Comments on Richard Pettigrew's *Accuracy and the Laws of Credence*

MATTHEW KOTZEN
University of North Carolina at Chapel Hill

Accuracy and the Laws of Credence is required reading for anyone interested in the foundations of epistemology. It is that rare philosophical work which serves both as a stunningly clear overview of a topic and as a cutting-edge contribution to that topic. I can't possibly address all of the interesting and philosophically rich components of *Accuracy and the Laws of Credence* here, so I will largely restrict my attention to pieces of Parts I, II, and III of the book, though I'll have some more general things to say about Pettigrew's accuracy-only approach to epistemology toward the end.

1. Measuring Inaccuracy

One of Pettigrew's core postulates in his account of inaccuracy is

> **Perfectionism**: The accuracy of a credence function at a world is its proximity to the ideal credence function at that world.

This notion of "proximity to the ideal credence function" is made precise in two steps. First, Pettigrew appeals to **Alethic Vindication**, according to which the ideal credence function at a world is the "omniscient" credence function that assigns 1 to all true (uncentered) propositions at that world and 0 to all false (uncentered) propositions at that world. Second, Pettigrew makes the notion of "proximity" precise by appealing to a *divergence* from one credence function to another; the inaccuracy of a particular credence function at a world, then, is just the divergence between that credence function and the ideal credence function at that world.

However, though Pettigrew appeals to the metaphor of "distance" in characterizing a divergence, he is explicit that he does *not* mean to assume that his preferred divergence has all of the same properties as a metric, which is the usual device that mathematicians use to formalize the notion of distance. In particular, whereas metrics satisfy the "triangle inequality" — according to which the distance from a to c is no greater than the sum of the distance from a to b and the distance from b to c — Pettigrew does not assume divergences to satisfy the triangle inequality. Indeed, as Pettigrew acknowledges, his preferred divergence — the Squared Euclidean Distance divergence — does not satisfy the

triangle inequality. For example, consider three credence functions a, b, and c that are each defined only on the single proposition P, and suppose that $a(P) = .6$, $b(P) = .8$, and $c(P) = 1$. The Squared Euclidean Distance divergence between a and b is $(.8-.6)^2 = .04$, the Squared Euclidean Distance divergence between b and c is $(1-.8)^2 = .04$, and the Squared Euclidean Distance divergence between a and c is $(1-.6)^2 = .16$. Pettigrew argues that

> "... the triangle inequality is only intuitively appealing when we are measuring physical distance. In those cases, its intuitive appeal arises from the thought that the distance between two points is the length of the shortest path between them (together with the thought that the length of a path that divides into two parts is obtained by summing the length of the first path with the length of the second path)."

However, I'm inclined to think that the triangle inequality is quite plausible even for lots of non-physical distances; for example, I have a hard time understanding a notion of distance between two colors in color-space, or two organisms in gene-space, or two different companies in financial-valuation-space, if the relevant notion of distance doesn't obey the triangle inequality. Notably, a lot of very helpful diagrams — including ones in Pettigrew's book — make use of an analogy between spacial distance and credal "distance" in order to make an argument or proof strategy more intuitively compelling. Of course, there are crucial similarities between the standard Euclidean Distance divergence and the Squared Euclidean Distance divergence which may underwrite this analogy; since squared Euclidean distance is a strictly increasing function of standard Euclidean distance, it follows that the two divergences will always agree on whether x is closer to y than x is to z, for any x, y and z. And *perhaps* that is enough to vindicate the metaphor of "proximity," and the graphical appeal to intuitions based on standard Euclidean distances, in the motivations for Pettigrew's account of inaccuracy. But I do wish that I had some more intuitive guidance here about how to think of an alleged notion of a proximity that violates the triangle inequality.

Pettigrew also addresses and rejects the **Weak Convexity** constraint on inaccuracy:

> **Weak Convexity**: If $\mathfrak{J}$ is a legitimate inaccuracy measure and c and c' are distinct credence functions that are equally inaccurate at w, then the equal mixture of c and c' is less inaccurate than either c or c'. That is: If $\mathfrak{J}(c,w) = \mathfrak{J}(c',w)$, then $\mathfrak{J}(\frac{1}{2}c + \frac{1}{2}c', w) < \mathfrak{J}(c,w) = \mathfrak{J}(c',w)$

This constraint has a somewhat complicated history. Joyce motivates **Weak Convexity** in his 1998 by appealing to the intuition that, if a change from c to $\frac{1}{2}c + \frac{1}{2}c'$ does not correspond to an increase in accuracy, then neither will a twice-as-extreme change in the same direction (i.e., from c to c'). Maher 2002 objects that **Weak Convexity** also implausibly rules out the situation in which the change from c to $\frac{1}{2}c + \frac{1}{2}c'$ is neutral in respect of accuracy and yet where the change from c to c' is neutral as well. Pettigrew concludes that "if Joyce's argument establishes anything, it establishes something weaker than Weak Convexity; it establishes that, for each world, the inaccuracy of a credence function is a *convex*, but not necessarily *strictly convex* function of it. But this is too weak to establish the central theorem of Joyce's paper, which underlies his accuracy argument for Probabilism." Still, it seems to me that Joyce *has* established this weaker claim, and that it can and should be assumed as a constraint on measures of inaccuracy, even if the

argument of Joyce 1998 doesn't survive the weakening. Moreover, Joyce 2009 provides an additional argument in favor of **Weak Convexity**; in brief, the main idea is that strict convexity discourages an epistemic agent from ignoring her evidence and employing a random belief-altering process; for example, if the credence in p that my evidence justifies is .5, a strictly convex measure discourages me from employing a random belief-altering process that switches my credence in p either to 1 or 0 with equal probability, since strict convexity guarantees that the inaccuracy of a credence of .5 will be less than the average of the inaccuracy of a credence of 1 in a world where p is false and the inaccuracy of a credence of 0 in a world where p is true. So, I would have been interested to see a more detailed discussion of convexity in Chapter 3, as well as the implications of assuming it for the main results of the book.

A few other thoughts on Pettigrew's account of inaccuracy:

Pettigrew assumes **Divergence Additivity**, according to which the total inaccuracy of a credence function is calculated by *summing* the individual inaccuracies of the credences that the function assigns to each proposition. This can seem natural enough, though Pettigrew perhaps overstates its naturalness somewhat. For one thing, this constraint immediately rules out the (non-squared) Euclidean Distance divergence; the Euclidean distance between 0 and 1, and between 1 and 0, is 1, and hence relative to the omniscient credence function (1,0), **Divergence Additivity** entails that the inaccuracy of credence function (0,1) is 2, whereas the Euclidean distance between (0,1) and (1,0) is $\sqrt{2}$. Of course, Pettigrew wants to rule out the Euclidean Distance divergence, and the *squared* Euclidean distance divergence clearly obeys **Divergence Additivity**; in the example above, $1^2 + 1^2 = (\sqrt{2})^2$. But this does have the feel of ruling out a class of prima facie plausible divergences more or less by fiat. Second, it can seem a bit strange to say that the overall inaccuracy of a credence function is always increased when a new highly-but-not-perfectly-accurate credence is added to it; I, for one, welcome credences of .999999 in many true propositions that I hadn't previously considered, notwithstanding the fact that **Divergence Additivity** entails that such new credences will increase the overall inaccuracy of my credence function. In fact, Pettigrew himself makes a similar point in his criticism of the calibrationist account of inaccuracy on p. 61; there, he argues that in certain situations the calibrationist approach implausibly entails that a credence function that assigns .5 to every proposition is perfectly calibrated and hence perfectly accurate. Pettigrew explains that "[w]hatever accuracy is, it is not a virtue that can be gained so easily. Accuracy is something for which one strives over the course of an epistemic life. It is a goal such that one collects evidence in order to achieve it better. On the calibrationist approach, accuracy is something many agents can acquire maximally at the beginning of their epistemic life, prior to acquiring any evidence." I agree, but I think that an analogous worry applies to **Divergence Additivity**; if nearly any consideration of new questions is bound to decrease my overall accuracy (and will never increase it), then I wonder how an accuracy-only epistemology like Pettigrew's is able to account for the epistemic value of such consideration.

2. Chance-credence Principles

In Part II, Pettigrew considers various chance-credence principles, beginning with the Principal Principle and culminating with his preferred **Evidential Temporal Principle**, according to which (roughly) a rational agent with total evidence E ought to have a conditional credence in proposition X, conditional on the claim that the current chances

are given by *ch*, equal to the conditional chance that *ch* assigns to *X*, conditional on *E*. Pettigrew's argument for this principle relies on **Current Chance Evidential Immodest Dominance**, according to which (again, roughly, and leaving out some details) if all of the chance functions that you can consider possible agree, once they have been updated on your total evidence, that credence function c_1 does better in terms of accuracy than credence function c_2 does, then it would be irrational for you to adopt c_2 over c_1.

Pettigrew is aware that there is a concern about circularity here: even if we agree that **Current Chance Evidential Immodest Dominance** is true, it's not obvious that **Current Chance Evidential Immodest Dominance** is more fundamental than or can be justified independently of the very sorts of chance-credence principles (such as **Evidential Temporal Principle**) that it is being used to derive. After all, it is natural to think that the whole project here is that of justifying the treatment of chances as a kind of "expert" in the sense that we ought to allow our opinions about the chances of propositions to constrain our credences in the propositions themselves. But **Current Chance Evidential Immodest Dominance** looks to simply *assume* that chance is just this sort of expert: it demands that you treat the chance functions that you consider possible to be sufficiently authoritative that you never go against their unanimous verdict with regard to accuracy.

Pettigrew's response to this worry is to concede that his argument from **Current Chance Evidential Immodest Dominance** to **Evidential Temporal Principle** "will not satisfy someone who is not already convinced that we should defer to the chances *in some way*." But he then argues that **Evidential Temporal Principle** (actually, a slightly stronger version of the principle that also follows from his assumptions) articulates a much more *precise* way of deferring to chances than **Current Chance Evidential Immodest Dominance** does. In particular, (the strengthened version of) **Evidential Temporal Principle** "demands that your credence in a proposition lies within the (closure of the) span of the possible current chances of that proposition conditional on your evidence," whereas **Current Chance Evidential Immodest Dominance** "merely says that, on the rare occasions where [the possible chance functions] all agree in their ordering of two credence functions with respect to accuracy (once they have been brought up to speed with your evidence), then you should adopt that ordering yourself on pain of irrationality."

Of course, Pettigrew is correct that there is *some* distance between these claims, but it is hard to feel fully satisfied by this argument, at least if what you were hoping for is the "reduction" promised in the Introduction of apparently evidentialist constraints like the Principal Principle to purely accuracy-oriented assumptions. Consider, for example, the case where there is only one chance function that I regard to be possible; suppose that I *know* that the chance that the coin will land heads is 50%. In this case, **Current Chance Evidential Immodest Dominance** says that, since the only chance function that I regard to be possible expects a credence of .5 in heads to be more accurate than any other credence, I should adopt a credence of .5; **Evidential Temporal Principle** says that, since the span of possible current chances of heads (conditional on my evidence) is the trivial span of [.5,.5], I should set my credence in heads to the only value within that trivial span – i.e., to .5. Even if both of those principles are true and deliver the correct verdict, it is hard to see clearly how **Current Chance Evidential Immodest Dominance** articulates a more basic or fundamental requirement on rationality, from which **Evidential Temporal Principle** is supposed to be illuminatingly derived. To my mind, the traditional problem of justifying the Principal Principle just is the problem of explaining why we should defer in our credences to the chances: what is so special about *chances* such

that credences in *chances of propositions* so tightly constrain our credences in the propositions themselves? To anyone in the grip of this sort of question, **Current Chance Evidential Immodest Dominance** sounds a lot like the claim that you should adopt the credences that the possible chances recommend, which is precisely what was at issue. Again, I agree with Pettigrew that there is some distance between **Current Chance Evidential Immodest Dominance** and (the stronger version of) **Evidential Temporal Principle**, and I think that Pettigrew has successfully shown that, *if* you ought to follow the advice of the possible chances in the case where they are unanimous, *then* you ought to adopt a credence in the span of the possible chances even in the case where they are not unanimous. That is an interesting result, but not the result that the rhetoric in the Introduction led me to expect.

3. The Principle of Indifference

One of the most novel elements of *Accuracy and the Laws of Credence* is Pettigrew's discussion of epistemic risk-aversion in Part III on the Principle of Indifference. In Lara Buchak's recent book *Risk and Rationality*, Buchak makes a compelling case that certain kinds of risk-aversion can be rational in the practical case and she develops a sophisticated decision theory to accommodate this insight. Here, Pettigrew aims to apply some of the same insights about practical decision theory to the epistemic case. The (rough) thought is that just as an interest in avoiding (or lowering) the risk of bad practical outcomes can be an additional factor in our practical decision theory that is separable from the interest in maximizing expected practical utility, so too can an interest in avoiding (or lowering) the risk of bad epistemic outcomes (i.e., inaccurate credences) be a relevant factor in our epistemic decision theory, separable from the interest in maximizing expected epistemic utility.

The main idea of Pettigrew's proposal is that the Principle of Indifference can be justified on the grounds that it minimizes the risk of inaccuracy in a certain way. More specifically, the idea is that if an agent "at the beginning of her epistemic life" were to follow a Maximin strategy, where she always chooses the option that maximizes minimum accuracy (alternatively: that minimizes maximum inaccuracy) — i.e., where she never risks greater than necessary inaccuracy — she will obey the Principle of Indifference. Obviously, Maximin is an *extremely* conservative policy; it directs an agent *never* to risk greater than necessary inaccuracy, even in situations where risking slightly greater than necessary inaccuracy would yield *huge* decreases in expected inaccuracy. Pettigrew acknowledges this, though he doesn't say much to defend Maximin beyond pointing out that "it is intended to apply only at the beginning of an agent's epistemic life — it governs her only at the point when she is setting her initial credences; that is, at a point when she has no credences to guide her decisions, epistemic or otherwise." Pettigrew also draws an analogy to Rawls's veil of ignorance, but unfortunately I just don't see why such extreme conservatism should be required even at the beginning of our epistemic lives, nor do I see why there should be different constraints on risk-aversion at the beginning of our epistemic lives than there are for the rest of our epistemic lives. That said, the possibility of rationally mandatory — or even permissible — risk aversion is a fascinating one, and one that certainly deserves further exploration. In particular, once we countenance permissible risk-aversion in our epistemic decision theory, there's a question about how this would impact decision-theoretic arguments for all kinds of epistemic constraints. Arguments which depend on *dominance* reasoning may still fare reasonably well,

since considerations of risk-aversion won't in general affect our reasons for preferring a dominating option to a dominated one. But arguments that appeal to expected epistemic utility — for example, the argument of Greaves and Wallace 2006, which is central to the argument for Plan Conditionalization offered in Section 14.1 of *Accuracy and the Laws of Credence* — might not fare as well, since the whole point of allowing risk-aversion into our epistemic decision theory is to leave room for an option not to be all-things-considered best even if it is among the options with the highest expected utility.

One final point about the Principle of Indifference: on Pettigrew's view, his version of the Principle of Indifference is language-dependent, but it is not inconsistent. The traditional concern here, arising from von Mises 1957 and van Fraassen 1989, is that the Principle of Indifference recommends inconsistent assignments depending on how the space of possibilities is partitioned; if I think of the space as being partitioned into {BLUE, YELLOW, RED}, then the Principle of Indifference requires me to assign a credence of $\frac{1}{3}$ to BLUE, whereas if I think of the space as being partitioned into {BLUE, ¬BLUE}, then the Principle of Indifference requires me to assign a credence of $\frac{1}{2}$ to BLUE. However, Pettigrew claims that "[s]ince it is never the case that the set of propositions to which [the agent] assigns credences is both {BLUE, ¬BLUE} and {BLUE, YELLOW, RED}, no inconsistency arises." Pettigrew goes on to explain that "[w]hat rationality demands of an agent is determined by the resources that are available to her," and that what rationality demands of an agent with an impoverished conceptual scheme might be different from what it demands of an agent with a more expansive conceptual scheme. But I think that much more needs to be said here. First, the issue isn't restricted to agents who are conceptually impoverished; I don't need to lack the concept of "yellow" or "red" in order to have or deploy the concept "not blue," any more than I need to lack the concept "volume of a cube" in order to deploy the concept "side length of a cube." Second, even if it were true (which I doubt) that I can't *simultaneously* assign credences to {BLUE, YELLOW, RED} and to {BLUE, ¬BLUE}, I certainly can "go back and forth" between these different conceptual schemes for all sorts of reasons, and it is implausible that rationality recommends switching my credence in BLUE every time that happens. So, even if Pettigrew has succeeded in deriving a version of the Principle of Indifference, I don't see that he has done anything to allay decades of legitimate fears about its consistency.

4. Accuracy-only Epistemology

There are some big implications of all of this material for epistemology more generally, and unfortunately there isn't space here to explore them all. The clearest statements of Pettigrew's overall epistemological aspirations are in the Introduction:

> I will argue that, in fact, accuracy is the *only* epistemic virtue. Or, more precisely, I will argue that it is the only *fundamental* epistemic virtue: all other epistemic virtues derive their goodness from their ability to promote accuracy.[1]

and in Chapter 2:

> Joyce's strategy here is in line with a broadly *accuracy-first epistemology*. For him, accuracy may not be the only source of value for credence functions—there is, perhaps, also the value that comes from respecting one's evidence—but it is the primary source of

---

[1]   p. 6.

value: it constrains how the other sources of value give rise to rational requirements and constraints. I favour a rather more radical accuracy-based epistemology. It might be better named an *accuracy-only* epistemology. It is embodied in Veritism. On this view, the only constraints that evidence can place on credence functions come from considerations of accuracy, together with decision-theoretic principles.[2]

This is extremely ambitious language, and while it is certainly a significant accomplishment to show how considerations of accuracy might be used to justify several central formal constraints on credence, that is far short of a reduction of all evidential constraints on credences to considerations of accuracy.

First, the putative constraints imposed on credences by considerations having to do with (for example) simplicity, coherence, knowledge, understanding, explanation, ad hocness, and the epistemic virtues all have to be either rejected or accounted for in terms of accuracy. And while Pettigrew is explicit that he "make[s] no claim that other doxastic states — full beliefs, imprecise credences, etc. — should be evaluated by the lights of veritism... [or] ... that other subjects of epistemic evaluation — belief-forming processes, rules of inference, collective doxastic states of groups, institutions — should be evaluated in that way," he is committed to the view that the relevance of any of those matters *to credences* can be accounted for in accuracy-only terms. So, for example, if our *knowledge* places any constraints on our credences, or if an *inference to the best explanation* can ever mandate a credal change in any proposition, then those too would need to be accounted for in terms of accuracy alone (together with decision theoretic principles).

Second, there are some very serious reasons to be doubtful about the possibility of accounting for some epistemic phenomena in accuracy-only terms. For instance, Selim Berker (building on work by Roderick Firth and Richard Fumerton) has argued in his 2013a and 2013b that veritist views can require us to form a belief in a particular proposition that goes *against* the evidence in cases where such a belief would lead to the accumulation of true beliefs (and the avoidance of false ones) in lots of *other* propositions. Berker uses the example of a person who has considered the question of whether God exists and has concluded that He doesn't, but is presented with a grant opportunity from a religious organization that will fund only theists. If we assume that any attempt to deceive the organization will be detected and that the grant will result in new true beliefs (and the avoidance of false beliefs) on a variety of matters, veritism looks to entail that it would be rational under these circumstances to form a belief in God's existence. This is a highly counterintuitive result; Berker puts the point by saying that veritism fails to respect the "epistemic separateness of propositions" by permitting trade-offs of the sort countenanced above. Pettigrew addresses this sort of worry very briefly, at the end of the Chapter 16 ("Where next for epistemic utility theory?"), but his response refers only to work by Konek and Levinstein on which dominance principle to use when adopting a credal state that affects the truth values of the propositions to which the state assigns credence. That's an interesting question, but I don't see how it's relevant to Berker's problem; Berker's cases aren't ones where the agent's credal states impacts the *truth values* of other propositions, but are rather cases where the agent's credal state impacts his *opportunity to form accurate credences* in other propositions. Of course, it's perfectly fair for Pettigrew to leave

---

important questions for future work, but it is a bit disappointing that such a central challenge to Pettigrew's epistemological project is addressed so quickly in the final pages of the book.

References

Berker, S. (2013a), "The Rejection of Epistemic Consequentialism", *Philosophical Issues* 23(1), pp. 363–387.

—— (2013b). "Epistemic Teleology and the Separateness of Propositions", *Philosophical Review* 122(3), pp. 337–393.

Greaves, H. and Wallace, D. (2006). "Justifying conditionalization: Conditionalization maximizes expected epistemic utility", *Mind* 115(459), pp. 607–632.

Joyce, J. (1998). "A nonpragmatic vindication of probabilism", *Philosophy of Science* 65 (4), pp. 575–603.

—— (2009). Accuracy & Coherence: Prospects for an Alethic Epistemology of Partial Belief. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*, pages 263–297. Springer.

Maher, P. (2002). "Joyce's Argument for Probabilism", *Philosophy of Science* 69(1), pp. 73–81.

van Fraassen, B. C. (1989). *Laws and Symmetry*. Oxford: Oxford University Press.

von Mises, R. (1957). *Probability, Statistics & Truth*. New York: MacMillan, revised English edition.