Why Evidentialists Need not Worry About the Accuracy Argument for Probabilism

James M. Joyce
Department of Philosophy
University of Michigan
jjoyce@umich.edu

In their (2012) paper "An Evidentialist Worry About Joyce's Argument for Probabilism" Kenny Easwaran and Branden Fitelson raise a "basic and fundamental" worry about the *accuracy argument* for probabilism of Joyce (1999) and (2009). The accuracy argument aspires to establish probabilistic coherence as a core normative requirement in an *accuracy-centered* epistemology for credences[1]. It does this by showing that any system of credences which violates the laws of probability will be *accuracy-dominated* by an alternative system that is strictly more accurate in every possible world. The argument relies on the key normative premise that accuracy-dominated credal states are categorically forbidden: no matter what other virtues they might possess, holding accuracy-dominated credences is an epistemic sin *in every evidential situation*. Easwaran and Fitelson object to this uncompromising position, alleging that the pursuit of accuracy can undermine the legitimate epistemic goal of having credences that are well-justified in light of the evidence. In short, they see a conflict between the following two norms:

> **Accuracy**: The cardinal epistemic good is *doxastic accuracy*, the holding of beliefs that accurately reflect the world's state. Believers have an unqualified epistemic duty to rationally pursue the goal of doxastic accuracy.
>
> **Evidence**: Believers have an unqualified epistemic duty to hold beliefs that are *well-justified* in light of their total evidence.

If Easwaran and Fitelson are right, then we will be forced to choose between accuracy-centered epistemology, which takes the first norm as fundamental, and evidence-centered epistemology, which gives the second pride of place.

---

[1] A believer's *credence* (a.k.a. degree of belief, partial belief) in a proposition *X* is her level of confidence in *X*'s truth. It reflects the extent to which she is disposed to presuppose *X*'s truth in her theoretical and practical reasoning. Credences are contrasted with *categorical* (a.k.a. full, all-or-nothing) beliefs which involve the unreserved acceptance of some proposition as true.

Fortunately, there is no tension between the accuracy and evidence-centered approaches to epistemology. Once we properly understand the workings and ambitions of the accuracy-centered framework it will become clear that Easwaran and Fitelson's worries are misplaced. The rational pursuit of accuracy never requires us to invest more confidence in propositions than our evidence warrants, and honoring our duty to hold well-justified beliefs never forces us to adopt credences that we take to be less than optimally accurate. **Accuracy** and **Evidence** are two sides of the same coin: epistemically rational believers will, in all circumstances, pursue the goal of accuracy by adopting the credences that are best justified in light of their evidence.

The paper has six sections. The first explains and motivates the general idea of an accuracy-centered epistemology for credences. Section 2 provides a brief sketch of the accuracy argument for probabilism and develops the modest formal apparatus that will be needed for the rest of the paper. Section 3 sketches Easwaran and Fitelson's objection, and the next section explains how it goes wrong. However, this is a minor skirmish since, as becomes clear in Section 5, a deeper worry remains. To assuage it one need to prove that no conflict between accuracy norms and legitimate evidential norms can ever arise. This is accomplished in Section 6, which makes it plain that the two sorts of norms will have a symbiotic relationship in any adequate accuracy-centered epistemology of credences. In such an epistemology, all legitimate norms of evidence will be consistent with the central requirement of accuracy-nondomination, and all reasonable measures of accuracy will reflect the epistemic values that norms of evidence codify. As this last section will make clear, while one of my aims in this paper is to explain where Easwaran and Fitelson's arguments go wrong, my larger and more important objective is to paint a compelling picture of an accuracy-centered epistemology in which norms of evidence and norms accuracy live in peace and harmony, and together ensure that believers are always encouraged to hold credences that are both well-justified and as close to the truth as the evidence allows.

1. The Idea of an Accuracy-Centered Epistemology for Credences

Accuracy-centered approaches should be familiar from traditional epistemology, where having true 'full' beliefs is frequently seen as the cardinal epistemic good and having false 'full' beliefs is the chief epistemic evil. This enshrines the *alethic commandment* – Believe truths and eschew falsehoods! – as the font of all doxastic normativity.[2] The resulting epistemology, which is in the business of telling us how best to pursue legitimate epistemic ends, values other aspects of beliefs – justification, safety, reliability, sensitivity,... – to the extent that they further the core alethic goal. Accuracy-centered epistemologies for credences have a similar structure, but with the traditionalist's black-and-white view of accuracy replaced by a more nuanced picture which reflects the fact that credences come in degrees. The categorical good of fully believing truths is

[2] Some traditionalists propose a 'truth-plus' state as the ultimate epistemic good, e.g., Williamson (2000) argues that *knowledge* plays this role. For current purposes, such views count as accuracy-based as long as the putative goal state is essentially truth-entailing.

replaced by the gradational good of investing high credence in truths (the higher the better); the categorical evil of fully believing falsehoods is replaced by the gradational evil of investing high credence in falsehoods (the higher the worse); and the overarching goal changes from that of fully believing truths and only truths to that of minimizing the *degree of divergence* between credences and truth-values. This graded alethic requirement serves as the source of all epistemic normativity, and epistemology has the job of explaining what believers must do to rationally pursue credal accuracy.[3]

The first challenge is to explain what it means for credences to 'diverge' from truth-values. While it is easy to define accuracy in traditional epistemology (accurate = true), the notion is less clear when credences are in play. Joyce (1999) and (2009) propose to measure divergence using the formal device of *epistemic scoring rules* or *inaccuracy scores*. An inaccuracy score is a function $\boldsymbol{I}$ that associates each credal state $\boldsymbol{b}$ and each 'possible world' ω with a non-negative real number, $\boldsymbol{I}(\boldsymbol{b}, \omega)$, which measures $\boldsymbol{b}$'s overall inaccuracy when ω is actual. Inaccuracy is graded on a scale were zero is perfection and larger numbers reflect greater divergence from actuality, so that $\boldsymbol{b}$'s credences are more accurate than $\boldsymbol{c}$'s at ω when $\boldsymbol{I}(\boldsymbol{b}, \omega) < \boldsymbol{I}(\boldsymbol{c}, \omega)$.

Following Joyce (2009), we assume that accuracy scores meet the following conditions:

Truth-Directedness. Moving credences uniformly closer to truth-values always improves accuracy. If $\boldsymbol{b}$ and $\boldsymbol{c}$ differ only in that $\boldsymbol{b}$ assigns higher/lower credences than $\boldsymbol{c}$ does to some truths/falsehoods, then $\boldsymbol{b}$ is more accurate than $\boldsymbol{c}$.

Extensionality.[4] The inaccuracy of a credence function $\boldsymbol{b}$ at a world ω is solely a function of the credences that $\boldsymbol{b}$ assigns and the truth-values that ω assigns.

Continuity. Inaccuracy scores are continuous.

Strict Propriety. If $\boldsymbol{b}$ satisfies the laws of probability, then $\boldsymbol{b}$ uniquely minimizes expected inaccuracy when expectations are calculated using $\boldsymbol{b}$ itself.[5]

A scoring rule that meets these conditions captures *a consistent way of valuing closeness to the truth*. **Truth-directedness** ensures that being close to the truth is lexically prior to any other value that might be incorporated into the score. **Extensionality** stipulates that features of propositions other than their truth-values do not figure into assessments of closeness to the truth.

---

[3] Alvin Goldman (2010) has recently endorsed a similar picture, writing that "just as we say that someone 'possesses' the truth categorically when she categorically believes something true, so we can associate with a graded belief [= credence] a degree of truth possession (n.b., not a degree of truth) as a function of the degree of belief and the truth-value of its content." citemmm

[4] This has the effect of identifying each possible world with a consistent truth-value assignment.

[5] As shown in Joyce (2009), the accuracy argument goes through as long as no coherent credal state is ever accuracy-dominated by another credal state.

This means, for example, that high credences are not worth more when they are invested in informative truths, or when they attach to 'verisimilar' falsehoods, or when they fall near known objective chances. Once $\boldsymbol{I}$ is specified, nothing affects accuracy except the numerical values of credences and truth-values. (Though, as we will see, $\boldsymbol{I}$'s *functional form* can reflect other aspects of epistemic value, like the value of having credences that track known chances.) **Continuity** says that small shifts in credence never cause large leaps in inaccuracy. This is a non-trivial assumption, but we will not discuss it further. **Strict Propriety** ensures that any probabilistically coherent credal state will seem optimal *from its own perspective*. Given a coherent $\boldsymbol{b}$ and any other credence function $\boldsymbol{c}$ (coherent or not) one can calculate $\boldsymbol{c}$'s *expected accuracy* according to $\boldsymbol{b}$ and can compare it to $\boldsymbol{b}$'s expected accuracy computed relative to $\boldsymbol{b}$ itself.[6] If $\boldsymbol{c}$'s expected accuracy exceeds $\boldsymbol{b}$'s in this comparison, then a $\boldsymbol{b}$-believer will judge that $\boldsymbol{c}$ strikes a better balance between the epistemic good of being confident in truths and the epistemic evil of being confident in falsehoods. Following Gibbard (2008), Joyce (2009) argues that believers have an unqualified epistemic duty to abandon such 'self-deprecating' credal states, and uses this fact that to provide a rationale for **Strict Propriety**. We will consider this rationale in §3 below.

While these four requirements rule out many potential inaccuracy scores, many others pass the test. Consider the score of Brier (1950), which identifies inaccuracy with the mean squared Euclidean distance from credences to truth-values. When $\boldsymbol{b}$ is defined on a set of N propositions, the Brier score defines $\boldsymbol{b}$'s inaccuracy at $\boldsymbol{\omega}$ as ${}^1/_N \cdot \Sigma_n (b_n - \omega_n)^2$, where b_n is the credence $\boldsymbol{b}$ assigns to the n^{th} proposition and ω_n is that proposition's truth-value at $\boldsymbol{\omega}$. Alternatively, the *logarithmic* score defines the inaccuracy of investing credence b in a true or false proposition, respectively, as $-log(1 - b)$ or $-log(b)$, and identifies $\boldsymbol{b}$'s total inaccuracy at $\boldsymbol{\omega}$ with ${}^1/_N \cdot \Sigma_n -log(|(1 - \omega_n) - b_n|)$, its mean logarithmic distance from the truth. One can think of these scores, and any others that satisfy the above requirements, as encoding a distinctive way of valuing 'closeness to the truth'.

Anyone who endorses a scoring rule $\boldsymbol{I}$ as the right way to value accuracy (in a context)[7] will rewrite *Accuracy* like this:

> ***Accuracy for Credences***: The cardinal epistemic good/evil is that of having credences with low/high $\boldsymbol{I}$-inaccuracy. Believers have an unqualified epistemic duty to rationally pursue the goal of minimizing $\boldsymbol{I}$-inaccuracy.

This sets up the minimization of gradational inaccuracy as the paramount epistemic end, and puts epistemologists in the business of telling believers how to most rationally pursue it.

---

[6] The definition is $Exp_{\boldsymbol{b}}(\boldsymbol{I}(c)) = \Sigma_\omega \boldsymbol{b}(\omega) \cdot \boldsymbol{I}(\boldsymbol{c}, \omega)$ where ω ranges over all consistent truth-value assignments. **Strict Propriety** says that, when $\boldsymbol{b}$ is a probability, one must have $Exp_{\boldsymbol{b}}(\boldsymbol{I}(c)) > Exp_{\boldsymbol{b}}(\boldsymbol{I}(b))$ for all $\boldsymbol{c} \neq \boldsymbol{b}$.

[7] There is a temptation to think that there is some single correct way to assess inaccuracy. I think, instead, that such assessments are highly contextual.

Now, it might seem that taking this strong stand on the value of accuracy forces us to say that people with more accurate credences are always doing better, all-epistemic-things-considered, than those with less accurate credences. Not so! While *de facto* accuracy is always the goal, the rational pursuit this goal often involves making trade-offs in which some level of guaranteed inaccuracy is tolerated as a means of avoiding the likelihood of even greater inaccuracy. Suppose you and I see a coin land heads 200 times in 1000 independent tosses. On the basis of this evidence you assign credence 0.2 to the proposition that the coin will land heads on its next toss, while I assign credence 1.0. If a head does comes up, does the accuracy-centered picture imply that you made a mistake? Does it hold me up as an ideal? Definitely not! My belief turned out to be more accurate than yours, but by luck. Since neither of us knew how the coin would fall, we both had to rely on data about previous tosses to settle on a credence that would strike the best balance between the epistemic good of being confident in truths and the epistemic evil of being confident in falsehoods. Ignoring the evidence, I took an epistemic risk and invested maximum credence in heads while you 'hedged your epistemic bets' by adopting a credence that the evidence suggested was likely to be highly, yet not perfectly, accurate. Which one of us did the right thing? The answer depends on the question. While I achieved higher accuracy, you better discharged your duty to rationally pursue accuracy since the evidence strongly suggested that my beliefs would be less accurate than yours. Indeed, if we ran the experiment many times, using the observed frequencies as our guide, my average (Brier) inaccuracy would be 0.64 while yours would be 0.16. Which of these considerations – actual accuracy or estimated accuracy in light of evidence – matters most to assessments of credence? Both do, but for different kinds of assessments! Epistemology should both specify the goals toward which believers should strive, and identify the practices and policies that characterize the rational pursuit of these goals. Since a person can attain a goal without having pursued it rationally, or can fail to secure a goal that was rationally pursued, success and failure must be assessed in both arenas. So, just as traditional epistemology draws a distinction between mere true beliefs, which may have been achieved by luck, and beliefs (true or false) that are well-justified by a believer's evidence, an accuracy-centered epistemology for credences should says that, while I had better luck achieving the overall goal of accuracy, you better fulfilled the epistemic duty to pursue this goal in a rational way. Accuracy is the cardinal epistemic virtue, but its rational pursuit is the primary epistemic duty.[8]

Most of the duties imposed by the requirement to rationally purse accuracy depend on the character of a believer's evidence. Believers are obliged to hold credences that, according to their best estimates in light of their evidence, are likely to strike the optimal achievable balance between the good of being confident in truths and the evil of being confident in falsehoods

[8] This has an exact parallel in moral philosophy. Conseqentialists say that the best acts cause the best *actual* outcomes, but also recognize that agents with imperfect information should strive to maximize *estimated* utility in light of their evidence. This can require people to behave in ways that they know will produce less than optimal results so as to avoid the high probability of even worse results.

(where the magnitudes of these goods and evils are measured by an appropriate scoring rule). Rational believers with different evidence will judge different credences to be optimal, and so will have duties to hold different beliefs. So, most epistemic duties are *hypothetical* imperatives. They says that *if one's evidence is such-and-such*, then it is permitted/prohibited/mandatory that one's credences be so-and-so. A fully developed accuracy-centered epistemology will identify such imperatives, and explain how they contribute to the overarching duty to rationally pursue epistemic accuracy. Here are two hypothetical imperatives of this sort:

> **Truth.** If your evidence conclusively shows that some proposition X is true, then you should be fully confident of X.
>
> **Principal Principle (PP)**. If you know that the current objective chance of X is x, and if you have no 'inadmissible' evidence regarding X,[9] then it is impermissible to assign any credence other than x to X, so that $\boldsymbol{b}(ch(X) = x) = 1$ only if $\boldsymbol{b}(X) = x$.

Accuracy-centered approaches unreservedly endorse **Truth** because, relative to *any* scoring rule that satisfies **Truth Directedness**, one always minimizes inaccuracy by being fully confident of truths. The **Principal Principle** is trickier. It places a value on having credences that agree with the objective chances rather than truth-values. If, say, you know that the coin about to be tossed is perfectly fair, then **PP** dictates ½ as the only allowable credence for heads (H). While aligning credences with known chances in this way seems optimal from the perspective of justification, it also puts a ceiling on your accuracy.[10] Indeed, any other credal assignment guarantees you a 50% chance of a better accuracy score (but also a 50% chance of a worse score). In light of this, one might wonder whether there any reason to think that $\boldsymbol{b}(H) = ½$ is the best credence to hold, on grounds of accuracy, when H's objective chance is known to be ½. To put it more bluntly, is there any reason to think that the rational pursuit of accuracy requires, or is even compatible with, **PP**'s demand that believers align their credences with known objective chances?

This is the question Easwaran and Fitelson want to press. They detect a tension between **PP** and the requirement of *accuracy-nondominance*, which sits at the very heart of the accuracy-centered framework. Say that one credal state $\boldsymbol{b}$ *accuracy-dominates* another $\boldsymbol{c}$ when $\boldsymbol{b}$ is sure to be more accurate than $\boldsymbol{c}$ *no matter what the world is like*, i.e., when $\boldsymbol{I}(\boldsymbol{b}, \boldsymbol{\omega}) > \boldsymbol{I}(\boldsymbol{c}, \boldsymbol{\omega})$ for every possible world $\boldsymbol{\omega}$. It is a non-negotiable tenet of accuracy-centered epistemology that accuracy-dominated credal states are rationally defective. The general principle, a *categorical* imperative, is this:

[9] For current purposes, that is direct evidence about X's chances at later times.
[10] With the Brier score your inaccuracy for H will be exactly 0.25.

Accuracy-Nondominance (**AN**). It is epistemically impermissible, whatever one's evidence might be, to hold credences that are accuracy-dominated by some available alternative.

In the same way that non-dominance principles are essential to the idea that pragmatic or moral value can be represented by utility functions, **AN** is essential to the idea that inaccuracy scores capture a coherent sense of 'epistemic (dis)value'. Unless we are willing to endorse **AN** for a given score ***I***, we cannot portray ***I*** as providing a coherent way of valuing 'closeness to truth'. If we do endorse **AN** for ***I***, however, then **Accuracy for Credences** commits us to saying that ***c*** is always worse than ***b*** *all-epistemic-things-considered* when ***b*** accuracy-dominates ***c***. This means, among other things, that *any* advantage that ***c*** might have over ***b*** in terms of justification (say because its values are uniformly closer than ***b***'s to the known objective chances) is trumped by the fact that ***b*** accuracy-dominates ***c***.

This is the aspect of the accuracy-centered approach that Easwaran and Fitelson worry about. The maintain that **AN** and **PP** can conflict, and that when they do the duty to conform one's credences to **PP** overrides the duty to avoid accuracy dominance. Before considering their argument in detail, it may help to first see how **AN** functions in the accuracy argument.

2. The Accuracy Argument for Probabilism

The gist of the accuracy argument can be conveyed by a simple example. Let H say that a head will come up on the next toss of a coin, and consider credence functions defined on the set $\{H \vee {\sim}H, H, {\sim}H, H \,\&\, {\sim}H\}$. The laws of probability require: $\boldsymbol{b}(H \vee {\sim}H) = 1$; $\boldsymbol{b}(H), \boldsymbol{b}({\sim}H) \geq 0$; and $b(H) + b({\sim}H) = 1$. The accuracy argument shows that believers who violate these laws pay a price in accuracy that probabilistically coherent believers can avoid. The key result is this:

Accuracy Theorem:[11] If accuracy is measured using a scoring rule ***I*** that satisfies the four conditions listed above, then

i. every credence function that fails to satisfy the laws of probability is accuracy dominated by some credence function (indeed by one that obeys the laws of probability), and

[11] There are a variety of versions of this theorem, each starting from slightly different premises about scoring rules and arriving at with slightly different conclusions. The differences between these results is not important here. It should be said, however, that the ideal version of the Theorem remains unproven. On this version, one would start with an arbitrary algebra of propositions (not a partition), and would show that the result holds for arbitrary decision rules that satisfy the four conditions above. See Joyce (2009) for further discussion. Interestingly different versions of the result and related results can be found in Joyce (1998), Lindley, D. (1982) and Predd, et. al., (2009).

ii. no credence function that obeys the laws of probability is dominated by anything.

When thinking about this result it helps to have a simple picture in mind. Let's represent credences by pairs ⟨*h*, *t*⟩, with $h = \boldsymbol{b}(H)$ and $t = \boldsymbol{b}(\sim H)$. Consistent truth-value assignments will correspond to the points $\boldsymbol{\omega}_1 = \langle 1, 0\rangle$ (the most accurate credences when *H* is true) and $\boldsymbol{\omega}_0 = \langle 0, 1\rangle$ (the most accurate credences when *H* is false). Probabilistically coherent credences sit on the line segment $\{\langle h, t\rangle: t = 1 - h \text{ and } 0 \leq h \leq 1\}$ running from $\boldsymbol{\omega}_0$ to $\boldsymbol{\omega}_1$. Readers should convince themselves that points which violate either of the first two laws are dominated. For the third law, *Additivity*, suppose *h* and *t* do not sum to one. Then, as FIGURE-1 indicates, there will be curves C_0 and C_1 which contain all the credence functions that are exactly as accurate as ⟨*h*, *t*⟩ when *H* is, respectively, true or false. As long as ***I*** satisfies the four conditions of §1, the **Theorem** shows that interior of the region bounded by C_0 and C_1 is non-empty and that it contains all and only points that accuracy dominate ***b***.

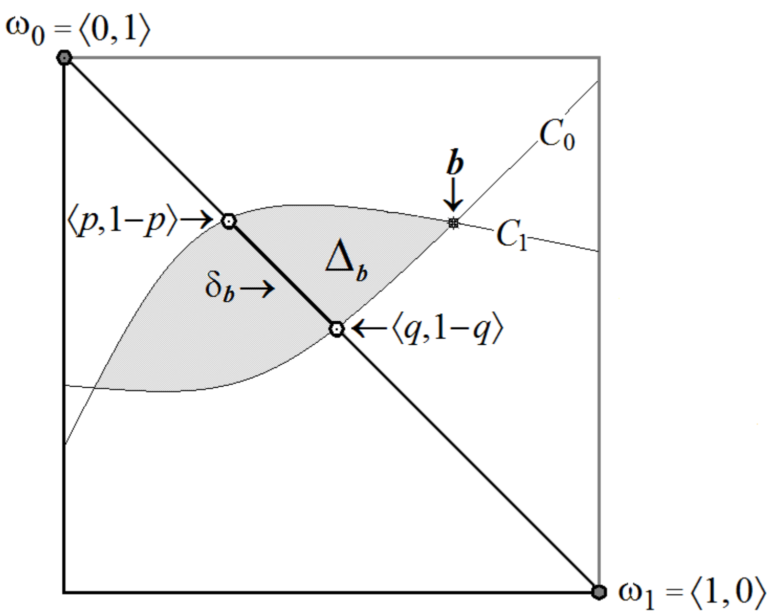

FIGURE-1
The Accuracy Theorem

$\boldsymbol{\omega}_0$ and $\boldsymbol{\omega}_1$ are consistent truth-value assignments: *H* is false in $\boldsymbol{\omega}_0$ and true at $\boldsymbol{\omega}_1$. The line segment between $\boldsymbol{\omega}_0$ and $\boldsymbol{\omega}_1$ contains all *coherent* credence functions. Curve $C_0 = \{\boldsymbol{c} : \boldsymbol{I}(\boldsymbol{c}, \boldsymbol{\omega}_0) = \boldsymbol{I}(\boldsymbol{b}, \boldsymbol{\omega}_0)\}$ passes through all points that are exactly as accurate as ***b*** when *H* is false, and points above and to the left of C_0 are strictly more accurate than ***b*** when $\boldsymbol{\omega}_0$ is actual. Curve $C_1 = \{\boldsymbol{c} : \boldsymbol{I}(\boldsymbol{c}, \boldsymbol{\omega}_1) = \boldsymbol{I}(\boldsymbol{b}, \boldsymbol{\omega}_1)\}$ passes through all points that are exactly as accurate as ***b*** is when *H* is true, and points below and to the right of C_1 are strictly more accurate than ***b*** when $\boldsymbol{\omega}_1$ is actual. The interior of the grey region $\Delta_{\boldsymbol{b}}$ contains all and only credence functions that accuracy-dominate ***b***. The constraints imposed on ***I*** ensure that $\Delta_{\boldsymbol{b}}$ is non-empty. The segment $\delta_{\boldsymbol{b}}$ is composed of coherent credence functions that accuracy-dominate

b. It contains all points $\langle h, 1-h\rangle$ with $p < h < q$, where $\langle p, 1-p\rangle$ lies on C_1 and $\langle q, 1-q\rangle$ lies on C_0. The constraints on $\boldsymbol{I}$ ensure that p and q are unique and that $p < q$.[12]

This should make the basic contours of the accuracy argument fairly clear. It starts by assuming both that inaccuracy scores must satisfy the four conditions in §1, and that accuracy-dominated credal states are categorically forbidden. The **Theorem** then ensures that credences are dominated if and only if they violate the laws of probability. Since it is forbidden to hold dominated credences, a categorical prohibition against probabilistic incoherent credences is thereby derived from the unqualified epistemic duty to rationally pursue the goal of doxastic accuracy. So, on an accuracy-centered picture, there can be no evidential situation in which it is rational to hold incoherent credences, e.g., no evidence can ever make it rationally permissible to assign credences of 0.2 and 0.7 to a proposition and its negation.

3. Easwaran and Fitelson's Evidentialist Worry

As already noted, within the accuracy framework believers have a general duty to hold credences that, in light of their evidence, strike the best balance between the epistemic good of being confident in truths and the epistemic evil of being confident in falsehoods. Achieving this balance often requires trading away the hope of perfect accuracy to obtain an optimal mix of epistemic risk and reward. A key challenge for accuracy-based epistemology is to explain how such tradeoffs are made.

A concrete example might be useful: Imagine a believer, Joshua, who has opinions about whether a certain coin will come up heads or tails when next tossed, and who also has evidence about the coin's bias. We may think of Joshua's credences as assigning real numbers to atomic events $[\pm H \;\&\; ch(\pm H) = x]$, where $\pm H$ might be H or $\sim H$ and where $[ch(\pm H) = x]$ says that the coin's objective chance of landing $\pm H$ is $x \in [0, 1]$.[13] Let's suppose further that Joshua knows that the coin's bias toward heads is 0.2, so that $\boldsymbol{b}(ch(H) = 0.2) = 1$, and that this is all the relevant evidence he has about the coin. According to the accuracy-centered approach, Joshua should use his evidence to find a credal pair $\langle h, t\rangle$ that strikes the best attainable balance between accuracy in the event of heads and accuracy in the event of tails. This forces him to undertake a kind of epistemic cost-benefit analysis in which the costs of holding $\langle h, t\rangle$ are given by $\boldsymbol{I}(\langle h, t\rangle, \langle 1, 0\rangle)$ when H is true and by $\boldsymbol{I}(\langle h, t\rangle, \langle 0, 1\rangle)$ when H is false. On the Brier score, these penalties work out to $½\cdot[(1-h)^2 + t^2]$ and $½\cdot[h^2 + (1-t)^2]$, respectively. The tradeoffs are clear: higher h-values lower the first cost but raise the second, while higher t-values raise the first cost but lower the second. Which credences offer just the right mix of epistemic risk and reward? **PP** provides

[12] The argument generalizes to credences defined over arbitrary finite partitions $\langle X_1, X_2, \ldots, X_N\rangle$, where each X_n is logically consistent, $(X_1 \vee X_2 \vee \ldots \vee X_n)$ is a logical truth, and $X_j \;\&\; X_n$ is a contradiction for each $j, n \le N$.

[13] Caution: We do *not* assume that $[ch(H) = x]$ and $[ch(\sim H) = 1 - x]$ are the same event, e.g., we do not identify a 1-to-4 (20%) bias toward heads with a 4-to-1 (80%) bias toward tails. This matters lot since the Easwaran/Fitelson argument only makes sense if these events are distinct.

a natural answer. It mandates $h = 0.2$ as the right credence for someone who knows $ch(H) = 0.2$. But, is this advice consistent with the accuracy-centered picture?

Easwaran and Fitelson say no. There is, they claim, a *general* conflict between **AN** and **PP**, a conflict that does not depend on what scoring rule is used or on any aspect of the accuracy-centered approach other than its commitment to **AN**. If they are right, then anyone who endorses **PP** as a norm of evidence (i.e., anyone who thinks it characterizes a part of an epistemic duty to hold well-justified credences) must repudiate **AN**, and with it any hope of an accuracy-centered epistemology for credences.

Easwaran and Fitelson reject **AN** on the grounds that (i) ***b***'s dominance of ***c*** only reflects badly on ***c*** only if is ***b*** is an *available* credal state, and (ii) a believer's evidence might make ***b*** unavailable. They write:

> "Joyce's argument tacitly presupposes that – for any incoherent agent *S* with credence function ***c*** – some (coherent) functions ***b*** that dominate ***c*** are always 'available' as 'permissible alternative credences' for *S*. But, there are various reasons why this may not be the case. The agent could have good reasons for adopting (or sticking with) *some* of their credences. And, if they do, then the fact that some accuracy-dominating (coherent) functions ***b*** 'exist' (in an abstract mathematical sense) may not be epistemologically probative."

Easwaran and Fitelson say surprisingly little about what it means for credal states to be available or unavailable as permissible alternatives.[14] This is unfortunate since, as we shall shortly see, their argument founders on an equivocation about the meaning of this central notion.

Easwaran and Fitelson contend that the combination of **Accuracy Non-dominance** and the **Principal Principle** leads to problematic "order-effects" in which serial application of **AN** then **PP** sanctions one set of credences while serial application of **PP** then **AN** sanctions another. To make their case, they read **PP** as a rule that makes any credal state with $\boldsymbol{b}(X) \neq x$ unavailable to epistemically rational believers who are certain that $ch(X) = x$. When **PP** is construed this way, 'order effects' do indeed arise. Here is an example (developed on the inessential assumption that inaccuracy is measured by is the Brier score):

> Joshua, who knows nothing about a coin except that $ch(H) = 0.2$, wants to obey **PP** by aligning his credences with the known chances, but also hopes to avoid accuracy-domination. To figure out which credences he may permissibly adopt, he might proceed in one of two ways:

[14] They do say (p. 430) that they are, "concerned with *evidential* reasons why [credences] may be unavailable to an agent," and add that "there may also be *psychological* reasons why some [credences] may be unavailable, but we are bracketing that possibility here."

> *Accuracy-then-Evidence*. Every credal state starts out as available. Joshua first satisfies **AN** by ruling out all ⟨*h*, *t*⟩ pairs that are accuracy dominated by any available pair. This leaves the coherent pairs ⟨*h*, 1 − *h*⟩ with $0 \le h \le 1$ as the only live options. Joshua can then apply **PP** to rule out every remaining pair except the one with *h* =0.2. So, when Joshua knows (only) that *H*'s objective chance is 0.2, *Accuracy-then-Evidence* says that ⟨0.2, 0.8⟩ is his only permissible credal state.
>
> *Evidence-then-Accuracy*. Here Joshua first invokes **PP** to rule out all ⟨*h*, *t*⟩ pairs with $h \neq 0.2$, leaving only pairs of the form ⟨0.2, *t*⟩ available as permissible credal states. But, since none of these pairs dominate any other relative to the Brier score, none is dominated by a *still available* credal state. So, when Joshua knows (only) that *H*'s objective chance is 0.2, *Evidence-then-Accuracy* says that the permissible credal states are the coherent pair ⟨0.2, 0.8⟩ and all incoherent pairs ⟨0.2, *t*⟩ with $0 \le t \le 1$.

This disparity between what is permitted by Accuracy-then-Evidence and by Evidence-then-Accuracy allegedly indicates "a conflict between *evidential* norms for credences and a certain (accuracy dominance) coherence norm for credences." (p. 430)

This argument hinges crucially on the claim that credal states made 'unavailable' by an application of **PP** may *not* be invoked in subsequent applications of **AN**. For example, the fact that ⟨0.25, 0.75⟩ dominates ⟨0.2, 0.7⟩[15] does not reflect badly on the latter credences in *Evidence-the-Accuracy* because **PP** has already made the former credences unavailable at the point when **AN** gets applied. Unfortunately, the idea that credal states made unavailable by **PP** may not be invoked in subsequent applications of **AN** is based on an equivocation 'unavailable'. As the next section shows, the term must mean one thing for **AN** to be true and another for **PP** to be true.

4. The 'Availability' Equivocation

It is surely true that accuracy-dominance only counts against a credal state when the dominating alternative is, in some sense, *available* for adoption. If physical or psychological limitations, lying beyond the agent's control, prevent her from holding the dominating credences *even if she thinks it advisable to do so*, then Easwaran and Fitelson's are entirely right that the dominating alternative's mere 'abstract' existence does nothing to make the dominated credences impermissible. But, norms like **PP** do *not* make credences unavailable in this strong way. When Joshua invokes **PP** to reject ⟨0.75, 0.25⟩ he does not erect some impenetrable psychological or

[15] This assumes the Brier score. Let $\boldsymbol{b}(H) = 0.2$ and $\boldsymbol{b}(\sim H) = 0.7$ and $\boldsymbol{c}(H) = 0.25$ and $\boldsymbol{c}(\sim H) = 0.75$. Then $\boldsymbol{I}(\boldsymbol{b}, \omega_0) = 0.065 > \boldsymbol{I}(\boldsymbol{c}, \omega_0) = 0.0625$ and $\boldsymbol{I}(\boldsymbol{b}, \omega_1) = 0.565 > \boldsymbol{I}(\boldsymbol{c}, \omega_1) = 0.5625$.

physical barrier that prevents him from holding those credences. On the contrary, he continues to see them as credences that he *could* hold *if he thought it wise to do so*. Adherence to **PP** leads Joshua to regard the adoption of ⟨0.75, 0.25⟩ as a *mistake*, not an *impossibility*. This distinction gets slurred over in Easwaran and Fitelson's "available as permissible alternative credences" phrasing which suggests that being impermissible, in the sense of contravening requirements of epistemic rationality, is something like being unavailable, in the sense of being a state that the agent could not adopt even if she thought that doing so was a good idea. To keep the distinction straight let's use the term *inaccessible* for credal states that a believer feels he would be unable to adopt even if, in light of his evidence, he deemed them to be among his best epistemic options. In contrast, a (merely) *impermissible* state is one the believer feels that he *could* adopt, but will not adopt because he does not rank it among his best epistemic options given his evidence.[16]

Let's note two things about this distinction. First, while accuracy-dominance may not reflect badly on a credal state when the dominating alternative is genuinely inaccessible, it *does* reflect badly on it when that alternative is merely impermissible. For if the dominant state is accessible yet impermissible then the dominated state is inferior all-epistemic-things-considered to a state that the believer thinks she could occupy if she saw it as her best option. Given that both states can be adopted, the fact that the dominant state looks bad makes the dominated state look even worse! So, on an accuracy-centered picture, domination by an *accessible* alternative is always a defect. It matters not a whit whether or not that alternative is itself permissible – what matters is that it is a superior system of credences that the believer could adopt *if her evidence warranted doing so*. The upshot is that **AN** is true when 'available' means 'accessible', but false when it means 'impermissible'.

The second point is that evidential norms do *not* make credal states inaccessible merely by ruling them impermissible. This is true of norms generally. A norm – be it practical, moral, social, epistemic or cultural – that prohibits some act or state does *not* thereby make that act or state inaccessible. We introduce norms only when we think it is possible to contravene them. (This is why, e.g., it would be superfluous and silly to introduce statutes to outlaw the creation of zombies by reanimation of the dead: we have no reason to prohibit such actions since we do not believe that anyone can actually perform them.) Additionally, we do not think that those who endorse a norm lose the ability to violate it. I know I should not lie, gossip, be easily angered, or eat more than recommended for my daily diet, but it is, alas, all too easy for me to do what is prohibited by the norms I endorse.

---

[16] When Easwaran and Fitelson give examples of unavailable alternatives they cite options that are clearly inaccessible. For example, in discussing the evaluation of practical alternatives, they write that "there is always some formally defined alternative that would be better – rather than betting a dollar at even odds on the outcome of a coin flip, I should choose the action that pays me a million dollars regardless of how the coin comes up! But this is no criticism of my action, or my utility function, since the alternative that is better is one that is not *available* to me." This is clearly a case of inaccessibility.

These points transfer straightforwardly to Joshua's situation. If Joshua were, for reasons beyond his control, *unable* invest credences of 0.25 and 0.75 in *H* and ~*H* even if he saw these as the best beliefs to adopt in light of his evidence, then their dominance of ⟨0.2, 0.7⟩ would indeed be an 'abstract mathematical' curiosity of no real consequence. But, when **PP** forbids the ⟨0.75, 0.25⟩ credences for Joshua it does not erect any barrier that prevents him from adopting those credences. **PP**, in other words, is false if it is interpreted as making credences inaccessible rather than merely impermissible. So, when Joshua uses **PP** to rule out ⟨0.25, 0.75⟩ he continues to see it as a credal state he *could* occupy, even though, in light his evidence, he does not think it is a state he *should* occupy. In short, Joshua's use of **PP** when he knows $ch(H) = 0.2$ has the effect of making ⟨0.75, 0.25⟩ impermissible, not inaccessible.

Once we understand this, it becomes clear that Easwaran and Fitelson's order effects will do not arise as long as each of **AN** and **PP** is interpreted in the way that makes it true. **PP** rules the credences ⟨0.75, 0.25⟩ impermissible but leaves them accessible, and **AN** rules ⟨0.7, 0.2⟩ impermissible because it is dominated by an accessible alternative. The order in which the norms are invoked is immaterial. Whether Joshua uses Accuracy-then-Evidence or Evidence-then-Accuracy he will end up with the same set of permissible credences: viz., {⟨0.2. 0.8⟩}. Easwaran and Fitelson see "order effects" here only because they conflate situations in which credences are rendered impermissible by evidential norms with situations in which they are made inaccessible by external contingencies. It is a general feature of evidential norms, however, that they forbid *without* foreclosing: they tell us what we *should or should not* believe, in light of our evidence, not what we *can* or *cannot* believe. It can be hard to keep this straight because it is so easy to slip into the habit of characterizing norms like **PP** by saying that they make certain belief states *impossible for an epistemically rational believer*. This makes it sound as if the states are impossible *per se* for a rational believer, but they remain possible – it is just that anyone who adopted them would not be counted as rational.[17] This may be what leads Easwaran and Fitelson into trouble. But whatever the cause, the fact is that, contrary to what they suppose, credal states deemed impermissible by evidential norms like **PP** (and not made inaccessible by independent external limitations) *can* be invoked in **AN** to show that other states are impermissible.

5. The Real Issue: Are There Conflicts Between Accuracy and Justification?

While the forging remarks show the flaw in Easwaran and Fitelson's reasoning, readers might not feel that the itch has been fully scratched. The problem of 'order effects' seems like a sideshow anyhow. The real issue is that an accuracy-centered epistemology is committed to the thesis that accuracy dominated credal states are inferior *all-epistemic-things-considered* to the states that dominates them, and it seems like this commitment might conflict with **PP** or other

---

[17] Compare: A devout Roman Catholic must defer to the Pope's teachings on matters of faith and morals. So, Jane, a devout Roman Catholic, lacks the power to reject the Pope's teachings. No! Jane is entirely free to reject them, though she would not count as a devout Roman Catholic if she did.

widely accepted evidential norms. It is, after all, a part of the accuracy-centered position that no matter how decisively the evidence might favor ***c***, this can *never* offset a dominant ***b***'s advantage in accuracy. This seems troubling. Accuracy considerations and evidential considerations seem like different sorts of beasts, and what guarantee do we have that they will 'play nicely' with one another in epistemology? For all we know, there might be some argument, other than the one Easwaran and Fitelson attempt, which proves that **PP**, or another legitimate norm of evidence, really does *conflict* with **AN**. If that happens, i.e., if evidential considerations point one way while considerations of accuracy point the other, why should accuracy prevail?

I suspect that this is the real source of Easwaran and Fitelson's concerns. When they stand back and describe their 'evidentialist worry' in general terms, they do not talk of order effects or availability. They focus, instead, on the possibility of direct conflicts between evidence norms and accuracy norms. In a revealing passage (pp. 430-431) they write that their worries "remain pressing, provided only that the following sorts of cases are possible" (lightly rewritten):

(a) Agent *S* has an incoherent credence function ***c***,
(b) ***c***(*X*) falls in the interval [*a*, *b*],
(c) *S* knows that epistemic rationality requires ***c***(*X*) to be in [*a*, *b*],
(d) but all credence functions ***b*** that dominate ***c*** place ***b***(*X*) outside of [*a*, *b*].

They go on to say, "to avoid our worry completely, one would need to argue that no examples satisfying (a)-(d) are possible. And, that is a tall order. Surely, we can imagine that an oracle concerning epistemic rationality has informed *S* that [the right credence for *X*] is in [*a*, *b*] – despite the fact that all (coherent) dominating functions ***b*** are such that ***b***(*X*) is not in [*a*, *b*]".

Easwaran and Fitelson are right to think that cases like (a)-(d) are possible, but wrong to think they pose problems for accuracy-centered epistemology. They would be big trouble if they entailed genuine conflicts in which evidential considerations forced believers to hold credences forbidden by the accuracy approach, but (a)-(d) do not entail any such thing. I suspect Easwaran and Fitelson think they do because they see (a)-(d) as describing a case in which epistemic rationality requires ***c***(*X*) to be in [*a*, *b*] while the accuracy framework requires it to be outside that interval. But, to extract the conclusion ***c***(*X*) ∉ [*a*, *b*] from (a)-(d) we need the further premise:

(e) The accuracy-centered approach recommends that *S* adopt a credence that dominates ***c***.

But, if (e) is false then no conflict need arise since it might (and will) turn out that, anytime (c) holds, the undominated credences that strike the best overall balance between the benefits of being confident /doubtful of truths/falsehoods and the risks of being confident/doubtful of falsehoods/truths will place *X*'s credence in [*a*, *b*]. Such a credal state would not dominate ***c***, of course, but proponents of the accuracy-centered approach will see it as being superior to ***c***, all-epistemic-things-considered.

And, (e) is definitely false! It is no part of accuracy-centered epistemology that believers with dominated credences should adopt a dominating alternative. For example, Joshua is *not* obliged to adopt a credence $\langle h, 1 - h\rangle$ with $0.24834 < h < 0.25495$ just because these dominate his own credences of ⟨0.2, 0.7⟩. It can seem plausible that the accuracy-centered view imposes this obligation on Joshua, and sanctions (e), because it is so tempting to read the statement '***b*** dominates ***c***' as a *recommendation* of ***b***. But, dominance arguments do *not* work this way. When we learn that one option dominates another we acquire a (conclusive) reason for rejecting the dominated option without acquiring any complementary reason for adopting the dominating one. By pointing out that ***b*** dominates ***c*** we denigrate ***c*** *without* commending ***b***. We affirm that ***b*** is *better* than ***c***, of course, but do not imply that ***b*** is best or even very good at all. I do not praise Franco when I say that Hitler was worse along every dimension of dictatorial evil. I am not recommending Northern Manitoba's climate when I tell you that the weather in Churchill is better than the weather in Vostok in every season. Likewise, when I point out that Joshua's credences are accuracy-dominated by ⟨0.25, 0.75⟩ I do not imply that he should adopt the latter beliefs. In fact, as we will see below, I should be quite certain that a person who knows what Joshua does about the chances should *not* adopt *any* of the credences that dominate his own. What he should do, instead, is to reflect more carefully on his total evidence with the goal of finding a credal state that strikes the optimal balance between the good of being confident in truths and the evil of doubting them, it being understood that this optimal state might well *not* be found among the dominating credences. When he does he will see that the optimal credences are ⟨0.2, 0.8⟩. So, proponents of accuracy-centered epistemology have nothing to fear from (a)-(d). Such cases pose no problem as long as we keep in mind that learning that a credal state is dominated shows that the dominated state is impermissible without implying that the dominating state is in any way permissible.

Now, one might object that I *do* recommend ***b***, at least a little bit, when I assert that ***b*** dominates ***c***. I imply, at least, that ***b*** is the best alternative in any context where it and ***c*** are the *only* accessible alternatives. If you are being banished to Churchill or Vostok, then I surely do mean to recommend Churchill when I say that its weather dominates Vostok's. Likewise, if Joshua is (for odd psychological reasons) is only capable of adopting one of the two credal states ⟨0.2, 0.7⟩ or ⟨0.25, 0.75⟩ then the accuracy-centered approach is committed to saying that he *should* adopt the latter credences and so violate the Principal Principle. Isn't this enough, all by itself, to show that there is a conflict between the accuracy norm **AN** and the evidence norm **PP**?

To see why this is a non-issue and, more generally, why conflicts between **AN** and legitimate evidential norms, like **PP**, can never arise, let's think about how Joshua might try to show that ⟨0.2, 0.7⟩ is superior to ⟨0.25, 0.75⟩. Appealing to **PP**, he might argue that 0.2 is better *justified* than 0.25 as credence for H since the former is closer to (indeed identical to) the known objective chance. Proponents of accuracy centered-epistemology will agree, and will even offer a (partial) analysis of justification that bears out Joshua's intuition. Suppose, temporarily, that objective chances are known to be probabilities (so that $ch(\sim H) = 1 - x$ when $ch(H) = x$), and that an appropriate accuracy score $\boldsymbol{I}$ has been identified. We can then define the *objective expected accuracy* of the credence $\boldsymbol{b}(H) = p$ when H's chance is known to be x as $E(\boldsymbol{I}(p)|ch(H) = x) = x\cdot\boldsymbol{I}(p, 1) + (1 - x)\cdot\boldsymbol{I}(p, 0)$, where $\boldsymbol{I}(p, 1)$ is the assignment's accuracy when H is true and $\boldsymbol{I}(p, 0)$ is its accuracy when H is false. The proposed theory of justification is this:

> **Justification by Chance (JBC)**. For a believer whose only relevant information about H's truth-value is $ch(H) = x$, the credence $\boldsymbol{b}(H) = h$ is better justified than the credence $\boldsymbol{b}(H) = h^*$ if and only if the *objective expected inaccuracy* of the second assignment exceeds that of the first.

To get the idea, imagine that one must settle on the same credence for each of a large series of independent events that are all known to have objective chance x (e.g., tosses of a coin of fixed bias). In this context, **JBC** says that the best justified credence is the one that produces the least total inaccuracy when frequencies align the known chances.[18] Likewise, **JBC** ranks $\boldsymbol{b}(H) = h$ as better justified than $\boldsymbol{b}(H) = h^*$ exactly if it is objectively likely that the h-assignment will produce less inaccuracy than the h^*-assignment over an indefinitely long run of trials.

This picture of justification dovetails nicely with the idea that all epistemic duties involve the rational pursuit of doxastic accuracy. In any context where objective chances are known to satisfy the laws of probability, proponents of accuracy-centered approaches will see the following as comprising an essential part of the duty to rationally pursue accuracy:

> **Accuracy by Chance** (**ABC**). An epistemically rational believer who knows that H's objective chance is x, and who has no other relevant evidence about H's truth-value, will see credences for H with higher/lower objective expected accuracies as striking better/worse balances between accuracy in the event of H and accuracy in the event of $\sim H$.

The upshot is that the rational pursuit of accuracy as detailed in **ABC** requires believers to hold credences that are well-justified by the lights of **JBC**. The two duties – to have a well-justified credence, and to have a credence that strikes the best overall accuracy balance –never clash. Moreover, when conjoined with **Strict Propriety**, **JBC** entails that a believer whose only relevant evidence about a proposition is its objective chance does best, justification-wise, by setting her credence for that proposition equal to the known objective chance. In this way, the

[18] One would anticipate this happening, with probability approaching one, as the series grows.

combination of **JBC** and **ABC** provides an accuracy-based rationale for **PP**![19] So, why should we embrace **PP**? It's not because there is anything especially virtuous, *per se*, about having credences that agree with known chances. It's because doing optimally balances the epistemic good of being confident in truths and the epistemic evil of being confident in falsehoods (but see below for caveats). You should use **PP** to regulate your credences because it's part of what is involved in the rational pursuit of accuracy! So, proponents of accuracy-centered epistemology will happily concede that Joshua's 0.2 credence for H is *perfectly* justified in light his knowledge of the chances, and that the h-values of any $\langle h, t\rangle$ pairs that dominate his credences are less well justified because they are farther away from the known chance value.

This would be the end of the story (and a bad end for the accuracy-centered view) if the justificatory impact of the data $ch(H) = 0.2$ were confined to its impact on Joshua's credence for H. However, since the logic of negation ensures that evidence for/against H is also evidence against/for $\sim H$, we cannot fully assess the degree to which Joshua's credences are justified until we consider the evidence's impact on his credence for $\sim H$. But, if Joshua knows that chances are probabilities he will know $ch(\sim H) = 0.8$, and so recognize that (by both his own criterion and **JBC**) the 0.75 credence for $\sim H$ is better justified than the credence 0.7. Since coordinate-wise comparison does not yield a uniform verdict (as it would if $\langle 0.2, 0.7\rangle$ were compared to $\langle 0.1, 0.6\rangle$), we need to figure out how Joshua's justification for the *pair* $\langle 0.2, 0.7\rangle$ compares with his justification for other pairs, like $\langle 0.25, 0.75\rangle$.

To make this determination, proponents of accuracy-centered epistemology will again invoke considerations of objective expected accuracy. The basic principle (still assuming that chances satisfy the laws of probability and that an adequate epistemic accuracy score has been identified) is this:

> **JBC (General)**.[20] Credal state $\boldsymbol{b}$ is better justified than credal state $\boldsymbol{b}^*$ in light of evidence about the objective chances (with no 'inadmissible' data) when the *objective expected inaccuracy* of the second assignment determinately exceeds that of the first. In particular, when $ch(H) = x$ is the only thing known, the objective expected accuracy of a pair $\langle p, q\rangle$ is given by
>
> $$E(\boldsymbol{I}(p, q)|x) = x\cdot\boldsymbol{I}(\langle p, q\rangle, \langle 1, 0\rangle) + (1 - x)\cdot\boldsymbol{I}(\langle p, q\rangle, \langle 0, 1\rangle),$$
>
> and $\langle h, t\rangle$ is better justified than $\langle h^*, t^*\rangle$ when $E(\boldsymbol{I}(\langle h, t\rangle)|x) < E(\boldsymbol{I}(\langle h^*, t^*\rangle)|x)$.

[19] This argument is similar to, and inspired by, one offered in Pettigrew (forthcoming). See, in particular, Pettigrew's Theorem 3.

[20] Caveats: (A) This is only a *sufficient* condition. (B) $\boldsymbol{b}$'s objective expected inaccuracy *determinately* exceeds $\boldsymbol{b}^*$'s just when the evidence is sufficiently informative to limit the possible chance functions to those that yield a higher expected inaccuracy for $\boldsymbol{b}$ than for $\boldsymbol{b}^*$.

There is a similarly generalized version of **ABC**:

> **ABC (General)**. An epistemically rational believer who knows that *H*'s objective chance is *x*, and who has no other relevant evidence about *H*'s truth-value, will see ⟨*h*, *t*⟩ pairs with higher/lower objective expected accuracies as striking better/worse balances between accuracy in the event of *H* and accuracy in the event of ~*H*.

It is then automatic that ⟨*h*, *t*⟩ is better justified than ⟨*h**, *t**⟩ if the former accuracy-dominates the latter. Moreover, *this will be true no matter what evidence about the chances a believer might have*! As in the case of ⟨0.2, 0.7⟩ and ⟨0.25, 0.75⟩, if the *h*-component of the dominated pair is better justified than the *h*-component of the dominant pair, this deficient will be more than offset by the dominant pair's justificatory advantage in the *t*-component.[21] So, even if Joshua were somehow only able to adopt one of these two credal states (⟨0.2, 0.8⟩ being inaccessible for some reason) there would still be no friction between **AN** and **PP**: the accuracy dominant pair ⟨0.25, 0.75⟩ is also the pair that is best justified in light of the total evidence. Note how the accuracy score, which assess *entire* credal states, is used to balance off the justificatory merits and defects of the various individual credences to produce an *aggregate* assessment. Easwaran and Fitelson obscure this point by framing their objections in terms of the impact of evidence on a *single* credence, like Joshua's credence of 0.2 to *H*, and this leads them to ignore its impact on other credences. For instance, when they wonder what will happen if epistemic rationality requires $\boldsymbol{c}(X) \in [a, b]$ they never notice that certain things will be required of $\boldsymbol{c}(\sim X)$ too, and that, as a result, the pair $\langle \boldsymbol{c}(X), \boldsymbol{c}(\sim X)\rangle$ might end up being *less* well justified in the aggregate than $\langle \boldsymbol{b}(X), \boldsymbol{b}(\sim X)\rangle$ even if $\boldsymbol{c}(X)$ is better justified than $\boldsymbol{b}(X)$, which is exactly what happens if **JBC** is correct.

Proponents of accuracy-centered epistemology will want to generalize **JBC** beyond evidence about objective chances so that *all* facts about justification are interpreted as facts about the rational pursuit of doxastic accuracy. This would make it a truism (on the level of 'evidence for *X* is evidence for *X*'s truth') that believers are justified in holding credences exactly to the extent that their evidence makes it reasonable for them to expect those credences to be accurate (in the aggregate). This picture of the relationship between justification and accuracy has a number of consequences *in re* dominance:

- Evidence that justifies a credal state ***b*** always provides even stronger justification for any state that accuracy-dominates ***b***.

[21] For Brier, we get (**a**) $E(\boldsymbol{I}(0.2)|ch = 0.2) = 0.08 < 0.08125 = E(\boldsymbol{I}(0.25)|\ ch = 0.2)$, (**b**) $E(\boldsymbol{I}(0.7)|\ ch = 0.8) = 0.085 > 0.08125 = E(\boldsymbol{I}(0.75)|\ ch = 0.8)$, and thus (**c**) $E(\boldsymbol{I}(0.25, 0.75)|\langle x, 1 - x\rangle) = 0.1625 < 0.165 = E(\boldsymbol{I}(0.2, 0.7)|\langle x, 1 - x\rangle)$. So, what ⟨0.25, 0.75⟩ loses in the first expected accuracy comparison it more than makes up in the second.

- Evidence that tells against ***b*** always tells even more strongly against any credal state that ***b*** dominates.

These points distill the core tenets of a theory of justification in which doxastic accuracy is the cardinal epistemic virtue, its pursuit is the fundamental epistemic duty, and in which accuracy-dominated credal states are inferior *all-epistemic-things-considered* to the (accessible) states that dominate them. They also answer the question of what would happen if the evidence were to favor ***c*** over ***b*** when ***b*** dominates ***c***, thereby generating a conflict between norms of evidence and norms of accuracy. These principles tell us that no such conflict will ever arise (as long as chances are probabilities and an acceptable accuracy score has been identified) because all legitimate norms of evidence are ultimately answerable to norms of accuracy.

This last point is worth emphasizing. On the account of justification sketched here, rules of evidence have no *independent* normative status. They are ancillary norms that regulate beliefs for the purpose of achieving doxastic accuracy. If a putative rule of evidence ever recommends accuracy-dominated credences we can safely repudiate it since it is not doing its job. Consider the Principal Principle. Within an accuracy-centered framework there is nothing admirable *per se* about holding credences that align with objective chances: such alignment is merely a means to the end of achieving high objective expected accuracy. **PP**'s status is entirely derived from its ability to recommend credences that rank among the best all-epistemic-things-considered, where it is understood that the optimal credences all-epistemic-things-considered are those that have the highest objective expected accuracy. Indeed, as we have seen, **PP** can be justified as a legitimate norm of evidence within the accuracy-based framework (as long as chances are probabilities) because it can be shown that followings its recommendations leads believers to hold credences that maximize objective expected accuracy.

Absent such an accuracy-based rationale there would be no reason for believers to defer to **PP** when settling on credences. To see why, consider a case in which the accuracy-based framework would repudiate **PP**. Suppose the objective chances are revealed to Joshua by an infallible 'oracle', like the one to which Easwaran and Fitelson allude. Let's call her Julika. When Joshua asks Julika for *H*'s chance he is told '0.2', and, invoking **PP**, he invests credence 0.2 in *H*. So far so good, but he still needs to fix a credence for ~*H*. He could, of course, settle on 0.8 since ⟨0.2, 0.8⟩ is the unique undominated credal pair whose *h*-component is 0.2. Instead of taking this option, however, suppose that Joshua asks Julika to reveal ~*H*'s chance directly, and she tells him '0.7'. He learns, to his shock, that the chances of *H* and ~*H* do not sum to one! Now **PP** really does contradict **AN**. Should Joshua stick with ⟨0.2, 0.7⟩ because **PP** tells him to or should he look for the undominated pair that is best justified in light of his odd evidence? He should do the latter! The reason is simple: **PP** should have normative standing for Joshua only to the extent that it helps him find an optimal credal state, all-epistemic-things-considered. Since no dominated state can have this feature, Joshua cannot both defer to **PP** and discharge his duty to rationally pursue accuracy. **PP** has to go. On the accuracy-centered approach, Joshua should recognize **PP** as a legitimate norm of evidence only if chances are probabilities (in which case he will adopt ⟨0.2, 0.8⟩, the credal pair with the highest objective expected accuracy). If chances are not probabilities (an outlandish notion),[22] then it would be an mistake for Joshua to defer to them because doing so would lead him to pay an unnecessary cost in accuracy. This point is general. Any 'oracle' who recommends probabilistically incoherent credences must be ignored since following its advice is inconsistent with the duty to rationally pursue doxastic accuracy.[23]

6. Accuracy and Epistemic Value: The Choice of *I*

We have now reached a delicate point in the dialectic. We have seen that, *once an accuracy score **I** that satisfies the four requirements laid down in* §1 *has been endorsed*, the norm of ***I*** non-dominance can never conflict with any legitimate evidential norm. But, this is because all epistemic norms recognized as *legitimate* by the accuracy-centered framework will entail that evidence which favors a credal state always favors any ***I***-dominant state even more strongly. We have also seen than the key evidential norm, **PP**, is legitimate relative to *any* accuracy score t (as long as chances are probabilities). Other norms might be legitimized in similar manner, though one might expect that the status of many norms would depend on the choice of an accuracy score. It's all a cozy picture.

A bit too cozy, perhaps. The whole house of cards depends on our endorsement of some score ***I*** as the right measure of doxastic inaccuracy. But, the choice of such a score is closely interwoven with hard questions about when credences are and are not justified. From a certain perspective, the endeavor seems circular. On the accuracy-centered picture, part of what it *is* to agree that the Brier score, say, is the right gauge of inaccuracy is to think ⟨0.25, 0.75⟩ is better justified than ⟨0.2, 0.7⟩ given any evidence. This might be palatable if there were free-standing, independent standards for identifying the ideal accuracy score for use in any context. In the absence of such a criterion, however, the various success listed above – the lack of conflicts between legitimate norms of accuracy and norms of evidence, and the rationale for **PP** as a recipe for minimizing objective expected inaccuracy – look to have been secured by an *ad hoc* choice of an 'accuracy' score that was motivated by the desire to secure these very successes.

---

[22] I am entertaining this possibility purely as devil's advocate. I do not see *any* plausibility to the idea that chances are not probabilities. Hypotheses about chances play two primary roles in our epistemic lives: (i) they are used to explain stable frequencies in large sets of independent trials; (ii) they are, in turn, confirmed by facts about the frequencies observed in such trials. Given the probabilistic structure of relative frequencies, it is hard to imagine anything but a probability playing either role.

[23] More generally, a mapping ***Q*** of propositions to real numbers will not be treated as an *epistemic expert* by a rational believer unless the believer is certain that ***Q*** obeys the laws of probability. Here, the believer regards ***Q*** as an epistemic expert just when her credences satisfy $\boldsymbol{b}(X|\boldsymbol{Q}(X) = x) = x$ for every X. To put is another way, rational believers will never defer to experts who recommend dominated credences.

To put a point on it, notice that there are plausible seeming ways of measuring 'closeness to truth' relative to which ⟨0.25, 0.75⟩ does *not* dominate ⟨0.2, 0.7⟩. Consider the *absolute-value score,* which sets $\boldsymbol{I}(\langle h, t\rangle, \langle 1, 0\rangle) = 1 - (h - t)$ and $\boldsymbol{I}(\langle h, t\rangle, \langle 0, 1\rangle) = 1 + (h - t)$. As is easy to see, ⟨0.2, 0.7⟩ and ⟨0.25, 0.75⟩ have the same absolute-value score whether *H* is true or false, which also means that their objective expected accuracies will coincide for any (probabilistic) chance distribution. So, if we developed the accuracy-centered framework using the absolute-value score, we would have to say that there is no epistemic difference between ⟨0.2, 0.7⟩ and ⟨0.25, 0.75⟩ relative to *any* body of data. Since this includes the data [$ch(H) = 0.25$ & $ch(\sim H) = 0.75$] we no longer have any rationale for **PP**: the absolute-value score makes it permissible to adopt ⟨0.2, 0.7⟩ as one's credences even when one knows that ⟨0.25, 0.75⟩ agrees *perfectly* with the chances. **PP** becomes optional.

There are even scores relative to which ⟨0.2, 0.7⟩ strictly dominates ⟨0.25, 0.75⟩. One is the *square-root score*: $\boldsymbol{I}(\langle h, y\rangle, \langle 1, 0\rangle) = ½\cdot[(1 - h)^{½} + y^{½}]$ and $\boldsymbol{I}(\langle h, y\rangle, \langle 0, 1\rangle) = ½\cdot[h^{½} + (1 - y)^{½}]$. It is easy to show that ⟨0.2, 0.7⟩ gets a better square-root score than ⟨0.25, 0.75⟩ whether *H* is true or false. So, if we developed an accuracy-centered epistemology based on this score, we would have to definitively prohibit believers from using **PP** at all. It is no longer even permissible to align one's credences with the known chances.

These examples make it clear that the success of accuracy-centered epistemology hinges crucially on the exclusion of certain scoring rules. If scores on which ⟨0.2, 0.7⟩ dominates ⟨0.25, 0.75⟩ are allowed, then the cozy relationship between accuracy and evidence breaks down. Now, it turns out that the accuracy-centered approach will not allow either of the scores just discussed because both violate **Strict Propriety**. Yet, it would be futile to argue for their exclusion on this basis, since the challenge would then be to justify **Strict Propriety**, and all the same issues will reemerge. The question, in most general terms, is this: Is there any compelling reason to think that the right epistemic accuracy score (for use in a given context) will recognize **PP**, and other familiar norms of evidence, as legitimate, so that the credences they recommend or permit are never accuracy-dominated?

Easwaran and Fitelson introduce a version of this worry in the last two paragraphs of their paper, writing that:

> One might think that violation of the Principal Principle doesn't make a credence function unavailable, but instead just represents some dimension of epistemic 'badness'. If this badness is different from the badness of inaccuracy, then it becomes clear that Joyce's arguments need to be modified – even if ***b*** dominates ***c*** with respect to inaccuracy, if ***c*** has less *overall* epistemic badness, then ***c*** may still be perfectly acceptable as a credence function. Thus, Joyce's arguments would need to consider overall badness rather than just inaccuracy.

The only way to save Joyce's arguments here seems to be to say that somehow the badness of violating the Principal Principle is *already included* when one has evaluated the accuracy of a credence function. Perhaps there is some way to argue for this claim. But this claim needs more support than it has been given. And nothing here turns on the use of the Principal Principle in particular – if there can be any epistemic norm whose force is separate from accuracy, then the same sort of problem will arise. Joyce's argument works only if *all* epistemic norms spring from accuracy." (pp. 432-433)

There is much right in this passage. Easwaran and Fitelson seem to recognize that their 'unavailability' worry might be resisted, and they rightly focus attention on the issue of potential conflicts between accuracy norms and evidence norms. They also are right that **AN** loses its bite if $\boldsymbol{b}$ can accuracy-dominate $\boldsymbol{c}$ when $\boldsymbol{c}$ is superior to $\boldsymbol{b}$ all-epistemic-things-considered (i.e., "$\boldsymbol{c}$ has less *overall* epistemic badness"). They even recognize that the solution is to show that norms of evidence are 'already included' in accuracy scores.

It is misleading, however, to claim that "Joyce's argument works only if *all* epistemic norms spring from accuracy" since the "spring from" locution suggests a hierarchical picture in which all legitimate epistemic norms are deduced from independently established principles that define doxastic accuracy and govern its rational pursuit. The relationship between epistemic norms and accuracy norms, however, is not hierarchical, but symbiotic. While it is true that, in a fully-articulated accuracy-based epistemology, all norms of evidence will be underwritten by rationales which show how they contribute to the rational pursuit of accuracy, this will not be because there is some free-standing theory of doxastic accuracy from which these norms can be derived. Rather than being autonomous, our concept of accuracy will be informed by, and highly dependent on our considered views about which epistemic norms are legitimate. Indeed, it is essential to the accuracy-centered picture that evidential considerations should factor into the choice of an inaccuracy score. These scores are, at bottom, ways of measuring 'closeness to the truth' that reflect our views about how such closeness should valued. Different scores will encourage different epistemic practices, and part of our goal in choosing among them will be to promote practices that promote our epistemic values. A few examples should make the point.

Consider first a streamlined version of an argument used in Joyce (2009) to dismiss the absolute-value score. Suppose you are about to toss a three sided die that you *know* to be fair. **PP** has you set $\boldsymbol{b}(side_1) = \boldsymbol{b}(side_2) = \boldsymbol{b}(side_3) = {}^1/_3$, which seems like the right thing to do. If you measure credal inaccuracy with the absolute-value score, then the credences $\langle b_1, b_2, b_3\rangle$ will produce scores of:

$$\boldsymbol{I}(\langle b_1, b_2, b_3\rangle, \text{side}_1) = 1 - b_1 + b_2 + b_3$$
$$\boldsymbol{I}(\langle b_1, b_2, b_3\rangle, \text{side}_2) = 1 + b_1 - b_2 + b_3$$
$$\boldsymbol{I}(\langle b_1, b_2, b_3\rangle, \text{side}_2) = 1 + b_1 + b_2 - b_3$$

$\langle {}^1/_3, {}^1/_3, {}^1/_3 \rangle$ receives a score of $1^1/_3$ in all circumstances, which is better than some assignments, but not as good as the 1s-across-the-board scores that go to $\langle 0, 0, 0 \rangle$. So, embracing the absolute-value score within the accuracy-centered framework requires you to think that it is better to be certain that each side will *not* come up than it is to make the uniform ${}^1/_3$ assignment, and this is true even when you *know* that each side has one-chance-in-three of coming up. Even worse, you must say that it is worse to invest ${}^1/_3$ credence in all three sides than it is to invest credence one in the disjunction ($side_1 \vee side_2 \vee side_3$) while investing zero credence in each disjunct. One might react to this by biting the bullet and arguing that $\langle 0, 0, 0 \rangle$ is a better set of credences in every evidential situation, including those in which $\boldsymbol{ch}(side_1) = \boldsymbol{ch}(side_2) = \boldsymbol{ch}(side_3) = {}^1/_3$. Or, one might retain the absolute-value score as one's measure of inaccuracy and reject **AN** (thereby giving up on the whole accuracy-centered approach). Or, one could say that the absolute-value score is a lousy measure of accuracy *partly because it ranks* $\langle 0, 0, 0 \rangle$ *above* $\langle {}^1/_3, {}^1/_3, {}^1/_3 \rangle$.

The last strategy is the right way to go. As surely as we know anything in epistemology, we know that $\langle {}^1/_3, {}^1/_3, {}^1/_3 \rangle$ is the right credal state in the imagined evidential situation, and (as an independent point) that $\langle 0, 0, 0 \rangle$ is wrong in *any* evidential situation. Since the absolute-value score ranks the latter credences above the former, it must go. We do not reject the score because it fails to be a way of measuring closeness to truth (it definitely is) or because it violates *a priori* insights we have about how such closeness should be measured. Instead, we reject it *because* it encourages epistemic practices that conflict with our considered normative judgments about the proper ways for beliefs to be influenced by evidence. The dominance of $\langle {}^1/_3, {}^1/_3, {}^1/_3 \rangle$ by $\langle 0, 0, 0 \rangle$ is a symptom of this failing, but, at root, the problem is that the absolute-value score encourages a kind of doxastic extremism in which one can only minimize inaccuracy by being *certain* of the falsity of propositions for which there is significant evidence of truth. To see the point, suppose that a believer has credences $0 < b_1 \leq b_2 \leq \ldots \leq b_N$ for a partition $\langle X_1, X_2, \ldots, X_N \rangle$ with $N \geq 3$, and imagine that she has excellent evidence for investing a positive credence in X_1, say because she knows that its chance exceeds ${}^1/_{2N}$. According to absolute-value score this person can make her beliefs more accurate, no matter which X_j is true, by switching to the credences $c_j = b_j - b_1$. So, relative to that score, the only undominated credences are those for which some X_j has credence zero. This is nuts! It entails that whatever your evidence about the bias of a die – maybe you tossed it 1000 and saw 107 ones, 154 twos, 167 threes, 127 fours, 201 fives, 244 sixes – the rational pursuit of accuracy should you lead you to be entirely certain that one particular side (presumably $side_1$) will *not* come up when the die is tossed. Such credences are not responding correctly evidence, and a scoring rule which encourages them should be rejected. (The problems are only worse for the square-root score.)

For another instance of evidential norms informing the choice of inaccuracy scores, consider the *overtly* epistemic rationale that Joyce (2009) offers for a weakened version of **Strict**

Propriety which bans scores that permit probabilistically coherent credal states to be dominated. The absolute-value and square-root scores fails this test, while the Brier and logarithmic scores pass. To see why passing is a plus, suppose we have a score $\boldsymbol{I}$ and a probability $\boldsymbol{b}$ defined over a partition of (non-contradictory) events $\Omega = \langle \omega_1, \omega_2, \ldots, \omega_N \rangle$. Imagine that $\boldsymbol{b}$ is dominated by $\boldsymbol{c}$,so that $\boldsymbol{I}(\boldsymbol{b},\omega) > \boldsymbol{I}(\boldsymbol{c},\omega)$ for any $\omega \in \Omega$. An accuracy-centered epistemology based on $\boldsymbol{I}$ will deem $\boldsymbol{b}$ impermissible in all evidential situations. Given Extensionality,[24] this means that assigning the b_j credences to *any* partition of events of length N is also impermissible in all evidential situations. So, to show that $\boldsymbol{I}$ is unacceptable score we need only find a partition $\langle X_1, \ldots, X_N \rangle$ and a possible evidential scenario in which $b_j = \boldsymbol{b}(X_j)$ is clearly the correct credal assignment. This is easy: use **PP**! Imagine an N-sided die, and suppose that X_j says that the jth side will come up when the die is next tossed. It is reasonable to assume that, for any non-negative real numbers $b_1, \ldots, b_N$ that sum to one, there will always be an *epistemically possible* situation in which a believer knows nothing about an N-sided die except that the objective chance of each X_j is b_j. In this situation, as **PP** recommends, the obviously right credal state is $\boldsymbol{b}$, which means that any purported measure of epistemic accuracy that makes $\boldsymbol{b}$ impermissible must be dismissed. Though it was not clear in the (2009) paper, the appeal to chances is inessential here – all that matters is the possibility of *some* evidential situation in which $\boldsymbol{b}(X_j) = b_j$ is the correct credal assignment. This evidence could involve knowledge of the chances, or long experience observing frequencies, or a well-confirmed physical theory of the die and rolling process that makes it reasonable to believe that each face will come up with a probability proportional to its area, or even one of Easwaran and Fitelson's "oracles" who specifies the credences to adopt. However it is managed, if there is a possible evidential situation in which $\boldsymbol{b}$ is clearly the right credal state, then any accuracy score that has $\boldsymbol{b}$ come out dominated must be rejected. So, scores that violate **Strict Propriety** should be rejected because they encourage believers engaged in the rational pursuit of accuracy to hold credences other than those that **PP** and other legitimate norms of evidence advocate.[25]

[24] *Extensionality*, which is not assumed in the (2009) paper, helps to answer an objection, found in Hájek (2008), involving propositions that cannot be assigned arbitrary credences. Hájek offers the example of a 'Moore proposition' like M = "It rains in Minsk today and my credence for that is below ½." Arguably, the only credences $\langle b, 1 - b \rangle$ that a self-aware believer may assigned to $\langle M, \sim M \rangle$ will have $b < ½$. The advantage of *Extensionality* is that it makes the content of propositions in the partition immaterial to the import of *Strict Propriety*. If you assign the credences $\langle 0.6, 0.3 \rangle$ to $\langle M, \sim M \rangle$ you are making the mistake of assigning too high a credence to a Moore proposition, but you are also making another mistake (which is sufficient, in itself, to show that your credences are irrational), and this second mistake is exactly the same one you would be making if you assigned $\langle 0.6, 0.3 \rangle$ to $\langle H, \sim H \rangle$.

[25] Let me assuage one concern that might arise about this reasoning. Since the chances in question are probabilities, it can seem as if probabilistic coherence for credences is being imposed by *fiat*. This is wrong. While we are stipulating that credences must be coherent in the *special* evidential circumstances in which one knows only that $ch(X_j) = b_j$, it does *not* follow that credences must be coherent in *all* such circumstances – that's a much larger and more substantive claim, which can only be established by the full accuracy argument. In effect, we use the fact that it is always possible to find an evidential situation

Didn't we just beg the question? We said both that norms of evidence are legitimate only if they never sanction ***I***-dominated credences for an appropriate inaccuracy score ***I***, and that ***I***'s credentials as an inaccuracy score rest partly on the fact that it never lets ***b*** dominate ***c*** when a legitimate norm of evidence recommends ***c*** over ***b***. This is indeed a circle, but not a vicious one. The circle would be vicious if the objective were to prove that "*all* epistemic norms spring from" some antecedently understood notion of epistemic accuracy, but this is not the goal. The goal is to show that all epistemic norms we hold dear can live happily together within a framework in which doxastic accuracy is the cardinal epistemic desiderata and its rational pursuit the primary epistemic duty. We do this by showing that there are ways of valuing closeness to truth that respect and (in some cases) rationalize our core epistemic values and judgments. We have seen, e.g., that an accuracy-centered epistemology which employs strictly proper scores can provide a rationale for **PP** based on the fact that aligning credences with chances maximizes objective expected accuracy (which is the best one can do without recourse or 'inadmissible' information). This both shows that **PP** is consistent with an accuracy-centered epistemology, and explains why satisfying the Principle is part of the duty to rationally pursue accuracy.

Thinking more broadly, we can view the choice of an accuracy score as a *consistency test* for epistemic principles. One might have various views about what it takes for credences to be epistemically rational – that they should be probabilistically coherent, obey the Truth Norm, satisfy the Principal Principle, and so on – and one may wonder whether these views are jointly consistent with the idea that doxastic accuracy is the paramount epistemic good and that its pursuit is the core epistemic duty. There is a straightforward answer: A set of epistemic norms for credences are mutually consistent with the accuracy norm just in case there is an accuracy score ***I*** satisfying the four requirements imposed in §1 such that:

- No norm in the set ever permits a believer to hold the credences ***b*** in any evidential situation if ***b*** is ***I***-dominated by an accessible credence function (even one that is itself impermissible in that evidential situation).

- No norm ever prohibits a believer from hold the credences ***b*** in any evidential situation unless it also prohibits the believer from holding any credences that ***b I***-dominates.

Some putative evidential norms, like **Coherence**, **Truth** and **PP**, pass this test for every ***I***, others pass for some ***I***'s but fail for others, and still others fail for any such ***I***.

Let me emphasize, that the accuracy argument is just the *start* of an accuracy-centered epistemology for credences. A fully articulated account will involve further constraints on credal

---

in which the chances are given by the probabilities $\langle b_j \rangle$ to justify **Strict Propriety**, and then use **Strict Propriety**, in connection with the other constraints imposed on accuracy scores, to show that credences must be probabilities in *all* evidential situations, in particular those in which the chances are not known.

states and their relationships to evidence. When faced with some proposed norm of epistemic rationality, proponents of the accuracy- centered approach have three options: (i) they can reject the norm as illegitimate because it fails to promote epistemic accuracy, (ii) they can show how that the requirement is consistent with the existing framework by showing that it never allows accuracy dominated credences, or (iii) they can *make* it consistent with the framework by placing additional restrictions on inaccuracy scores that incorporate the norm's insights. For a case of (i), consider the claim that believers should aim to hold credences that are as well *calibrated* as possible (so that, on average, the proportion of truths among propositions assigned credence x is as close as possible to x). Joyce (1998) shows that this rule is inconsistent with the accuracy-centered approach because it is possible for $\boldsymbol{b}$ to be better calibrated than $\boldsymbol{c}$ even when $\boldsymbol{c}$'s credences are uniformly closer than $\boldsymbol{b}$'s are to the actual truth-values. That is fatal: since the unbridled pursuit of calibration conflicts with the pursuit of accuracy, it has to go.[26]

For an example of (ii) consider Alan Hájek's (2008) example of a 'Moore proposition' like M = "It will rain in Minsk today but my credence for that is below ½." Investing a high credence $m > ½$ in M is clearly irrational because such an assignment provides the believer with conclusive evidence of M's falsity (on the perhaps debatable assumption that a rational agent will know her own credences). Now, it might seem that we need a new norm to eliminate this sort of 'Moore incoherence', but it can be done within the accuracy-centered framework. Just notice that, whatever other credences $\boldsymbol{b}$ may assign, if it sets $\boldsymbol{b}(M) > ½$ it will be dominated by the credal state $\boldsymbol{c}$ defined by $\boldsymbol{c}(X) = \boldsymbol{b}(X)$ for $X \neq M$ and $\boldsymbol{c}(M) = ½\cdot(\boldsymbol{b}(M) + ½)$. (When $\boldsymbol{b}(M) = ½$ continuity guarantees a dominating $\boldsymbol{c}$ as well.) It does not matter that the dominating credal state $\boldsymbol{c}$ is probabilistically incoherent, since there will always be a coherent state that dominates $\boldsymbol{c}$ and thus also $\boldsymbol{b}$. So, the prohibition against 'Moore incoherent' credences follows from **AN**.

Finally, for an example of (iii), consider someone who thinks that the process of *Entropy Maximization* (MaxEnt) is the right way to settle on "prior probabilities".[27] To keep it simple, suppose one has symmetrical, but rather uninformative evidence about the propositions in some partition $\langle X_1, X_2, \ldots X_N \rangle$. (Think, say, of our current evidence about the last digit of the decimal expression for the number of humans alive at 12:00am GMT on 1 January 2000.) MaxEnt says that when choosing priors one should always select the credal state with maximum *Shannon entropy* $H(\boldsymbol{b}) = -\Sigma_n b_n \cdot log(b_n)$ from among those not directly contradicted by the data. If you believe this is the rationally mandated way to choose priors (which I don't!), then you may want

[26] This does not mean that calibration is immaterial to questions of epistemic rationality. As is well know, the so-called *calibration index* is a component of the quadratic score, along with something called the *discrimination index*. In contexts where the quadratic inaccuracy is used and where it is possible to increase calibration without decreasing discrimination by a larger amount, the pursuit of calibration is epistemically virtuous *because it increases accuracy*. See Joyce (2009) for details.
[27] See, for example, Jaynes (2003).

to incorporate your commitment into an inaccuracy score. Here is one way to do it, merely for purposes of illustration. Suppose you subscribe to the following two ideas:

A. The optimal credal state to have, among those not contradicted by the data, is the one that is the least committal about truth-values not entailed by the data (since these credences do the least amount of 'jumping to conclusions').

B. The relative degree to which two credence functions $\boldsymbol{b}$ and $\boldsymbol{c}$ 'jump to conclusions' is the difference in their Shannon entropies, $H(\boldsymbol{b}) - H(\boldsymbol{c})$.

Though it would take us too far afield to justify it here, there are reasons to think that someone who uses $\boldsymbol{I}$ to measure epistemic inaccuracy is thereby committed to thinking that one coherent credence function $\boldsymbol{b}$ is less committal than another $\boldsymbol{c}$ *in re* truth-values just when $Exp_{\boldsymbol{b}}(\boldsymbol{I}(\boldsymbol{b})) > Exp_{\boldsymbol{c}}(\boldsymbol{I}(\boldsymbol{c}))$, i.e., just when $\boldsymbol{b}$ expects its inaccuracy to be higher than $\boldsymbol{c}$ expects its inaccuracy to be. As a result, a person who accepts (A) and (B) will want to measure inaccuracy using a score with $Exp_{\boldsymbol{b}}(\boldsymbol{I}(\boldsymbol{b})) = H(\boldsymbol{b})$. It turns out that the logarithmic score has this property! So, one can incorporate the MaxEnt norm – maximize Shannon entropy (among those not contradicted by the initial data) – within the accuracy-centered framework by adopting the logarithmic score.

Let me emphasize that I am *not* endorsing this maneuver. In fact, I think (B) is entirely up for grabs. There are many ways to assess the degree to which a system of credences 'jumps to conclusions', and different ways of doing it have disparate effects on inaccuracy scores. For example, if, instead of using $\boldsymbol{H}$ (= self-expected Shannon information), one identifies the degree to which a credence function goes beyond the data with its *self-expected variance,* then the Brier score would turn out to the right way to measure inaccuracy. There are many, many other ways that (B) could be interpreted as well, and each will lead to its own $\boldsymbol{I}$. So, the point to take away, here, is not that an accuracy-centered approach should commit to the logarithmic score. Rather, it is that the approach is very flexible.

Let me close by reiterating the basic morals of this section:

- The relationship between epistemic norms and accuracy norms, however, is not hierarchical, but symbiotic. Rather than being autonomous, our concept of accuracy will be informed by, and highly dependent on our considered views about which epistemic norms are legitimate.

- Evidential considerations should factor into the choice of an inaccuracy score because these scores are ways of measuring 'closeness to the truth' that reflect our considered 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 that properly reflects our epistemic values, including the value we place on holding well-justified beliefs.

- The choice of an accuracy score is a *consistency test* for epistemic principles. A group of norms for credences are mutually consistent just in case there is an accuracy score such that: (i) no norm in the group ever permits a believer to hold credences that are accuracy-dominated; no norm ever prohibits a holding a system of credences in any evidential situation unless it also prohibits any credences that that system dominates.

- Some familiar evidential requirements, the Principal Principle for example, can be incorporated straightforwardly into the accuracy-centered framework by placing restrictions on the allowable accuracy measures.

- Some others can be shown to follow from the framework.

- Some important aspects of the process of settling on prior probabilities can be under understood as deciding about the (informational) values that we want our accuracy measures to exhibit.

Works Cited

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

Brier, G. W. (1950) "Verification of Forecasts Expressed in Terms of Probability,"

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.

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

E. T. Jaynes (2003). *Probability Theory: The Logic of Science*. Cambridge: Cambridge University Press.

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 (forthcoming) "A New Epistemic Utility Argument for the Principal Principle," *Episteme*