theorem** (de Finetti). b is *non*-probabilistic *if and only if* there exists a *probabilistic* credence function b' such that (a) *b' has a strictly lower Brier Score than b at some worlds*, and (b) *b' never has a greater Brier Score than b at any world*.
- And, the proof of de Finetti's theorem is *constructive* — it tells us *precisely which functions b'* "Brier-dominate" b.

- Joyce [6, 5] uses de Finetti's Theorem (and generalizations of it) to ground an (epistemic) *probabilistic coherence norm*.

(PC) S's credences b should be *probabilistic — on pain of being Brier-dominated* by (*specific*) credence functions b'.

- Because Joyce thinks that Brier Score is a good measure of "credal inaccuracy", he thinks this provides incoherent agents with some "*epistemic* reason" to be Pr-coherent.
- Maher [10] points out that other *prima facie* plausible measures of "inaccuracy" do *not* undergird (PC). I'll return to that issue below. But, first, a *concrete* toy example.

- Suppose S adopts the Brier Score as their I-measure, and that S's b is non-probabilistic. Then, there are alternative (coherent) credence functions b' that accuracy-dominate b.
- Intuitively, these b' functions should "look epistemically better" (in a precise sense) than S's current credences b.
- But, our "evidentialist" ("Kolodny's revenge") worry lingers.
- Consider a very simple toy agent S with one sentence P in their language. And, suppose S's credence function assigns $b(P) = 0.2$ and $b(\sim P) = 0.7$. So, S's b is *non*-probabilistic.
- It follows from de Finetti/Joyce's theorems that there is *a specific set of* credence functions b' that Brier-dominate b.
- The figure on the next slide depicts this situation. The red dot is S's credence function b. And, the shaded region depicts the credence functions b' that *Brier-dominate* b. [The black dot at $\langle 0.2, 0.8\rangle$ depicts the *only probabilistic* credence function that is compatible with $b(P) = 0.2$.]

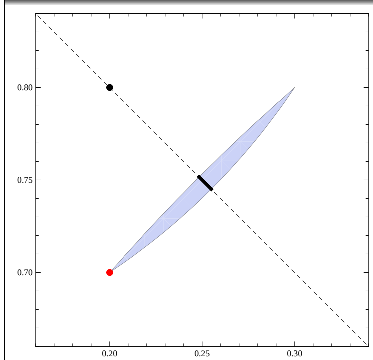

- Suppose that S *has good reason* to assign $b(P) = 0.2$ (*e.g.*, *S knows that the objective chance of P* is 0.2).
- Here, *all* the Brier-dominating functions b' are s.t. $b'(p) \neq 0.2$.
- So, *all the Brier-dominating functions b' may be "ruled-out" by S's evidence.*
- *Perhaps*, *b' needn't* "look better" than b.

- I don't have the space to delve into the various other worries I have about Joyce's argument(s) for probabilism. [But, in my lecture next week, I will discuss *another* worry.]
- For now, I have a *suggestion* re the quantitative case.
- Based on our experience from the qualitative and comparative cases, *we should not expect* an AD-justification of the *full* probabilistic norm(s) in the quantitative case…

- Rather than trying to "justify" the use of $\mathfrak{s}$ (or some other "distance measure" that yields the *full* probabilistic norms), why not start with *desiderata* for distance measures d? *E.g.*,
  - $d(x, x) = 0$.
  - $d(x, y) = d(y, x)$.
  - $d(x, y) \leq d(x, z) + d(z, y)$.
- Once we settle on desiderata ($\mathfrak{D}$) for adequate measures of distance (in this context), then we could ask the following:

(Q) What accuracy-dominance norms are entailed by $\mathfrak{D}$?

- In other words, (Q) is asking what accuracy-dominance norms are *agreed upon* by *all* inaccuracy measures $I_d(b, w)$, where *all* we assume about d is that it satisfies desiderata $\mathfrak{D}$.
- I don't have an answer to (Q). But, I conjecture that this *will* lead to norms for b that are similar to those we saw in the *comparative* case — *e.g.*, if $p \vDash q$, then $b(p) \leq b(q)$, *etc.*
  - Idea: start with $\mathfrak{s}(x, y)$ and Maher's $\delta(x, y) = |x - y|$.

[1] B. de Finetti, *The Theory of Probability*, Wiley, 1974.

[2] ______, *Foresight: Its Logical Laws, Its Subjective Sources*, in H. Kyburg and H. Smokler (*eds.*), *Studies in Subjective Probability*, Wiley, 1964.

[3] T. Fine, *Theories of Probability*, Academic Press, 1973.

[4] P. Fishburn, *The Axioms of Subjective Probability*, *Statistical Science*, 1986.

[5] J. Joyce, *Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief*, in F. Huber and C. Schmidt-Petri (*eds.*), *Degrees of Belief*, 2009.

[6] ______, *A Nonpragmatic Vindication of Probabilism*, *Philosophy of Science*, 1998.

[7] J.M. Keynes, *A Treatise on Probability*, MacMillan, 1921.

[8] N. Kolodny, *How Does Coherence Matter?*, *Proc. of the Aristotelian Society*, 2007.

[9] C. Kraft, J. Pratt and A. Seidenberg, *Intuitive Probability on Finite Sets*, *The Annals of Mathematical Statistics*, 1959.

[10] P. Maher, *Joyce's Argument for Probabilism*, *Philosophy of Science*, 2002.

[11] D. Scott, *Measurement Structures and Linear Inequalities*, *Journal of Mathematical Psychology*, 1964.

- If S violates *Monotonicity* (4), then S is accuracy-dominated.
  - (4) If p entails q, then $p \nsucc_S q$.

	P	Q	$P \succ_S Q$	$Q \succ_S P$
w_1	**T**	**T**	B	B
w_2	**T**	**F**	—	—
w_3	**F**	**T**	C	A
w_4	**F**	**F**	B	B

- Indeed, as this table shows, *any scoring scheme that satisfies our desiderata* [*viz.*, (†) $\Longrightarrow A < C$] entails Monotonicity.
- To see that de Finetti's additivity axiom (3) does *not* have a dominance justification, one must look at *all* the possible ways of "fixing" a violation of (3), and show that *none* of these lead to a comparison set that dominates the original.
- There aren't that many cases to check. [I won't show them.]
- On the next slides, I'll discuss the Scott Axiom…

- Before stating the Scott Axiom, we'll need one definition:

  Definition. For each state description $\mathfrak{s}$ and each *sequence* (n-tuple) of propositions $\mathbf{Z} = \langle z_1, \ldots, z_n \rangle \in \prod_n \mathcal{B}$, let $\mathfrak{c}(\mathfrak{s}, \mathbf{Z})$ be the number of elements of $\mathbf{Z}$ that are entailed by $\mathfrak{s}$.

- OK, here's the (dreaded) Scott Axiom:

  (SA) Let $\mathbf{X}, \mathbf{Y} \in \prod_n \mathcal{B}$ be (arbitrary) sequences of propositions, each having length $n > 0$. Let $\langle x_1, \ldots, x_n \rangle$ denote the members of $\mathbf{X}$, and $\langle y_1, \ldots, y_n \rangle$ denote the members of $\mathbf{Y}$.

  If the following two conditions are satisfied

  i. For every state description $\mathfrak{s}$, $\mathfrak{c}(\mathfrak{s}, \mathbf{X}) = \mathfrak{c}(\mathfrak{s}, \mathbf{Y})$.

  ii. For all $i \in (1, n]$, $x_i \succ_S y_i$.

  then, the following must also be the case

  iii. $y_1 \succ_S x_1$.

- Not only is (SA) *unintuitive*, it is also *quite strong*. It entails *both* de Finetti's "additivity" (3) *and* (full) transitivity of $\succ_S$.

- I think the best way to grasp the content of (SA) is *via* the following illuminating theorem of Fishburn [4, Ch. 4].

  Theorem (Fishburn). (SA) is true *if and only if* there exists a mass function $\mathfrak{m}$ on $\mathcal{B}$ such that, for all propositions p and q in $\mathcal{B}$, the following *real-valued representation* holds:

  $$(\star) \qquad p \succ_S q \textit{ if and only if } \sum_{\mathfrak{s}_p \vDash p} \mathfrak{m}(\mathfrak{s}_p) > \sum_{\mathfrak{s}_q \vDash q} \mathfrak{m}(\mathfrak{s}_q).$$

  And, given de Finetti's axiom (2), there will always be a *probability* mass function $\mathfrak{m}$ satisfying ($\star$).

- Fishburn's Theorem reveals that (SA) *alone* ensures a real-valued representation ($\mathcal{R}_{\succ_S}$) of the $\succ_S$-ordering.
- Not only does this imply de Finetti's additivity axiom (3), but it also implies axiom (1) as well ($>_{\mathbb{R}}$ is a strict total order).
- Thus, once we have (SA) on board, the only axiom of de Finetti that can do *any* work is his axiom (2), which just ensures that $\mathcal{R}_{\succ_S}$ is a *probabilistic* representation of $\succ_S$.