CALIBRATION AND PROBABILISM

MICHAEL CAIE

Syracuse University

In this paper, I consider an argument due to Bas van Fraassen that attempts to show that considerations of calibration can justify the claim that a rational agent ought to have probabilistically coherent credences. I develop a case that shows that this argument fails. I argue, further, that if a rational agent ought to have credences that are as close as possible to relative frequencies, then there are situations in which an agent ought to have probabilistically incoherent credences.

“The one paradigm rule for the reasonableness of judgment . . . is to see whether the axioms of probability are not violated. Let the frequentist either justify this rule or show why it should be rejected or restricted.” van Fraassen (1983: 299)

BELIEFS, we will assume, come in degrees. As a shorthand, we will refer to these graded doxastic attitudes as *credences*, and to the totality of an agent’s credences as their *credal state*. Granting this assumption, a natural question is: What general synchronic normative constraints, if any, are there on an agent’s credal state? Many have thought that the following provides, at the very least, part of the answer to this question:

Probabilism: A rational agent’s credences ought to be probabilistically coherent.¹

Contact: Michael Caie <caiemike@gmail.com>

1. Let $Cr(\cdot)$ be a function representing an agent’s credences. We say that $Cr(\cdot)$ is *probabilistically coherent* just in case it satisfies the following constraints:

Normalization: For any logical truth T , $Cr(T) = 1$

Non-Negativity: For any proposition ϕ , $0 \leq Cr(\phi)$

Finite Additivity: If ϕ and ψ are logically incompatible propositions, then $Cr(\phi \vee \psi) = Cr(\phi) + Cr(\psi)$

A number of putative justifications for this principle have been offered. In this paper, I will consider an argument for a restricted version of Probabilism provided in van Fraassen (1983).²

Call *credal frequentism* the view that credences constitutively aim at being close to *relative frequencies*. According to van Fraassen, credal frequentism can be used to provide a justification for, at the very least, a restricted version of Probabilism.

One may, I think, be reasonably skeptical of the claim that credences constitutively aim at being close to relative frequencies.³ However, for present purposes, I will grant this assumption. One reason for doing so is the following. In Caie (2013), I argue that the accuracy-dominance argument, presented in Joyce (1998) and Joyce (2009), as well as the Dutch-Book argument, presented in Ramsey (1931), both fail to provide adequate justifications for Probabilism. Indeed, I argue that the normative considerations that motivate the accuracy-dominance and the Dutch-Book arguments in fact support the claim that, in certain cases, an agent may be rationally required to have probabilistically incoherent credences. Thus, the normative considerations that motivate the accuracy-dominance and Dutch-Book arguments in fact provide us with positive reason to reject Probabilism. Assuming that the arguments presented in Caie (2013) are persuasive, a proponent of Probabilism, then, has reason to look elsewhere for a justification of this putative norm. And a natural candidate, once we have excluded the accuracy-dominance and Dutch-Book arguments, is van Fraassen's argument.

In what follows, however, I will argue that, even granting the assumption that credences constitutively aim at being close to relative frequencies, van Fraassen's argument for a restricted version of Probabilism fails. Indeed, I will argue that it is a consequence of this credal frequentist claim that there are cases (which are in the domain of van Fraassen's restricted version of Probabilism) in which one ought to have credences that are probabilistically incoherent. The sorts of cases that make trouble for van Fraassen's argument are those in which the truth-value of a certain proposition depends in a certain manner on an agent's credence in that very proposition. The conclusion to be drawn is that if one endorses the claim that credences constitutively aim at being close to relative frequencies, then one should, in fact, reject Probabilism.

2. A similar putative justification for Probabilism is provided in Shimony (1988). The points that follow apply to both van Fraassen and Shimony's arguments. However, in order to streamline the discussion, I will focus on the argument as it is presented by van Fraassen.

3. See, e.g., Joyce (1998) for an expression of such skepticism.

1. The Calibration Argument

According to van Fraassen, we can think of the Dutch-Book argument as having the following structure. First, there is a claim about what it is for an agent to have a credal state that is *vindicated*. According to van Fraassen, the proponent of the Dutch-Book argument will maintain:

Betting Vindication: An agent has a vindicated credal state just in case her credal state does not sanction as fair a losing set of bets.⁴

Second, there is an appeal to the following norm:

Possible Vindication: It is irrational for an agent to have a credal state, if it can be determined *a priori* that if the agent has that credal state, then she will fail to have a vindicated credal state.⁵

Given Betting Vindication and Possible Vindication, we can provide an argument for Probabilism, then, if we can establish, for any probabilistically incoherent credal state and any agent *S*, that it is *a priori* that if *S* has such a probabilistically incoherent credal state, then she will sanction as fair a losing set of bets.

Call a *Dutch-Book* a set of simultaneous bets on a set of propositions Γ that, for any possible truth-value distribution over Γ , has negative value. We can show that the following holds:

Dutch-Book Theorem: An agent will sanction as fair a Dutch-Book just in case her credences are probabilistically incoherent.⁶

The Dutch-Book Theorem, then, establishes that, for any probabilistically incoherent credal state and any agent *S*, it is *a priori* that if *S* has such a probabilistically incoherent credal state, then she will sanction as fair a losing set of bets. And so, given Betting Vindication and Possible Vindication, we have an argument for Probabilism.⁷

4. Thus: "The Bayesian equates a probabilistic expression of opinion with an announcement of betting odds the person is willing to accept. Vindication consists clearly in gaining, or at least not losing, as a consequence of such bets." van Fraassen (1983: 297)

5. Thus: "Looking again at the parallel of practical and moral decisions, we see one minimal criterion of reasonableness that connects it with vindication. A decision is unreasonable if vindication is *a priori* precluded." van Fraassen (1983: 297)

6. For an elegant proof of this see Paris (2001).

7. Let me briefly note why I think that this argument ultimately fails. While the Dutch-Book Theorem establishes that, for any agent *S* and any probabilistically incoherent credal state, it is *a priori* that if *S* has such a credal state, then she will sanction as fair a losing set of bets, it does not es-

Van Fraassen's strategy is to provide an alternative argument for a restricted version of Probabilism, one that shares the same structure as the Dutch-Book argument, so construed, but that relies on a different notion of credal vindication.⁸ Thus, while van Fraassen accepts Possible Vindication, he rejects Betting Vindication.

How should a credal frequentist understand the notion of credal vindication?

One natural thought is that an agent's credal state should count as vindicated just in case, for every proposition ϕ over which it is defined, the agent's credence in ϕ perfectly matches the frequency of truths for propositions appropriately like ϕ .

There is, however, a serious worry for this way of trying to understand vindication. For, while we have not said what it is for some proposition to be appropriately like ϕ , given any reasonable way of making this notion precise, there could in principle be an infinite number of propositions in this class. But if there are an infinite number of propositions appropriately like ϕ , then the frequency of truths amongst this class need not be well-defined.

In response to this worry, it is natural to focus instead on *finite* classes of propositions over which the credal state is defined. One may, for example, be tempted to say that an agent's credal state should count as vindicated just in case, for every finite class of propositions Q over which it is defined and every $\phi \in Q$, the agent's credence in ϕ matches the frequency of truths that are appropriately like ϕ in Q .

But a little thought shows that this cannot possibly work. For, in concert with Possible Vindication, this would rule out as irrational any credal state that failed to assign to every proposition ϕ credences of either 1 or 0. For consider the set $\{\phi\}$. The frequency of truths amongst the class of propositions that are appropriately like ϕ in $\{\phi\}$ will be either 1 or 0 depending on whether ϕ is true or false. But, then, it is a priori that one's credence cannot match the frequency of truths in this class if one does not assign to ϕ credence 1 or 0.

To deal with this problem, one may say, instead, that an agent's credal state should count as vindicated just in case, for every finite class of propositions Q over which it is defined and every $\epsilon > 0$, there is some $Q' \supseteq Q$ and some suitably constrained extension of the agent's credal state, defined over Q' , such that, for every

establish that *only* probabilistically incoherent credal states have this property. In Caie (2013), I argue that there are cases in which, for some agent S, every credal state defined over some algebra $\mathcal{A}$ is such that it is a priori that if S has that credal state, then she will sanction as fair a losing set of bets. And so, Betting Vindication and Possible Vindication lead to the claim that if this agent has a credal state defined over $\mathcal{A}$, then she will be irrational no matter what credal state she has. Betting Vindication and Possible Vindication, then, are incompatible with a plausible ought-implies-can principle that tells us that for any algebra of propositions an agent should always have some rational credal state available to her. This, I think, gives us sufficient reason to reject at least one of Betting Vindication and Possible Vindication.

8. Thus: "[T]he possibility of vindication is taken as a requirement of reasonableness. This general insight and strategy are open to all contestants. Let the frequentist equate probabilistic expression of opinion with something else; and let him investigate the conditions under which such vindication is not *a priori* excluded." van Fraassen (1983: 297)

$\phi \in Q'$, were the agent to have the extended credal state, the difference between her credence in ϕ and the frequency of truths amongst the class of propositions that are appropriately like ϕ in Q' would be less than ϵ .

Vindication for an agent's credal state, then, does not consist in matching the actual frequencies, since this may be precluded for a variety of reasons that do not impeach the rationality of that state. Instead, vindication for an agent's credal state consists in it being such that there are natural extensions of that state that can be made to arbitrarily approach certain frequencies.

That, at least, is the intuitive gloss. Let us now try to make these ideas more precise. To do so, we first need to lay down a number of definitions.

We will let $\mathcal{D}$ be a set of individuals and $\mathcal{F}$ an algebra of monadic properties, i.e., a class of monadic properties that is closed under conjunction, disjunction, and negation. We will let X be the class of propositions of the form $A(x)$ for $x \in \mathcal{D}$ and $A \in \mathcal{F}$. Throughout, Q and Q' will be finite sets. A model M will assign truth-values to the members of X in the standard way. We will let $C(\cdot)$ rigidly denote a function mapping members of X to values in $\mathbb{R}$. We will call these *credal functions*. And we will let $Cr_s(\cdot)$ serve as an abbreviation for S 's *credal state*, $Cr_s(\phi)$ as an abbreviation for S 's *credence in ϕ* and $Cr_s(\cdot) = C(\cdot)$ as an abbreviation for *For every ϕ , S 's credence in ϕ is $C(\phi)$* . Note that while $C(\cdot)$ rigidly denotes a certain function from X to $\mathbb{R}$, $Cr_s(\cdot)$ does not. Thus, even if $Cr_s(\cdot) = C(\cdot)$, this need not hold of necessity.

Def. We say that the *frequency of truths in Q* is the ratio of truths in Q to the total number of propositions in Q .

Def. We say that the *frequency of truths in Q given M* is the ratio of truths in Q given M to the total number of propositions in Q .

Def. Let $Q \subseteq X$. For each proposition $A(x) \in Q$, we say that the *reference class for $A(x)$ in Q given $C(\cdot)$* is the set of propositions $A(z) \in Q$ such that $C(A(x)) = C(A(z))$.

Def. We say that $C(\cdot)$ is *calibrated to within $\epsilon > 0$ over Q* just in case, for every $\phi \in Q$, the absolute difference between $C(\phi)$ and the frequency of truths in the reference class for ϕ in Q is less than or equal to ϵ . If $Cr_s(\cdot) = C(\cdot)$, then we say that $Cr_s(\cdot)$ is calibrated to within $\epsilon > 0$ over Q just in case $C(\cdot)$ is.

Def. We say that $C(\cdot)$ is *calibrated to within $\epsilon > 0$ over Q given M* just in case, for every $\phi \in Q$, the absolute difference between $C(\phi)$ and the frequency of truths in the reference class for ϕ in Q given M is less than or equal to ϵ . Again, if $Cr_s(\cdot) = C(\cdot)$, then we say that $Cr_s(\cdot)$ is calibrated to within $\epsilon > 0$ over Q given M just in case $C(\cdot)$ is.

Def. For every $x, z \in \mathcal{D}$, we say that x and z are $C(\cdot)$ -*alike* just in case $C(A(x)) = C(A(z))$, for every $A \in \mathcal{F}$.

Def. We say that $\mathcal{Q}'$ is a $C(\cdot)$ -*alike extension* of $\mathcal{Q}$ just in case

- (i) $\mathcal{Q} \subseteq \mathcal{Q}'$.
- (ii) $C(\cdot)$ is defined for every member of $\mathcal{Q}'$.
- (iii) If $A(z) \in \mathcal{Q}'$, then there is some y such that y and z are $C(\cdot)$ -alike, and for every $B \in \mathcal{F}$, $B(z) \in \mathcal{Q}'$ just in case $B(y) \in \mathcal{Q}$.

Def. Let $C'(\cdot)$ be defined on the class of propositions of the form $A'(x')$ for $x' \in \mathcal{D}'$ and $A' \in \mathcal{F}'$. We say that $C'(\cdot)$ is an *extension* of $C(\cdot)$ just in case $\mathcal{D} \subseteq \mathcal{D}'$, $\mathcal{F} \subseteq \mathcal{F}'$, and for every $x \in \mathcal{D}$ and $A \in \mathcal{F}$, $C(A(x)) = C'(A(x))$.

Def. We say that $C(\cdot)$ is *potentially calibratable* on $\mathcal{Q}$ just in case for every $\epsilon > 0$, there is an extension of $C(\cdot)$, $C'(\cdot)$, and a $C'(\cdot)$ -alike extension $\mathcal{Q}'$ of $\mathcal{Q}$ such that $C'(\cdot)$ is calibrated to within ϵ on $\mathcal{Q}'$, given some model M .

Def. Let $Cr_s(\cdot) = C(\cdot)$. We say $Cr_s(\cdot)$ is *calibratable* on $\mathcal{Q}$ just in case, for every $\epsilon > 0$, there is an extension of $C(\cdot)$, $C'(\cdot)$, and a $C'(\cdot)$ -alike extension $\mathcal{Q}'$ of $\mathcal{Q}$ such that were it to be the case that $Cr_s(\cdot) = C'(\cdot)$, then S 's credal state would be calibrated to within ϵ on $\mathcal{Q}'$.

Def. For each $x \in \mathcal{D}$, we let $C^x(\cdot) : \mathcal{F} \rightarrow \mathbb{R}$ be such that, for each $A \in \mathcal{F}$, $C^x(A) = C(A(x))$.

Having presented these definitions, we can now present the calibration argument for a restricted version of Probabilism.⁹ First, we note the following theorem:

Calibration Theorem: $C(\cdot)$ is potentially calibratable, for every finite set $\mathcal{Q}$ over which it is defined, just in case, for every $x \in \mathcal{D}$, $C^x(\cdot)$ is a probability function over $\mathcal{F}$.¹⁰

9. A method of generalizing this argument to provide a justification for an unrestricted version of Probabilism is suggested in an appendix to van Fraassen (1983). However, since, as I will argue, this putative justification for the restricted version of Probabilism fails, I will not consider how this argument may be generalized.

10. See Theorems 5.2 and 6.1 in van Fraassen (1983). Let $C^x(\cdot)$ be a function mapping members of $\mathcal{F}$ to real numbers. We say that $C^x(\cdot)$ is a *probability function* over $\mathcal{F}$ just in case it satisfies the following constraints:

Property Normalization: For every $A(\cdot) \in \mathcal{F}$ such that it is logically guaranteed that every object satisfies $A(\cdot)$, $C^x(A(\cdot)) = 1$.

Property Non-Negativity: For every $A(\cdot) \in \mathcal{F}$, $0 \leq C^x(A(\cdot))$.

Property Finite Additivity: For every $A(\cdot), B(\cdot) \in \mathcal{F}$ such that it is logically guaranteed that no object satisfies both $A(\cdot)$ and $B(\cdot)$, $C^x(A(\cdot) \vee B(\cdot)) = C^x(A(\cdot)) + C^x(B(\cdot))$.

Note that, given the Calibration Theorem, it follows that if $C(\cdot)$ does not satisfy the condition that, for every $x \in \mathcal{D}$, $C^x(\cdot)$ is a probability function over $\mathcal{F}$, then, for any agent S , it is a priori that if $Cr_s(\cdot) = C(\cdot)$, then $Cr_s(\cdot)$ is not calibratable.

Next we assume that credal vindication may be characterized as follows:

Frequency Vindication: An agent S has a vindicated credal state just in case her credal state is calibratable for every finite Q over which it is defined.¹¹

Given Frequency Vindication, the Calibration Theorem then entails that, if $C(\cdot)$ does not satisfy the condition that, for every $x \in \mathcal{D}$, $C^x(\cdot)$ is a probability function over $\mathcal{F}$, then, for any agent S , it is a priori that if $Cr_s(\cdot) = C(\cdot)$, then $Cr_s(\cdot)$ is not vindicated. And so, given Potential Vindication, it follows that every agent S is rationally required to be such that if $Cr_s(\cdot) = C(\cdot)$, then, for every $x \in \mathcal{D}$, $C^x(\cdot)$ is a probability function over $\mathcal{F}$.

Now let $\mathcal{A}$ be an algebra of propositions of the form $A(x)$, where each member of $\mathcal{A}$ concerns the same object x . Call $\mathcal{A}$ a *single-object algebra*. The claim that it is rationally required that an agent have a credal state such that, for every $x \in \mathcal{D}$, $Cr^x(\cdot)$ is a probability function over $\mathcal{F}$, then entails the following norm:

Restricted Probabilism: It is a rational requirement that an agent have a credal state that is probabilistically coherent for every single object algebra over which it is defined.

That, then, is the calibration argument for Restricted Probabilism. I will now argue that we should reject this argument. In particular, I will argue that we should reject at least one of Frequency Vindication and Potential Vindication. I will further argue that *if* one ought to try to have credences that can be extended in a suitably constrained manner to line up as close as possible to appropriate frequencies, then in certain cases one ought to have probabilistically incoherent credences.

11. It should be noted that although van Fraassen is explicit about the form that his argument for a restricted version of Probabilism is supposed to take, he is not completely explicit about how credal vindication is to be characterized on his view. What van Fraassen says is: "Vindication I shall explicate in terms of *calibration*, a measure of how reliable one's judgments have been as indicators of actual frequencies." van Fraassen (1983: 300). However, 'calibration' here cannot mean being perfectly calibrated over every finite Q over which the credal state is defined. For, as we have seen, such calibration can be ruled out a priori for almost every credal state. I suggest, then, that we should interpret 'calibration' here as meaning *calibratability*. Given this, it is clear how the Calibration Theorem together with Potential Vindication entail a restricted version of Probabilism.

At the end of §2, having argued that we should reject at least one of Frequency Vindication and Potential Vindication, I will, however, consider two alternative ways of characterizing credal vindication. I will argue that neither alternative provides an adequate frequentist characterization of credal vindication.

2. Impermissible Probabilistic Coherence

While Frequency Vindication and Potential Vindication do entail that an agent ought to have probabilistically coherent credences over single-object algebras, I will now show that they in fact impose much stricter requirements. In particular, I will show that in certain cases Frequency Vindication and Potential Vindication rule out as rationally impermissible a large class of probabilistically coherent credences over certain single-object algebras. I will argue that these restrictions are manifestly implausible, and, so, conclude that we should reject either Frequency Vindication or Potential Vindication.

I will begin by showing that there are cases in which $C(\cdot)$ is a probability function defined over a single-object algebra and yet, for some agent S , it is a priori that if $Cr_s(\cdot) = C(\cdot)$, then $Cr_s(\cdot)$ will not be calibratable over some finite $\mathcal{Q}$ over which it is defined.¹²

To see this, consider the following case. Let ‘(*)’ refer to the following interpreted sentence:

Annie’s credence that () is true isn’t greater than or equal to 0.5.*

The above can, then, be represented as:

$$(*) \neg Cr_a(T(*)) \geq 0.5$$

Note that, as an instance of the T-schema, we have the following:

$$(T) T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$$

Let $\mathcal{D} = \{(*)\}$ and let $\mathcal{F}$ be the algebra of monadic properties generated by the following set of atoms: $\{T(\cdot), Cr_a(T(\cdot)) \geq 0.5\}$. We again let X be the set of propositions of the form $A(x)$, for $x \in \mathcal{D}$ and $A \in \mathcal{F}$. Note that X , so defined, is a single-object algebra.¹³

12. Note that this suffices to show that there are cases in which $C^x(\cdot)$ is a probability function over $\mathcal{F}$, for every $x \in \mathcal{D}$, and yet, for some agent S , it is a priori that if $Cr_s(\cdot) = C(\cdot)$, then $Cr_s(\cdot)$ will not be calibratable over some finite $\mathcal{Q}$ over which it is defined.

13. Might one reasonably reject the claim that X , so defined, is an algebra of propositions? No, I do not think so. If there is a worry here, it presumably stems from the fact that $(*)$ is a self-referential sentence. However, despite the self-referential nature of $(*)$, it is perfectly clear which object $(*)$ is. And it is perfectly clear which properties $T(\cdot)$ and $Cr_a(T(\cdot)) \geq 0.5$ and their Boolean combinations are. And so it is perfectly clear which propositions result from the application of such properties to $(*)$.

To see this, note that there are various possible scenarios in which it is natural to say, for example, that the proposition that $Cr_a(T(*)) \geq 0.5$ is true or false. So, for example, imagine that Annie falsely believes that any self-referential sentence is true and, furthermore, that she knows that $(*)$ is

We will assume the following:

- (i) Annie's credal state, $Cr_a(\cdot)$, is defined over X .
- (ii) $Cr_a(\cdot)$ is a probability function over X .
- (iii) $Cr_a(T(*)) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5 = 1$.

Note that there are many possible credal states that satisfy these conditions. For example, let $Cr_a(\cdot)$ meet the following conditions:

- (1) $Cr_a(T(*)) = 1$
- (2) $Cr_a(Cr_a(T(*)) \geq 0.5) = 0$
- (3) $Cr_a(\neg\phi) = 1 - Cr_a(\phi)$
- (4) $Cr_a(\phi \wedge \psi) = \min\{Cr_a(\phi), Cr_a(\psi)\}$
- (5) $Cr_a(\phi \vee \psi) = \max\{Cr_a(\phi), Cr_a(\psi)\}$.

Then, given that $(\phi \leftrightarrow \psi) =_{df} (\phi \wedge \psi) \vee (\neg\phi \wedge \neg\psi)$, it follows that $Cr_a(\cdot)$ is a probability function over X such that $Cr_a(T(*)) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5 = 1$. We assume that it is possible for Annie to have such a credal state.¹⁴

self-referential. Assuming that Annie puts these two beliefs together, it seems that she will have a very high credence that $(*)$ is true, and so the proposition $Cr_a(T(*)) \geq 0.5$ will be true. Or, perhaps Annie has the equally mistaken view that any self-referential sentence is false. If, then, Annie knows that $(*)$ is self-referential, and if she puts these two beliefs together, she will have very low credence that $(*)$ is true, and so the proposition $Cr_a(T(*)) \geq 0.5$ will not be true.

Given these sorts of considerations it seems completely implausible to deny that there is a proposition that $Cr_a(T(*)) \geq 0.5$. And similar considerations apply to the proposition that $T(*)$ as well as the other propositions in X .

14. Might one reasonably reject this assumption? No, I do not think so. First, note that it is clearly possible for Annie to have a credal state defined over an algebra generated by a pair of distinct propositions p_1 and p_2 such that $Cr_a(p_1) = 1$, $Cr_a(p_2) = 0$, $Cr_a(\neg\phi) = 1 - Cr_a(\phi)$, $Cr_a(\phi \wedge \psi) = \min\{Cr_a(\phi), Cr_a(\psi)\}$, and $Cr_a(\phi \vee \psi) = \max\{Cr_a(\phi), Cr_a(\psi)\}$. If, then, it is impossible for Annie to have such a credal state defined over X , this must be because of the particular propositions $T(*)$ and $Cr_a(T(*)) \geq 0.5$.

Now, given how $(*)$ is defined, it follows that necessarily $T(*)$ is true just in case $\neg Cr_a(T(*)) \geq 0.5$ is true. However, the strongest constraint that this fact could enforce is the following: necessarily, $Cr_a(T(*)) = Cr_a(\neg Cr_a(T(*)) \geq 0.5)$. This would follow, for example, if we assumed that propositions that are true in all the same worlds are identical. But this constraint is perfectly compatible with the assumption that $Cr_a(\cdot)$ meets conditions (1)–(5). The definition of $(*)$, then, does not preclude the possibility of Annie having a probabilistically coherent credal state over the single-object algebra X , such that $Cr_a(T(*)) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5 = 1$.

Now there are further assumptions that we could make that *would* rule out the possibility of Annie having a credal state meeting conditions (1)–(5). In particular, this could be ruled out as impossible if we assumed that there were certain necessary connections between Annie's first-order credence in $T(*)$ and her second-order credence in $Cr_a(T(*)) \geq 0.5$. For example, if it were metaphysically necessary that if $Cr_a(T(*)) \geq 0.5$ then $Cr_a(Cr_a(T(*)) \geq 0.5) \geq 0.5$, then it would not be possible that both $Cr_a(T(*)) = 1$ and $Cr_a(Cr_a(T(*)) \geq 0.5) = 0$. Of course, as a matter of contingent fact, there may be such connections between Annie's first-order credence in $T(*)$ and her second-order credence in

We can now show that it is a priori that if Annie has a credal state satisfying conditions (i)–(iii), then her credal state will not be calibratable over $Q = \{T(*), T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5\}$. The key fact here is that, given that a proposition of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ holds, it follows that the truth-value of $T(x)$ will depend on Annie's credence in $T(x)$. In particular, given that $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ holds, it follows that if $Cr_a(T(x)) \geq 0.5$, then $T(x)$ is false, while if $Cr_a(T(x)) < 0.5$, then $T(x)$ is true. And this fact, as we will see, entails that there are values of ϵ such that there is no appropriate Q' over which Annie's credences may be extended so that, were she to have such credences, propositions of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ would be within ϵ of 1, while propositions of the form $T(x)$ would be within ϵ of Annie's actual credence in $T(*)$.

Claim: It is a priori that if $Cr_a(\cdot)$ satisfies conditions (i)–(iii), then $Cr_a(\cdot)$ is not calibratable over $Q = \{T(*), T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5\}$.

Justification: Let $Cr_a(T(*)) = r$. First note that for Annie to be calibratable over Q , it must be the case that for every $\epsilon > 0$, there is some finite superset of $\mathcal{D}$, $\mathcal{D}'$, such that the following are all compossible:

- ($\mathcal{I}_a$) For every $x \in \mathcal{D}'$, $Cr_a(T(x)) = r$ and $Cr_a(T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5) = 1$.
- ($\mathcal{I}_b$) The frequency of truths amongst the class of propositions of the form $T(x)$, for $x \in \mathcal{D}'$, is within ϵ of r .
- ($\mathcal{I}_c$) The frequency of truths amongst the class of propositions of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$, for $x \in \mathcal{D}'$, is within ϵ of 1.

To see this, let $Cr_a(\cdot) = C(\cdot)$. For Annie to be calibratable over Q , it must be the case that, for every $\epsilon > 0$, there is an extension of $C(\cdot)$, $C'(\cdot)$, and a $C'(\cdot)$ -alike extension Q' of Q , such that were it to be the case that $Cr_a(\cdot) = C'(\cdot)$, Annie would be calibrated within ϵ over Q' . Now, given that $C(T(*)) = r$ and $C(T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5) = 1$, it follows that if $C'(\cdot)$ is an extension of $C(\cdot)$ and Q' is a $C'(\cdot)$ -alike extension of Q , then there will be a largest $\mathcal{D}' \supseteq \mathcal{D}$ such that there is a non-empty subset of Q' , Q'_s , such that:

$$Q'_s = \{T(x), T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5 \mid x \in \mathcal{D}' \text{ and } C'(T(x)) = r \text{ and } C'(T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5) = 1\}.$$

$Cr_a(T(*)) \geq 0.5$. For example, Annie may have remarkably good introspective powers. It is, however, completely implausible to maintain that such connections hold necessarily. We have no good reason, then, to reject the assumption that Annie may have a credal state satisfying conditions (1)–(5) and so no good reason to deny that Annie may have a credal state satisfying assumptions (i)–(iii).

In order for it to be the case that Annie would be calibrated to within ϵ over $\mathcal{Q}'$, were $Cr_a(\cdot) = C'(\cdot)$, it must therefore be the case that, were $Cr_a(\cdot) = C'(\cdot)$, the frequency of truths amongst the class of propositions of the form $T(x)$, for $x \in \mathcal{D}'$, would be within ϵ of r , and the frequency of truths amongst the class of propositions of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$, for $x \in \mathcal{D}'$, would be within ϵ of 1. If, however, there is no finite superset of $\mathcal{D}$, $\mathcal{D}'$, for which $(I_a)-(I_c)$ are compossible, given this value of ϵ , then this condition cannot be met, and so Annie will not be calibratable over $\mathcal{Q}$.

To show, given (i)–(iii), that Annie is not calibratable, it suffices, then, to show that, no matter what value r is, there are values of ϵ such that there is no $\mathcal{D}'$ that extends $\mathcal{D}$ for which $(I_a)-(I_c)$ are compossible.

First, assume that $Cr_a(T(*)) = r \geq 0.5$. Further, assume that (I_a) holds. Given these two assumptions, it follows that if the frequency of truths amongst the class of propositions of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$, for $x \in \mathcal{D}'$, is at least q , then the frequency of truths amongst the class of proposition of the form $T(x)$, for $x \in \mathcal{D}'$, is *at most* $1 - q$. The reason for this is the following. Given that $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ is true, it follows from the fact that $Cr_a(T(x)) = r \geq 0.5$ that $T(x)$ must be false. Thus, for every proposition that contributes to the frequency of truths amongst the class of propositions of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$, for $x \in \mathcal{D}'$, there is a proposition that contributes to the frequency of *falsehoods* amongst the class of propositions of the form $T(x)$, for $x \in \mathcal{D}'$. And so, as the number of truths of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ approaches 1, the number of truths of the form $T(x)$ must approach 0, and so move away from r . But this assures us that there are values of ϵ for which at least one of (I_b) or (I_c) must fail.

To pick a somewhat arbitrary example to illustrate this point, if the frequency of truths amongst the class of propositions of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$, for $x \in \mathcal{D}'$, is within 0.2 of 1, it follows that the frequency of truths amongst the class of propositions of the form $T(x)$, for $x \in \mathcal{D}'$, can be at most 0.2. But, then, since $Cr_a(T(*)) = r \geq 0.5$, it follows that the frequency of truths amongst the class of propositions of the form $T(x)$, for $x \in \mathcal{D}'$, is not within 0.2 of r . And so, given that $Cr_a(T(*)) = r \geq 0.5$, it follows that there are values of ϵ such that there is no $\mathcal{D}'$ that extends $\mathcal{D}$ for which $(I_a)-(I_c)$ are compossible.

Next, assume that $Cr_a(T(*)) = r < 0.5$. And, again, assume that (I_a) holds. Given these two assumptions, it follows that if the frequency of truths amongst the class of propositions of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$, for $x \in \mathcal{D}'$, is at least q , then so too is the frequency of truths amongst the

class of propositions of the form $T(x)$, for $x \in \mathcal{D}'$. The reason for this is the following. Given that $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ is true, it follows from the fact that $Cr_a(T(x)) = r < 0.5$ that $T(x)$ must be true. Thus, for every proposition that contributes to the frequency of truths amongst the class of propositions of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$, for $x \in \mathcal{D}'$, there is also a proposition that contributes to the frequency of truths amongst the class of propositions of the form $T(x)$, for $x \in \mathcal{D}'$. But this assures us that there are values of ϵ for which at least one of (I_b) or (I_c) must fail.

To again pick a somewhat arbitrary example to illustrate this point, if the frequency of truths amongst the class of propositions of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$, for $x \in \mathcal{D}'$, is within 0.2 of 1, it follows that the frequency of truths amongst the class of propositions of the form $T(x)$, for $x \in \mathcal{D}'$, is also within 0.2 of 1. But, then, since $Cr_a(T(*)) = r < 0.5$, it follows that the frequency of truths amongst the class of propositions of the form $T(x)$, for $x \in \mathcal{D}'$, is not within 0.2 of r . And so, given that $Cr_a(T(*)) = r < 0.5$, it follows that there are values of ϵ such that there is no $\mathcal{D}'$ that extends $\mathcal{D}$ for which $(I_a)-(I_c)$ are compossible.

It follows that, whatever value r is, there are values of ϵ such that there is no $\mathcal{D}'$ that extends $\mathcal{D}$ for which $(I_a)-(I_c)$ are compossible. This suffices to establish that, given (i)–(iii), Annie is not calibratable.

Note that no empirical assumptions were required for this proof. Thus, we have established that it is in fact a priori that if Annie has a credal state satisfying conditions (i)–(iii), then she will not be calibratable.

We have shown, then, that there are cases in which an agent may have a probabilistically coherent credal state defined over some single-object algebra and yet it is a priori that if the agent has such a credal state, then she will not be calibratable. Given Frequency Vindication and Potential Vindication, then, it follows that it is rationally impermissible for such an agent to have such a probabilistically coherent credal state.

Now, in and of itself, this conclusion is not necessarily a bad thing. Many proponents of Probabilism think that there are in fact stronger norms that serve to rule out certain probabilistically coherent credal states as irrational. However, upon closer inspection, I think that we can see that the class of probabilistically coherent credal states that Frequency Vindication and Potential Vindication rule out as irrational is implausibly restrictive.

To see this, note the following.

Claim: It is a priori that if $Cr_a(T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5) > 0.5$, then $Cr_a(\cdot)$ is not calibratable over $Q = \{T(*), T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5\}$.

Justification: Again, we let $Cr_a(T(*)) = r$. Given that $Cr_a(T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5) = m > 0.5$, for Annie to be calibratable over $Q = \{T(*), T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5\}$, it must be the case that for every $\epsilon > 0$, there is some finite superset of $\mathcal{D}, \mathcal{D}'$, such that the following are all compossible:

- (2_a) For every $x \in \mathcal{D}'$, $Cr_a(T(x)) = r$ and $Cr_a(T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5) = m$.
- (2_b) The frequency of truths amongst the class of propositions of the form $T(x)$, for $x \in \mathcal{D}'$, is within ϵ of r .
- (2_c) The frequency of truths amongst the class of propositions of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$, for $x \in \mathcal{D}'$, is within ϵ of m .

To see that these conditions are not compossible, first assume that $Cr_a(T(*)) = r \geq 0.5$. Then for each proposition of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ that is true, there is a corresponding proposition of the form $T(x)$ that is false. Thus, as the number of truths of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ approaches m , the upper bound for the number of truths of the form $T(x)$ will approach $1 - m$. Since $m > 0.5$ and $Cr_a(T(*)) = r \geq 0.5$, it follows, then, that there are values of ϵ such that Annie cannot be calibrated to within ϵ for both $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$ and $T(*)$.

Next, assume that $Cr_a(T(*)) = r < 0.5$. Then for each proposition of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ that is true, there is a corresponding proposition of the form $T(x)$ that is also true. Thus, as the number of truths of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ approaches m , the lower bound for the number of truths of the form $T(x)$ will also approach m . Since $m > 0.5$ and $Cr_a(T(*)) = r < 0.5$, it follows, then, that there are values of ϵ such that Annie cannot be calibrated to within ϵ for both $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$ and $T(*)$.

Since (2_a)–(2_c) are not compossible, it follows that Annie is not calibratable over Q , given that $Cr_a(T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5) = m > 0.5$. Once again, no empirical assumptions were required for this proof. Thus, we have established that it is a priori that if $Cr_a(T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5) > 0.5$, then $Cr_a(\cdot)$ is not calibratable over Q .

Given Frequency Vindication and Potential Vindication, then, it follows that Annie is rationally required to have credence less than or equal to 0.5 in $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$. This is quite a strong and, I think, quite an implausible constraint

on Annie's credal state. For given how $(*)$ was introduced, $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$ is, I submit, an obvious truth. Moreover, it is a truth in which, I take it, *we* ourselves have a high credence. However, given Frequency Vindication and Potential Vindication, it follows that it would be irrational for Annie to have greater than 0.5 credence in this truth, even if her epistemic situation with respect to this proposition were the same as ours. This is a rather unbelievable claim.¹⁵ The implausibility of this putative rational restriction, then, gives us good reason to be skeptical of the pair Frequency Vindication and Potential Vindication, and so to be skeptical of this putative justification of Restricted Probabilism.

Having presented this argument, let me now consider a few responses.

First, the proponent of the calibration argument may suggest that there is in fact good reason to have a low credence in $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$. For, as is well-known, if we want to endorse classical logic, we must reject at least some instances of the T-schema. In particular, if λ is the sentence that says of itself that it is not true, then we must reject the instance of the T-schema for this sentence. The proponent of the calibration argument may suggest, then, that since both λ and $(*)$ are self-referential, the fact that we must reject the instance of the T-Schema for the former sentence, gives us reason to reject the instance of the T-schema for the latter.

This argument, however, has little force. For there is an important difference between $(*)$ and λ . For, while endorsing the instance of the T-schema for λ leads to inconsistency given classical logic, the same is not true if one endorses: $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$. The latter bi-conditional is perfectly consistent with classical logic.¹⁶ And given the intuitive plausibility of the T-biconditionals, if there is no logical reason to reject an instance of this schema, we should endorse it. Indeed, there are perfectly general non-ad-hoc treatments of the truth predicate that allow us to endorse various instances of the T-schema, including instances involving self-referential sentences, as long as there is no inconsistency with classical logic.¹⁷ We do not, then, have good reason to reject the assumption that $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$ holds. And so, we do not have good reason to maintain that, if Annie is rational, then she will have a low credence in this claim.

Second, the proponent of the calibration argument may suggest that in characterizing Frequency Vindication I have picked the wrong credal frequentist account

15. See Caie (2013) for an extensive defense of the claim that low credence in claims such as $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$ cannot be rationally required.

16. This can be shown by generating a classical model for a language involving a truth-predicate T as well as an operator $Cr_a(\cdot) \geq 0.5$ by using a Kripke-style fixed-point construction. The basic idea is that one will identify the extension of T in the model with the so-called positive extension that results from the fixed-point construction. There are models of this type in which ' $T(*)$ ' and ' $\neg Cr_a(T(*)) \geq 0.5$ ' both come out true, and ones in which both such sentences come out false. The basic idea of how to generate such models for languages without credence operators comes from Kripke (1975). See Caie (2011) for details on how to extend such models to languages containing credence operators.

17. See, e.g., Field (2008) for a description of theories of this type.

of credal vindication. Indeed a natural response to the forgoing argument, which sought to establish that the constraints that Frequency Vindication and Potential Vindication impose on Annie's credal state are implausibly restrictive, is to try to find some other frequentist-friendly notion of credal vindication that might be used to justify Restricted Probabilism in a similar manner, but which is not subject to the same sort of criticism.

Two alternative notions of vindication suggest themselves.

To characterize the first, we need an additional piece of terminology.

Def. Let $Cr_s(\cdot) = C(\cdot)$. We will say that S's credal state is *weakly calibratable* just in case, for every $\epsilon > 0$, there is an extension of $C(\cdot)$, $C'(\cdot)$, and a $C'(\cdot)$ -alike extension Q' of Q such that $C'(\cdot)$ is calibrated to within ϵ on Q' .

The first suggestion, then, is the following.

Frequency Vindication₁: An agent S has a vindicated credal state just in case her credal state is weakly calibratable for every finite Q over which it is defined.

Now it is a consequence of the Calibration Theorem that if, for some $x \in \mathcal{D}$, $C^x(\cdot)$ is not a probability function over $\mathcal{F}$, then it is a priori that if $Cr_s(\cdot) = C(\cdot)$, then $Cr_s(\cdot)$ is not weakly calibratable. It follows, then, that Frequency Vindication₁ and Potential Vindication entail Restricted Probabilism.

Frequency Vindication₁ and Potential Vindication do not, however, entail that Annie is rationally prohibited from having credence greater than 0.5 in $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$. To see this, first recall why it is that Annie is not calibratable over $Q = \{T(*), T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5\}$, given that $Cr_a(T(*)) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5 > 0.5$. Let $Cr_a(\cdot) = C(\cdot)$. As we saw earlier, for Annie to be calibratable over Q , it must be the case that:

For every $\epsilon > 0$, there is some $Q' = \{T(x), T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5 \mid x \in \mathcal{D}' \supseteq \{(*)\}\}$ and some $C'(\cdot)$ such that $C'(T(x)) = C(T(*))$ and $C'(T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5) = C(T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5)$, for every $x \in \mathcal{D}'$, such that were it to be the case that $Cr_a(\cdot) = C'(\cdot)$ then $Cr_a(\cdot)$ would be calibrated to within ϵ over Q' .

The following facts, however, preclude this from obtaining, given that $Cr_a(T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5) > 0.5$:

- (i) If $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ is true, then it follows that if $Cr_a(T(x)) \geq 0.5$, then $T(x)$ is false.

- (ii) If $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ is true, then it follows that if $Cr_a(T(x)) < 0.5$, then $T(x)$ is true.

For, given (i), it follows that if $Cr_a(T(*)) \geq 0.5$, then the number of truths of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ will impose an upper bound on the number of truths of the form $T(x)$. While, given (ii), it follows that if $Cr_a(T(*)) < 0.5$, then the number of truths of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ will impose a lower-bound on the number of truths of the form $T(x)$. And these bounds are such that, if $Cr_a(T(*)) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5 > 0.5$, then there are values of ϵ such that Annie's credence cannot be calibrated to within ϵ for $T(*)$, given that she is calibrated to within ϵ for $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$.

Unlike calibration, however, weak calibration does not require that it be the case that Annie would be calibrated to within ϵ over Q' , were it to be the case that $Cr_a(\cdot) = C'(\cdot)$. It simply requires that there is, in fact, some extension $C'(\cdot)$ of $C(\cdot)$ that is calibrated to within ϵ over Q' . In assessing whether $C'(\cdot)$ is calibrated to within ϵ over Q' , we need not, then, assume that $Cr_a(\cdot) = C'(\cdot)$. And without this assumption, the number of truths of the form $T(x) \leftrightarrow \neg Cr_a(T(x)) \geq 0.5$ will not impose any upper bound on the number of truths of the form $T(x)$, given that $C'(T(x)) \geq 0.5$, or any lower bound, given that $C'(T(x)) < 0.5$. There is not, then, the same in-principle barrier to Annie's credal state being weakly calibratable over Q , given that $Cr_a(T(*)) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5 > 0.5$, as there is to her credal state being calibratable.

The second suggestion for an alternative frequentist characterization of credal vindication is the following.

Frequency Vindication₂: Let $Cr_s(\cdot) = C(\cdot)$. An agent S has a vindicated credal state just in case $C(\cdot)$ is *potentially calibratable*, for every finite Q over which it is defined.

If we accept Frequency Vindication₂, then the Calibration Theorem entails that it is a *necessary and sufficient* condition for $Cr_s(\cdot)$ to be vindicated, that, for every $x \in \mathcal{D}$, $C^x(\cdot)$ is a probability function over $\mathcal{F}$. Moreover, since this is a priori, it follows that Frequency Vindication₂ and Possible Vindication rule out as irrational all and only those credal states violating this condition. Thus, Frequency Vindication₂ and Possible Vindication entail Restricted Probabilism, but do not rationally preclude Annie from having credence greater than 0.5 in $T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5$.

Both Frequency Vindication₁ and Frequency Vindication₂, then, in concert with Potential Vindication, entail Restricted Probabilism without also imposing the same implausible restrictions as those that are imposed by Frequency Vindication. Despite these facts, however, I do not think that Frequency Vindication₁ or Frequency

Vindication₂ are particularly plausible ways for a frequentist to characterize credal vindication.

To see this, let us ask why it might be an epistemic failing for a credal state to be such that it is not weakly calibratable, or to be such that it is not potentially calibratable? I maintain that the most plausible account of why such features of a credal state count as epistemic defects will have as a consequence that it is also an epistemic defect for a credal state to not be calibratable.

Let us focus first on the failure to be weakly calibratable. If one endorses Frequency Vindication₁, then, at a minimum, one must claim that it is an epistemic failure for one to have a credal state that is not weakly calibratable. Why though should this be thought to be an epistemic failure?

Here is, I think, a very natural way of justifying this claim. Let Q be a finite set of propositions and $\epsilon > 0$. We start with the idea that it is a good-making feature of an agent's credal state $Cr_s(\cdot)$ that $Cr_s(\cdot)$ is calibrated to within ϵ over Q . Now, as we noted earlier, an agent's credal state may fail to have such a good-making feature in virtue of the paucity of propositions within Q . For example, if Q is a singleton set consisting of a true proposition, then any credal state that fails to assign to the unique proposition in Q the value 1 will fail to be calibrated to within ϵ , for infinitely many such values. We will not, then, want to say that in these sorts of cases it is an epistemic failure for such a credal state to exhibit this type of good-making feature. However, so the thought goes, we should say that it is an epistemic failure if such a credal state precludes the agent from having this good-making feature when one abstracts away from the limitations imposed by the number of propositions in the class Q . That is, it is an epistemic failure if an agent's credences over Q are such that she would still fail to have this good-making feature were she to have these credences over arbitrarily larger classes of suitably similar propositions.

Now this line of thought certainly supports the claim that it is an epistemic failure if an agent has a credal state that is not weakly calibratable. For, if one's credal state is not weakly calibratable, then there will be some finite Q such that, even when we abstract away from the limitations imposed by the number of propositions in this class, the agent's credences over Q preclude her from having the good-making feature of being calibrated to within ϵ . However, it is plain to see that this way of justifying the claim that it is an epistemic failure to be weakly calibratable does not serve to support Frequency Vindication₁. For it is also true that if an agent is not calibratable over Q , then there is some feature of her credal state that precludes her from having the good-making feature of being calibrated to within ϵ , even when we abstract away from the number of propositions in Q . The preceding line of thought, then, not only supports the claim that it is an epistemic failure to have a credal state that is not weakly calibratable, it also supports the stronger

claim that it is an epistemic failure to have a credal state that is not calibratable. And this latter claim is not compatible with Frequency Vindication₁.

What the foregoing shows, then, is that if one is to provide a justification for the claim that it is an epistemic failure to have a credal state that is not weakly calibratable, which does not also serve to justify the claim that it is an epistemic failure to have a credal state that is not calibratable, then one must provide some account of why there is something epistemically bad about having a credal state for which there is no appropriate extension that is calibrated within ϵ that does not appeal to the idea that it is a good-making feature of one's credences to in fact be calibrated to within ϵ . I, at least, can see no plausible story meeting this description. And so, I cannot see any plausible justification for Frequency Vindication₁.

The same worries will apply, *mutatis mutandis*, to Frequency Vindication₂. For, first, while we can justify the claim that it is an epistemic defect to have a credal state that is not potentially calibratable by appealing to the idea that a credal state is epistemically defective if features of the credal state preclude an agent from having certain epistemic goods even when we abstract away the limitations that are imposed by the cardinality of the sets over which the state is defined, this, again, will not serve to justify Frequency Vindication₂. For as we have seen, this thought will also support the claim that it is an epistemic defect to have a credal state that is not calibratable over some finite Q , and this is not compatible with Frequency Vindication₂. And, second, as with weak calibratability, it is not at all obvious how one might support the claim that a credal state is epistemically defective if it is not potentially calibratable, without appealing to the fact that it is a good-making feature of a credal state to be calibrated to within ϵ over some set Q .

The prospects, then, for justifying Restricted Probabilism by an appeal to either of these alternative frequentist characterizations of credal vindication look quite dim.

3. Required Probabilistic Incoherence

As it happens, I think that a frequentist should reject Frequency Vindication as a characterization of credal vindication. However, as we will see, the reason for this will provide little succor for the proponent of Restricted Probabilism.

The motivating idea behind Frequency Vindication is that a credal state should count as vindicated just in case, for every finite set Q over which it is defined and every $\epsilon > 0$, the credal state can be extended to some $Q' \supseteq Q$ so that were the agent to have the extended credal state, the agent's credence, for every $\phi \in Q'$, would be within ϵ of the relative frequency of truths amongst the propositions that are relevantly similar to ϕ in Q' . This motivating idea, however, is not perfectly precise. In particular, it is not clear what is required for some propositions to be

relevantly similar to one another. Frequency Vindication, then, provides a particular *precisification* of this motivating idea. For the formulation of Frequency Vindication involves a precise characterization of which reference classes are relevant for determining how close an agent's credences are to the relative frequencies. In principle, then, one could endorse this motivating idea, while rejecting the particular precisification of this idea that Frequency Vindication provides. For one could hold that Frequency Vindication depends on an incorrect characterization of the reference classes that are appropriate for assessing the closeness of an agent's credences to the relative frequencies.

In fact, I think that there is good reason to hold that, at least in certain cases, the reference classes appealed to in Frequency Vindication are not the right classes for assessing how close an agent's credences are to the relative frequencies. I think, then, that even one who is attracted to the idea that credences constitutively aim at closeness to relative frequencies should not accept Frequency Vindication. This fact, however, will not help the proponent of Probabilism. For, as I will now argue, given the most natural way of demarcating the reference classes for $T(*)$, it follows that if credences constitutively aim at closeness to relative frequencies, then Annie ought to have a probabilistically *incoherent* credal state.

According to our earlier definition, borrowed from van Fraassen (1983), given a set of propositions Q and a credal function $C(\cdot)$, the reference class for $A(x)$ is the set of propositions $A(z) \in Q$ such that $C(A(x)) = C(A(z))$. I think that this way of demarcating reference classes gets two things right. First, in assessing how well-calibrated an agent's credence in the proposition $A(x)$ is, the propositions that are *relevantly similar* to $A(x)$ should all involve the application of the same property to some object. Second, the propositions that are *relevantly similar* to $A(x)$ should all be such that the agent has the same credence in them as she has in the proposition $A(x)$. The conditions that van Fraassen cites, then, are both *necessary* conditions for a proposition being in the reference class for $A(x)$. According to van Fraassen (1983), however, they are also sufficient. And this, I think, is not true, at least not in general.

To see this, consider the following three sentences:

(*) $\neg Cr_a(T(*)) \geq 0.5$.

(#) $\neg Cr_a(T(\#)) \geq 0.5$.

(S) Snow is white.

Let $Q = \{T(*), T(\#), T(S)\}$ and let us assume that $C(T(*)) = C(T(\#)) = C(T(S))$. The question that we want to consider is: What is the appropriate reference class, given Q and $C(\cdot)$, for $T(*)$? According to van Fraassen (1983), the correct answer is: $\{T(*), T(\#), T(S)\}$. But a much more natural answer, I submit, is: $\{T(*), T(\#)\}$. Of

course, $T(*)$, $T(\#)$, and $T(S)$ all share the property of being assigned the same credence by Annie. And they all share the property of ascribing the property of truth to a sentence. However, $T(*)$ and $T(\#)$ share an additional property that $T(S)$ lacks. For both of these propositions ascribe truth to a sentence that has the property of being such that, necessarily, the sentence is true just in case Annie does not have credence greater than or equal to 0.5 that that very sentence is true. $T(S)$, on the other hand, ascribes the property of truth to a sentence whose truth is, at least in principle, independent of Annie's credence in its truth. This commonality in the nature of the sentence to which the property of truth is ascribed strikes me as being just as important an aspect of similarity between $T(*)$ and $T(\#)$ as the fact that they both ascribe the same property. For this reason, then, $\{T(*), T(\#)\}$ is, I think, a better candidate for being the appropriate reference class for $T(*)$ than $\{T(*), T(\#), T(S)\}$.

It is, of course, a vexed issue how, in general, we should determine the appropriate reference class for a given proposition ϕ in order to determine the relative frequency of truths amongst propositions like ϕ . However, for present purposes, we do not need a general account in hand. All that is required are the following two claims.

First, given a class of propositions Q containing $T(*)$ and a possible credal function $C(\cdot)$, there is a natural answer about what the reference class for $T(*)$ is. It is the set of $\phi \in Q$ satisfying the following three conditions:

- (3_a) ϕ of the form $T(x)$.
- (3_b) $C(T(x)) = C(T(*))$.
- (3_c) x is a sentence such that, necessarily, x is true just in case Annie does not have credence greater than or equal to 0.5 that that very sentence is true.

Second, given a class of propositions Q containing $\neg T(*)$ and a possible credal function $C(\cdot)$, there is a natural answer about what the reference class for $\neg T(*)$ is. It is the set of $\phi \in Q$ satisfying the following three conditions:

- (4_a) ϕ is of the form $\neg T(x)$.
- (4_b) $C(\neg T(x)) = C(\neg T(*))$.
- (4_c) x is a sentence such that, necessarily, x is true just in case Annie does not have credence greater than or equal to 0.5 that that very sentence is true.

We will assume that the correct account of the reference class for some proposition ϕ given some class of propositions Q and some credal function $C(\cdot)$ satisfies conditions (3_a)–(3_c) and (4_a)–(4_c). Call this the *real reference class* for ϕ in Q , given $C(\cdot)$. Given our earlier series of definitions, a change in our characterization of reference classes will result in a corresponding change in the notion of calibratability.

Call the notion of calibratability that results from this amended characterization of reference classes *real calibratability*. I suggest, then, that the correct frequentist characterization of credal vindication should take the following form.

Real Frequency Vindication: An agent S has a vindicated credal state just in case her credal state is real calibratable for every finite Q over which it is defined

We can show, however, that real calibratability, while an epistemic good, may be impossible for certain agents and certain classes of propositions.

Claim: Whatever Annie's credence is in $T(*)$, her credal state is not real calibratable over $Q = \{T(*)\}$.

Justification: Let $Cr_a(T(*)) = r$. Given $(3_a)-(3_c)$, if Annie's credal state is to be real calibratable, then it must be the case that for every $\epsilon > 0$, there is some $\mathcal{D}' \supseteq \{(*)\}$ satisfying the following two conditions:¹⁸

- (5_a) Every $x \in \mathcal{D}'$ is a sentence satisfying condition (3_c) .
- (5_b) The following two claims are compossible:
 - (i) For every $x \in \mathcal{D}'$, $Cr_a(T(x)) = r$
 - (ii) The absolute difference between r and the number of truths of the form $T(x)$, for $x \in \mathcal{D}'$, is less than or equal to ϵ .

It is easy to see, however, that no matter what value r takes, if (5_a) holds, then there are values of ϵ for which (i) and (ii) are not compossible.

Let $Cr_a(T(*)) = r \geq 0.5$. Then it follows, given (5_a), that if (i) holds, then each proposition of the form $T(x)$, for $x \in \mathcal{D}'$, is false. Thus, the absolute difference between the number of truths of the form $T(x)$, for $x \in \mathcal{D}'$ and r must be greater than or equal to 0.5 .

Next, let $Cr_a(T(*)) = r < 0.5$. Then it follows, given (5_a), that if (i) holds, then each proposition of the form $T(x)$, for $x \in \mathcal{D}'$, is true. Thus, the absolute difference between the number of truths of the form $T(x)$, for $x \in \mathcal{D}'$, and r must be greater than 0.5 .

It follows, then, that for any value of ϵ less than 0.5 , if (4_a) holds, then (i) and (ii) are not compossible. And so it follows that Annie's credence in $T(*)$

18. The justification for this claim parallels the justification, given in §2, for the claim that in order for Annie to be calibratable over $Q = \{T(*), T(*) \leftrightarrow \neg Cr_a(T(*)) \geq 0.5\}$, conditions $(1_a)-(1_c)$ must be compossible.

cannot be made to be arbitrarily close to the relative frequency of truths amongst propositions similar to $T(*)$.

I have argued that the correct frequentist characterization of credal vindication is given by Real Frequency Vindication. This result, then, gives us good reason to maintain that, pace van Fraassen, a frequentist should not accept Possible Vindication. For what the forgoing result shows is that Real Frequency Vindication and Possible Vindication are jointly incompatible with the following plausible general constraint on principles of rationality:¹⁹

Ought-Can: It must always be possible for an agent to meet the requirements imposed by rationality.

This, however, does not mean that one cannot hold that closeness to relative frequencies is normative for credences. It does, however, mean that if there is a normative role that closeness to relative frequencies plays, then it must be characterized differently than van Fraassen suggests. Let us explore, then, how we might characterize this normative role.

While a frequentist should maintain that to be fully vindicated one must be real calibratable, it seems quite plausible to me that she should also allow that the notion of credal vindication is not all or nothing.

With credal vindication, a miss is not as good as a mile. For it seems quite plausible that there may be two possible credal states that an agent could have, neither of which would make the agent fully vindicated, and yet were the agent to have one such credal state, she would be more vindicated than if she were to have the other.

There are a number of tricky questions about how we should think about relations of relative vindication. Luckily, for present purposes, all that we require is a very plausible *sufficient* condition for it to be the case that, for some agent S , were she to have one credal state, she would be more vindicated than she would be were she to have some other credal state.

First, however, some definitions.

19. See Caie (2013) for further discussion of this principle in a similar context. It is worth highlighting that the following is a very weak ought-implies-can principle. What it demands is that it be *metaphysically* possible for an agent to meet the requirements imposed by rationality. Now, there are perhaps good reasons for a proponent of Probabilism to reject stronger ought-implies-can principles that require that it be possible, in some more demanding sense, for an agent to meet the requirements of rationality. For, given how actual agents are constituted, there are good reasons to hold that such agents cannot meet all of the requirements imposed by Probabilism. However, the proponent of Probabilism does not have similar reason to reject the principle Ought-Can. For this principle allows that we may abstract away from the contingent cognitive limitations of actual agents.

Def. Let $Cr_s(\cdot) = C(\cdot)$. We say $Cr_s(\cdot)$ is *real calibratable to within $\epsilon > 0$ on Q* just in case there is an extension of $C(\cdot)$, $C'(\cdot)$, and a $C'(\cdot)$ -alike extension Q' of Q such that were it to be the case that $Cr_s(\cdot) = C'(\cdot)$, then S 's credal state would be real calibrated to within ϵ on Q' .

Def. We say that α is the *real calibratability value of $Cr_s(\cdot)$ on Q* just in case α is the greatest lower bound of the set of ϵ such that $Cr_s(\cdot)$ is *real calibratable to within ϵ on Q* .

Def. Let $C(\cdot)$ and $C'(\cdot)$ be credal functions defined over the same set of propositions. Let α_c^q be the real calibratability value that would result for $Cr_s(\cdot)$ on Q , were it the case that $Cr_s(\cdot) = C(\cdot)$. And let $\alpha_{c'}^q$ be the real calibratability value that would result for $Cr_s(\cdot)$ on Q , were it the case that $Cr_s(\cdot) = C'(\cdot)$. We say that $C(\cdot)$ *calibration dominates $C'(\cdot)$ for S* just in case:

- (i) For every finite Q , $\alpha_c^q \leq \alpha_{c'}^q$.
- (ii) For some finite Q , $\alpha_c^q < \alpha_{c'}^q$.

If one accepts Real Frequency Vindication, then I suggest that one should accept:

Graded Vindication: If $C(\cdot)$ calibration dominates $C'(\cdot)$ for S , then S 's credal state would be more vindicated were it to be the case that $Cr_s(\cdot) = C(\cdot)$ than it would be were it the case that $Cr_s(\cdot) = C'(\cdot)$.

Now consider the following normative principle:

Vindication Maximization: If S has credences defined over some set X , and there is some $C(\cdot)$ defined over X such that it is a priori that were it to be the case that $Cr_s(\cdot) = C(\cdot)$, then S would be more vindicated than she would be were she to have any other credal state defined over X , and it is possible that $Cr_s(\cdot) = C(\cdot)$, then S is rationally required to be such that $Cr_s(\cdot) = C(\cdot)$.

I suggest that if one thinks that credal states constitutively aim at closeness to relative frequencies, then one should endorse, in addition to Real Frequency Vindication, both Graded Vindication and Vindication Maximization.²⁰ The latter two principles, however, entail that, if Annie's credal state is defined over certain algebras, then Annie ought to have probabilistically incoherent credences.

To see this, consider the single-object algebra $\mathcal{A} = \{T(*), \neg T(*), T(*) \vee \neg T(*)\}$,

20. Note that these principles will be compatible with Ought-Can. For the antecedent of Vindication Maximization ensures that an agent is only required to have some particular credal state if it is indeed possible for her to have such a credal state.

$T(*) \wedge \neg T(*)$.²¹ We will first show that there is some probabilistically incoherent credal function $C(\cdot)$, defined over $\mathcal{A}$, such that it is a priori that $C(\cdot)$ calibration dominates all other such credal functions for Annie.

Claim: Let $C(\cdot) = \{C(T(*)) = 0.5, C(\neg T(*)) = 1, C(T(*) \vee \neg T(*)) = 1, C(T(*) \wedge \neg T(*)) = 0\}$. $C(\cdot)$ calibration dominates $C'(\cdot)$ for Annie, for every other $C'(\cdot)$ defined over $\mathcal{A}$.

Justification: The first point to note is that if $Cr_a(T(*) \vee \neg T(*)) = 1$, then Annie's credence in this proposition perfectly matches the frequency of truths in the real reference class for $T(*) \vee \neg T(*)$ in $\mathcal{A}$. And given that the real reference class for $T(*) \vee \neg T(*)$ will consist of propositions of the form $T(x) \vee \neg T(x)$ in which Annie has credence 1, it follows that Annie's credence in this proposition will continue to perfectly match the frequency of truths in its real reference class for any extension of her credal state to some larger set of propositions.

Similar considerations show that if $Cr_a(T(*) \wedge \neg T(*)) = 0$, then Annie's credence in this proposition will perfectly match the frequency of truths in its real reference class for any extension of her credal state to some larger set of propositions.

The next point to note is that $Cr_a(T(*)) = 0.5$ is the unique credence that can be made to most closely approximate the relative frequency of truths amongst propositions similar to $T(*)$. To see this, note that the following two claims are consequences of $(3_a)-(3_c)$. First, if $Cr_a(T(*)) = r \geq 0.5$, then every member in the real reference class for $T(*)$ will be false. Second, if $Cr_a(T(*)) = r < 0.5$, then every member in the real reference class for $T(*)$ will be true. From these two claims it follows that the difference between Annie's credence in $T(*)$ and the number of truths in the real reference class is always greater than or equal to 0.5 and that it is equal to 0.5 just in case $Cr_a(T(*)) = 0.5$. Thus, it follows that 0.5 is the unique credence in $T(*)$ that can be made to most closely approximate the relative frequency of truths amongst propositions similar to $T(*)$.

21. In what follows, I will take logically equivalent propositions to be identical. Thus, for example, $T(*)$ and $\neg\neg T(*)$ will be taken to be the same proposition. This identification is useful in simplifying the following proof. In principle, however, one could prove, mutatis mutandis, that the same conclusion holds given an algebra in which propositions such as $T(*)$ and $\neg\neg T(*)$ are not identified. Note that this assumption is consistent with our assumption that propositions have as constituents properties and objects. We simply need to assume that properties that are logically guaranteed to be coextensive, e.g., $T(\cdot)$ and $\neg\neg T(\cdot)$, are, in fact, identical. This actually seems quite plausible to me.

Finally, we note that, given $(4_a)-(4_c)$, it follows from the fact that $Cr_a(T(*)) = 0.5$ that every member of the real reference class for $\neg T(*)$ is true. Thus, it follows that if $Cr_a(T(*)) = 0.5$, then if $Cr_a(\neg T(*)) = 1$, then Annie's credence in this proposition will perfectly match the frequency of truths in its real reference class for any extension of her credal state to some larger set of propositions.

Given these facts, it follows that:

- (i) For every $Q \subseteq \mathcal{A}$ such that $T(*) \notin Q$, the real calibratability value of $Cr_s(\cdot)$ on Q , given that $Cr_s(\cdot) = C(\cdot)$, is 0.
- (ii) For every $Q \subseteq \mathcal{A}$ such that $T(*) \in Q$, the real calibratability value of $Cr_s(\cdot)$ on Q , given that $Cr_s(\cdot) = C(\cdot)$, is 0.5.
- (iii) For every $Q \subseteq \mathcal{A}$ such that $T(*) \in Q$, the real calibratability value of $Cr_s(\cdot)$ on Q , given that $Cr_s(\cdot) = C'(\cdot)$, for some $C'(\cdot) \neq C(\cdot)$, is greater than 0.5.

And so it follows that:

- (iv) For every $Q \subseteq \mathcal{A}$, the real calibratability value of $Cr_s(\cdot)$ on Q , given that $Cr_s(\cdot) = C(\cdot)$, is less than or equal to the real calibratability value of $Cr_s(\cdot)$ on Q , given that $Cr_s(\cdot) = C'(\cdot)$, for every other $C'(\cdot)$ defined over $\mathcal{A}$.
- (v) There is some $Q \subseteq \mathcal{A}$, e.g., $\{T(*)\}$, such that the real calibratability value of $Cr_s(\cdot)$ on Q , given that $Cr_s(\cdot) = C(\cdot)$, is strictly less than the real calibratability value of $Cr_s(\cdot)$ on Q , given that $Cr_s(\cdot) = C'(\cdot)$, for every other $C'(\cdot)$ defined over $\mathcal{A}$.

We have shown, then, that there is some probabilistically incoherent $C(\cdot)$, defined over $\mathcal{A}$, such that it is a priori that $C(\cdot)$ calibration dominates $C'(\cdot)$ for Annie, for every other $C'(\cdot)$ defined over $\mathcal{A}$. Given Graded Vindication, it therefore follows that it is a priori that Annie would be more vindicated were it to be the case that $Cr_a(\cdot) = C(\cdot)$ than she would be were she to have any other credal state defined over $\mathcal{A}$. Since there is no reason to deny that it is possible for Annie to have this particular credal state, it follows that if Annie has a credal state defined over $\mathcal{A}$, then Graded Vindication and Vindication Maximization entail that Annie ought to have a probabilistically incoherent credal state. Note, furthermore, that $\mathcal{A}$ is a single-object algebra. Thus, it follows that considerations of calibration, instead of motivating Restricted Probabilism in fact provide us with reason to reject this principle.

Acknowledgments

Thanks go to two anonymous referees for extremely helpful comments on an earlier draft of this material.

References

- Caie, Michael (2011). *Paradox and Belief* (Doctoral dissertation, University of California, Berkeley).
- Caie, Michael (2013). Rational Probabilistic Incoherence. *The Philosophical Review*, 122 (4), 527–575. <http://dx.doi.org/10.1215/00318108-2315288>
- Field, Hartry (2008). *Saving Truth from Paradox*. Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199230747.001.0001>
- Joyce, James (1998). A Non-Pragmatic Vindication of Probabilism. *Philosophy of Science*, 65, 575–603. <http://dx.doi.org/10.1086/392661>
- Joyce, James (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In Franz Huber and Christoph Schmidt-Petri (Eds.), *Degrees of Belief* (263–297). Synthese Library.
- Kripke, Saul (1975). Outline of a Theory of Truth. *Journal of Philosophy*, 72 (19), 690–716. <http://dx.doi.org/10.2307/2024634>
- Paris, Jeffrey (2001). A Note on the Dutch Book Method. In Gert De Cooman, Terrence Fine, and Teddy Seidenfeld (Eds.), *Proceedings of the Second International Symposium on Imprecise Probabilities and their Applications* (301–306). Shaker.
- Ramsey, Frank (1931). Truth and Probability. In Richard Braithwaite (Ed.), *Foundations of Mathematics and Other Essays* (156–198). Routledge.
- Shimony, Abner (1988). An Adamite Derivation of the Calculus of Probability. In J. H. Fetzer (Ed.), *Probability and Causality* (79–90). D. Reidel.
- van Fraassen, Bas C. (1983). Calibration: A Frequency Justification for Personal Probability. In R. S. Cohen and L. Laudén (Eds.), *Physics, Philosophy, and Psychoanalysis: Essays in Honor of Adolf Grünbaum* (295–315). D. Reidel.