ACCURACY-CENTERED EPISTEMOLOGY

August 16, 2013
Bristol Summer School
Bristol, UK

Jim Joyce
Department of Philosophy
University of Michigan
jjoyce@umich.edu

Two Approachs to Epistemology

Accuracy-centered. The cardinal epistemic good is the holding of beliefs that accurately reflect the world's state. Believers have a duty to rationally pursue doxastic accuracy. (Compare A. Goldman's "veritistic value".)

Evidence-centered. Believers have an epistemic duty to hold beliefs that are *well-justified* in light of their evidence.

My Aim: To paint a compelling picture of accuracy-centered epistemology in which evidential considerations play a central role, and to refute a recent objection to the accuracy-centered approach.

Basic Line: The rational pursuit of accuracy never requires us to invest more confidence in any proposition than our evidence warrants. Honoring our duty to hold well justified beliefs never interferes with the rational pursuit of accuracy.

ACCURACY SCORES AND EPISTEMIC VALUES

- An **inaccuracy score** $\mathfrak{I}$ associates each credal state ***b*** and world ω with a non-negative real number, $\mathfrak{I}$(***b***, ω), which measures ***b*'s** overall inaccuracy when ω is actual (where 0 = perfection).

  Truth-Directedness. Moving credences closer to truth-values always improves accuracy.

  Extensionality. ***b***'s inaccuracy at $\boldsymbol{\omega}$ is solely a function of the credences ***b*** assigns and the truth-values $\boldsymbol{\omega}$ assigns.

  Continuity. Inaccuracy scores are continuous (for each $\boldsymbol{\omega}$).

  Propriety. If ***b*** is a probability then ***b*** itself uniquely minimizes expected inaccuracy when expectations are calculated using ***b***.

- A score meeting these conditions captures ***a consistent way of valuing closeness to the truth***.

What kinds of values? EPISTEMIC VALUES!

Useful Examples of Accuracy Scores

$\langle h, t\rangle$ is the state in which a believer assigns credence h to H and t to $\sim H$.

- *Brier*: $\mathfrak{B}_1(h, t) = \frac{1}{2}\,[(1 - h)^2 + t^2]$

  $\mathfrak{B}_0(h, t) = \frac{1}{2}\,[h^2 + (1 - t)^2]$

- *Log*: $\mathfrak{L}_1(h, t) = -\frac{1}{2}\,[ln(h) + ln(1 - t)]$

  $\mathfrak{L}_0(h, t) = -\frac{1}{2}\,[ln(1 - h) + ln(t)]$

- *Absolute Value**: $\mathfrak{A}_1(h, t) = \frac{1}{2}\,[1 + (t - h)]$

  $\mathfrak{A}_0(h, t) = \frac{1}{2}\,[1 + (h - t)]$

- *Square Root**: $\mathfrak{S}_1(h, t) = \frac{1}{2}\,[(1 - h)^{1/2} + t^{1/2}]$

  $\mathfrak{S}_0(h, t) = \frac{1}{2}\,[h^{1/2} + (1 - t)^{1/2}]$

*not proper

ACCURACY FOR CREDENCES

- Believers have an epistemic duty to hold credences that minimize *estimated inaccuracy* in light of their evidence.

- These are the credences they will see as likely to strike the optimal balance between the good of being confident in truths and the evil of being confident in falsehoods (where the magnitudes of the goods and evils are measured by an appropriate scoring rule).

  - To hedge against even greater inaccuracy, believers often must take 'epistemic gambles' by holding credences that are certain to be less than perfectly accurate.

    - Example: If you know that a coin is biased 3:1 for heads, then when investing a credence of h in heads you must balance off a 75% chance of having an inaccuracy of $(1 - h)^2$ against a 25% chance of having an inaccuracy of h^2 (assuming Brier).

    - The best thing to do in such circumstances is to set $h = 0.75$. You are sure to be less than perfectly accurate, but (as we will see) you do minimize your estimated inaccuracy.

JUSTIFICATION IN ACCURACY-CENTERED EPISTEMOLOGY

Justification Principle. A body of evidence $\mathcal{E}$ provides more justification for one credal state than for another when $\mathcal{E}$ requires a rational believer to fix a higher estimated accuracy for the first state than for the second.

- Justification is a matter of holding credences that have low estimated inaccuracy in light of the available evidence.

NOTE: Just as in traditional 'full belief' epistemology, in accuracy-centered epistemology we can distinguish two notions of doxastic: being accurate and being well-justified justification.

- If I set h = 1 and you set h = 0.75 and the coin lands HEADS, who did better, epistemically speaking?

- The answer depends on the question! My belief turned out to be more accurate, but I was lucky. Your belief turned out to be less accurate, but was better justified in light of the evidence that *chance*(HEADS) = 0.75.

ESTIMATION AND EXPECTATION

A Key Question: When does evidence $\mathcal{E}$ requires a rational believer to fix a lower estimated inaccuracy for credal state ***c*** than for credal state ***c****?

- For ***probabilistically coherent*** believers with credences ***b*** there is an easy answer

$$\mathit{Estimate}(\mathfrak{I}(\boldsymbol{c}^*)) > \mathit{Estimate}_{\boldsymbol{b}}(\mathfrak{I}(\boldsymbol{c})) \text{ iff } \Sigma_\omega\, \boldsymbol{b}(\omega)\cdot\mathfrak{I}(\boldsymbol{c}^*, \omega) > \Sigma_\omega\, \boldsymbol{b}(\omega)\cdot\mathfrak{I}(\boldsymbol{c}, \omega)$$

So, estimates are just ***expectations*** for coherent believers.

- The evidence $\mathcal{E}$ is factored into the probability ***b***.
- Propriety ensures that $\mathit{Estimate}(\mathfrak{I}(\boldsymbol{c})) > \mathit{Estimate}(\mathfrak{I}(\boldsymbol{b}))$ for all $\boldsymbol{b} \neq \boldsymbol{c}$, i.e., every coherent credence function sees itself as best justified.
- This leaves an important open question: What coherent credence function should one hold in light of any given body of evidence $\mathcal{E}$?

?? What about incoherent believers? Not really my problem!

ACCURACY-NONDOMINANCE

Joyce (1998): Whether a believer is coherent or not, her estimates must obey:

Nondominance. If ***b* dominates *b***, so that $\mathfrak{I}(\boldsymbol{c}, \omega) > \mathfrak{I}(\boldsymbol{b}, \omega)$ for all worlds ω, then, ***whatever one's evidence might be***, one is required to have an inaccuracy estimate for ***c*** that exceeds one's inaccuracy estimate for ***b***.

- It follows from **JP** that one is *always* better justified in holding ***c*** than in holding ***c***** when *c* dominates ***c*****.

  ★ **Evidence that supports a credal state always provides even more support for any state that accuracy-dominates it.**

  ★ **Evidence that tells against a state always tells even more strongly against anything that state dominates**.

Important! Choosing an inaccuracy score commits us to views about which credal states are better justified in light of the available evidence.

THE ACCURACY ARGUMENT

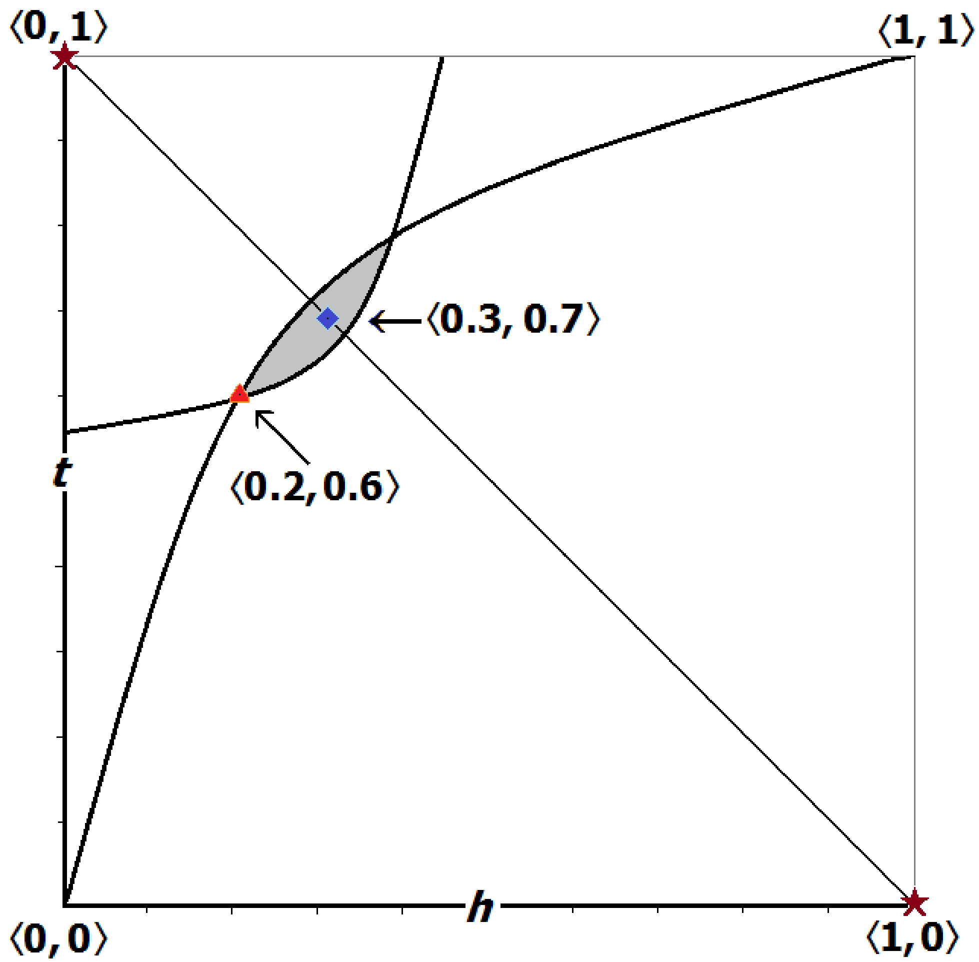

Points ⟨*h*, *t*⟩ are credences for *H* and ~*H*. The points ⟨0, 1⟩ and ⟨1, 0⟩ represent consistent truth-value assignments. The incoherent pair ⟨0.2, 0.6⟩ is accuracy-dominated by the coherent pair ⟨0.3, 0.7⟩.

CATEGORICAL VERSUS NON-CATEGORICAL NORMS FOR CREDENCES

The accuracy argument establishes probabilistic coherence as a *categorical* epistemic norm for credences: it holds good for ***any*** scoring rule (meeting the conditions above), and in ***any*** evidential situation.

- *Hypothetical* norms depend on the evidential situation.
  - **Truth Norm**: If your evidence entails that *H* is true, then you should be certain of *H*. (Good for any accuracy score)
  - **Principal Principle**. If your evidence is that the current objective chance of *H* is *h*, and if you have no 'inadmissible' information about *H*, then your evidence requires you to assign a credence of *h* to *H*. (Good for any accuracy score, see below.)

- *Qualified* norms depend on the choice of a scoring rule.
  - **Split the difference**. If the totality of your evidence justifies the credal states $\langle x, y\rangle$ and $\langle 1 - y, 1 - x\rangle$ equally well, then you should adopt the credences $\langle ½ + ½(x - y), ½ + (y - x)\rangle$. Good for Brier, but not all scores.

KEY TASKS IN AN ACCURACY-CENTERED EPISTMOLOGY

Theory of Justification. Identify the norms that describe the ways in which evidence constrains rational estimates of inaccuracy for credal states.

Example (**Chance Estimation)**: A believer who knows that *chance*(*H*) = x = 1 – *chance*(~*H*), and has no 'inadmissible' information, should use $Exp_x(\mathscr{I}(h, t)) = x \cdot \mathscr{I}_1(h, t) + (1 - x) \cdot \mathscr{I}_0(h, t)$ to estimate $\langle h, t \rangle$'s inaccuracy.

- When combined with **Propriety** this entails the **Principal Principle**.

Theory of Accuracy. Identify and justify the (further) properties that epistemic scoring rules should possess.

Example (**0/1 Symmetry**): $\langle h, t \rangle$ should be as accurate when *H* is true as $\langle 1 - h, 1 - t \rangle$ is when *H* is false.

- This is where ***epistemic values*** come into play.

JUSTIFICATION EXAMPLE-1: CHANCE ESTIMATION

A believer who knows only that *chance*(H) = x should use $\langle h, t\rangle$'s *objective expected inaccuracy* as her estimate of its inaccuracy, so that

$$Estimate(\mathscr{I}(h, t)) = Exp_x(\mathscr{I}(h, t)) = x\cdot\mathscr{I}_1(h, t) + (1 - x)\cdot\mathscr{I}_0(h, t)$$

Rationale: Suppose you must fix credences for each of a series of independent events $\pm H_1, \pm H_2, \ldots, \pm H_N$ in which each H_n has the same objective chance x.

- Since trials are IID, you should settle on the same credences for each, so that $\langle h_n, t_n\rangle = \langle h_1, t_1\rangle$ for $n \leq N$.
- Then, the state that minimizes objective expected inaccuracy is also the one that produces the least *total* inaccuracy in the most typical scenario in which long-run frequencies line up with chances.
- So, to minimize inaccuracy over the long run, your objectively best bet is to hold credences that minimize $Exp_x(\mathscr{I}(h, t))$.
- This is a situation were the strategy that works in the long run should also be used in the single case.

JUSTIFYING THE PRINCIPAL PRINCIPLE

(Compare the "Chance Dominance Argument" in Pettigrew, 2013)

- **Chance Estimation** implies that a person who knows only that *chance*(H) = x should use $Exp_x(\mathscr{I}(h, t))$ as her estimate of $\langle h, t\rangle$'s inaccuracy.

- **Propriety** then implies that $Exp_x(\mathscr{I}(h, t)) > Exp_x(\mathscr{I}(x, 1 - x))$ for all $\langle h, t\rangle$.

- So, by the **Justification Principle**, $\langle x, 1 - x\rangle$ is the best justified credal state in light of the evidence that *chance*(H) = x.

NOTE: The 'no inadmissible evidence' clause of **PP** can be explained in terms of accuracy.

- An item of data is inadmissible relative to the information $\boldsymbol{ch}(H) = x$ just in case it requires believers to assign a lower estimated inaccuracy to some *specific* alternative credence $\boldsymbol{b}(H) = h \neq x$.
- **Accuracy information always trumps chance information**!

JUSTIFICATION EXAMPLE-2: PETTIGREW ON INDIFFERENCE
(as retold or mistold by Joyce)

- **PoI**. If you are entirely ignorant about H, you should set $h = ½$ and $t = ½$.

- **MiniMax Estimation.** If you are entirely ignorant about H, your estimate of $\langle h, t\rangle$'s accuracy should be its maximum potential inaccuracy, so that

$$Estimate(\mathscr{I}(h, t)) = \max\{\mathscr{I}_1(h, t), \mathscr{I}_0(h, t)\}.$$

- So, by **Justification Principle**, the best justified credal state is the one that minimizes $\max\{\mathscr{I}_1(h, t), \mathscr{I}_0(h, t)\}$.

- For a wide range of scoring rules, including the Brier score, $\langle ½, ½\rangle$ uniquely minimizes $\max\{\mathscr{I}_1(h, t), \mathscr{I}_0(h, t)\}$.

I like the *style* of this argument, but not the substance since I reject both PoI and MinMax Estimation.

Choosing an Accuracy Score, Example-1

Choosing an inaccuracy score commits us to definite judgments about which credal states are better justified in light of the available evidence.

- They decide how inaccuracies of individual credences are to be weighed off against one another in overall assessments of total credal states.

- Let's compare the support the data ***chance***(H) = 0.2 provides for $\langle 0.2, 0.6\rangle$ and for $\langle 0.3, 0.7\rangle$ by asking which is better justified.

- **The answer depends on how we measure inaccuracy**!

Brier Score: ***Exp***($\mathcal{B}$(0.2, 0.6)) = 0.18 > ***Exp***($\mathcal{B}$(0.3, 0.7)) = **0.17**
Root Score: ***Exp***($\mathcal{R}$(0.2, 0.6)) = **0.5988 < *Exp***($\mathcal{R}$(0.3, 0.7)) = 0.6055

A WRONG WAY TO THINK ABOUT THE PROBLEM

- It is tempting to think that we should try to develop a free-standing theory of accuracy, *untainted by evidential considerations*, that rules out scores like Square Root and rules in requirements like Propriety.

- This is the wrong. Evidential considerations are in the picture from the start!

  - The relationship between epistemic norms and accuracy norms is symbiotic, not hierarchical. Not all epistemic norms are will be *derived* from accuracy norms (though some well), they *cohere* with them.

  - Inaccuracy scores are ways of measuring ‘closeness to truth’ that reflect our considered views about how such closeness should valued. Part of our goal in choosing a score is to promote correct epistemic values.

  - Methodological Point: It ***is*** legitimate to reject some putative accuracy score merely because, in the presence of principles of justification we accept, it entails that ***c*** is better justified than ***b*** when we know ***b*** is better justified than ***c***.

EXAMPLE: 💭💭 ABSOLUTE VALUE AND SQUARE ROOT SCORES 💭💭

You have rolled a die 10000 times and have observed these outcomes:

ONE	TWO	THREE	FOUR	FIVE	SIX
1000	1000	1000	2000	2000	3000

Evidential Norm: Based on this evidence, $\langle 0.1, 0.1, 0.1, 0.2, 0.2, 0.3\rangle$ is better justified than $\langle 0, 0, 0, 0.1, 0.1, 0.2\rangle$, i.e., your inaccuracy estimate for the former credal state should be lower than you estimate for the latter.

- Fact: If we measure inaccuracy using the absolute value score the latter credal state dominates the former! For this score the only undominated credences are those for which some *N* has credence zero. Same for the Square Root score.

Here our norms of evidence inform and constrain our accuracy scores.

The Accuracy Score as Consistency Test

In developing an "accuracy-centered" epistemology for credences, the goal is to show that various evidential norms – **Truth**, **PP**,… – are jointly consistent with the idea that accuracy is the cardinal epistemic good.

- To achieve this goal one must prove that there exists an inaccuracy score $\mathcal{I}$ such that:

  ★ No norm permits $\mathcal{I}$-dominated credences in any evidential situation.

  ★ No norm prohibits the credences ***b*** in any evidential situation unless it also prohibits any credences that are $\mathcal{I}$-dominates by ***b***.

  ★ More generally, no norm permits ***c*** when the available evidence requires the believer to fix a higher estimate for the accuracy of another state ***b***.

Key Point: In proving that the required $\mathcal{I}$ exists one, in effect, shows that there is at least one way of valuing accuracy that reflects the epistemic values that the evidential norms incorporate.

EXAMPLE: ACCURACY AND THE HIERARCHY OF EPISTEMIC EXPERTS

- Imagine a believer, with a probabilistically coherent credence function ***b***, who treats some source of (probabilistically coherent) information ***q*** as an *expert* about a proposition *H*.

  Def. ***q*** is an **expert** for ***b*** about *H* just in case $\boldsymbol{b}(H \mid \boldsymbol{q}(H) = x) = x$.

  - This means that the believer will defer to ***q***'s values when she knows what those values are. (Note: the JP and the Accuracy Argument require ***q*** to be a *probability*: never defer to an incoherent 'expert'!)

- An **expert principle** is an evidential norm of the form:

  If you know $\boldsymbol{q}(H) = x$, and if this is all your relevant evidence about *H*, then *x* should be your credence for *H*.

  Examples:

  - $\boldsymbol{q}(H) = @(H)$ = *H*'s actual truth value (always an expert)
  - $\boldsymbol{q}(H) = \boldsymbol{chance}_{now}(H)$ = *H*'s current chance
  - $\boldsymbol{q}(H) = \boldsymbol{chance}_{later}(H)$ = *H*'s chance in an hour

EXAMPLE: ACCURACY AND THE HIERARCHY OF EPISTEMIC EXPERTS

Expert Principle: If you know $\boldsymbol{q}(H) = x$, and if this is all your relevant evidence about X, then your $Exp_x(\mathcal{I}(h, t)) = x \cdot \mathcal{I}_1(h, t) + (1 - x) \cdot \mathcal{I}_0(h, t)$ should be your estimate for the inaccuracy for $\langle h, t \rangle$.

- **Propriety** and **JP** then entails that the $\langle x, 1 - x \rangle$ are the best justified credences given the evidence.
-

The rub in any expert principle is the "this is all your evidence about *X*" clause, which can be problematic when experts give conflicting advice.

Def. $\boldsymbol{q}$ *trumps* $\boldsymbol{r}$ for $\boldsymbol{b}$ exactly if $\boldsymbol{b}(X \mid \boldsymbol{q}(X) = x \;\&\; \boldsymbol{r}(X) = y) = x$ for all x, y.

E.g., @ trumps everything, later chances trump earlier chances.

Question: Under what conditions will $\boldsymbol{q}$ trump $\boldsymbol{r}$ for $\boldsymbol{b}$?

- Intuitively, the trumping expert should be the one that ***b*** expects to be most accurate. This turns out to be true if we assume **Propriety**!

EXAMPLE: ACCURACY AND THE HIERARCHY OF EPISTEMIC EXPERTS

Fact: Given any $\mathscr{I}$-score that satisfies Truth-directedness and Propriety, ***q*** *trumps* ***r*** for ***b*** only if ***b***'s estimate of ***r***'s inaccuracy exceeds ***b***'s estimate of ***q***'s inaccuracy.

See the Appendix for the Proof

- So, this *one* aspect of the theory of epistemic experts can be subsumed into the accuracy-based framework: the expert 'pecking order' goes by increasing expected accuracy.

- That the theory gets this right is a further reason to accept Propriety!

EXAMPLE: PRIORS AS EXPRESSIONS OF EPISTEMIC VALUE

- On an accuracy-centered picture the choice of prior probabilities can be seen as an expression of the way we **value** closeness to the truth.

    - $\mathcal{I}$ has ***0/1-symmetry*** when $\mathcal{I}_0(1 - h) = \mathcal{I}_1(h)$.
        - Brier and Log have this property.
        - The score $\mathcal{I}_1(h) = (1 - h)^3$ and $\mathcal{I}_0(h) = ½\cdot(3h^2 - 2h^3)$ lacks this proprerty. Call this the ***Skew Score***.

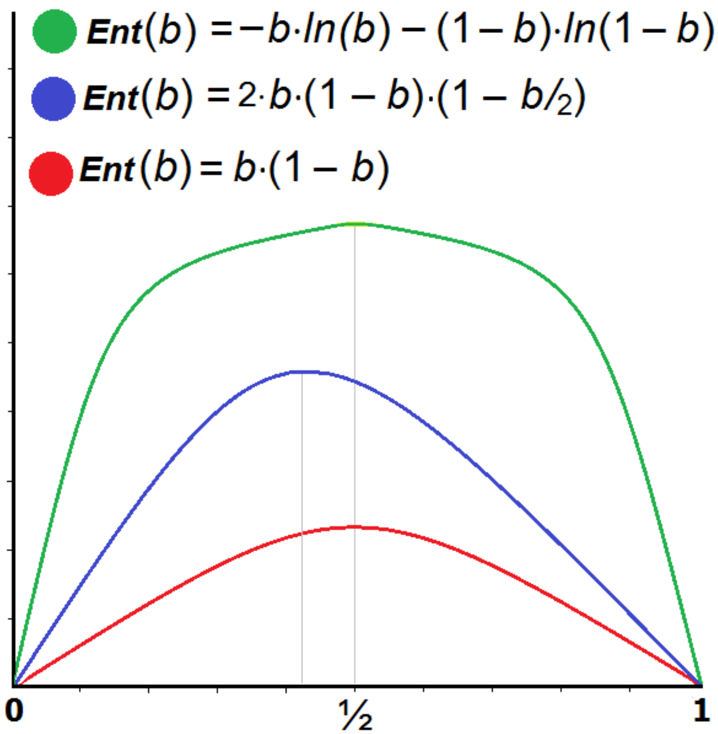

The **green** ***Ent***-function, the *Shannon entropy*, generates the Log score.

The **blue** ***Ent***-function generates the Skew score.

The **red** ***Ent***-function, the *variance*, generates the Brier score.

- The Log and Brier self-scores are symmetrical and achieve maxima at ½, whereas Skew has its maximum at about 0.423.
- ***Ent***'s maximum indicates the credal assignment that sees itself as being "farther from the truth" than any other assignment sees itself being.

Max-Ent AS A RECIPE FOR CHOOSING PRIORS

- When one is committed to an inaccuracy score $\mathscr{I}$ as a codification of one's epistemic values there is a kind of mechanism for translating one's values into prior probabilities.

- When ***Ent***(***b***) > ***Ent***(***c***) a believer with credences ***c*** will take herself to be more accurate than a believer with credence ***b*** does (holding $\mathscr{I}$ fixed). But, one might think that someone with no evidence at all should be minimally confident about the accuracy of her credences.

Max-Ent: An believer with no information should hold the prior that maximizes ***Ent***:

- If $\mathscr{I}$ = Log, this E.T. Jaynes' *MaxEnt* recipe for choosing priors (which I see as an expression of one kind of epistemic value.)

*Let me stress: I am **not** endorsing Max-**Ent**. This is just one example of how an accuracy-based epistemology might go.*

IS HAVING THE MAX-ENT PRIOR REASONABLE? (I'M NOT SURE)

Here are some (fairly lame) considerations in favor of Max-*Ent*:

- It seems that anything more than a minimal level of epistemic self-confidence should be earned (by having evidence).

- An 'epistemic peer' argument: If your credal state has an ***Ent*** below the maximum, then there are always other states with that same ***Ent***-value. Perhaps neither of you deserves to be so confident in your accuracy.

- There is a sense in which ***Ent*** (***b***) is a kind of measure of informativeness (e.g., in the cases of Brier and Log), and in which Max-***Ent*** tells one to have the least informative priors.

But, I am not defending Max-***Ent***. My point is only this: **IF** we accept it, then we have a way of choosing priors that reflects our views about the value of doxastic accuracy.

HOW TO BE LIKE LAPLACE (IF THAT FLOATS YOUR BOAT)

For scoring 0/1-symmetric scores Max-***Ent*** entails Laplace's ***Principle of Indifference***: given no data and a partition $X_1, X_2, \ldots, X_N$, one ends up assigning the uniform probabilities $\boldsymbol{b}(X_n) = {}^1/_N$.

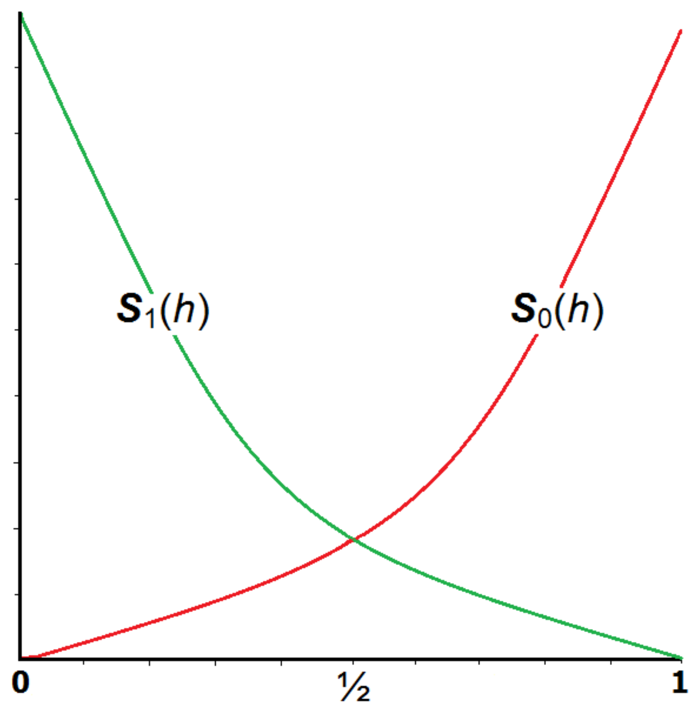

The person values improvements accuracy for $\boldsymbol{b}(H)$ at any point in just the same way that she values improvements in accuracy for $\boldsymbol{b}(\sim H)$ at that point.

HOW NOT TO BE LIKE LAPLACE

Rules that are not 0/1-symmetric can generate non-uniform priors.

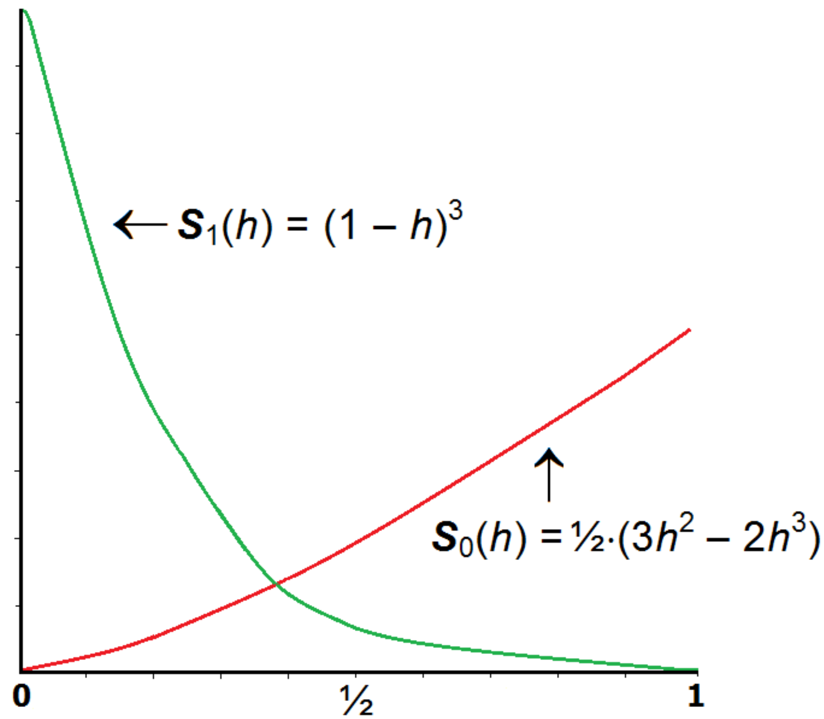

This person worries a more about being way off about *H* than about being way off about ~*H*, but she also cares a little more about making small shifts close to the truth when ~*H* than when *H*.

The result is a prior with ***b***(*H*) = 0.423.

Conclusions

- Well justified credeneces minimize estimated inaccuracy.
- Substantive epistemic norms tell us when the evidence mandates or permits various accuracy estimates.
- The relationship between epistemic norms and accuracy norms is not hierarchical, but symbiotic.
- Evidential considerations factor into the choice of an inaccuracy score because these scores are ways of measuring ‘closeness to the truth’ that reflect our views about how such closeness should valued.
- Conflict between norms of accuracy and norms of evidence should never arise as long as our inaccuracy score properly reflects our epistemic values, including the value we place on well-justified beliefs.
- A set of norms for credences are mutually consistent just in case there is an accuracy score such that: (i) no norm in the set permits believers to hold that do not minimize estimated accuracy.

APPENDIX

Imagine a believer, with a probabilistically coherent credence function ***b***, who treats two sources of probabilistically coherent information ***q*** and ***r*** as experts[1] about proposition *H*. For example, *H* might say that it will rain Friday, and ***r*** and ***q*** might be the chance functions for *H* evaluated, respectively, on Wednesday and Thursday. To say that the believer treats ***q*** and ***r*** as experts means that she will defer to their opinions if she knows what those opinions are, so that for any q, $r \in [0, 1]$ her conditional probabilities satisfy $\boldsymbol{b}(H \mid \boldsymbol{q}(H) = q) = q$ and $\boldsymbol{b}(H \mid \boldsymbol{r}(H) = r) = r$. Say that ***q*** *trumps* ***r*** for ***b*** if the believer aligns her credences with ***q***'s values whenever it and ***r*** deliver different results about *H*, so that $\boldsymbol{b}(H \mid \boldsymbol{q}(H) = q \;\&\; \boldsymbol{r}(H) = r) = q$. Intuitively, the trumping expert should be the one that is expected to be more accurate. And, this turns out to be true as long as scoring rules satisfy Truth-directedness, Extensionality, Normality and Propriety.

> **Theorem**: ***q*** *trumps* ***r*** for ***b*** only if ***b***'s estimate of ***r***'s inaccuracy exceeds its estimate of ***q***'s inaccuracy (relative to any inaccuracy score that satisfies the criteria set down above).

Proof: Use the abbreviations $\boldsymbol{b}(Y, q, r)$ for $\boldsymbol{b}(Y \;\&\; \boldsymbol{q}(H) = q \;\&\; \boldsymbol{r}(H) = r)$ and $\boldsymbol{b}(q, r)$ for $\boldsymbol{b}(\boldsymbol{q}(H) = q \;\&\; \boldsymbol{r}(H) = r) = r$. Now, if ***q*** trumps ***r***, then (because ***b*** is coherent) we will have

[1] Gaifman, Haim. 1988. A Theory of Higher Order Probability. In *Causality, Chance and Choice*, edited by B. Skyrms and W. Harper, pp. 191-219. Dordrecht: Kluwer.

both $\boldsymbol{b}(H, q, r) = q\cdot\boldsymbol{b}(q, r)$ and $\boldsymbol{b}(\sim H, q, r) = (1 - q)\cdot\boldsymbol{b}(q, r)$. As before, $\boldsymbol{I}_1(r)$ is the accuracy that results from investing credence r in H when H is true, and $\boldsymbol{I}_0(r)$ is the accuracy that results from that investment when H is false. Then, $\boldsymbol{b}$'s expectation for $\boldsymbol{r}$'s inaccuracy is

$$
\begin{aligned}
Exp(\boldsymbol{I}(\boldsymbol{r})) &= \textstyle\sum_r \boldsymbol{b}(H, r)\cdot\boldsymbol{I}_1(r) + \boldsymbol{b}(\sim H, r)\cdot\boldsymbol{I}_0(r) \\
&= \textstyle\sum_{q,\, r} \boldsymbol{b}(H, q, r)\cdot\boldsymbol{I}_1(r) + \boldsymbol{b}(\sim H, q, r)\cdot\boldsymbol{I}_0(r) && \text{Law of Total Probability} \\
&= \textstyle\sum_{q,\, r} q\cdot\boldsymbol{b}(q, r)\cdot\boldsymbol{I}_1(r) + (1 - q)\cdot\boldsymbol{b}(q, r)\cdot\boldsymbol{I}_0(r) && \text{Trumping} \\
&= \textstyle\sum_{q,\, r} \boldsymbol{b}(q, r)\cdot[q\cdot\boldsymbol{I}_1(r) + (1 - q)\cdot\boldsymbol{I}_0(r)] \\
&< \textstyle\sum_{q,\, r} \boldsymbol{b}(q, r)\cdot[q\cdot\boldsymbol{I}_1(q) + (1 - q)\cdot\boldsymbol{I}_0(q)] && \text{when } q \neq r \text{ by Propriety.} \\
&= \textstyle\sum_{q} \boldsymbol{b}(q)\cdot[q\cdot\boldsymbol{I}_1(q) + (1 - q)\cdot\boldsymbol{I}_0(q)] && \text{Law of Total Probability} \\
&= \textstyle\sum_{q} q\cdot\boldsymbol{b}(q)\cdot\boldsymbol{I}_1(q) + (1 - q)\cdot\boldsymbol{b}(q)\cdot\boldsymbol{I}_0(q) \\
&= \textstyle\sum_{q} \boldsymbol{b}(H, q)\cdot\boldsymbol{I}_1(q) + \boldsymbol{b}(\sim H, q)\cdot\boldsymbol{I}_0(q) \\
&= \textstyle\sum_{q} \boldsymbol{b}(H, q)\cdot\boldsymbol{I}_1(q) + \boldsymbol{b}(\sim H, q)\cdot\boldsymbol{I}_0(q) \\
&= Exp(\boldsymbol{I}(\boldsymbol{q}))
\end{aligned}
$$

So, $Exp(\boldsymbol{I}(\boldsymbol{r})) > Exp(\boldsymbol{I}(\boldsymbol{q}))$ when $\boldsymbol{q}$ trumps $\boldsymbol{r}$.

USEFUL REFERENCES

K. Easwaran and B. Fitelson (2012) “An ‘Evidentialist’ Worry about Joyce’s Argument for Probabilism,” *Dialectica* **66**: 425-433.

Gibbard, A. (2008) “Rational Credence and the Value of Truth,” in T. Gendler and J. Hawthorne, eds., *Oxford Studies in Epistemology*, vol. 2. Oxford: Clarendon Press.

Goldman, Alvin (2010) “Epistemic Relativism and Reasonable Disagreement,” in Richard Feldman & Ted Warfield (eds.), *Disagreement*. London and New York: Oxford University Press.

Lindley, D. (1982) “Scoring Rules and the Inevitability of Probability,” *International Statistical Review* 50: 1–26.

Joyce, J. (1998) “A Non-Pragmatic Vindication of Probabilism,” *Philosophy of Science* 65: 575–603.

Joyce, J. (2009) “Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief,” in Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of Belief*. Synthese (2009)

Lewis, D. (1980) “A Subjectivist’s Guide to Objective Chance,” in *Studies in Inductive Logic and Probability*, edited by R. Jeffrey, vol. 2. Berkeley: University of California Press: 263–94.

J. Predd, R. Seiringer, E. H. Lieb, D. Osherson, V. Poor, and S. Kulkarni (2009). “Probabilistic Coherence and Proper Scoring Rules.” *IEEE Transactions on Information Theory*, **55**(10): 4786–4792, 2009.

Pettigrew, Richard (2013) “A New Epistemic Utility Argument for the Principal Principle,” *Episteme* **10***: 19-35.*