Chance Recalibrated

Michael Townsen Hicks

Abstract

Work on chance has, for some time, focused on the *normative* nature of chance: the way in which objective chances constrain what partial beliefs, or credences, we *ought to* have.

According to me, an agent is an expert if and only if their credences are maximally accurate; they are an analyst expert with respect to a body of evidence if and only if their credences are maximally accurate conditional on that body of evidence. I argue that the chances are maximally accurate conditional on *local, intrinsic* information. This matches nicely with a requirement that Schaffer (2003, 2007) places on chances, called at different times (and in different forms) the *Stable Chance Principle* and the *Intrinsicness Requirement*. I call my account the Accuracy-Stability account.

I then show how the Accuracy-Stability account underlies some arguments for the New Principle, and show how it revives a version of Van Fraassen's calibrationist approach. But two new problems arise: first, the Accuracy-Stability account risks collapsing into simple frequentism. But simple frequentism is a bad view. I argue that the same reasoning which motivates the Stability requirement motivates a *continuity* requirement, which avoids at least some of the problems of frequentism. I conclude by considering an argument from Briggs (2009) that Humean chances aren't fit to be analyst experts; I argue that the Accuracy-Stability account overcomes Briggs' difficulties.

Introduction

Objective chances determine what partial beliefs, or credences, we ought to have. If I am playing a dice game using a fair die, I should expect each side to come up with the same degree of confidence; this is because they are each equally likely, in an objective sense, to be facing upwards after a roll. If I am playing poker and drawing for an inside straight, I should be considerably more confident than not that I will fail, because if the deck was shuffled in a sufficiently random way the objective chance of success is quite low: $\frac{4}{remaining\ cards}$.

Precisely characterizing and justifying the link between objective chances and partial belief is a philosophical project that has met with failure on all fronts: there is no agreement about what exactly the correct chance-credence linking principle is, there is no fully satisfying normative justification for any of the proposed chance-credence principles, and there is no widely accepted metaphysical account of what objective chances are–in fact, all of the extant accounts face apparently insurmountable objections. In this paper, I aim to make progress on all three fronts. I start by briefly reviewing the discussion surrounding the Principal Principle and the New Principle.

Then, I take up a new proposal: following Hicks (2017), I argue that to be an expert is to be maximally conditionally accurate with respect to some subject matter. An accuracy measure represents how close–or far–a credence function is from the truth. I take up the intuitive proposal that to be an expert with respect to some subject matter is to be accurate with respect to that subject matter; to be an analyst expert with respect to some body of evidence is to have accurate credences conditional on that evidence. The question then is, what sort of evidence are the chances suited to analyze–what inputs do they give us output information about? I argue that the chances are maximally accurate conditional on *local, intrinsic* information. This departs from Hicks's view and matches nicely with a requirement that Schaffer (2004, 2007) places on chances, called at different times (and in different forms) the *Stable Chance Principle* and the *Intrinsicness Requirement.* I call my account the Accuracy-Stability account.

I then show how the Accuracy-Stability account underlies some arguments for the New Principle. First, I'll argue that there is a consistency-based argument for the New Principle, given the fact that the chances are stipulated to be maximally accurate. Second, I'll review van Fraassen's (1983) calibrationist approach to chance. I'll show how a calibrationalist norm motivates the New Principle on the Accuracy-Stability account.

If the Accuracy-Stability Account is correct, two new problems arise and one old problem remains. I'll conclude the paper by responding to these problems. First, the accuracy-stability account risks collapsing into simple frequentism. But simple frequentism is a bad view. I argue that the same reasoning which motivates the Stability requirement motivates a *continuity* requirement, which avoids at least some of the problems of simple frequentism. Second, the Accuracy-Stability Account looks suspiciously like an indifference-based view. But indifference principles are also bad. I argue that despite appearances my arguments do not appeal to any hidden indifference principle. I conclude by considering a problem for the older Humean view that seems to remain a problem for the Accuracy-Stability Account: Briggs (2009) argues that Humean chances aren't fit to be analyst experts; I argue that the Accuracy-Stability account overcomes Briggs' difficulties.

1 Setting the Stage: the Principal Principle and the New Principle

Since the majority of this paper will be focused on the connection between credence and chances, it's worth our time to more formally lay out what that connection is. Lewis (1980) argued that credences should match the chances in a form given by the Principal Principle:

> Principal Principle: If $b(*|*)$ is a reasonable initial credence function, T is the true objective chance theory, A is any proposition, E is a proposition admissible with respect to A, and $P_T(A)$ is the probability T assigns A, then
>
> $$b(A|E\&T) = P_T(A)$$

The Principal Principle tells us to match our credences to the objective chances, and tells us that information about the chances screens off other evidence, provided that evidence is *admissible*; Lewis held that any information about the chance function and the past is admissible. If we know what the chances are, and are reasonable, we will have credences which match the chances.

On any reasonable chance function, the outcomes of different chance setups are probabilistically independent. So the outcome of my tossing this coin now and your tossing that coin tomorrow don't depend probabilistically on one another. This implies that the chance function allows any sequence and the chance of every coin ever tossed, in the past and future, coming up heads is small, but it is nonzero.

Unfortunately, so long as the chance theory is admissible, theories of chance on which chances are grounded in the non-chancy facts of the world, facts about things like symmetries and frequencies, are in trouble. This is because if the chances metaphysically depend on the frequencies, single-case coin flips cannot depart too far from the total coin-flip frequency. So if some chance-theory assigns a chance of 0.5 to the coin coming up heads, then the all-heads sequence is metaphysically incompatible with that theory being the chance theory. And–here's how we get an explicit contradiction–we should give credence zero to anything that's metaphysically impossible. And the all heads sequence is not the only "undermining" sequence of flips: any total sequence which is very far from the actual chances would have produced another chance theory, and so is metaphysically incompatible with the actual chances.

The contradiction then is this: conditional on chance 0.5 for heads, the Principal Principle instructs us to have a small, but nonzero, credence in any finite string of heads. But because those sequences are metaphysically incompatible with the chance being 0.5, our conditional credence in such a sequence should be precisely zero. Either we give up the Principal Principle or we think the impossible might happen. This contradictory advice is called 'The Big Bad Bug'.

This has lead a number of authors to replace the Principal Principle with the New Principle (Tau (1994), Lewis (1994), Hall (1994)). The New Principle looks like this:

> New Principle: If $b(*|*)$ is a reasonable initial credence function, T is the true objective chance theory, A is any proposition, E is any proposition (both in the domain of the chance theory), and $P_T(A|E)$ is the probability T gives A conditional on E, then
>
> $$b(A|E\&T) = P_T(A|E)$$

There are a few things to note about the New Principle. The first is that there is no mention of admissibility, nor is there a time-index on the chances. With the Principal Principle, we had to rule out, say, future information because if you knew that tomorrow the a flipped coin would land heads, you would be crazy to have a credence of 0.5 in its landing heads because that was its chance today. The PP got out of this by holding that the PP did not apply when you have information about the future. The NP assumes that the objective chance of heads *given heads* is one, not 0.5. So the NP tells you, correctly, to have a credence of 1 in the coin coming up heads, given that it will come up heads. This feature is precisely how the NP gets out of the Big Bad Bug. Presumably, the chances are no more willing than you are to have nonzero credence in impossible propositions. So, conditional on the information that T is the chance theory, the probability of an arbitrarily long all-heads run is zero, allowing you to obey NP and rule out impossible futures.

This feature, though, looks weird to many philosophers. Since $P_T(AllHeads|T) = 0 \neq P_T(AllHeads)$, the chance of some particular events change in a surprising way when they–and you–learn that they are the chances. It seems odd that (a) the chances can have this sort of self-knowledge[1], and (b) that the chance of some particular event will change on the basis of it. I'll say a little about this in §4.3. In the meantime, keep the New Principle in the back of your head as an expression of the claim that our credences should match the chances.

2 The Accuracy-Stability Account

The account of chance I give here mirrors Lewis's Best System Account (1980, 1994) of laws and chances, but is different in crucial ways. On Lewis's account, the chances aim to *fit* the world by making the actual history of the world very probable. This is the *matching condition*: it tells us how the chances match the world. Of course, they could fit the world maximally well by assigning it probability one; if they did, then the chance of every outcome would either be zero or one. To rein the chances back from these extremes, the chances have to feature in a *simple* set of axioms, which includes any deterministic laws the world might have. Because a system which assigned different outcome-chances to otherwise identical physical interactions would be horrendously complex, the chances must be non-maximal if they are to be simple. Simplicity acts as a limiting factor on the fit of a chance theory.

The view I defend has a similar structure: the chances are determined by a correctness condition, which evaluates how well they match the world, and

[1] One way of expressing this first worry is to just claim that the chances do not have this sort of proposition in their domain. But remember that according to reductionist theories of chance, the chance theory supervenes on the actual sequence. Taking advantage of this, many presentations of the Bug substitute some actual sequence for the chance theory; surely the chance function has the sequence as part of its domain. Of course to make the argument work, there must also be a "that's all" clause, saying that no more e.g. coins are ever flipped; it's harder to see the "that's all" proposition as in the domain of the theory, but we won't look too deeply at that now.

are prevented from being maximal by a limiting factor. However the view I defend departs from Lewisian orthodoxy in important respects. First, both the matching condition and the limiting factor are different on my view. Both, I argue, are more closely connected to norms which already govern partial belief: the matching condition, on my view, is the same as the correctness condition for a partial belief. And the limiting factor is not based on simplicity, but instead arises from the epistemic limitations of agents embedded in the world and learning from it.[2] Finally, I do not require that the chances fit into a system of laws and chances, although they are compatible with some views of laws: the limiting factor I defend here works well with the Humean approaches of Dorst (2018)and Jaag and Loew (2018).

The matching condition here is inspired by Joyce (1998, 2009)'s claim that partial beliefs are estimates of truth. Unlike guesses, estimates are better when they are closer to being correct: so for example, if I estimate the number of jelly beans in a jar, my estimate is better if I'm within ten jelly beans of the correct answer than if I'm only within a few hundred. Joyce argues that taking partial beliefs to be estimates of truth gives us reason to introduce a measure of their distance from truth: the closer, the better. The measure Joyce introduces, which has been defended by Leitgeb and Pettigrew (2010a, 2010b) and Pettigrew (2016) among others, is the Brier Score:

> Let b be a credence function defined on a set of propositions $\mathcal{F}$ and w is a possible world relative to $\mathcal{F}$. Let v_w be the *truth-function at* w, where $v_w(A) = 1$ if and only if A is true at w and $v_w(A) = 0$ otherwise. Then
>
> $$I(b, w) = \Sigma_{A \in \mathcal{F}}(b(A) - v_w(A))^2$$
>
> is the *Brier Inaccuracy* of b at w.

The Brier score measures how far a credence function is from truth (or falsity); the lower a partial belief's Brier innaccuracy, the closer it is to the truth: the more accurate it is. Joyce, Leitgeb, and Pettigrew argue that just as full beliefs aim for truth, partial beliefs aim to be close to the truth, as measured by this score (for a thorough defence of Brier-inaccuracy, see Pettigrew 2016).

I hold that the Brier scores should also be used to measure how well a chance-function matches the actual world. My reasoning goes by way of analogy: just as propositions match the world by being true or false, chances match the world by being close (or far) from the truth function. Accuracy is constitutive of correctness for partial belief and chance functions, just as truth is constitutive of correctness for full belief and propositions.

Of course, the Brier score is minimized by the credence function that assigns 1 to each truth and 0 to each falsehood. But it also has features in common with

[2] This gives my view the tools to respond to criticisms like that of Sturgeon (1998: 332): "Once simplicity and strength are objectified, it's genuinely unclear that knowledge of them constrains rational credence. It's unclear, for instance, that systems employing less quantifiers are more rational than those employing more..."

van Fraassen and Shimony's notion of calibration (see §3.2): as will become important shortly, if a credence function assigns the same value to a number of distinct events, its Brier innaccuracy is lowest if that value matches the frequency of truths within that set of events.

In non-deterministic worlds, the chances are not maximal: they do not assign 1 (or zero) to each truth (or falsehood). So the chances do not minimize inaccuracy, and so do not fully satisfy the matching condition. Why not? This is where the limiting factor kicks in. Chances are meant to be used by agents embedded in the world; such agents have to learn the chances by observing relatively local subsystems of the world; they have to determine the chance of future events happening based on the observable features of their locality, or of the present.

Consequently, the chance function must assign equal probabilities to future events conditional on the same local, qualitative features of the present. This is also suggested by the chance's role as an analyst expert: analysts are good at making estimates based on a specific kind of evidence: a meteorological expert makes accurate predictions about the weather based on evidence concerning measurably atmospheric conditions, like windspeed and pressure; an economic expert makes accurate predictions about the economy based on evidence concerning measureable economic factors, such as unemployment and asset liquidity. The fact that these experts make their predictions based on a certain sort of evidence is part of what makes them analyst experts, and it requires them to make the same estimations about future behavior given the same evidence-based input. A good metereological expert will give the same estimate of hurricane behavior if the atmospheric conditions are the same; a good economic expert will give the same estimates for market behavior if the economic conditions are the same. This gives rise to a limiting factor: the chances must assign the same outcome probabilities to situations with the same qualitative character.

Principles of this sort have been proposed under a few names, most recently by Jonathan Schaffer. Schaffer defends both a Stable Chance Principle (SCP) (Schaffer 2003) and an intrinsicness requirement (IR) (Schaffer 2007). These principles are similar, but I will focus on the SCP. Here's the SCP: "Consider the STABLE TRIAL PRINCIPLE (STP): If (i) A concerns the outcome of an experimental setup E at t, and (ii) B concerns the same outcome of a perfect repetition setup E at a later time t', then $P_{tw}(A) = x = P_{t'w}(B)$." (Schaffer 2003: 37).

Inspired by the SCP and the IR, I claim that chances must be *stable*:

> STABILITY: For all propositions describing outcome events *A, A'*, and all propositions describing chance setups E, and E*, if there are no qualitative differences between A and *A'*, or between E and E*, then $ch(A|E) = ch(A^*|E^*)$.

One thing I'd like to note here is that the Stability requirement is not an indifference principle, nor is it obviously equivalent to any indifference principle. An indifference principle requires agents to assign equal credence to a partition of mutually incompatible possibilities based on the agents' ignorance of which is actual, and the argument for indifference principles is often presaged on the (in my opinion, false) notion that this equal credence is representative of or required by the ignorance of the agents. Stability is not like this: a credence function which obeys Stability is ignorant of which outcome will occur in a particular situation (that is, it does not assign them zero or one), but it does not react to that ignorance by assigning equal credence to each possible outcome. In fact, given the accuracy requirement on chances, the chances typically will assign some outcomes a much higher chance than others; it is because they do this that they are worth thinking of as an analyst expert. Nonetheless authors regularly conflate principles like Stability with indifference principles; I'll discuss the difference in greater detail in §4.2.

Another thing I'd like to say about the stability requirement is motivational: objective chances are a feature of the world, and one which we discover empirically. Empirical discovery only works if we can observe the behavior of a variety of systems which behave in the same way. Consequently, some level of similarity and uniformity between systems is a precondition for us to empirically discover the chances. I think that the Stability requirement states one way in which the chance function marks out uniform behavior between distinct systems, and so meeting the Stability requirement, or something like it (see §4.1), is a precondition for inductive discovery. Since if we are to know about objective chances at all, we must learn about them through induction, the chances must be Stable for us to know about them. Making Stability a requirement for chance-hood makes this Humean account of chances fit well with Humean accounts of law: Humean accounts of law build uniformity into lawhood, in part to make them fit with inductive practice; here we've similarly built it into chancehood.

This then is the account I favor, and the one which I think is best positioned to motivate the New Principle as an expert function. With a bit of ominous foreshadowing, I'll here refer to it as the Naïve Accuracy-Stability Account:

> Naïve Accuracy-Stability Account (NASA): The chances are the least Brier-inaccurate Stable probability function.

If you're looking for a more sophisticated view, you'll have to wait for §4, where I slightly modify things to respond to some objections. In the meantime here are two arguments–not for this view of chances, but instead for credence-chance principles given this view of chances.

3 Two Arguments

In this section I will argue that the Accuracy-Stability Account is well-suited to underwrite an argument for the New Principle. Rather than arguing directly for any chance-credence matching principle, I will argue that our credences should aim for the chances. This general claim can then be used to support a specific chance-credence matching principle when supplemented by a richer formal framework for partial belief. My own favorite here is Pettigrew's (2012) accuracy-based argument for the Principal Principle, which can be slightly modified to support the New Principle. I argue that, by building accuracy directly into the account, the Accuracy-Stability Account is best suited to underwrite a key premise in Pettigrew's account, namely, that credence should aim at chance[3].

The first of my arguments is based on arguments given in Howson & Urbach (1993) and Hoefer (forthcoming) and has two parts: first, it has a characterization of partial belief, or credence. Then, I argue that, given the account of chance under discussion, knowingly having credences which depart from the chances involves a sort of internal incoherence.[4].

The second argument is slightly more complicated, and builds on van Fraassen (1983) and Shimony's (1988) notion of *calibration*. First, in §3.2, I'll describe their account, and explain both what is attractive about it and why it is not widely held today. Then, I'll argue that the NASA captures the intuitive plausibility of their account and solves its problems.

Before digging into this section, its worth discussing the value of an argument for the New Principle. Showing that, given some specific normative assumptions, a particular metaphysical account of chance can vindicate NP shows that, for agents normatively constrained by those assumptions, the structure mentioned in the account plays the role of chance. I take this to be the primary function of a proof of a chance-credence principle: these arguments are not attempts to show from *a priori* assumptions that disobeying the principle is irrational; they are instead meant to show what other normative notions the principal depends upon.

Some normative chance-credence principle is often presupposed in formal (and traditional) epistemology; but without an account of chance, and connected proof, epistemologists cannot evaluate how that principle coheres with other epistemic norms. For example, if the proof of NP requires an indifference principle, accepting NP without accepting indifference in other areas is unmotivated; if the proof of NP rests on some *a posteriori* facts, for example a posit that future events will roughly match some frequency, then adding NP to a set of purely internalist formal norms undercuts an internalist position on epistemic rationality.

3.1 The Argument from Consistency

In this section I will argue that knowing the chances but failing to match one's credence to them gives rise to a certain form of internal tensions in one's partial beliefs. The inconsistency is not a strict contradiction; rather, if you know what the chances are but have credences that depart from them, I will argue, you are knowingly violating the constitutive norms governing credence.

It's worthwhile examining a similar case involving full belief. If you assert "it's raining, but I don't believe it" or "I believe P, but P is false", you haven't

[3]Note that Pettigrew gives slightly different arguments for PP and NP in his (2013) and (2016). The differences aren't important for us here, but I do think the reasoning below best underwrites the key premise of the argument in the earliest paper (2012).

[4]This argument is also suggested by Loewer (2004:1123).

asserted any logical contradiction; in fact it's perfectly normal to have false beliefs. But someone who makes these assertions is knowingly violating the constitutive norm governing belief: beliefs aim at truth. To believe something is to take it to be true, and anyone making the assertion takes themselves to believe something false.

Similarly, the norm governing partial belief is accuracy. We should accept this because partial beliefs are beliefs; beliefs of all kinds aim to be true, and accuracy measures the truth of partial beliefs. Partial beliefs are estimates of truth. What, then, is a conditional credence? It's the estimate of the truth of one proposition on the basis of, or estimated given, another. We can express this counterfactually: $P(A|B)$ is the estimate we *would* give of the truth of A if we learned B.

Now recall what, on my view, the chances are: they are the most accurate stable credence function, where stability requires them to provide the same conditional credence to outcomes of qualitatively identical situations.

Recall also that the stability condition was motivated by the sorts of epistemic constraints we believe ourselves to be subject to. We must make decisions about what to do on the basis of local, qualitative information and our experience of situations similar to the one in which we find ourselves.

Consequently, an agent who knows that what the chances are but fails to match her credences to them is either (a) knowingly diverging from the truth more than she needs to or (b) distinguishing between situations which she has no way to distinguish. But her conditional credence was her estimate of the truth of some proposition on the basis of some hypothetical evidence. This is the sense in which an agent whose credences diverge from the chances is internally inconsistent: she first tells us her estimate of the truth of some proposition based on some evidence, and then she tells us either that she herself believes that there is another estimate closer to the truth (the chances), or that her estimate was based on features of the evidence that she couldn't herself recognize.

If the chances are the most accurate stable credal function, then, agents who don't match their credences to the chances are internally incoherent: they don't meet their own standards for rationality.

3.2 Calibration

Soon, I will present an argument for the NASA account based on the notion of *calibration*. But first, I need to explain what calibration is, and show both what motivated it and why it was abandoned.That's what I'll do in this section. Then, in the next section I'll show how it connects to an argument for the NASA.

Van Fraassen (1983) and Shimony (1988) argue that our credences, to be correct, must be *well-calibrated*: that is, a credence should match the proportion of beliefs which are true. More formally:

> Suppose b is a credence function defined on a set of propositions $\mathcal{F}$ and w is a possible world relative to $\mathcal{F}$. Then b is well calibrated

at w if, for each x in the range of b,

$$x = \frac{|\{X \in \mathcal{F} : b(X) = x \quad and \quad v_w(X) = 1\}|}{|\{X \in \mathcal{F} : b(X) = x\}|}$$

The idea is this: take all of your partial beliefs in which you have the same credence–say, 0.7. Now look at the proportion of those beliefs which are true. If exactly 70% of them are true, you are well calibrated.

Of course, most of us are not well calibrated. But calibrationism allows us to define an epistemic ideal for ourselves; just as in the full belief case, we can define ideal agents as those who are logically consistent, or have only true beliefs, so here we can define an ideal for any particular agent: her *well-calibrated counterpart*(this discussion follows Pettigrew 2016):

> Suppose b is a credence function defined on a set of propositions $\mathcal{F}$ and w is a possible world relative to $\mathcal{F}$. Then b's well calibrated counterpart at w, b^w, is the credence function such that, for all propositions $Z \in \mathcal{F}$
>
> $$b^w(Z) = \frac{|\{X \in \mathcal{F} : b(X) = b(Z) \quad and \quad v_w(X) = 1\}|}{|\{X \in \mathcal{F} : b(X) = b(Z)\}|}$$

Your well-calibrated counterpart agrees with you about what propositions deserve equal credence, but she disagrees about what those credences should be. Her credences, unlike yours, match the actual ratio of truths to falsehoods, and so are well calibrated.

Calibrationism provides correctness conditions for partial belief without falling into the pitfalls which plague some accounts of objective probability. Specifically, calibrationism is meant to recover the intuitive appeal of actual frequentism without succumbing to the reference class problem. Actual frequentism says that the probability of some outcome O is exactly equal to the proportion of events of that type T which have that outcome: $\frac{T}{O}$. Actual frequentism lends itself easily to an argument for matching our credences to the probabilites, because if we act, or bet, based on the actual frequencies we are guaranteed to break even at some point (if we don't die first, of course). But actual frequentism faces a number of apparently insurmountable difficulties. The first is the reference class problem. Because any event can be seen as an outcome of any number of types T of event, the probability of any outcome is hopelessly undefined. The second is the problem of single-case events: according to frequentism, if something happens exactly once, it can only have a probability of 1 or zero. But many things aren't like that: so it seems I could have a credence of .8 in the proposition that Croatia will win the world cup final though I believe that they will never again enter it, and so the actual proportion of their wins must be one or zero.

Calibrationism solves the reference class problem by ignoring it: according to calibrationism, you get to decide which events deserve the same credence, and you need not do so on the basis of similarity or types. It solves the problem of the single-case by allowing us to assign the same credence to dissimilar events, so my confidence in Croatia's win will be correct just in case it is in a group of propositions with the .8 credence of which $\frac{8}{10}$ are true, even if these propositions are all dissimilar.

Unfortunately, calibrationism's strengths are also its weaknesses. By allowing agents to group beliefs in any way they want, it provides too weak of a constraint on partial belief. This can be illustrated by a simple example (from Pettigrew (2016)):

	P	$\neg P$
b	0.5	0.5
b^w	0.5	0.5
b_ϵ	$0.5+\epsilon$	$0.5-\epsilon$
b_ϵ^w	1	0

Where b and b_ϵ are credence functions and b^w, b_ϵ^w are their well calibrated counterparts. The first row reminds us that if an agent suspends belief about all of the propositions she considers and their negations, she is automatically well-calibrated. But this is too easy: surely we are required to do more than just be skeptics. But imagine such an agent changes her beliefs slightly, so that she becomes more confident in a truth and less confident in a falsehood. Intuitively, she is now epistemically better off. But according to calibrationism, she has gone from being perfectly calibrated to being quite far from ideal, just because her ideal has changed.

It's this sort of worry which, I believe, provides the way forward for those of us who thinks that there is something correct about the calibrationist correctness condition for partial belief. We should not arbitrarily group propositions just to ensure that we are well-calibrated; this smacks too much of "teaching to the test." Rather, we should make distinctions when we're capable of it: if it is within our perceptual or cognitive abilities to distinguish one state of affairs from another, we should do so but of course, not otherwise. It's this epistemic constraint which we should add to the calibrationist norm; and of course, this is the constraint which motivated stability requirement. Below, I'll show that stability combined with calibrationism yields a strong argument for chance-credence matching.

3.3 The Argument from Calibration

The account of chance given above provides an interesting link to the notion of Calibration discussed in §3.2. The link goes by way of an interesting feature of the Brier score, which according the Accuracy-Stability account is minimized by the chance function. The Brier score can be decomposed into two terms (see Pettigrew (2016: ch. 4), Joyce (2009), and Murphy (1974)):

Decomposition:

$$I(b,w) = \Sigma_{A\in\mathcal{F}}(b(A) - v_w(A))^2$$

$$= \Sigma_{A\in\mathcal{F}}(b^w(A) - v_w(A))^2 + \Sigma_{A\in\mathcal{F}}(b(A) - b^w(A))^2$$

Where $b^w(*)$ is just the well-calibrated counterpart from §3.2. In other words, the Brier-innaccuracy of a credal function can be divided into the distance between that credal function and its well-calibrated counterpart, and the innaccuracy of the well-calibrated counterpart on its own. I can use this to show that, if stability determines which well-calibrated counterpart an agent ought to have, accuracy will insure that the agent should match her credences to the chances.

How could stability determine which well-calibrated counterpart an agent ought to have? Well, recall that Stability requires the agent to assign equal conditional credence to qualitatively identical outcomes of qualitatively identical chance setups. This requirement does not by itself ensure that two Stable agents will partition the relevant probabilities into the same equal-credence classes: one such agent could hold that two propositions which describe qualitatively distinct outcomes nonetheless are equally likely, while another would disagree; this would lead them to disagree about which outcomes are equally likely despite both being Stable, and so would result in their having distinct well-calibrated counterparts.

The easy patch for the argument here would be to strictly rule this out, and require agents to assign different credence to qualitatively distinguishable outcomes. Call an agent who does this *maximally discerning.*

> Maximally Discerning: A credence function b if maximally discerning if and only if, for all propositions describing outcome events *A, A'*, and all propositions describing chance setups E, and E*, if there are local qualitative differences between A and *A**, or between E and E*, then $b(A|E) \neq \underline{(}A^*|E^*)$.

Such an agent's credences treat propositions as differently as she can, given her epistemic limitations. If she can tell the difference between any two outcomes or any two chance processes, she has different expectations concerning them.

Requiring agents to be maximally discerning would be a mistake. For suppose that two discernible outcome events occur just as often as one another: say, heads and tails outcomes from an unbiased coin flip. A maximally discerning agent would be required to have different credences in these two, and so could not treat heads and tails as equally likely. But it is obviously not irrational to treat heads as just as likely as tails. So maximum discernment is not a requirement of rationality.

Nonetheless it's interesting that, in cases where multiple distinguishable outcomes happen just as frequently as one another, the well-calibrated counterpart of a maximally discerning agent will not be maximally discerning. The well-calibrated counterpart will instead assign equal credence to discernible outcomes if and only if doing so is required for her to be well-calibrated; that is, if and only if doing so maximizes her accuracy (without violating stability). We can call such an agent *Satisfactorily Discerning.*

> Satisfactorily Discerning: A credence function b is satisfactorily discerning if and only if, for all propositions describing outcome events *A, A'*, and all propositions describing chance setups E,

and E^*, if there are no local qualitative differences between A and A^*, or between E and E^*, then $b(A|E) = b(A^*|E^*)$ if and only if $b^m(A|E) = b^w(A^*|E^*)$, where $b^m(*)$ is the well-calibrated counterpart of a *maximally discerning* credence function.

A satisfactorily discerning agent makes all the distinctions allowed by her epistemic limitations except when doing so will require her to depart from calibrationist standards.

It's easy to see that any agent that is Stable and at least Satisfactorily Discerning will have the same well-calibrated counterpart. Moreover, on the Accuracy-Stability Account, this counterpart is guaranteed to be the chance function. This is because the chance function is defined to be the most accurate stable credence function; but given Decomposition, this can only be accomplished if the distance between the chance function and its well calibrated counterpart is minimized. Of course this distance will be minimal if and only if the chance function is its own well-calibrated counterpart, in which case it will be zero.

Similarly, given the argument above, the chance function will be the well-calibrated counterpart of any Stable agent which is at least Satisfactorily Discerning. This is because the chance function will be Satisfactorily Discerning, as a Satisfactorily Discerning agent is one which maximizes accuracy while obeying Stability, which by definition the chance function does. Since every Stable agent which is at least Satisfactorily Discerning has the same well-calibrated counterpart, the chance function will have the same well-calibrated counterpart of every Satisfactorily Discerning Stable agent. Since the chance function is its own well-calibrated counterpart, it is also the well-calibrated counterpart of every Satisfactorily Discerning Stable agent.

The upshot of this is that any Satisfactorily Discerning, Stable, well-calibrated agent automatically matches her credences to the chances, if the chances are those given by the NASA. What's interesting about this the that the calibrationalist norm is an independently motivated correctness condition for credences; if the view I'm defending here is correct, then the chances explain what's right about calibrationism without its less savory consequences.

4 Avoiding the Collapse

In this section I'll consider a number of objections to the NASA view presented in §2. First, I'll consider the objection that the account collapses into naïve frequentism by showing how it can be slightly modified to overcome two arguments against frequentism. Next, I'll argue that the Stability constraint is not an indifference principle, and that reductionists about chance are not committed to indifference. Finally, I'll respond to an argument from Rachel Briggs that reductionist chances are simply not the sorts of things that could ground the New Principle and are better suited for the regular old Principal Principle, with which they are inconsistent.

4.1 The Collapse into Frequentism

The NASA view, as presented in §2, risks collapse into naïve finite frequentism. This is because the Brier score is only defined for credence functions that range over finitely many propositions. Stability requires those propositions to be partitioned into equal-credence sets. There will therefore be a ratio of truths to falsehoods within the set, and the Brier score will be maximized if the credence given for each member of the set matches that ratio. This, of course, is naïve frequentism, a horrible view. For a thorough accounting of the problems with naïve frequentism, see Hájek (1996). Here, I'll consider two: naïve frequentism can only deliver rational chances, but we know from QM that some chances are irrational; and in cases where there is only one event in a class, naïve frequentism can only deliver extremal chances (1 or 0)[5].

Because its possible for there to be qualitatively unique events–in fact, there are probably few or no actual events that are intrinsic duplicates–this second problem is particularly pressing for the NASA defended here. My response to it will expand the NASA into a sophisticated Accuracy-Stability based account, which diverges sufficiently from Naïve frequentism to overcome the first objection without undermining the arguments given in §2.

In §2, I motivated the Stability requirement by appealing to the epistemic state of embedded agents. Such agents must discover the chances via inductive generalization: they do so by observing frequencies of event outcomes in controlled environments where the precise qualitative facts are accessible to them, and extend these discovered frequencies to estimate the results of qualitatively similar situations. I believe that this expresses an important feature of chance: chances are *a posteriori* information which impacts credence, rather than full belief.

However the Stability requirement is a relatively weak requirement given these limitations of embedded agents. For such agents are fallible, and are not able to discern the precise qualitative state of their environment, even in carefully controlled situations. Consequently, their predictive tools must be able to accommodate some error: if an agent is slightly wrong about what qualitative situation she is examining, the predictions she makes must be at most slightly wrong. This motivates a continuity-like requirement on the chances: the chance of outcome events must vary continuously with small changes in the chance setup. The requirement is stronger than a mere continuity requirement: continuity would merely require infinitesimal changes in the chance setup to result in at most infinitesimal changes in outcome chances. But many such changes would still be imperceptible to realistic agents; the barrier at which changes become perceptible is vague, and so I will use vague language to characterize it[6]:

[5]I don't have space here to respond to all of Hájek's problems for frequentism, but that's okay, as the NASA view is already more sophisticated than naïve finite frequentism and so doesn't have all of the problems Hájek attributes to frequentism. For example, the NASA view already has a built-in solution to the reference class problem; the relevant reference class for a proposition is determined by the local qualitative state that proposition represents.

[6]This may strike some as insufficiently objective: if what the chances are depends too

Smoothness: For all propositions describing outcome events A, A', and all propositions describing chance setups E, and E', if there are only small differences between A and A', and between E and E', then $ch(A|E) \approx ch(A'|E')$.

Adding Smoothness[7] to the Nïave Accuracy-Stability Account yields what I call the Sophisticated Accuracy-Stability-Smoothness Account, or SASSA. Like NASA, SASSA is motivated by the epistemic goals and constraints of embedded agents learning the chances inductively and applying them to make predictions about their environment. So the arguments in §3 are undisturbed by the addition of Smoothness. But Smoothness allows us to avoid some of the pitfalls which beset nïave frequentism.

Sophisticated Accuracy-Stability-Smoothness Account (SASSA): The chances are the least Brier-inaccurate Smooth and Stable credence function.

First, given Smoothness, not all chances need be rational. For the best way of achieving high accuracy while minimizing the distance in outcome-chances between very similar situations is to have the chances of outcomes vary continuously as a function of the qualitative quantities of the chance-setup. To achieve this, this function will need to provide outcome chances for a continuum of similar chance-setups.

Similarly, chances for unique events need not be one or zero, if those events are sufficiently similar to other events with different–but similar–outcomes. Consider two very similar but qualitatively distinct situations, E and E' which have different outcomes A and $\neg A'$, where A is very similar to A' (so the outcome of E' is almost–but not quite–the opposite of the outcome of E). Realistic agents may not be able to determine with certainty which sort of situation they are in, or tell the difference between these two situations. So in order or a chance function to be useful to such limited agents, it cannot assign a maximal

closely on agents, then the difference between objective chance and subjective credence might seem too small, and the view might become susceptible to 'ratbag idealism'. But I think this vagueness is an advantage to the view: with some perceptive apparatus (e.g. our eyes) two events might be quite similar, but with another (e.g. a slow-motion film camera) they might be quite different. So an ordinary person in a casino might regard each card as equally likely to come up, but someone with a device allowing them to precisely observe how cards fall in the shuffle could predict with accuracy what the draw will be. In my view this difference is often objective–but context-sensitive–rather than subjective, and leads to a hierarchy of chances linked on the information available to agents, similar to that described by Handfield and Wilson. The stratification into levels comes not from variability in the chance function, but instead from variability in the available evidence propositions E.

[7] Smoothness looks a lot like a continuity condition. As I'll argue below, in the right situations it leads to one. But Smoothness is more closely connected with our epistemic states, and so diverges from a continuity condition in two ways: first, if there are only a few sufficiently similar events, the argument I give later for continuity won't go through. So smoothness is in some respects weaker than continuity. Second, continuity permits chance functions to vary smoothly beneath the level perceptible to agents in such a way that imperceptibly different setups have wildly different chances. I aim to rule out such continuous functions with Smoothness.

value to either. This is why we require that it satisfies Smoothness. If it assigns $ch(A|E) = 1$, then by Smoothness $ch(A'|E') \approx 1$, but since we have stipulated $\neg A'$, this assignment will drag down the accuracy of the chance function as a whole. So instead the chance function will assign $ch(A|E) \approx 0.5 \approx ch(A'|E')$ (of course, the more situations there are which are similar to E and E', the wider the range of values that this can take), despite the fact that we have stipulated that E and E' are one-off chance setups.

This example also illustrates how the chances may take irrational values. For the requirement here is only that $ch(A|E) \approx 0.5 \approx ch(A'|E')$, not that $ch(A|E) = 0.5 = ch(A'|E')$. A chance function which assigns a slightly higher value to A given E than A' given E' will score higher on accuracy than one which assigns the same value to both. So Accuracy will require $ch(A|E) > ch(A'|E')$; the easiest way to satisfy both this and $ch(A|E) \approx ch(A'|E')$ for a wide range of situations is to use a chance function defined over an arbitrarily small changes, that is, the continuum.

I believe that some of Hájek's other criticisms of frequentism can be successfully resisted by the SASSA account. But not all: Hájek worries that frequentism operationalizes chance, and that it doesn't allow chances to explain frequencies– a objection advanced forcefully by Nina Emery (2017). These worries seem to me to be connected to similar worries about the explanatory power of Humean laws, and I think Humeans have responses to them. It is my hope the success (or failure) of Humean responses to the explanatory worries surrounding laws will extend to the present account of chance. See [redacted for anonymity] for my views on this question.

4.2 Indifference?

I have already argued that the Stability principle is not an indifference principle. But similar principles are regularly confused for indifference principles, or, worse, accused of being based on inexplicit indifference reasoning. The idea is this: I have claimed that agents should have the same outcome-credences for situations which they cannot distinguish. This general idea is meant to motivate both Stability and Smoothness. But why should we accept it? Plausibly, we should accept it because if we cannot distinguish two situations, we have no evidence which tells us which we are in. And in absence of evidence differentiating two possibilities we should assign those possibilities equal credence (this is the indifference step). So we should take it to be equally likely that we are in either situation. And given that we take it to be equally likely that we are in either situation, the only way we can maximize our accuracy about outcomes is by matching our credence to the ratio of outcomes of certain types. This seems to be the sort of reasoning that Strevens (1999: 259) is arguing against when he accuses short-run frequentist views of being based on an illicit indifference principle, and it is explicitly embraced by Wolfgang Schwarz's interesting argument for the Principal Principle (Schwarz 2014: 94).

The first thing I'd like to say about this sort of reasoning is that it is not clear that it can lead to Stability. If this argument doesn't lead to Stability, it's hard to see how Stability could be based on a secret indifference principle. Recall that Stability concerns my conditional credences: if A and A' are qualitatively identical, and E and E' are qualitatively identical, Stability requires $ch(A|E) = ch(A'|E')$. But this doesn't follow from the argument above. The argument above claims that if E and E' are indistinguishable, $ch(E) = ch(E')$, and then seeks to show that $ch(A \vee A'|E \vee E')$ should equal the fraction of events qualitatively identical to E and E' that result in events qualitatively indistinguishable from A and A'. Perhaps, with the right background constraints on A and A', this will follow from Stability, but I don't see how it could lead to it.

The second thing worth pointing out is that it's not at all clear whether the indifference principle could really apply here, as you could well know what you're supposedly ignorant of. Schwarz, for example, argues that you can get the Principal Principle out of a frequentist-like view if you are indifferent to where in the relevant chance sequence you are; a similar argument could be drummed up by claiming you should be indifferent to which sequence is actual amongst those sequences with the actual frequency. But I think it's quite clear that you may well want to reason using the chances even if you have information about where you are in the sequence or about which of the equal-frequency sequences is actual. You could, and should, take yourself to know as much as you want about the past part of the sequence, which could tell you where you are (at the *nth* place in the sequence) and lots of information about which sequence is actual (you could know everything about the sequence up to n, for example; you could also know that the sequence doesn't have any recursively-definable patterns).

What's motivating Stability–and Smoothness–isn't your ignorance of where you are or of which sequence is actual; it's a different feature of our shared epistemic situation. We have to use the past to shape our expectations about the future. We do this by taking the past behavior of some systems to be the same, in some ways, as the future behavior of similar systems. This requires the chances to treat similar situations at different times in roughly the same way. That requirement is meant to be embodied by Stability and Smoothness. But that requirement is just an occurrence of the Principle of the Uniformity of Nature (PUN), and as Hume showed, PUN cannot be founded on any *a priori* principle, including the Principle of Indifference. It's a mistake to think that it can be, because it confuses a precondition of *a posteriori* reasoning with an unmotivated *a priori* principle.

To their credit, both Strevens and Schwarz recognize the connection between chance and induction, although they both conflate this with a reliance on a certain sort of indifference principle. Here's Schwarz (2014: 94) "[t]he principle of indifference required here is closely related to the principle of induction." Similarly, Strevens claims that a major theme of his paper is "that there is a resemblance between rules for probability coordination and the rule of enumerative induction and a corresponding resemblance in the difficulties encountered in pursuit of their justification" (Strevens 1999: 248-9); in arguing that most attempts to prove the Principal Principle rely on an indifference principle Strevens seeks to show that they rest on unmotivated *a prioristic* reasoning, just like bad defences of the Principle of Uniformity of Nature.

The aim of the view here is to bypass the indifference principle altogether, and motivate constraints on credences directly from inductive practice. Of course, in order for this to lead to successful predictions, the world must cooperate: the behavior of past and future systems must be uniform enough for us to learn the chances by observing past instances. It is a feature of Humean accounts that though they rely on this uniformity they make no attempt to justify our belief in it, and the Humean account of chances here is no different.

4.3 Briggs' Challenge

Briggs (2009) has argued that reductionist accounts cannot justify the New Principle, rather than the Principal Principle. Recall that the Principal Principle asks us to match our credences to the *unconditional* chances rather than the *conditional* chances, as the New Principal advises. In motivating the move from PP to NP, Ned Hall (2004) argues that PP and NP are both expert principles, but of different sorts. PP tells us how to defer to an expert who has strictly more evidence than we do; NP tells us how to defer to an expert who has no more evidence than we do, but is better at evaluating that evidence. Hall calls the former a "database expert" and the latter an "analyst expert."

Briggs' argument is that, because the reductionist's definition of chance makes the chances encode information about the future, it's wrong to regard reductionist chances as analyst rather than database experts. Here's Briggs: "If HS [reductionism] is true, then in addition to 'knowing' about the past, the chance function 'knows' about the distribution of outcomes among future trials [...] We are justified in trusting the chance function only because it is sensitive to information about the future that we ourselves lack. Thus, although consistent with HS, NP is poorly motivated." (Briggs, 2009: 440). The idea is this: any reductionist chance function will depend both on past and future information. Its dependence on the future means that we trust it because it has information or evidence we do not; but on Hall's way of dividing things up this describes a database, not an analyst, and so should be deferred to via the PP, not the NP. This of course puts the Humean in a bind, as the Big Bad Bug shows that reductionism is inconsistent with PP. For Briggs, reductionism makes Humean chances databases, and the Big Bad Bug shows that they cannot be databases: together, this amounts to a reductio of the view.

I think Briggs is correct to be worried about the future-dependence of the chance function, but that both Hall and Briggs are wrong to think that typical analyst-experts have no more information than we do. Typical analyst-experts are good at evaluating evidence because they have rather general information regarding situations of the sort under consideration; they differ from database experts by not having specific information about the exact situation under consideration.

Let me illustrate this by analogy. I recently visited Baltimore with some British friends. None of us had been there before, so none of us had specific information about the age of houses in Baltimore. Nonetheless, my British friends deferred to me concerning the age of Baltimorean houses because of my experience with qualitatively similar houses in Richmond, Washington, and Philadelphia, and our shared expectation that the houses of Baltimore were similar in design and age to houses in those cities. To be safe, we double-checked our estimates of house age with our host, a Baltimore native, who was directly familiar with the houses and neighborhoods we had been looking at. I claim that in this situation I acted as an analyst expert and our host acted as a database expert.

Despite the fact that I was a mere analyst expert, my expectations of Baltimore were based on some facts which depended on the houses of Baltimore: specifically, my friends and I had to assume that the Baltimorean houses were sufficiently similar to Richmonder, Philadelphian, and Washingtonian houses, and specifically that correlations between their quality and age were pretty uniform amongst these cities. This gave me information about correlations between the qualitative properties of houses and their ages. By contrast, our host had information concerning specific houses and neighbourhoods which did not relate them to any information she had about houses in other cities or to broader correlations between their qualitative properties and ages. It was her possession of specific information unrelated to similar situations or to the qualitative features of the houses that made her a database, rather than an analyst.

The chances, on a reductionist picture, depend on the future in the same way that my knowledge of Baltimorean houses depends on Baltimore. They rest on quite general features of the future, and ways in which it resembles the parts of the world we've already observed–namely, the past. They are not like our host; they do not encode specific information about future events in ways which do not relate those events to similar events in the past or to their qualitative features. And this sort of general information concerning similar situations isn't in conflict with their role as an analyst expert: if it was, there would be few or no analyst experts at all.

5 Concluding Ontological Ramblings

There's a bit of an ontological mystery surrounding reductionist theories of chance. Most reductionists follow Lewis (1980) in taking chances to be connected to laws of nature, and to fit somewhere in the Best System Account of laws (and chances). But the chances, on a reductionist picture, are fundamentally different from the laws. Laws, for a reductionist, are just everyday true universal generalizations which some philosopher has put nicer clothes on. They were there already; the best systematization just differentiates them from other true generalizations. Chances, on this view, are like that except without the true universal generalizations. Lewis's view makes it look like the chances come into being along with the clothes; they're a shirt with nobody wearing it. This makes Lewisian chances much stranger than laws.

Reductionists since Lewis have spent some time looking for something to dress up as the chances. One option is the frequencies, and this is the sort of move favored by Schwarz (2014). But the tack I follow here is closer to what

Loewer suggests in (Loewer 2004: 1122). For Loewer, the chances are like other laws because they provide information about the world. Loewer holds that the sort of information provided by the chances is determined by the Principal Principle: there's really nothing to ask about their nature beyond this characterization of how they link the world to our beliefs: "Fit can be understood as a kind of *informativeness*–the information that probabilistic propositions provide concerning the propositions they attribute probability to. [...] But theses probabilities are informative only to someone who is willing to let them constrain her degrees of belief."

The accuracy-based accounts given here–the NASA and the SASSA–are meant to elucidate what this probabilistic information is. Just as the informational content of propositions is given by their truth conditions, the informational content of a chance function is determined by its accuracy at various worlds. By connecting the chances to truth, we can see what we're dressing up, and how what we're dressing up is like enough to the true universal generalizations we take to be laws.

Now, there are a lot of informative–and so accurate–credence functions, just as there a lot of true universal generalizations. What separates the chances and laws, on the view advocated here, is their compatibility with inductive discovery and their applicability in the sort of local, quasi-isolated situations in which realistic embedded agents find themselves. This makes the reductionist account of chance presented here play well with recent Humean views of laws of nature provided by Dorst (2018), and Jaag and Loew (forthcoming). I think this account brightens what Lewis saw dimly, that such chances can explain why we employ chance-credence linking principles. And I hope its ability to overcome objections, which depends crucially on how we understand the role of inductive discovery, lends support to this general approach to Humean laws and chances.

References

Dorst, Chris (2018). "Toward a Best Predictive System Account of Laws of Nature." *British Journal for the Philosophy of Science.* https://doi.org/10.1093/bjps/axy016.

Briggs, Ray (2009). "The anatomy of the big bad bug." *Noûs* 43 (3):428-449.

Emery, Nina (2017). "A Naturalist's Guide to Objective Chance." *Philosophy of Science* 84 (3):480-499.

Hájek, Alan (1996). "'Mises Redux' – Redux: Fifteen Arguments against Finite Frequentism." *Erkenntnis* 45(2/3): 209-227.

Hall, Ned (1994). "Correcting the Guide to Objective Chance." *Mind* 103 (412): 505-518.

Handfield, Toby & Wilson, Alastair (2014). "Chance and Context." In

Alastair Wilson (ed.), *Chance and Temporal Asymmetry.* Oxford University Press.

Hicks, Michael Townsen (2017). “Making Fit Fit.” *Philosophy of Science* 84 (5):931-943.

Hoefer, Carl (manuscript). *Chance in the World: a Humean Guide to Objective Chance.*

Jaag, Siegfried and Christian Loew (2018). “Making best systems best for us.” *Synthese*. https://doi.org/10.1007/s11229-018-1829-1.

Joyce, James M. (1998). “A nonpragmatic vindication of probabilism.” *Philosophy of Science* 65 (4):575-603.

Joyce, James (2009). “Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief.” In Franz Huber & Christoph Schmidt-Petri (eds.), *Degrees of Belief. Synthese* 263-297.

Murphy, Allan (1973). “A New Vector Partition of the Probability Score.” *Journal of Applied Meteorology* 12: 595-600.

Leitgeb, Hannes & Richard Pettigrew (2010a). “An Objective Justification of Bayesianism I: Measuring Inaccuracy.” *Philosophy of Science* 77 (2):201-235.

Leitgeb, Hannes & Richard Pettigrew (2010b). “An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy.” *Philosophy of Science* 77 (2):236-272.

Lewis, David (1980) “A Subjectivist’s Guide to Objective Chance.” In Richard C. Jeffrey (ed.), *Studies in Inductive Logic and Probability.* University of California Press: 83–132

Lewis, David (1994). “Humean Supervenience Debugged.” *Mind* 103 (412): 473–490.

Loewer, Barry (2004). “David Lewis’s Humean Theory of Objective Chance.” *Philosophy of Science* 71 (5):1115–25.

Pettigrew, Richard (2012). “Accuracy, Chance, and the Principal Principle.” *Philosophical Review* 21:2 241-275.

Pettigrew, Richard (2013). “A New Epistemic Utility Argument for the Principal Principle.” *Episteme* 10 (1):19-35.

Pettigrew, Richard (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Schaffer, Jonathan (2003). “Principled chances.” *British Journal for the Philosophy of Science* 54 (1):27-41.

Schaffer, J. (2007). “Deterministic Chance?” *British Journal for the Philosophy of Science* 58 (2):113-140.

Shimony, Abner (1988). “An Adamite Derivation of the Calculus of Probability.” In J. Fetzer (ed.) *Probability & Causality* Dordrecht: D. Reidel.

Sturgeon, Scott (1998). “Humean Chance: Five Questions for David Lewis.” *Erkenntnis* **49**: 321-335.

Strevens, Michael (1999). “Objective Probability as a Guide to the World.” *Philosophical Studies* 95 (3):243-275.

Schwarz, Wolfgang (2014). “Proving the Principal Principle.” In Alastair Wilson (ed.), *Chance and Temporal Asymmetry*, Oxford: Oxford University Press, 2014: 81-99.

Thau, Michael (1994). “Undermining and Admissibility.” *Mind* 103 (412): 491-504.

van Fraassen, Bas (1983). “Calibration: A Frequency Justification for Personal Probability.” In In R. Cohen & L. Laudan (eds.), *Physics, Philosophy, and Psychoanalysis*. D. Reidel.