J. Joyce: "Some Remarks on Easwaran & Fitelson's 'An Evidentialist Worry About Joyce's Argument for Probabilism"

E & F allege that the pursuit of epistemic accuracy can conflict with other legitimate epistemic goals, like the goal of having credences that reflect known objective chances. In particular, they claim that applying legitimate epistemic norms in conjunction with the norm of accuracy can lead to order effects, so that we get one answer by minimizing inaccuracy and then applying the norm, while we get an incompatible answer by applying the norm and then minimizing in accuracy. Here is an example (using the Brier Score to measure epistemic inaccuracy):

> Joshua knows that the coin about to be flipped is biased 4:1 in favor of tails, so that Ch(H) = 0.2. This is all he knows. His (incoherent) credences are ***b***(H) = 0.2 and ***b***(T) = 0.7. Joshua wants to obey the Principal Principle by aligning his credences with the known chances, but he also does not want to hold accuracy-dominate credences. He might proceed in two ways:

1. **PP-then-AD**: Joshua first makes sure his credences satisfy PP, which they do since ***b***(H) = 0.2, and then minimizes inaccuracy. According to E&F, my approach recommends that Joshua carry out the second step by moving to credences that accuracy dominate ***b***. This requires him to adopt some credence function ***b'*** with ***b'***(H) = x and ***b'***(T) = $1 - x$ where $0.248 \leq x \leq 0.36$. Joshua cannot possibly satisfy PP after this move since 0.2, the known chance of heads, is not in this range of remaining credence functions.

2. **AD-then-PP**: Joshua first makes sure his credences are not accuracy dominated. There are two ways he might do so:

   - He might eliminate all non-coherent credences from consideration, and so resolve to have some credence ***b'***(H) = x and ***b'***(T) = $1 - x$ with $0 \leq x \leq 1$.

   - He might eliminate all credences except those that accuracy dominate his current credences, so that all the permissible credences have form ***b'***(H) = x and ***b'***(T) = $1 - x$ with $0.248 \leq x \leq 0.36$.

   E&F think that I am committed to the second of these recommendations. But, if Joshua does the second thing, then he is left with an "available" set of credences that does not contain any element that assign heads a probability of 0.2. When he applies PP and eliminates the credences for which ***b'***(H) $\neq$ 0.2 he thereby eliminates everything.

I have two main misgivings about this reasoning.

Misgiving-I: As E&F admit in Fn 8, **PP-then-AD** relies on a misapplication of the PP. If Joshua knows Ch(H) = 0.2 (and understands what chances are) then he also knows Ch(T) = 0.8. So, if he starts by making sure his credences conform to PP, then the first thing he should do is to eliminate every credence function except ***b'***(H) = 0.2 and ***b'***(T) = 0.8. But then, applying AD leaves him right where he is (as long as the "epistemic scoring rule" is proper). So, **PP-then-AD** will take Joshua from ***b***(H) = 0.2 and ***b***(T) = 0.7 to ***b'***(H) = 0.2 and ***b'***(T) = 0.8.

> Note: This argument assumes that chances are probabilities. I am happy to defend this assumption (see below). I am even happy to defend the claim that any function ***q*** that takes propositions to real numbers must obey the laws of probability if a rational agent is to treat it as an "expert". Roughly, ***q*** is an expert relative to a set of propositions $\mathcal{X}$ just in case you are committed to aligning your credences to *q*'s values when you know these values, so that $\boldsymbol{b}(X \mid q(X) = x) = x$ for all $X \in \mathcal{X}$ and all real numbers *x*. (This must be qualified with a "no undermining" clause which precludes other experts ***q****, which you see as more reliable than ***q***, that make recommendations at odds with those of ***q***. That is, it should never happen both that $\boldsymbol{b}(q(X) = x \,\&\, q^*(X) = y) > 0$ for some $x \neq y$ and that $\boldsymbol{b}(X \mid q(X) = x \,\&\, q^*(X) = y) = y$. Here, think of *q* as giving the chances at one time and of *q** as giving the chances at a later time. The *q** expert trumps the *q* expert.)

Misgiving-II: The "AD then PP" argument relies on a misapplication of AD, albeit a subtle one. Let's start with the two way of interpreting the requirement to avoid accuracy dominated credences:

- AD-WEAK: An incoherent agent is obliged to adopt some credence function that is not accuracy dominated (i.e., some coherent credence function, if I am right about the character of epistemic scoring rules), but is ***not*** necessarily obliged to move to a credence function that accuracy dominates *her own* credence function.

- AD-STRONG: An incoherent agent is obliged to adopt some credence function that is not accuracy dominated and which dominates *her own* credence function.

If AD-WEAK is right, then the set of credence functions that remain after applying AD will include all coherent ***b'*** with $0 \leq$ ***b'***(H) ≤ 1. There is then no problem with **AD-then-PP**. Since ***b'***(H) = 0.2 and ***b'***(T) = 0.8 is not eliminated by AD, when you apply PP to the available credence set you end up at ***b'***(H) = 0.2 and ***b'***(T) = 0.8, just where you did when you applied **PP-then-AD**.

On the face of things, it seems like there are problems for AD-STRONG. For it forces you to first eliminate all coherent credence functions except those with $0.248 \leq$ ***b'***(H) ≤ 0.36. Since ***b'***(H) = 0.2 and ***b'***(T) = 0.8 is no longer "available" it seems like we cannot apply PP. In fact, since PP eliminates all remaining credence functions from consideration, you are left without rational credence to hold. This does conflict with answer one gets from **PP-then-AD**.

I would resist this argument on two fronts. First, I do *not* endorse AD-STRONG. While an incoherent agent will do ***better***, from the perspective of epistemic accuracy or epistemic utility, if she adopts coherent credences that dominate her own, it does not follow that she does ***best*** (*or even particularly well*), given her overall epistemic situation, by adopting such credences. When ***b*** is dominated by ***b'*** this shows that there is something wrong with ***b***, not necessarily that there is something right about ***b'***. Here is a decision-theoretic analogy. Suppose a die will be rolled, and consider the following wagers on its outcome:

	One, Two, Three	Four, Five	Six
A	$1	$1	$1
B	$2	$2	$2
C	$10	$5	$0

Suppose that you own *A*, and someone offers you the option of keeping it or trading it for *B* or for *C*. You might recognize that keeping *A* is definitely irrational since it is dominated by *B*, but it does *not* follow that you should trade *A* for *B*. What you should do depends entirely on your evidence about the die's chances of coming up in various ways. If you know it to be fair, then it would be crazy to choose anything but *C*! It's the same here. Often it is an epistemic *mistake* to remedy incoherence by moving to a dominating alternative. It might be, after all, that none of the dominating alternatives is remotely justified in light of the evidence. The E&F example is just such a case. Since none of the agent's dominating alternatives is justified in light of her evidence about the objective chances, she should *not* to adopt *any* of them. Rather, upon seeing that there are ***better*** credences for her to have, she should consult her evidence to find the ones that are ***best*** given her total evidential situation. If, as here, the incoherence is the result misusing evidence (specifically of failing to align credences with the known chances), then the agent should move to the credence function that is best supported by her evidence or, as I like to say, the credal state that best reflects her total evidential situation. This will always be a coherent credence function. How do I know? Because the credences are justified to the extent that the evidence makes it reasonable to think (a) that the credences are accurate and (b) that the credences are, in terms of overall epistemic value, the best one can do. (See below for more on this point.) Thus, I deny that there can be "conflict between evidential norms for credence and an (accuracy-dominance) coherence norm for credences." (Fn 5, p. 4) Insofar as we take our "epistemic scoring rule" to correctly reflect considerations of accuracy and other epistemic virtues, the policy of proportioning our beliefs to our evidence will automatically ensure that we adopt some set of credences that we can regard as having the best "epistemic score". Of course, no dominated system of credences can ever be so regarded.

As a second related point, suppose (contrary to what I just said) that I am somehow committed to AD-STRONG. Even here, I deny that "order effects" arise. This is because, independent of my views about AD-STRONG, I reject the way that E&F apply PP in the **AD-then-PP** scenario. E&F restrict applications of PP to "available" credence functions. More precisely, they assume:

PP^A. If you know Ch(H) = *x* (and Ch(T) = 1 – *x*) and if ***b'***(H) = *x* and ***b'***(T) = 1 – *x* is available (i.e., has not already been ruled out before you apply PP), then PP requires you to adopt it as your credence function.

PP^B. If you know that Ch(H) = *x* (and Ch(T) = 1 – *x*) and if ***b'***(H) = *x* and ***b'***(T) = 1 – *x* is ***not*** available (i.e., has already been ruled out), then PP prevents you from rationally holding any credence function (all credences are made "unavailable").

PP^A is fine, but PP^B is wrong -- indeed, it contradicts the Principal Principle. To see why, notice that in the **AD-then-PP** scenario E&F invoke AD-STRONG to move straight from the incoherent assignment ***b***(H) = 0.2 and ***b***(T) = 0.7 to the set of "available alternatives" with 0.248 ≤ ***b'***(H) ≤ 0.36. Moreover, they do so in a way that does ***not*** make use of the information Ch(H) = 0.2. If Ch(H) were 0.0 or 0.5 or 0.75 or 1.0, the credence functions available after AD is applied would be the same. This is something AD-STRONG requires. So, if AD-STRONG is right, we can ignore chance information when applying AD. But, it then makes no difference whether we imagine the agent having the chance information at the start or imagine her acquiring it after only AD is applied. In the later case, however, no one (I hope) will say that an agent with a credence function in the range 0.248 ≤ ***b'***(H) = 1 – ***b'***(T) ≤ 0.36 who learns Ch(H) = 0.2 should continue to hold a credence in that range. Unless the agent has undermining information (not envisioned here), **PP** says that the agent should move to new credences with ***b''***(H) = 0.2 even if, prior to factoring in the chance information, such a function was "unavailable". This contradicts PP^B, which requires us to despair of finding any posterior probability consistent with the chance information Ch(H) = 0.2 unless we start with a prior credal state that is already consistent with the ***b***(H) = 0.2 assignment. It should not be surprising that PP^B fails this way. The Principal Principle does not care how you came to hold your prior credences. It says that, however you got where you are, when you acquire information about chances, you should align your new credences with known objective chances (barring undermining information). Likewise, it says that if the credences "available" to you are misaligned with the known chances, then (barring undermining information) you are obliged adopt some "unavailable" credences.

To step back a bit, consider E&F's general presentation of their objection. They maintain that my view approach will have problems whenever:

(★1) An agent has an incoherent credence function ***b***.

(★2) The agent has good reason to believe (or knows) that epistemic rationality requires her to adopt some credence in set $\mathcal{I}$.

(★3) $\mathcal{I}$ contains no credence function that accuracy dominates ***b***.

I am supposed to be committed to the inconsistent conclusion that this agent must adopt a credence function that both accuracy dominates *b* and is in $\mathcal{I}$. E&F write that, to avoid this,

> "one would need to argue that no examples satisfying (★) are possible. And that is a *tall* order. Surely, we can imagine that an oracle concerning epistemic rationality has informed us that [some credence function] in $\mathcal{I}$ is required -- despite the fact that all (coherent) Brier dominating functions ***b'*** are [outside] $\mathcal{I}$."

Given my previous remarks, it should be clear why I'll disagree. Examples satisfying (★) pose no problem for me a*s long as we cannot have good reason to believe that epistemic rationality requires us to adopt **incoherent** credences*. To make things specific, suppose the oracle gives us probabilities for heads and tails in the example we have been considering, so that, for *x* and *y* in the unit interval, she tells us something of the form *O*(*x*, *y*) = "set ***b'***(H) = *x* and set ***b'***(T) = *y*". To the extent that we regard the oracle as reliable when it comes to rational belief we will treat her as an "expert", i.e., our prior credences will satisfy ***b***(H | *O*(*x*,*y*)) = *x* and ***b***(T | *O*(*x*,*y*)) = *y*. Now, if we are incoherent and the oracle recommends a coherent credence assignment, one in which $y = 1 - x$, and if we really do take her to be an expert about what credences it would be rational for us to adopt (given our evidence), then nothing in my view prevents us from saying that we should adopt ***b'***(H) = *x* and ***b'***(T) = *y*, even when this does not dominate our current incoherent assignment. Even if all the dominating credence functions are ***better*** than the one we hold, the oracle's function might still be ***best***.

On the other hand, if it is possible for the oracle to recommend ***in***coherent credences *and legitimately retain her status as an oracle/expert*, then I'm in trouble given my commitment to the view that it is epistemically irrational to hold incoherent credences. But, there is no danger. Insofar as we think that our "epistemic utility" correctly captures the accuracy of credences (perhaps together with other epistemic virtues), we should reject any "expert opinions" that recommend dominated credences. We know that such an expert cannot be recommending the very best options since all the dominating options are better (from the perspective of accuracy and the other epistemic virtues). So, if I am right about what epistemic scoring rules look like, an oracle who recommends incoherent credences should be ignored (or at least should not be treated as an epistemic expert). Thus, *a necessary condition for it to be rational to treat an alleged "oracle" as an expert about epistemic rationality is that her recommendations should be probabilistically coherent*.

A fly in the ointment? While I generally endorse the argument just given, I really use it to argue that Chance, which PP says is an epistemic expert, must obey the laws of probability. This is because I appeal to the fact that chances are probabilities to justify a central constraint on epistemic utility functions, namely

> COHERENT ADMISSIBILITY: An acceptable epistemic scoring rule ***S*** will never allow a coherent credence function ***b*** to be even weakly dominated by an incoherent credence function ***c***, i.e., it will never be the case that $\boldsymbol{S}(\boldsymbol{b}, \boldsymbol{v}) \geq \boldsymbol{S}(\boldsymbol{c}, \boldsymbol{v})$ for every truth-value assignment *v*.

I argue, roughly, that allowing for violations of CA would prevent people from satisfying the Principal Principle. The best form of the argument (which, unfortunately, does not coincide yet with one that appears in print) depends on the idea that for any probability ***p*** defined on any finite abstract Boolean algebra there is a Boolean algebra of propositions and a possible world in which the objective chances of those propositions are given by ***p***.

One worry: Chances might not be probabilities. Wrong!

There are (overwhelming) reasons to think that chances satisfy the laws of probability. First, am not convinced by any arguments I have seen for non-probabilistic chances. Second, I think everything we know about the theoretical role of chances in science and in epistemology points to a probabilistic interpretation. Here is pretty close to everything I know about chances:

- Chances *explain* relative frequencies in long series of independent trials. *Ceteris paribus*, a good way to explain why the frequency of *A*'s among *B*'s in such a series is *f* is by showing that the conditional chance of observing *A* given *B* is approximately *f*.

- Facts about frequencies in long series of independent trials *confirm* hypotheses about chances. Observing a long series of independent trials in which the frequency of *A*'s among *B*'s is *f* is*, ceteris paribus*, good evidence for believing that the conditional chance of *A* given *B* is approximately *f*.

- The chances at time *t screen off* information about the history of events up to *t* (this included information about the chances at previous times). This makes it legitimate to use the chance of *X* at *t* as a kind of "summary statistic" when predicting *X*'s occurrence on the basis of data about the state of the world up to *t*. Barring information about future chances, knowing the value of $Ch_t(X)$ is the best evidence one can have about *X*. This is what justifies us in treating Chance_t as an epistemic expert as required by the Principal Principle.

- The chances at *t* are a function of past chances and the resolution of chance events that occur before *t*.

The first two facts convince me that chances are probabilities. If we use them to explain frequencies (which satisfy the laws of probability) and if we confirm them by appealing to frequencies, then chances should be probabilities themselves. (That's too quick, of course, but I'll elaborate if people want to hear more.)