Finding Balance in Uncertain Times

Adrian Liu, adrian.liu@rutgers.edu
Formal Epistemology Workshop, University of Iceland, Reykjavík, May 27th, 2025

§1. The Promises and Perils of Reflection

This handout also at adrianliu.net/balance

Let π be the credence function that describes the credences you in fact have and let R be a description of the rational credences to have in your situation, *whatever they might be*.

Let $[R(q) = t]$ be the proposition that the rational credence to have in proposition q is t. Then you **reflect** the rational opinions on a q if your credence in q, conditional on the indicative supposition that the rational credence to have in proposition q is t, is t.

The credences described needn't be the *rational* ones: this is just an important example. See footnote 3.

Fact 1. *If some prior credence* π reflects *a posterior credence* P*, definitely described as above, then three other features also hold:*

1. *π **values** P: π expects P to always give better estimates than π.*
2. *π **balances** R: π expects P to have the same estimates as π on average.*
3. *P is **higher-order certain**: P is certain about the actual value of P*

If an agent with credences π expects to rationally respond to some evidence, but isn't sure what the evidence will be, then they won't know what the rational credences upon getting new evidence are. But P^+ can describe the opinions the agent with prior π will have after updating on their evidence, whatever it is.[1]

[1] It follows from Fact 1 that if π is rationally required to **reflect** P^+,
1. π is rationally reqd to **value** P^+.
2. π is rationally reqd to **balance** P^+.
3. P^+ is rationally required to be **higher-order certain**.

Thus if any of (1–3) is too strong to be a rational requirement on updating, then reflection is also too strong.

Some think that *reflection* is too strong to be a rational requirement on updating. Williamson (1997), Elga (2013) and Dorst (2020) think *reflection* is too strong because *higher-order certainty* is too strong.

1. If *reflection* is a rational requirement, then an agent who updates on evidence E must always know that their evidence is E.
2. But it can be rational to be unsure what your evidence is.
3. So *reflection* cannot be a rational requirement.

If we suppose that *reflection* is too strong because *higher-order certainty* is too strong, then we think that both *reflection failures* and *higher-order uncertain posteriors* can be rational. What about **value** and **balance**?

× π is rationally required to **reflect** P^+.
1. ? π is rationally reqd to **value** P^+.
2. ? π is rationally reqd to **balance** P^+.
3. × P^+ is rationally reqd to be **higher-order certain**.

This Talk:
Question: how much balance can you get in conditions of higher-order uncertainty? **Answer**: Not much at all.

1. We should want conditional as well as unconditional formulations of balance, but these are hard to get.
2. It's hard for higher-order uncertain priors to balance higher-order uncertain posteriors, which may mean that balanced updates from higher-order uncertain priors are hopeless.

Upshot: It will be really hard to both allow higher-order uncertainty and require balance in one's theory of rational epistemic updating.

§2. Formalities

Definition 1. *Let $\mathbb{C}(\Omega)$ be the set of probabilistic credence functions*[2] *over Ω. A **credence family** is a function $P : \Omega \to \mathbb{C}(\Omega)$, mapping each world to a credence function. P_w, the value of P at w, is a credence function.*[3]

[2] A credence function over a universe Ω of possible worlds is a function $\mathcal{P}(\Omega) \to \mathbb{R}$.For finite cases, $\mathcal{P}$ is the powerset. For infinite cases, it is a σ-algebra over Ω. I will deal with probabilistic credence functions.

[3] Informally, we can think of a credence family as a definite description of a credence function, like 'the rational credences', or 'Gabrielle's credences', or 'the credences I will have after I get new evidence'.

Definition 2. *A credence function π **reflects** a credence family P if for any proposition q and any $t \in [0,1]$,*

$$\pi(q \mid [P(q) = t]) = t.$$

Definition 3. *A credence function family $P : \Omega \to \mathbb{R}$ is **higher-order certain***[4] *if for every $w \in \Omega$,*

$$P_w([P = P_w]) = 1,$$

*where $[P = P_w] =_{def} \{w' \in \Omega \mid P_{w'}) = P_w\}$, and **higher-order uncertain** if for some $w \in \Omega$, $P_w([P = P_w]) < 1$.*

[4] In this formalism (Dorst 2019), higher-order uncertainty is always equivalent to uncertainty about worlds already in the underlying algebra. So P is not uncertain about what values P takes on at different worlds. P is uncertain *which world* is actual, and thus what the actual value of P_w is.

Definition 4. *Let a **random variable** over a set Ω be any function $\Omega \to \mathbb{R}$. The **expectation***[5] *of X over Ω relative to π is given by*

$$\mathbb{E}_\pi[X] =_{def} \sum_{w \in \Omega} \pi(w) \cdot X(w).$$

[5] Informally, this is the *average value* of X over the possible worlds $w \in \Omega$, weighted by how likely each world w is, according to π.

Definition 5. *π **balances** P^+ on a random variable*[6] *X if*

$$\mathbb{E}_\pi[X] = \mathbb{E}_\pi[\mathbb{E}_{P^+}[X]],$$

where $\mathbb{E}_{P^+}[X]$ is itself a random variable $w \mapsto \mathbb{E}_{P^+_w}[X]$, so the double expectation is well-defined.

[6] π **balances** P^+ on a proposition q if

$$\pi(q) = \mathbb{E}_\pi[P^+(q)].$$

Since credence families are functions $P : \Omega \to \mathbb{C}(\Omega)$, credence families 'saturated' with propositional inputs are random variables $P(q) : \Omega \to \mathbb{R}$ that give the value of $P_w(q)$ at different w. So the expectation is well defined:

$$\mathbb{E}_\pi[P(q)] =_{\text{def}} \sum_{w \in \Omega} \pi(w) \cdot P_w(q).$$

Definition 6. *A **general local accuracy scoring rule** is a function $\mathbf{A}$ that takes as input a credence function $\pi \in \mathbb{P}(\Omega)$, a random variable $X : \Omega \to \mathbb{R}$, and a world $w \in \Omega$ and evaluates the accuracy at w of π's estimation of the value of X.*

Definition 7. *A general local accuracy scoring rule $\mathbf{A}$ is **strictly proper** for a random variable X if for any probability functions π, ρ,*[7]

$$\mathbb{E}_\pi[\mathbf{A}(\pi, X, w)] \geq \mathbb{E}_\pi[\mathbf{A}(\rho, X, w)].$$

[7] Informally, $\mathbf{A}$ is strictly proper for X if every probability function evaluates itself as more accurate in expectation using $\mathbf{A}$ than any other (rigidly-designated) credence function. The expectation is well defined because $\mathbf{A}(\pi, X, \cdot)$ is a function $\Omega \to \mathbb{R}$.

Definition 8. *π **values***[8] *a credence family $P^+ : \Omega \to \mathbb{P}(\Omega)$ relative to some $X : \Omega \to \mathbb{R}$ if for every general local accuracy scoring rule $\mathbf{A}$ that is strictly proper for X, we have:*

$$\mathbb{E}_\pi[\mathbf{A}(P^+, X, w)] \geq \mathbb{E}_\pi[\mathbf{A}(\pi, X, w)].$$

*π **values** P^+* simpliciter *if π values P^+ for every random variable X.*

[8] Informally, π values P^+ relative to X if π evaluates P^+ as more accurate in expectation than itself on estimations about the value of X.

Fact 1, Formalized. *Let π be a credence function and P be a credence family, and suppose that π **reflects** P: for any proposition q and any $t \in [0,1]$, $\pi(q \mid [P(q)=t]) = t$. Then*

1. *π **values** P: for any random variable X and any general local scoring rule $\mathbf{A}$ that is strictly proper for X, $\mathbb{E}_\pi[\mathbf{A}(P^+, X, w)] \geq \mathbb{E}_\pi[\mathbf{A}(\pi, X, w)]$*
2. *π **balances** P: for any random variable X, $\mathbb{E}_\pi[X] = \mathbb{E}_\pi[\mathbb{E}_{P^+}[X]]$.*
3. *P is **higher-order certain**: for any world $w \in \Omega : P_w([P = P_w]) = 1$.*

§3. *The Possibility of Balance*

In any "good case / bad case" skeptical scenario, propositional evidence can be characterized so that conditionalizing on it (as well as other plausible updates) leads to a balance failure.

> **Old Friend / Stranger**: Either a stranger walks in (w_1) or your old friend you haven't seen in ages walks in (w_2). It is stipulated that if it's your friend, you'll be sure it's your friend. But if it's a stranger, you'll be unsure whether it's your friend or a stranger. It is thus alleged that your propositional evidence can be modeled as follows:
>
> 1. In w_1, your evidence is $\{w_1, w_2\}$: "it's either my friend or a stranger.'
> 2. In w_2, your evidence is $\{w_2\}$: "it's my friend for sure."

If you're antecedently 50/50 and conditionalize on your evidence, $\pi(\textsc{friend}) = \frac{1}{2}$ but $\mathbb{E}_\pi(P^+(\textsc{friend})) = \frac{3}{4}$.

$\mathbb{E}_\pi(P^+(\textsc{friend})) = \pi(w_1)P^+_{w_1}(w_2) + \pi(w_2)P^+_{w_2}(w_2) = \frac{1}{2}\cdot\frac{1}{2} + \frac{1}{2}\cdot 1 = \frac{3}{4}$.

The conditionalizing update in Old Friend / Stranger satisfies value: π values P^+. In general, under conditions of higher-order uncertainty, value and balance come apart: there can be many valuable updates that don't satisfy balance.

Dorst (2023) argues from this that it's *value* that's important for epistemic rationality, not balance; and that balance failures can be rational when evidence warrants higher-order uncertainty.

But these particular conditions of higher-order uncertainty never *require* agents to have balance failures:

Fact 2. *Suppose that in each world an agent gets evidence in the form of some proposition e. Then there is a function $U : \mathbb{P}(\Omega) \times \mathcal{P}(\Omega) \to \mathbb{P}(\Omega)$ such that an agent with prior π who adopts credence function $U(\pi, e)$ upon getting evidence e satisfies balance in that π balances $U(\pi, e)$.*

Idea: For any proposition e that might be an agent's evidence, define $\mathcal{L}(e) = \{w \in \Omega \mid \text{the agent's evidence is } e\}$. Then define $U(\pi, e) = \pi(\cdot \mid \mathcal{L}(e))$. Then π balances $U(\pi, e)$. (Also, $U(\pi, e)$ maximizes expected accuracy).

Theorem 1. *Suppose that in each world an agent gets evidence in the form of some proposition e. Then for any $\alpha \in [0,1]$ there is a function $U : \mathbb{P}(\Omega) \times \mathcal{P}(\Omega) \to \mathbb{P}(\Omega)$ s.t. an agent with prior π who adopts credence function $U(\pi, e)$ upon getting evidence e satisfies balance, and the credence family describing the agent's credences has $P_w([P = P_w]) = \alpha$ for all w.*

Idea: define $U(\pi, e) =$

$$\alpha \cdot \pi(\cdot \mid \mathcal{L}(e)) + (1-\alpha)\pi(\cdot \mid \Omega \setminus \mathcal{L}(e)).$$

Then π balances $U(\pi, e)$ and $U(\pi, e_w([U(\pi, e) = U(\pi, e_w)] = \alpha)$ for all w. (And $U(\pi, e)$ maximizes expected accuracy among credence families with $P_w([P = P_w]) \leq \alpha \ \forall w$).

§4. *The Tenability of Balance*

How much balance is compatible with higher-order uncertainty? Not much.

- **§4.1**: Some higher-order uncertain posteriors have no priors that balance it both *unconditionally* and *conditionally*.
- **§4.2**: Except in special cases, higher-order uncertain posteriors cannot be balanced by higher-order uncertain priors.

§4.1. *Conditional Balance*

> **Coin Tosses**: I flip a fair coin 1000 times. My prior expectation of the number of heads is 500. Suppose I rationally update on veridical evidence about how each coin toss landed, and write down my credences.

It seems like my prior expectation of the sum of my credences in heads over the 1000 trials in the notebook should also be 500.

This is the intuition for balance.

But by this logic, if I were to learn that the coin, which I thought was fair, is actually 3/4 heads-biased, my new prior expectation of the number of heads should be 750. Then my prior expectation of the sum of my credences in heads over 1000 trials, assuming that I take into account the bias of the coin, should also be 750.

If we want balance to hold even when both the prior and the planned posterior are conditionalized on some proposition, we want priors to *conditionally-balance* posteriors.

Definition 9. *With π, P as above and $q \subset W$ a proposition, say that π* ***balances*** *P* ***conditional on*** *q*[9] *if for any variable X,*

$$\mathbb{E}_\pi(X \mid q) = \mathbb{E}_\pi\left(\mathbb{E}_P(X \mid q) \mid q\right). \qquad \text{(Conditional Balance)}$$

[9] When P has no higher-order uncertainty, conditional balance is satisfied by any valuable update for any condition q. Basically this is because the only valuable updates without higher-order certainty are conditionalizing ones, and

$$\begin{aligned}&\mathbb{E}_\pi\left(\mathbb{E}_{\pi(\cdot|q)}(X \mid q) \mid q\right)\\ =&\mathbb{E}_\pi\left(\mathbb{E}_{\pi(\cdot|q)}(X) \mid q\right)\\ =&\mathbb{E}_\pi\left(\mathbb{E}_{\pi(\cdot|q)}(X)\right)\\ =&\mathbb{E}_\pi(X \mid q).\end{aligned}$$

Unfortunately, conditional balance is hard to come by for higher-order uncertain P.

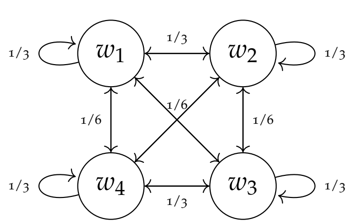

$$\pi = \left(\begin{array}{c|cccc} & w_1 & w_2 & w_3 & w_4 \\ \hline \pi & 1/4 & 1/4 & 1/4 & 1/4 \end{array}\right)$$

$$P^+ = \left(\begin{array}{c|cccc} & w_1 & w_2 & w_3 & w_4 \\ \hline \rho_{w_1} & 1/3 & 1/3 & 1/6 & 1/6 \\ \rho_{w_2} & 1/3 & 1/3 & 1/6 & 1/6 \\ \rho_{w_3} & 1/6 & 1/6 & 1/3 & 1/3 \\ \rho_{w_4} & 1/6 & 1/6 & 1/3 & 1/3 \end{array}\right)$$

Balance but not Conditional Balance: a pair $\langle \pi, P^+ \rangle$ over four worlds with the following features: (1) π trusts P^+, (2) π does **not** conditional-balance P^+ on $\neg w_1$, and (4) π balances P^+.

Alleged Interpretation of the Last Figure: Two coins will be flipped. You'll get ambiguous evidence about the first coin (red/left), warranting a 2/3 credence in h if h is true, and 2/3 credence in t if h is true. You get no evidence about the second coin (blue/center).

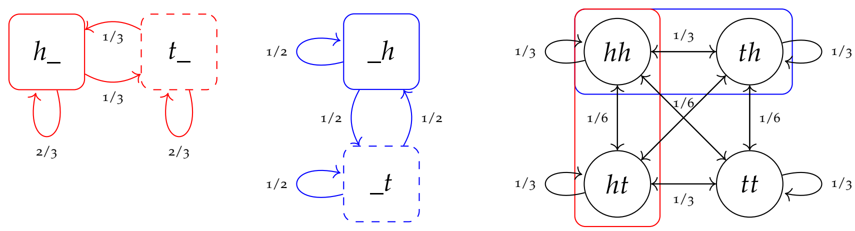

The prior is uniform between four worlds corresponding to the four possibilities (hh, ht, th, tt). The posterior credence family is balanced by this uniform prior. But *conditional* on the proposition that both coins landed heads (hh = w_1), the posterior is no longer reflected by the prior. Here, e.g., $\pi(w_2 \mid \neg w_1) = 1/3$, but

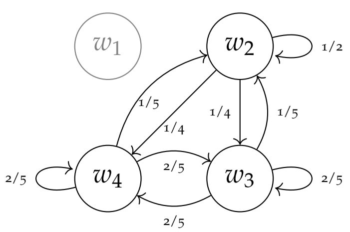

$$\left(\begin{array}{c|cccc} & w_1 & w_2 & w_3 & w_4 \\ \hline \pi & 0 & 1/3 & 1/3 & 1/3 \end{array}\right) \qquad \pi(\cdot \mid \neg w_1)$$

$$\left(\begin{array}{c|cccc} & w_1 & w_2 & w_3 & w_4 \\ \hline \rho_{w_2} & 0 & 1/2 & 1/4 & 1/4 \\ \rho_{w_3} & 0 & 1/5 & 2/5 & 2/5 \\ \rho_{w_4} & 0 & 1/5 & 2/5 & 2/5 \end{array}\right) \qquad P^+(\cdot \mid \neg w_1)$$

$$\begin{aligned} \mathbb{E}_\pi(P^+(w_2 \mid \neg w_1) \mid \neg w_1) &= \pi(w_2)P^+_{w_2}(w_2) + \pi(w_3)P^+_{w_3}(w_2) + \pi(w_4)P^+_{w_4}(w_2) \\ &= \tfrac{1}{3} \cdot \tfrac{1}{2} + \tfrac{1}{3} \cdot \tfrac{1}{5} + \tfrac{1}{3} \cdot \tfrac{1}{5} = 3/10 < 1/3. \end{aligned}$$

Fact 3. *There are credence function families where no credence function balances it conditionally and unconditionally.*

For example, P^+ in the double coin-toss example is unconditionally balanced only by $\pi = (1/4, 1/4, 1/4, 1/4)$ and balanced conditional on $\neg w_1$ only by $\pi = (1-\lambda, \frac{2\lambda}{7}, \frac{5\lambda}{14}, \frac{5\lambda}{14})$ for $\lambda \in [0, 1]$.

§4.2. *Higher-Order Uncertain Priors*

We can interpret P^+ as itself a prior that could be further updated into a P^{++}. For instance, suppose you gain the ambiguous evidence about the first coin, updating from π to P^{++}, and then gain evidence about whether w_1 is true, thus updating into P^{++}.

Definition 10. *Where $P, R : W \to \mathbb{P}(W)$ are credence families and $P[W]$ is the image of W under P, say that P **balances** R if for any variable X,*

$$\forall \pi \in P[W] : \mathbb{E}_\pi (X) = \mathbb{E}_\pi (\mathbb{E}_R(X)). \qquad \text{(Balance (for Families))}$$

Question: when does P^+ balance P^{++}? **Answer**: only in a very restricted set of circumstances.

Theorem 2. *If P, R are credence families, P balances R only if there is a partition $\Pi = \{T, C_1, C_2, \ldots\} \subseteq \mathcal{P}(\Omega)$ of propositions where,*
1. *for all $C \in \Pi$, for all $w, w' \in C$, $P_w = P_{w'}$.*
2. *for all $C \in \Pi$, for all $w \in C$, $R_w(C) = 1$, and so $R_w([R_w \in R[C]]) = 1$.*
3. *for all $\pi \in P[\Omega]$, $\pi(T) = 0$.*

Informally, there must be a partition of conditions where,
1. Within each condition, P is constant, and thus higher-order certain. In other words, *conditional* on any condition C in the partition, P is always certain about the actual value of P_w, even though P may not know which condition is the actual condition
2. R is always certain about which condition is the actual condition.
3. There is a set of 'pathological' possibilities that P is certain will never obtain.

I.e. the higher-order uncertainty in P and R must be 'orthogonal'.

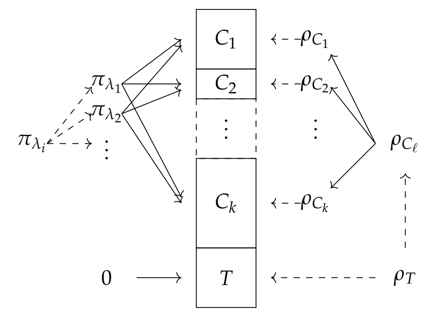

§5. *On Balance*

Consequences for Rational Updating
1. From §4.1: When we're considering a single prior credence function and a higher-order uncertain posterior credence family, there may be no way for the prior to think the updating will be unbiased conditional on different hypotheses.
2. From §4.2: When we're considering both a higher-order uncertain posterior credence family and a higher-order uncertain prior credence family, there may be no way for the prior to even *unconditionally* balance the posterior. So updating into a higher-order uncertain state can mean giving up balance in the future.

Upshot: balanced updates are really hard to achieve in conditions of higher-order uncertainty.

Consequences for Higher-Order Uncertainty There may really be a tradeoff between balance and higher-order uncertainty. If the latter is rationally permissible, it's hard to see how the former, in desirable generalities, could be rationally required.

→ Dorst (2023), Williamson (1997), Zendejas Medina (2024): keep higher-order uncertainty, do away with balance.

→ Gallow (2021), Isaacs and Russell (2023, implicitly): keep both. (I think this talk poses problems for their position).

→ Me, tentatively: Keep balance, and thus *do away with higher-order uncertainty*. Or, at least: find a *different way* to model it.