An accuracy-based approach to quantum conditionalization

Alexander Meehan and Jer Steeger

Abstract

A core tenet of Bayesian epistemology is that rational agents update by conditionalization. Accuracy arguments in favour of this norm are well-known. Meanwhile, scholars working in quantum probability and quantum state estimation have proposed multiple updating rules, all of which look *prima facie* like analogues of Bayesian conditionalization. The most common are Lüders conditionalization and Bayesian mean estimation (BME). Some authors also endorse a lesser-known alternative that we call retrodiction. We show how one can view Lüders and BME as complementary rules, and we give expected accuracy and accuracy dominance arguments for both. By contrast, we find that retrodiction is accuracy-dominated, at least on many measures of accuracy.

1 Introduction

How should rational agents revise their beliefs in response to new information? A core tenet of Bayesian epistemology is the norm of conditionalization, which says (roughly) that if one begins with a rational prior C and learns E, then one should revise one's credences by conditionalizing on E. One particularly influential strategy for justifying this norm involves appealing to the accuracy of one's beliefs. It turns out that non-conditionalizing plans for credal revision lead agents to adopt credences that do relatively poorly, accuracy-wise. In particular, one can show that conditionalization has higher expected accuracy than every other plan (Greaves and Wallace [2006]) and accuracy-dominates every other plan (Briggs and Pettigrew [2020]).

What does this updating norm look like for agents situated in a quantum world? The answer is not immediately apparent. In the quantum setting, subjective and objective probabilities are typically expressed indirectly, via operators on a Hilbert space. Standard formulations of

Bayesian norms do not immediately carry over to this setting. (Here and in what follows, we assume that there are objective quantum probabilities, and that quantum systems have true, objective quantum states. The idea that there are both subjective and objective probabilities, and we can discuss their relationship, is sometimes known as Lewisian dualism (Wallace [2012]).)

While it is unclear how to translate standard Bayesian norms into this quantum operator-theoretic setting, we can draw a natural analogy. Usually, the Bayesian represents both credences and states of the world as real-valued functions over an algebra of propositions; we can view credences as 'estimates' of the truth values assigned by a given state.[1] In quantum theory, one represents states with bounded, linear operators on a Hilbert space. It turns out that one can also use such operators to encode estimates of quantum states. We adopt this paradigm to construct natural analogues of the accuracy arguments, affording novel Bayesian justifications for several quantum updating rules.

It is controversial whether such a non-classical approach to probability is necessary once we fix an appropriate solution to the measurement problem (Loewer, Healey [1994, 2020]). However, even if it is not necessary, this alternate setting is immensely popular. Experimenters typically work within the framework of Hilbert-space operators, as this framework houses most approaches to quantum state estimation (that is, the reconstruction of quantum states based on measurement outcomes; see Jones, Blume-Kohout, Granade *et al.* [1991, 2010, 2016]). So, the status of Bayesian norms in this setting seems well worth exploring.

Interestingly, authors have proposed multiple updating norms in this framework, all of which look prima facie like analogues of Bayesian conditionalization:

1. The rule of Lüders conditionalization (Bub, Bobo, Earman [1977, 2013, 2019]) says that if an agent with a prior estimate ρ, where ρ is some density operator, learns E, the outcome of some projective measurement (represented by a projection operator), then they should adopt the posterior estimate

$$\mathcal{L}_E(\rho) := \frac{E\rho E}{\mathrm{Tr}(E\rho E)}. \tag{1}$$

   More generally, if an agent with prior estimate ρ learns the result m of some POVM measurement M, then they should adopt the posterior estimate

$$\mathcal{L}_{M_m}(\rho) := \frac{M_m \rho M_m^\dagger}{\mathrm{Tr}(M_m \rho M_m^\dagger)}, \tag{2}$$

   where M_m is the operator describing the effective evolution of a system that yields the outcome m when measured by M.[2] We will call this generalization of Lüders conditionalization the credal state reduction rule.

2. The retrodiction rule (Leifer and Spekkens [2013], Eqn. (126)) says that if an agent with prior estimate ρ learns m by measuring M, then, defining $E := M_m^\dagger M_m$, they should adopt the posterior estimate

$$\frac{\rho^{1/2} E \rho^{1/2}}{\mathrm{Tr}(\rho^{1/2} E \rho^{1/2})}, \tag{3}$$

[1] This idea has precedent in (de Finetti [1974]).

[2] We follow the notation of (Nielsen and Chuang [2010], pp. 84-93) and (Paris [2012]). The only constraint on the measurement operators is that the products $M_m^\dagger M_m$ are positive and sum to the identity, and so they form a POVM. A measurement is projective if the $M_m^\dagger M_m$ are projections. Note that we can obtain Equation 1 as a special case of Equation 2 by taking $M_m = E^{1/2} = E = M_m^\dagger M_m$.

which represents the agent's updated beliefs about the early region (Leifer and Spekkens [2013, 2014]), that is, their updated beliefs about the system (or world) as it was before the learning experience took place.

3. The rule of Bayesian mean estimation (BME) (Blume-Kohout [2010], Eqn. 11), similarly to retrodiction, yields a posterior estimate which represents the agent's updated belief about the initial state of the system or world. However, unlike retrodiction, BME's recommendation depends on the choice of decomposition of ρ, a particular probability function p on density operators such that $\sum_k p(k)\rho_k = \rho$.[3] Given a decomposition p, if the agent learns m by measuring M, then, defining $E := M_m^\dagger M_m$, BME recommends the posterior
$$\frac{\sum_k p(k)\,\mathrm{Tr}(\rho_k E)\,\rho_k}{\mathrm{Tr}(\rho E)}. \tag{4}$$

This paper aims to evaluate these three rules through the lens of accuracy. Accuracy arguments for probabilism in the quantum setting have been given (Steeger [2019]), but not yet for quantum conditionalization. The main result of the paper is that both BME and Lüders (more generally, credal state reduction) do the best accuracy-wise and are, in a sense, complementary rules. The retrodiction rule, on the other hand, is the odd one out: there is a precise sense in which it is accuracy-dominated.

Our approach of evaluating quantum updating rules with respect to their accuracy is premised on two background assumptions, which are important to make clear from the outset. First, a metaphysical commitment, mentioned above: we assume that quantum systems have true, objective quantum states and that we can distinguish the objective quantum state of a system from the credal state of an agent. This is crucial as we will evaluate the accuracy of a credal state (at least in part) in terms of its closeness to the true quantum state.

Second, we assume a standard (effective) collapse dynamics applies in cases of measurement. In particular, we assume that, upon measurement, the quantum state of a system (effectively) collapses. That is, if ψ is the true quantum state of a system before measurement, and M is performed on the system with outcome m, then the effective final state is given by the standard dynamical state reduction rule (Paris [2012]):
$$\mathcal{L}_{M_m}(\psi) = \frac{M_m\,|\psi\rangle}{\langle\psi|M_m^\dagger M_m|\psi\rangle}. \tag{5}$$

This assumption of collapse or reduction, or at least effective collapse or reduction—according to which although the quantum state does not collapse, the relevant conditional or relative

[3]In the main text, we will assume that decompositions are discrete (and thus that $\sum_k p(k)\rho_k$ is a countable convex sum of states) for ease of discussion. Although we do not treat them explicitly, continuous distributions also define density operators. Let $\Omega = \{\rho_k\}_{k\in K}$ be any subset of the set of density operators. Given a distribution $\pi(\rho_k)\,\mathrm{d}\rho_k$ on Ω (equipped with the Hilbert-Schmidt topology), where $\mathrm{d}\rho_k$ is some Borel measure, $\pi \in L^1(\Omega)$, and $\int_\Omega \pi(\rho_k)\,\mathrm{d}\rho_k = 1$, one can define the probabilistic mixture
$$\rho = \int_K \pi(\rho_k)\,\rho_k\,\mathrm{d}\rho_k,$$
where ρ itself is a density operator. For this mixture, BME recommends the posterior
$$\frac{\int_K \rho_k\,\pi(\rho_k)\,\mathrm{Tr}(\rho_k E)\,\mathrm{d}\rho_k}{\mathrm{Tr}(\rho E)}.$$
We conjecture that our results will straightforwardly generalize to continuous mixtures, but we leave the task of spelling out this generalization for future work.

quantum state is given by these rules—is motivated by familiar experimental findings. We won't argue for the assumption here, except to note that it follows from a natural commitment to a tight connection between quantum states and the objective probabilities of experimental outcomes. However, one might worry whether this assumption begs the question in favour of the credal state reduction rule. Formally, Equation 2 is just a generalization of Equation 5 to density operators. If the physical dynamics is one of collapse or reduction, shouldn't the credal dynamics be one of collapse or reduction, too?

While we agree that this proposal is exceptionally natural, insofar as one seeks to justify the credal state reduction rule, there remains a gap between the descriptive premise about physical states and any normative conclusions concerning how an agent should update their beliefs. As we will see, even on an accuracy-based approach, there are several complications to fleshing out this argument. For example, while it may be obvious—given the assumption that epistemic rationality is tied to accuracy and that accuracy is measured by closeness to the true state—that an agent who knows the true initial state ψ should adopt credences concerning post-measurement outcomes that match the true post-measurement state $\mathcal{L}_{M_m}(\psi)$, it is far less obvious what to say about an agent who is uncertain about the true initial state. Why should their estimate ρ update by reduction? There is also the issue of the agent's new estimate of the system's initial state—can that update also be viewed in terms of this rule? And if so, in what sense is that rule the best, accuracy-wise? The results below address these questions.

2 A Brief Review of the Synchronic Case

Before tackling these questions concerning conditionalization, it is worth reviewing the synchronic accuracy argument for probabilism in this quantum setting. This argument takes a particularly simple and general form when the background Hilbert space $\mathcal{H}$ is finite-dimensional, so we restrict our attention to this case in what follows.[4]

To start, we need the notion of a state, which corresponds to a possible way the world could be at a time. In light of our dualism, we suppose that this notion at least involves some stipulation of a true quantum state. It is standard to restrict attention to pure vector states ψ. But in what follows, we will only require that the set of possible states $\Omega = \{\rho_k\}_{k\in K}$ (where K is some set of indices) is some subset of the set of density operators on $\mathcal{H}$ that contains the pure states. This leaves room for the possibility that the true quantum state $\rho_k \in \Omega$ is impure.

We aim to assess an agent's estimate of the true quantum state. Let A denote some such estimate; we require that A is a bounded linear operator on $\mathcal{H}$, but nothing further. We say that A is probabilistic if A is a probabilistic mixture of elements of Ω, that is,

$$A = \sum_k p(k)\rho_k, \tag{6}$$

where $p(k)$ gives the probability that ρ_k is the true state (here and hereafter, when assigning probabilities p to operators on Hilbert space, we use only the index of the associated operator in the argument of p for brevity).[5] Following convention, we call such an assignment p a decomposition of A.

[4]See (Steeger [2019]) for a treatment in terms of C*-algebras. While this approach allows for a more flexible choice of quantum states, it requires a specific accuracy measure (namely, the Brier score). However, we conjecture that one can relax this requirement.

[5]For ease of exposition, we suppose that p is a discrete probability function. Rigorously, we define p on $\mathcal{B}(\Omega)$, the Borel sets of Ω equipped with the Hilbert-Schmidt topology, and we suppose that $p(B) = \sum_{\rho_k\in\Omega'} p(k)\delta_{\rho_k}(B)$ for all $B \in \mathcal{B}(\Omega)$, for some countable subset Ω' of Ω that we call the support of p (Çınlar [2011], p. 14). We implicitly take the sum in Equation 6 to range over the support Ω', yielding a countable convex sum of operators (which we

How do we measure the inaccuracy of an estimate A at true state ρ_k? Let $\mathfrak{I}_{\rho_k}(A)$ denote this inaccuracy. One natural choice for the inaccuracy measure $\mathfrak{I}$ is squared Hilbert-Schmidt distance, a kind of mean squared error that is analogous to the classical Brier score:

$$\mathfrak{I}_{\rho_k}(A) = \mathrm{Tr}[(\rho_k - A)^2].$$

Other natural choices are available, however. A common measure in quantum information theory, for example, is quantum relative entropy, which is given by $\mathfrak{I}_{\rho_k}(A) = \mathrm{Tr}[\rho_k(\log \rho_k - \log A)]$.

It turns out that both of these common measures are examples of Bregman divergences. Explicitly, each takes the form $\mathfrak{I}_{\rho_k}(A) = d_f(\rho_k, A)$ for some Bregman divergence d_f:

> Bregman Divergence: Let X be a real or complex topological vector space, and let $f : X \to \mathbb{R}$ be continuously differentiable and strictly convex. The Bregman divergence for f is
>
> $$d_f(x, y) := f(x) - f(y) - \langle \nabla f(x), x - y \rangle,$$
>
> where the last summand is the first term of the Taylor expansion of f at y evaluated at x.

In our case, X is the set $\mathcal{B}(\mathcal{H})$ of bounded linear operators on $\mathcal{H}$ equipped with the Hilbert-Schmidt topology. Squared Hilbert-Schmidt distance corresponds to $f(A) = \mathrm{Tr}[A^2]$, and relative entropy corresponds to $f(A) = \mathrm{Tr}[A \log A]$.

For our arguments, we stipulate that some Bregman divergence gives the measure of inaccuracy. Given that this notion of distance canvasses several intuitive and popular options, it is already of great interest to see what results from this stipulation. It might be interesting to give a principled argument for this stipulation in the quantum context, along the lines of the argument that Pettigrew ([2016]) provides in the classical case. Still, not all familiar quantum distance measures are Bregman divergences; for instance, infidelity, given by

$$\mathfrak{I}_{\rho_k}(A) = 1 - \left(\mathrm{Tr}\sqrt{\sqrt{\rho_k} A \sqrt{\rho_k}}\right)^2,$$

is not. So, for now, we prefer to treat our restriction to Bregman divergences as a natural stipulation rather than a principled choice.

Now that we have a general way to fix a measure $\mathfrak{I}$, we are interested in whether there is a sense in which a probabilistic estimate does better accuracy-wise than a non-probabilistic one. Intuitively, if an agent adopts estimate A and there is some other estimate A' that would be closer to the true state no matter what—in which case we say that A' strongly dominates A—then it seems like there is something wrong with the agent's estimate. Something still seems amiss if A' is always just as close but closer for at least one possible state, in which case we say that A' weakly dominates A. Precisely:

> Accuracy Dominance for Estimates: We say that A' accuracy-dominates A if
>
> $$\mathfrak{I}_{\rho_k}(A') \leq \mathfrak{I}_{\rho_k}(A) \quad \text{for all } k \in K$$
>
> and this inequality is strict for at least one k. If it is strict for all k, we say the dominance is strong; otherwise, we say it is weak.

evaluate as a limit of finite sums in the Hilbert-Schmidt topology). For a discussion of continuous distributions, see Footnote 3.

However, while being accuracy-dominated is a *prima facie* problem, it might not always make sense to demand that agents avoid estimates that are (weakly or strongly) accuracy-dominated by other estimates. Consider a variable that has one possible value, and that value is irrational, but an agent is constrained to choosing a rational number to estimate it. In this setup, every possible estimate is strongly dominated. Yet, it is still permissible for the agent to adopt some estimate in this scenario. Arguably, what this case illustrates is that we should only demand agents avoid dominated estimates if it is clear they can do better.

In the current setting, however, it does turn out that an agent can avoid estimates that are even weakly dominated by other estimates, and they can do so precisely by having an estimate that is probabilistic. In particular, in the current setting, we obtain the following analogue of Joyce ([1998])'s dominance result:

> Theorem 1: An estimate A is not weakly dominated by any other estimate if and only if A is probabilistic.

It will be helpful to get a sense for why this result holds (for the full proof, see the appendix). The first crucial observation is that A is probabilistic—that is, it has a decomposition—if and only if it lies in $\overline{\text{co}}\,\Omega$, the closed convex hull of Ω. The second crucial observation is that all Bregman divergences have a nice projection property: for any closed, convex, and compact subset S of X and any point x outside it, there is a unique point π_x in S that is closest to x, according to the divergence (Predd *et al.* [2009], Proposition 3). This point π_x is, indeed, closer to all points in S than x:

$$\forall y \in S, \quad d_f(y, \pi_x) < d_f(y, x). \tag{7}$$

Combining these observations for $S = \overline{\text{co}}\,\Omega$ eventually yields the result.

With this background in mind, we can move on to the next and more interesting step of quantum updating: revising our estimates in light of evidence. Can accuracy arguments go further and vindicate these revision strategies, too?

3 Unifying Credal State Reduction and BME

When assessing quantum revision strategies, one faces an immediate difficulty: agents learn about physical systems by measuring them, but measuring a quantum system changes its true (effective) state. As mentioned in the introduction, we assume that the usual effective collapse dynamics applies for a given outcome. According to the standard operational description of quantum measurement, this dynamics has precisely the mathematical form of the credal state reduction map $\mathcal{L}_{M_m}$.[6] So, if an agent learns m by measuring M for a system that starts in state ρ_k, the system's final state is $\mathcal{L}_{M_m}(\rho_k)$.

As discussed in the introduction, we only assume that $\mathcal{L}_{M_m}$ gives an effective description of the true final state. In other words, the reader is free to deny that the wavefunction collapses and take $\mathcal{L}_{M_m}(\rho_k)$ to describe a final relative or conditional quantum state. Thus, many approaches to measurement in quantum theory are compatible with our considerations.[7] Furthermore, we only assume that $\mathcal{L}_{M_m}$ gives an effective description of the true final state. In other words, while we stipulate that the actual states of affairs change in the way that $\mathcal{L}_{M_m}$ prescribes, we assume

[6] It might look odd to those new to quantum mechanics that the final state depends, not only on whether a measurement M is performed, but also on which outcome m is obtained. This is arguably a peculiarity of the quantum case, though see the discussion of effective collapse below.

[7] Note, however, that we judge accuracy by closeness to the true relative or conditional state (as opposed to, say, the universal wavefunction).

nothing about how agents should update their beliefs. So the question of whether reduction is the right updating rule for agents remains, at this stage, open.

With these clarifications in mind, note that agents might well be (and usually are) interested in revising their beliefs about both the initial and the final quantum state. So far, we have been using the vector space $\mathcal{B}(\mathcal{H})$ to keep track of the true state and an agent's estimate of it at a given time. However, this space might not be rich enough to describe an agent's beliefs about what the system is like at two different times. Indeed, the reader may have noticed that retrodiction and BME rules have a specific gloss: their posteriors represent the agent's updated beliefs 'about the early region', understood as the system (or world) as it was before the measurement took place. In this sense, one can view retrodiction and BME as merely partial rules—rules that specify how to revise only a certain subset of the agent's beliefs.

It will be useful to model states and beliefs in a richer space. One promising way to do so is to introduce an analogy between our situation and one common in experimental practice: namely, the case where we investigate two systems that we have good reason to believe are initially in the same state. Let

$$\mathcal{A} := \mathcal{B}(\mathcal{H})_{\mathcal{I}} \otimes \mathcal{B}(\mathcal{H})_{\mathcal{F}}$$

denote our new vector space, where we call $\mathcal{I}$ the 'initial' or 'early' system and $\mathcal{F}$ the 'final' or 'late' system (which labels we will justify shortly). Since both systems are in the same state by hypothesis, the set of possible true states Ω contains only symmetric tensor products; explicitly, $\Omega = \{\rho_k \otimes \rho_k\}_{k \in K}$. Let $A \in \mathcal{A}$ be an agent's initial estimate of the two-part system's before-measurement state. Suppose the agent performs some measurement on this two-part system with state $\rho_k \otimes \rho_k$ that yields M_m for the 'final' system but keeps the 'initial' system fixed. By the standard dynamics, the true after-measurement effective state is given by $\mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k)$. One might then ask whether one's posterior estimate a_{M_m} of the two-part system's after-measurement state should be given by the credal reduction rule, applied to the agent's initial estimate, A:

$$a_{M_m} \stackrel{?}{=} \mathcal{L}_{I\otimes M_m}(A). \tag{8}$$

Before addressing this question, we propose that its answer will also resolve our initial query. In particular, this two-system setup is strictly analogous to an agent learning about a single system that starts in an unknown state ρ_k. In the single-system case, the system $\mathcal{I}$ becomes counterfactual: it expresses what would have happened to a second, identical system untouched by a measurement of the first. The counterfactual system's after-measurement state is always identical to the actual system's before-measurement state. Thus, whatever posterior estimate one would adopt of $\mathcal{I}$'s after-measurement state is precisely the posterior estimate one should adopt of the actual system's before-measurement state (hence the label 'initial').

There is good reason for taking this analogy seriously, as it demonstrates a precise sense in which credal state reduction and BME are two sides of the same coin. To start, suppose that A is some probabilistic mixture of elements of Ω, $A = \sum_k p(k)\, \rho_k \otimes \rho_k$, and that this estimate is indeed updated via credal state reduction, as in Equation 8. From the assumption that $A = \sum_k p(k)\, \rho_k \otimes \rho_k$, one can check that:

$$\mathcal{L}_{I\otimes M_m}(A) = \frac{\sum_k p(k)\, \rho_k \otimes M_m \rho_k M_m^\dagger}{\mathrm{Tr}(\sum_k p(k)\, \rho_k \otimes M_m \rho_k M_m^\dagger)}.$$

Then to obtain BME, we simply trace out the final region's degrees of freedom from the estimate that $\mathcal{L}_{I\otimes M_m}$ yields: letting $\rho = \sum_k p(k)\rho_k = \mathrm{Tr}_{\mathcal{F}}[A]$, we have

$$\mathrm{Tr}_{\mathcal{F}}[a_{M_m}] = \mathrm{Tr}_{\mathcal{F}}[\mathcal{L}_{I\otimes M_m}(A)] = \frac{\sum_k p(k)\, \mathrm{Tr}(\rho_k E)\, \rho_k}{\mathrm{Tr}(\rho E)}, \tag{9}$$

where $E = M_m^\dagger M_m$. Recalling Equation 4, we note this is just what BME prescribes. Similarly, we obtain the single-region expression of credal state reduction, Equation 2, by tracing out the initial region's degrees of freedom:

$$\mathrm{Tr}_{\mathcal{I}}[a_{M_m}] = \mathrm{Tr}_{\mathcal{I}}[\mathcal{L}_{I\otimes M_m}(A)] = \frac{M_m \rho M_m^\dagger}{\mathrm{Tr}(M_m \rho M_m^\dagger)}. \tag{10}$$

This unification is powerful in its own right, and it might lead one to suspect that $\mathcal{L}_{I\otimes M_m}$ is the correct credal updating rule.

Applying accuracy considerations strengthens this suspicion. By Theorem 1, an estimate A of the two-part system's before-measurement state is not weakly accuracy-dominated by any other estimate at that time if and only if it is a probabilistic mixture. Thus, A is coherent if and only if it has some decomposition $A = \sum_k p(k)\rho_k \otimes \rho_k$. Likewise, an estimate a_{M_m} of its after-measurement state (for a measurement with outcome $I \otimes M_m$) is not weakly dominated by any other estimate at that time if and only if it is some probabilistic mixture $a_{M_m} = \sum_k p'(k)\rho_k \otimes \mathcal{L}_{M_m}(\rho_k)$. In short, any a_{M_m} that avoids being synchronically accuracy-dominated must be of the form $a_{M_m} = \mathcal{L}_{I\otimes M_m}(A')$ for some probabilistic A'.

These considerations do not, on their own, treat the diachronic case. In other words, they do not demonstrate that an agent should update their estimate A to the specific a_{M_m} that the credal state reduction rule recommends. However, along with the unification of BME and the reduction rule, these considerations do give us strong reason to pursue a diachronic argument for the reduction rule in the case of the two-part system. We can, in fact, give two justifications for this rule in terms of accuracy: one using expected inaccuracy and the other using accuracy dominance.

4 Two Approaches to the Diachronic Case

4.1 Expected inaccuracy

To state our results, we introduce a tool to track an agent's contingency plans for each alternative outcome of a given measurement. Let M denote a POVM measurement, that is, a collection $\{M_m\}$ of operators in $\mathcal{B}(\mathcal{H})$ such that $E_m := M_m^\dagger M_m$ is positive for each m and $\sum_{m=1}^N E_m = 1_{\mathcal{H}}$. (We restrict attention to measurements with a finite number of possible outcomes, that is, such that N is finite.)

We say that an updating plan a for M is a map that specifies, for each possible outcome m of M, the after-measurement estimate a_{M_m} that the agent should adopt according to the plan. Note that each of the three updating rules proposed in the literature naturally yields an updating plan.

How do we evaluate the expected inaccuracy of a plan a relative to a prior A? We need a way of weighting the inaccuracy contribution at each possible state and outcome according to A. Note that for an outcome m of M to be compatible with the effective collapse dynamics, it must be that $\mathrm{Tr}(M_m \rho_k M_m^\dagger) > 0$; otherwise the evolution $\mathcal{L}_{M_m}(\rho_k)$ is ill-defined. Thus we call pairs (ρ_k, M_m) where this condition holds dynamics-compatible pairs, and we define a state-outcome decomposition as a weighting of such pairs:

> State-Outcome Decompositions: Given M and $A \in \overline{\mathrm{co}}\{\rho_k \otimes \rho_k \ : \ k \in K\}$, we define a state-outcome decomposition of A as a probability function p over state-outcome pairs such that $\sum_k p(k) = 1$, $\sum_m p(m) = 1$, and $A = \sum_k p(k)\, \rho_k \otimes \rho_k$, with $p(k, m) = 0$ if (ρ_k, M_m)

is not dynamics-compatible.[8] We say that a state-outcome decomposition is a Born-rule decomposition if, furthermore,

$$p(m \,|\, k) = \mathrm{Tr}(M_m \rho_k M_m^\dagger),$$

for all m and k with $p(k) > 0$.

A Born-rule decomposition weighs each result M_m according to its objective probability (for a particular true state ρ_k; in the case where ρ_k is a pure state corresponding to a vector ψ_k, $\mathrm{Tr}(M_m \rho_k M_m^\dagger) = \langle \psi_k | M_m^\dagger M_m | \psi_k \rangle$). We give Born-rule decompositions special attention in what follows, but we put aside, for now, the question of whether they compel any special normative status. We return to this issue at the end of Section 4.

Now, fix a before-measurement estimate A with some (not necessarily Born-rule) state-outcome decomposition p of A. We define the expected inaccuracy of some updating plan a relative to that decomposition as

$$exp_p(\mathfrak{I}(a)) := \sum_{k,m} p(m \,|\, k)\, p(k)\, \mathfrak{I}_{\mathcal{L}_{I \otimes M_m}(\rho_k \otimes \rho_k)}(a_{M_m}),$$

where we adopt the convention of taking the terms in sums over k and m involving an incompatible pair of ρ_k and M_m to be zero, noting that $p(m \,|\, k) = 0$ in such cases, and similarly for k such that $p(k) = 0$. (We assume these conventions for all subsequent sums over k and m without comment, for ease of exposition.) Moreover, we say that the plan a is conditionalizing for a particular decomposition p just in case

$$a_{M_m} = \sum_k p(k \,|\, m)\, \mathcal{L}_{I \otimes M_m}(\rho_k \otimes \rho_k) \tag{11}$$

for each outcome m of M such that $p(m) > 0$.

We then obtain the following expected inaccuracy result (Blume-Kohout [2010], Sect. 2.3.3), which vindicates both the credal state reduction rule and BME (recall Equations 9 and 10).

> Theorem 2: Fix a before-measurement estimate $A \in \overline{\mathrm{co}}\{\rho_k \otimes \rho_k \;:\; k \in K\}$ with a state-outcome decomposition p. An updating plan a minimizes expected inaccuracy for p if and only if it is a conditionalizing plan for p (regardless of the choice of p). Suppose, moreover, that p is a Born-rule decomposition. Then a minimizes expected inaccuracy for p if and only if
>
> $$a_{M_m} = \mathcal{L}_{I \otimes M_m}(A)$$
>
> for all m such that $p(m) > 0$.

The proof of this theorem is quick, so let us walk through it. The crucial observation is that since $\mathfrak{I}$ is a Bregman divergence, it is strictly proper—that is, according to $\mathfrak{I}$, a probabilistic A has the lowest expected inaccuracy according to any of its decompositions. More formally, we say that $\mathfrak{I}$ is strictly proper if, given any density matrix A, for any decomposition p of A, every estimate A' satisfies

$$\sum_k p(k)\, \mathfrak{I}_{\rho_k}(A') \geq \sum_k p(k)\, \mathfrak{I}_{\rho_k}(A)$$

with equality if and only if $A' = A$. Then we have the following lemma (proof in the appendix):

[8] Rigorously, we define p on $\mathcal{B}(\Omega) \times \sigma(M)$, the product σ-algebra of the Borel sets of Ω and the σ-algebra generated by M. We assume that p is discrete with support on $\Omega' \times M$, where $\Omega' = \{\rho_{k'}\}_{k' \in K'}$ is some countable subset of Ω. We let k denote the disjunction $\bigcup_{M_m \in M}(\rho_k, M_m)$ and m denote the disjunction $\bigcup_{\rho_{k'} \in \Omega'}(\rho_{k'}, M_m)$. Finally, we require that the probability of dynamics-incompatible pairs is zero, that is, that $p(k, m) = 0$ when $\mathrm{Tr}(M_m \rho_k M_m^\dagger) = 0$. Once again, we implicitly take sums in k to range over the support.

Lemma 1: If $\Im$ is a Bregman divergence on $\mathcal{B}(\mathcal{H})$, then $\Im$ is strictly proper.

Like in the classical case (Greaves and Wallace [2006]), it is a short step from strict propriety to the result that the credal state reduction plan minimizes expected inaccuracy. Let p be any state-outcome decomposition of A. Then the p-expected inaccuracy of a is:

$$\begin{aligned} exp_p(\Im(a)) &= \sum_{k,m} p(m\,|\,k)\, p(k)\, \Im_{\mathcal{L}_{I\otimes M_m}(\rho_k\otimes\rho_k)}(a_{M_m}) \\ &= \sum_{k,m:\, p(m)>0} p(k\,|\,m)\, p(m)\, \Im_{\mathcal{L}_{I\otimes M_m}(\rho_k\otimes\rho_k)}(a_{M_m}) \\ &= \sum_{m:\, p(m)>0} p(m) \sum_k p(k\,|\,m)\, \Im_{\mathcal{L}_{I\otimes M_m}(\rho_k\otimes\rho_k)}(a_{M_m}). \end{aligned}$$

Note that for each m such that $p(m) > 0$, $A_m := \sum_k p(k\,|\,m)\, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k)$ is a density matrix. So it follows from strict propriety that this quantity is minimized if and only if $a_{M_m} = A_m$. If p is a Born rule decomposition, then

$$p(k\,|\,m) = \frac{p(k)\,\mathrm{Tr}(M_m\rho_k M_m^\dagger)}{\sum_{k'} p(k')\mathrm{Tr}(M_m\rho_{k'} M_m^\dagger)}.$$

And so by

$$\mathcal{L}_{I\otimes M_m}(A) = \sum_k \frac{p(k)\,\mathrm{Tr}(M_m\rho_k M_m^\dagger)}{\sum_{k'} p(k')\mathrm{Tr}(M_m\rho_{k'} M_m^\dagger)}\, \rho_k \otimes \mathcal{L}_{M_m}(\rho_k), \tag{12}$$

we get that $A_m = \mathcal{L}_{I\otimes M_m}(A)$, which completes the proof of our first accuracy-based argument for the credal state reduction rule.

However, one might harbour suspicions about this argument's reliance on particular decompositions to assess inaccuracies. Briggs and Pettigrew ([2020]) criticize the original argument due to Greaves and Wallace ([2006]) on similar grounds. Roughly, they ask: what is so special about a particular agent's credences at the initial time? Are those the right sort of things to justify a normative updating rule? One might think that facts about the world ought to justify such a rule, and that facts about an agent's subjective state should not play such a central role. If the reader shares this concern, they should note that the same problem hounds the quantum version—but twice over! Recall that the decomposition of A is not, in general, unique. So, the expected inaccuracy argument for the reduction rule depends not only on an agent's initial estimate but also on that agent's choice of decomposition for that estimate.[9]

Moreover, one might doubt the claim that rational agents should always minimize expected inaccuracy. In particular, while this claim follows from the familiar expected utility approach to decision theory (assuming that we view accuracy as a sort of epistemic utility), this general approach has been called into question by thought experiments like the Allais paradox, which aim to demonstrate that rational agents who are risk-averse might reasonably violate it.[10]

One way to address these concerns is to assess the agent's inaccuracies at the two different times directly, so that one is no longer forced to mediate the latter inaccuracy through the agent's expectations. That is precisely the idea behind the dominance approach.

[9] An agent might well have good reason to prefer a particular decomposition. Say, for instance, they flip a coin to decide the orientation of a Stern-Gerlach magnet for the preparation of pure spin states. They ought to pick the A that results from weighing each possible pure state by one-half, the particular decomposition that describes their experimental setup. The point is that there are many other ways of preparing a system well-estimated by A, and our updating rules are functions of the estimate A and the outcome M_m alone.

[10] See (Steele and Stefánsson [2020]) for further discussion of the challenge that the Allais paradox poses to expected utility theory.

4.2 Accuracy dominance

Intuitively, the accuracies of an agent's beliefs before and after measurement matter equally to their rationality. To capture this intuition, we define a strategy for M as a pair (A, a) consisting of a before-measurement estimate A and an updating plan a, and we suppose that a strategy's inaccuracy is the sum of the inaccuracies of its two components. Following Briggs and Pettigrew ([2020]), we call this assumption temporal separability.

> Temporal Separability: For each dynamics-compatible state-outcome pair (ρ_k, M_m), let
>
> $$\mathfrak{I}_{\rho_k\otimes\rho_k}(A) := d_f(\rho_k \otimes \rho_k, A)$$
> $$\mathfrak{I}_{\mathcal{L}_{I\otimes M_m}(\rho_k\otimes\rho_k)}(a_{M_m}) := d_f(\mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k), a_{M_m})$$
>
> for d_f a Bregman divergence for $\mathcal{A}$. The inaccuracy $\mathfrak{I}_{(\rho_k,M_m)}(A, a)$ of a strategy (A, a) given a true initial state $\rho_k \otimes \rho_k$ and outcome $I \otimes M_m$ is given by
>
> $$\mathfrak{I}_{(\rho_k,M_m)}(A, a) := \mathfrak{I}_{\rho_k\otimes\rho_k}(A) + \mathfrak{I}_{\mathcal{L}_{I\otimes M_m}(\rho_k\otimes\rho_k)}(a_{M_m}).$$

Given this assumption, we can appeal to our previous definition of accuracy dominance for estimates to get a precise definition of accuracy dominance for strategies. All we do is substitute each estimate A in that definition with a strategy (A, a), while making sure to quantify over both states ρ_k in Ω and outcomes M_m of M. Explicitly:

> Accuracy Dominance for Strategies: We say that a strategy (A', a') accuracy-dominates another strategy (A, a) if
>
> $$\mathfrak{I}_{(\rho_k,M_m)}(A', a') \leq \mathfrak{I}_{(\rho_k,M_m)}(A, a) \quad \text{for all dynamics-compatible pairs } (\rho_k, M_m),$$
>
> and this inequality is strict for at least one such (ρ_k, M_m). If it is strict for all (ρ_k, M_m), we say that the dominance is strong; otherwise, we say it is weak.

It turns out that we can give a nice characterization of desirable strategies in terms of accuracy dominance: a strategy (A, a) is not strongly dominated by any other strategy if and only if A is probabilistic and a is conditionalizing with respect to A (that is, a satisfies Equation 11 with respect to some state-outcome decomposition p of A).

> Theorem 3: A strategy (A, a) is not strongly dominated by any other strategy if and only if A is probabilistic and a is conditionalizing with respect to A. Moreover, no strategy that fails to satisfy these conditions even weakly dominates any strategy that does.

We relegate the details of this theorem's proof to the appendix, but the reader can get a broad sense of it by noting that we can represent strategies as vectors in a larger space. In particular, by tweaking the strategy of Briggs and Pettigrew ([2020]), we can represent each strategy (A, a) as the vector

$$(A, a) = A \oplus a_{M_1} \oplus \ldots \oplus a_{M_m} \oplus \ldots \oplus a_{M_N} \tag{13}$$

living the $N + 1$-fold direct sum of $\mathcal{A}$ (where N is the number of possible outcomes of the measurement M). Likewise, we can associate a dynamics-compatible state-outcome pair (ρ_k, M_m) with the vector

$$(A, a)_{(\rho_k,M_m)} = (\rho_k \otimes \rho_k) \oplus a_{M_1} \oplus \ldots \oplus \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k) \oplus \ldots \oplus a_{M_N}. \tag{14}$$

It turns out that one can characterize strategies with conditionalizing plans as members of the closed convex hull of these latter vectors. The result follows quickly from this characterization.[11]

Just as Nielsen ([2021]) finds for the classical case, we find that some conditionalizing and probabilistic strategies might weakly dominate others. Recall our earlier qualification, however: this case might well be one where it does not make sense to demand that agents avoid plans that are weakly dominated. For now, we will remain agnostic about the best way to characterize the normative upshot of Theorem 3. Note, however, that allowing an agent to switch amongst various different conditionalizing plans is a fairly cheap price to pay: plans conditionalizing for some prior p can only differ on events that prior assigns zero probability (that is, where $p(m) = 0$).

Still, even if one thinks that Theorem 3 straightforwardly shows that agents must adopt a conditionalizing plan, it does not follow that this plan must be Lüders conditionalization or, more generally, the credal state reduction plan. If a is the reduction plan and A is probabilistic, then a is conditionalizing (recall Equation 12), and so the strategy (A, a) is not strongly accuracy-dominated. But what about the converse? Just as in the case of minimizing expected inaccuracy, the converse holds when we restrict attention to Born-rule decompositions. As long as agents are independently required to avoid state-outcome decompositions that do not respect the Born rule, the above accuracy dominance argument implies that they must update their estimates according to the credal state reduction rule.

We are inclined to view this Born rule requirement as a chance-credence norm akin to the Principal Principle (Lewis, Wallace, Earman, Meehan [1980, 2012, 2018, 2021]). But what justification, if any, can be provided for it? One might try to adapt derivations of the Born rule found elsewhere in the literature. However, they tend to be housed within particular approaches to measurement—a matter on which we want to remain neutral.[12] Thus, for now, we remain agnostic about the best way to justify a restriction to Born-rule-respecting decompositions.

While the status of the Born rule is an interesting subtlety in these results, it is notable that accuracy-based tools go so far in vindicating the credal reduction rule and BME as the uniquely rational strategies for revising estimates in light of evidence. This result is of interest not just to accuracy-firsters but to any theorist who thinks there is a tight connection between optimizing accuracy and epistemic rationality.

5 Retrodiction

Let's take stock. We began with three candidate rules—reduction, retrodiction, and BME—and sought to evaluate these rules in terms of their accuracy performance. We found a precise sense in which accuracy considerations support reduction and BME, which can be viewed as complementary rules (recall Equations 9 and 10). Notably, our results are a bit subtler than the classical theorems. We only obtain the converse directions of the results—that following BME and reduction is necessary—if we restrict attention to decompositions that respect the Born rule. However, one might reasonably treat this restriction as a chance-credence norm.

[11] It is also worth noting that our analogue of the expected inaccuracy argument plays a crucial role in this characterization, just as it does for Briggs and Pettigrew ([2020]) and Nielsen ([2021]). See the appendix for more details.

[12] For example, Everettians often justify the Born rule by appealing to symmetries of the quantum state space (Deutsch, Wallace, Sebens and Carroll [1999, 2012, 2018]). However, it is a matter of some controversy whether this strategy requires the Everett interpretation: Wallace ([2012]) argues that it does, while Steeger ([2022]) argues that Bohmians can use it, too. Conditional on your stance in this debate, symmetry might offer one measurement-neutral way to complete our justification of the credal reduction rule.

What about retrodiction? Leifer and Spekkens ([2013], p. 27) suggest that this rule is a more faithful analogue of Bayesian conditionalization. However, it turns out that accuracy considerations speak decidedly against this rule. Indeed, retrodiction can be accuracy-dominated, even if one sets the Born rule aside. We end by noting this result. Fixing M and $A \in \overline{\mathrm{co}}\{\rho_k \otimes \rho_k \ :\ k \in K\}$, say that (A, a) is a retrodiction strategy given A if, for all m,

$$\mathrm{Tr}_{\mathcal{F}}[a_{M_m}] = \frac{A_{\mathcal{I}}^{1/2} E_m A_{\mathcal{I}}^{1/2}}{\mathrm{Tr}(A_{\mathcal{I}} E_m)},$$

where $E_m := M_m^\dagger M_m$ and $A_{\mathcal{I}} := \mathrm{Tr}_{\mathcal{F}}(A)$; in other words, the posterior estimate of the initial state is given by Equation 3. Then we have the following result:

> Proposition 1: There are cases where no retrodiction strategy given A has any decomposition (Born-rule or not) that avoids dominance. In other words, there are cases where every retrodiction strategy given A is strongly accuracy-dominated.

Of course, it is still open to proponents of the rule to dispute various assumptions of our accuracy setup here. They might, for example, challenge our restriction of accuracy scores to Bregman divergences—or they might deny a tight connection between accuracy and epistemic rationality altogether.[13] Perhaps most dramatically, they might reject the assumption of objective quantum states (that is, Lewisian dualism) at the centre of our paradigm. QBism, for example, famously denies the objectivity of quantum states, and Fuchs ([2003], §6) gives an earlier presentation of the retrodiction rule within that framework. QBists like Fuchs will not be perturbed by our results. Nonetheless, it is illuminating to find that the rule's tenability might well hinge on the QBist's strong subjectivism.

6 Conclusion

The landscape of quantum updating rules is complex, and several interesting formal and normative issues arise when we examine their accuracy performance—including issues concerning the role of extra-probabilistic constraints—that do not directly arise in the familiar classical setting. We do not take this paper as the final word on these issues. Rather, we hope it serves as a starting point for the application of accuracy-based tools to quantum conditionalization.

Appendix

This appendix collects the proofs we suppress in the main text; we group them together by the sections to which they pertain.

A.1 A brief review of the synchronic case

Proof of Theorem 1: First, we prove that no member of $\overline{\mathrm{co}}\,\Omega$ is even weakly dominated, which suffices to prove the 'if' direction. Then we prove the other direction.

[13] Leifer and Spekkens ([2013]) note that retrodiction is related to the protocol of pretty-good measurement (Hausladen and Wootters [1994]). Our accuracy considerations do not target pretty-good measurement, which is a protocol for deciding which POVM measurement to perform rather than a rule for how to update one's estimates given the result of a fixed POVM. Still, one might wonder how retrodiction fares specifically in such situations. One can check that all pretty-good measurement situations admit retrodiction strategies with dominance-avoiding decompositions. However, it may still be the case that none of those dominance-avoiding decompositions respect the Born rule.

(⇐) Suppose $A \in \overline{\text{co}}\,\Omega$. Now suppose towards a contradiction that some other $A' \in \mathcal{B}(\mathcal{H})$ weakly dominates it. In particular, for every ρ, $d_f(\rho, A') \leq d_f(\rho, A)$. Note that $d_f(\rho, A) - d_f(\rho, A')$ depends linearly on ρ, and so convex sums over ρ preserve the inequality. A is one such sum; thus, $d_f(A, A') \leq d_f(A, A)$. But now note that the strict convexity of f implies that $d_f(A, A') \geq 0$ with equality if and only if $A = A'$. So it must be that $A = A'$, for a contradiction.

(⇒) Suppose $A \notin \overline{\text{co}}\,\Omega$. Exploit the projection property of d_f to define π_A, the closest point in $\overline{\text{co}}\,\Omega$ to A. By Equation 7, π_A strongly dominates A, and by the first step of this proof, π_A is not even weakly dominated by any other estimate.

□

A.2 Expected inaccuracy

Proof of Lemma 1: Suppose $\mathfrak{I}$ is given by $\mathfrak{I}_{\rho_k}(A) := d_f(\rho_k, A)$ for a Bregman divergence for $\mathcal{B}(\mathcal{H})$. One can check by the strict convexity of f that it is truth-directed in the sense $d_f(\rho_k, A) \geq d_f(\rho_k, \rho_k)$ with equality if and only if $A = \rho_k$. Define

$$G(\rho_k, A) := d_f(\rho_k, A) - f(\rho_k) = -f(A) - \text{Tr}\left[(\rho_k - A)\nabla f(A)\right].$$

Now fixing a density matrix A and a decomposition p of A, note that, for all estimates A',

$$\sum_k p(k) G(\rho_k, A') = G\left(\sum_k p(k)\rho_k,\ A'\right) = G(A, A')$$

by the linearity of trace. So we have:

$$\begin{aligned}\sum_k p(k)\mathfrak{I}_{\rho_k}(A') &= \sum_k p(k)(G(\rho_k, A') + f(\rho_k)) \\ &= G(A, A') + \sum_k p(k)\, f(\rho_k) \\ &= d_f(A, A') - f(A) + \sum_k p(k)\, f(\rho_k).\end{aligned}$$

By truth-directedness, the first term is minimized if and only if $A' = A$. The remaining terms are independent of the estimate A'. So it follows that $\sum_k p(k)\,\mathfrak{I}_{\rho_k}(A') \geq \sum_k p(k)\,\mathfrak{I}_{\rho_k}(A)$ with equality if and only if $A' = A$. □

Theorem 2 is proved in the main text.

A.3 Accuracy dominance

To start, recall the vectors (A, a) and $(A, a)_{(\rho_k, M_m)}$ defined in Equations 13 and 14, respectively. With this structure in place, we prove the accuracy dominance theorem (Theorem 3) in two lemmas, following the basic strategy of Briggs and Pettigrew ([2020]) and Nielsen ([2021]). Like Nielsen, we do not require the transfinite induction step in Briggs and Pettigrew's original proof.

> Lemma 2: A strategy (A, a) lives in the closed convex hull of the $(A, a)_{(\rho_k, M_m)}$s if and only if $A \in \overline{\text{co}}\{\rho_k \otimes \rho_k \ :\ k \in K\}$ and there is a state-outcome decomposition p of A such that $a_{M_m} = \sum_k p(k\,|\,m)\, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k)$.

Proof: The proof follows a similar format to (Briggs and Pettigrew [2020], Lemma 2):

(⇒) Suppose (A, a) lies in the closed convex hull of the $(A, a)_{(\rho_k,M_m)}$s. That means that there is some sequence of non-negative real numbers $\lambda_{k,m}$ such that $(A, a) = \sum_{k,m} \lambda_{k,m} (A, a)_{(\rho_k,M_m)}$ and $\sum_{k,m} \lambda_{k,m} = 1$. It follows that $A = \sum_{k,m} \lambda_{k,m} \rho_k \otimes \rho_k$ and so $A \in \overline{\text{co}}\{\rho_k \otimes \rho_k \,:\, k \in K\}$. Now define the unique probability function $p : \mathcal{B}(K) \times \sigma(M) \to [0, 1]$ by extension from $p(k, m) := \lambda_{k,m}$ where defined and zero otherwise. It is straightforward to check that p is a state-outcome decomposition of A. It also follows that, for each m,

$$\begin{aligned} a_{M_m} &= \sum_{k,n\neq m} \lambda_{k,n}\, a_{M_m} + \sum_k \lambda_{k,m}\, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k) \\ &= \left(1 - \sum_k \lambda_{k,m}\right) a_{M_m} + \sum_k \lambda_{k,m}\, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k) \\ &= a_{M_m} - p(m)\, a_{M_m} + \sum_k p(k, m)\, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k) \end{aligned}$$

and so $a_{M_m} = \sum_k p(k \,|\, m)\, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k)$ whenever $p(m) > 0$, as claimed.

(⇐) Suppose p is a state-outcome decomposition of A such that

$$a_{M_m} = \sum_k p(k \,|\, m) \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k)$$

whenever $p(m) > 0$. It suffices to show

$$(A, a) = \sum_{k,m} \lambda_{k,m}\, (A, a)_{(\rho_k,M_m)},$$

where $\lambda_{k,m} := p(k, m)$. It is immediate that $A = \sum_{k,m} \lambda_{k,m} \rho_k \otimes \rho_k$. For each m, we note that our assumption entails

$$a_{M_m} = a_{M_m} - a_{M_m} \sum_k \lambda_{k,m} + \sum_k \lambda_{k,m}\, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k),$$

where in the case $p(m) := \sum_k \lambda_{k,m} = 0$, this equality holds automatically. Running the above equalities backward we obtain $a_{M_m} = \sum_{k,n\neq m} \lambda_{k,n}\, a_{E_m} + \sum_k \lambda_{k,m}\, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k)$, and the conclusion follows.

□

Now consider a Bregman divergence d_F on the $N + 1$-fold direct sum that is additive, that is, such that

$$\begin{aligned} d_F\left((A, a), (A, a)_{(\rho_k,M_m)}\right) &= d_f(a, \rho_k \otimes \rho_k) + d_f(a_{M_1}, a_{M_1}) + \ldots \\ &\quad + d_f(a_{M_m}, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k)) + \ldots + d_f(a_{M_N}, a_{M_N}) \\ &= d_f(a, \rho_k \otimes \rho_k) + d_f(a_{M_m}, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k)). \end{aligned}$$

where d_f is a Bregman divergence for $\mathcal{A}$. The squared norm of $\bigoplus^{N+1} \mathcal{A}$ provides one suitable F. In fact, any Bregman divergence d_f on $\mathcal{A}$ yields such an F: simply let

$$F(A, a_M) = f(A) + f(a_{M_1}) + \ldots + f(a_{M_N}),$$

recalling that strict convexity and continuous differentiability are preserved by sums, as well as the fact $\langle \nabla F(A, a_M), \cdot \rangle$ is linear.

Proposition 2: Let d_f be some Bregman divergence on $\mathcal{A}$. There exists a Bregman divergence d_F on $\bigoplus^{N+1}\mathcal{A}$ that is additive in the sense above.

Proof: Define $F(A, a_M) := f(A) + f(a_{M_1}) + \ldots + f(a_{M_N})$. To see that F is strictly convex, pick (A, a) and (A', a') in $\bigoplus^{N+1}\mathcal{A}$ such that $(A, a) \neq (A', a')$. Then pick some $0 < \lambda < 1$. We have

$$\begin{aligned} F(\lambda(A, a) + (1-\lambda)(A', a')) &= f(\lambda A + (1-\lambda)A') + \sum_i f(\lambda a_{M_m} + (1-\lambda)a'_{M_m}) \\ &< \lambda f(A) + (1-\lambda)f(A') + \sum_i \lambda f(a_{M_m}) + (1-\lambda)f(a'_{M_m}). \end{aligned}$$

Note, too, that F inherits continuous differentiability from the summed fs. Then we exploit the fact that the derivative $\langle \nabla F(A, a), \cdot \rangle$ is linear to recover

$$\begin{aligned} d_F\left((A, a), (A', a')\right) &= F(A, a) - F(A', a') - \langle \nabla F(A', a'), (A, a) - (A', a') \rangle \\ &= f(A) - f(A') - \langle \nabla f(A'), A - A' \rangle \\ &\quad + \sum_i f(a_{M_m}) - f(a'_{M_m}) - \langle \nabla f(a'_{M_m}), a_{M_m} - a'_{M_m} \rangle \\ &= d_f(A, A') + \sum_i d_f(a_{M_m}, a'_{M_m}), \end{aligned}$$

as desired. □

Given this proposition (and our earlier synchronic arguments), one suspects that the geometry of $\bigoplus^{N+1}\mathcal{A}$ might suffice, on its own, to characterize when a strategy avoids accuracy dominance. The following lemma shows that this is indeed the case.

Lemma 3: A strategy (A, a) avoids strong accuracy dominance if and only if it lives in the closed convex hull of the $(A, a)_{(\rho_k, M_m)}$s. Moreover, no strategy in that hull is weakly accuracy-dominated by any strategy outside it.

Proof: To start, we suppose that the Bregman divergence d_f fixes the inaccuracy score $\mathfrak{I}_{(\rho_k, M_m)}$ for strategies. Using Proposition 2, we let d_F be a Bregman divergence on $\bigoplus^{N+1}\mathcal{A}$ that is additive for d_f.

We first show (i) that any (A, a) outside the closed convex hull of the $(A, a)_{(\rho_k, M_m)}$s is strongly accuracy-dominated, demonstrating the 'only if' direction of the first part of the lemma. Then we show (ii) the 'if' direction, from which it quickly follows that strategies outside the hull cannot weakly accuracy-dominate strategies inside it.

(i) Suppose that (A, a) lives outside the closed convex hull of the $(A, a)_{(\rho_k, M_m)}$s. By the projection property, there is a point (A', a') in the hull that is closest to it, by the lights of d_F. That is, for each possible (ρ_k, M_m),

$$d_F\left((A, a)_{(\rho_k, M_m)}, (A', a')\right) < d_F\left((A, a)_{(\rho_k, M_m)}, (A, a)\right). \tag{15}$$

Moreover, note that

$$\begin{aligned} d_F\left((A, a)_{(\rho_k, M_m)}, (A, a)\right) &= d_f(\rho_k \otimes \rho_k, A) + d_f(\mathcal{L}_{I \otimes M_m}(\rho_k \otimes \rho_k), a_{M_m}) \\ &\quad + \sum_{n \neq m} d_f(a_{M_n}, a_{M_n}) \\ &= d_f(\rho_k \otimes \rho_k, A) + d_f(\mathcal{L}_{I \otimes M_m}(\rho_k \otimes \rho_k), a_{M_m}) \\ &= \mathfrak{I}_{(\rho_k, M_m)}(A, a) \end{aligned}$$

and

$$
\begin{aligned}
d_F\left((A,a)_{(\rho_k,M_m)},(A',a')\right) &= d_f(\rho_k\otimes\rho_k, A) + d_f(\mathcal{L}_{I\otimes M_m}(\rho_k\otimes\rho_k), a'_{M_m}) \\
&\quad + \sum_{n\neq m} d_f(a_{M_n}, a'_{M_n}) \\
&\geq d_f(\rho_k\otimes\rho_k, A') + d_f(\mathcal{L}_{I\otimes M_m}(\rho_k\otimes\rho_k), a_{M_m}) \\
&= \Im_{(\rho,M_m)}(A',a'),
\end{aligned}
$$

and so the strict inequality in (15) yields

$$\Im_{(\rho_k,M_m)}(A',a') < \Im_{(\rho_k,M_m)}(A,a),$$

and so (A',a') strongly dominates (A,a).

(ii) We now show that no (A,a) in the closed convex hull of the $(A,a)_{(\rho_k,M_m)}$s is strongly accuracy-dominated. Suppose towards a contradiction that it were, by some (A',a'). Let the expected inaccuracy of a strategy (A,a) according to a state-outcome decomposition p be given by

$$
\begin{aligned}
exp_p(\Im(A,a)) &:= \sum_{k,m} p(m\,|\,k)\,p(k)\,\Im_{(\rho_k,M_m)}(A,a) \\
&= \left(\sum_{k,m} p(m\,|\,k)\,p(k)\,\Im_{(\rho_k\otimes\rho_k)}(A)\right) + exp_p(\Im(a)), \\
&= \left(\sum_{k} p(k)\,\Im_{(\rho_k\otimes\rho_k)}(A)\right) + exp_p(\Im(a)),
\end{aligned}
$$

where the last equality follows from temporal separability. By Lemma 1, d_f is strictly proper. Then, by its strict propriety (directly for the first summand, and by Theorem 2 for the second), we can conclude that

$$exp_p(\Im(A,a)) \leq exp_p(\Im(A',a')).$$

However, on the assumption that (A',a') strongly dominates (A,a), we have that

$$\Im_{(\rho_k,M_m)}(A',a') < \Im_{(\rho_k,M_m)}(A,a)$$

for all possible (ρ_k, M_m). It would then follow that $exp_p(\Im(A',a')) < exp_p(\Im(A,a))$, for a contradiction.

Furthermore, no strategy outside the closed convex hull even weakly dominates (A,a). Suppose towards a contradiction that (A',a') did so. Then by part (i) of this proof, some (A'',a'') would strongly dominate (A',a'). But (A'',a'') would also strongly dominate (A,a), contradicting what we just proved.

□

Proof of Theorem 3: Immediate from Lemmas 2 and 3. □

A.4 Retrodiction

Proof of Proposition 1: Letting $P_x = |\uparrow_x\rangle\langle\uparrow_x|$ and $P_x^\perp = |\downarrow_x\rangle\langle\downarrow_x|$, we can consider

$$A = \frac{2}{3}\, P_x \otimes P_x + \frac{1}{3}\, P_x^\perp \otimes P_x^\perp \tag{16}$$

and, letting $P_z = |\uparrow_z\rangle\langle\uparrow_z|$ and $P_z^\perp = |\downarrow_z\rangle\langle\downarrow_z|$, we can consider M given by

$$M_1 = P_z,\ M_2 = P_z^\perp.$$

Note that A is a separable symmetric state, so it has a unique decomposition into symmetric pure states (Qian *et al.* [2020], Lemma 17), namely the one given by Equation 16. Additionally, one can check that, in this case, in this case, for both possible results $m = 1$ and $m = 2$ of M (and letting $\mathrm{Tr}_{\mathcal{I}} A = A_{\mathcal{I}}$),

$$P_m := \frac{A_{\mathcal{I}}^{1/2} M_m A_{\mathcal{I}}^{1/2}}{\mathrm{Tr}(A_{\mathcal{I}}^{1/2} M_m A_{\mathcal{I}}^{1/2})},$$

is a pure state that do not equal P_x or $P_x^\perp$. Let (A, a) be any retrodiction strategy for this A given M, so $\mathrm{Tr}_{\mathcal{F}}[a_{M_1}] = P_1$ and $\mathrm{Tr}_{\mathcal{F}}[a_{M_2}] = P_2$. Suppose toward contradiction that there is a state-outcome decomposition p of A such that, for both $m = 1$ and $m = 2$,

$$p(m)\, a_{M_m} = \sum_k p(m\,|\,k)\, p(k)\, \mathcal{L}_{I\otimes M_m}(\rho_k \otimes \rho_k).$$

It follows that, for all m,

$$p(m)\, \mathrm{Tr}_{\mathcal{F}}[a_{M_m}] = \sum_k p(m\,|\,k)\, p(k)\, \rho_k$$

and so, for all m,

$$p(m)\, P_m = \sum_k p(m\,|\,k)\, p(k)\, \rho_k.$$

Since each P_m is a pure state, this is only possible if, for all m, $p(m\,|\,k)\, p(k) = p(m)$ when $\rho_k = P_m$ and 0 otherwise. And so it follows that

$$A = \sum_{k,m} p(m\,|\,k)\, p(k)\, \rho_k \otimes \rho_k = \sum_m p(m)\, P_m \otimes P_m,$$

which contradicts the fact that Equation 16 is the only decomposition into symmetric pure states. One can find additional cases by using a spin measurement along a different direction for the measurement $M = \{M_1, M_2\}$. □

Acknowledgements

The authors would like to thank Kade Cicchella, Benjamin Feintzeig, Hans Halvorson, Mark Maxwell, Laura Ruetsche, the audience at the workshop Non-Classical Probability: Formal Epistemology Meets Foundations of Physics (on July 5 and 6, 2022, at the University of Bristol), audience members at PSA 2022, and two anonymous referees for helpful discussions of early versions of this work. Jer Steeger was supported during the completion of this work by the National Science Foundation under grant #2043089.

Alexander Meehan
Department of Philosophy
Yale University
City, USA
alexander.meehan@yale.edu

Jer Steeger
Department of Philosophy
University of Washington
Seattle, WA, USA
jsteeger@uw.edu

References

Blume-Kohout, R. [2010]: 'Optimal, reliable estimation of quantum states', *New Journal of Physics*, **12**(4) 043034.

Bobo, I. G. [2013]: 'On quantum conditional probability', *Theoria: An International Journal for Theory, History and Foundations of Science*, **28**(1), pp. 115–137.

Briggs, R. and Pettigrew, R. [2020]: 'An Accuracy-Dominance Argument for Conditionalization', *Noûs*, **54**(1), pp. 162–181.

Bub, J. [1977]: 'von Neumann's projection postulate as a probability conditionalization rule in quantum mechanics', *Journal of Philosophical Logic*, **6**(1), pp. 381–390.

Çınlar, E. [2011]: *Probability Theory and Stochastics*, New York: Springer.

de Finetti, B. [1974]: *Theory of Probability*, New York: Wiley.

Deutsch, D. [1999]: 'Quantum Theory of Probability and Decisions', *Proceedings: Mathematical, Physical and Engineering Sciences*, **455**(1988), pp. 3129–3137.

Earman, J. [2018]: 'The Relation between Credence and Chance: Lewis' "Principal Principle" Is a Theorem of Quantum Probability Theory', available at <http://philsci-archive.pitt.edu/14822/>.

Earman, J. [2019]: 'Lüders conditionalization: Conditional probability, transition probability, and updating in quantum probability theory', available at <http://philsci-archive.pitt.edu/15914/>.

Fuchs, C. A. [2003]: 'Quantum mechanics as quantum information, mostly', *Journal of Modern Optics*, **50**(6-7), pp. 987–1023.

Granade, C., Combes, J. and Cory, D. [2016]: 'Practical Bayesian tomography', *New Journal of Physics*, **18**(3) 033024.

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

Hausladen, P. and Wootters, W. K. [1994]: 'A 'pretty good' measurement for distinguishing quantum states', *Journal of Modern Optics*, **41**(12), pp. 2385–2390.

Healey, R. [2020]: ‘Is Quantum Mechanics a New Theory of Probability?’, in M. Hemmo and O. Shenker (*eds*), *Quantum, Probability, Logic: The Work and Influence of Itamar Pitowsky*, Jerusalem Studies in Philosophy and History of Science, Springer International Publishing, pp. 317–336.

Jones, K. R. W. [1991]: ‘Principles of quantum inference’, *Annals of Physics*, **207**(1), pp. 140–170.

Joyce, J. M. [1998]: ‘A nonpragmatic vindication of probabilism’, *Philosophy of science*, **65**(4), pp. 575–603.

Leifer, M. S. and Spekkens, R. W. [2013]: ‘Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference’, *Physical Review A*, **88**(5) 052130.

Leifer, M. S. and Spekkens, R. W. [2014]: ‘A Bayesian approach to compatibility, improvement, and pooling of quantum states’, *Journal of Physics A: Mathematical and Theoretical*, **47**(27) 275301.

Lewis, D. K. [1980]: ‘A Subjectivist’s Guide to Objective Chance’, in R. C. Jeffrey (*ed.*), *Studies in Inductive Logic and Probability*, University of California Press, pp. 83–132.

Loewer, B. [1994]: ‘Probability and quantum theory’, in P. Suppes and P. Humphreys (*eds*), *Patrick Suppes: Scientific Philosopher: Volume 2. Philosophy of Physics, Theory Structure, and Measurement Theory*, Springer Netherlands, pp. 3–28.

Meehan, A. [2021]: ‘States of ignorance and ignorance of states: examining the Quantum Principal Principle’, *Studies in History and Philosophy of Science Part A*, **90**, pp. 160–167.

Nielsen, M. [2021]: ‘Accuracy-dominance and conditionalization’, *Philosophical Studies*, **178**(10), pp. 3217–3236.

Nielsen, M. A. and Chuang, I. L. [2010]: *Quantum Computation and Quantum Information: 10th Anniversary Edition*, Cambridge: Cambridge University Press.

Paris, M. G. A. [2012]: ‘The modern tools of quantum mechanics’, *The European Physical Journal Special Topics*, **203**(1), pp. 61–86.

Pettigrew, R. [2016]: *Accuracy and the Laws of Credence*, Oxford: Oxford University Press.

Predd, J. B., Seiringer, R., Lieb, E. H., Osherson, D. N., Poor, H. V. and Kulkarni, S. R. [2009]: ‘Probabilistic Coherence and Proper Scoring Rules’, *IEEE Transactions on Information Theory*, **55**(10), pp. 4786–4792.

Qian, L., Chen, L. and Chu, D. [2020]: ‘Separability of symmetric states and vandermonde decomposition’, *New Journal of Physics*, **22**(3) 033019.

Sebens, C. T. and Carroll, S. M. [2018]: ‘Self-locating uncertainty and the origin of probability in Everettian quantum mechanics’, *The British Journal for the Philosophy of Science*, **69**(1), pp. 25–74.

Steeger, J. [2019]: ‘Probabilism for Stochastic Theories’, *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, **66**, pp. 34–44.

Steeger, J. [2022]: ‘One world is (probably) just as good as many’, *Synthese*, **200**(2) 97.

Steele, K. and Stefánsson, H. O. [2020]: ‘Decision Theory’, in E. N. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, Winter 2020 edition.

Wallace, D. [2012]: *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*, Oxford University Press.