

The Art of Learning

Jason Konek

Abstract

Confirmational holism is at odds with Jeffrey conditioning—the orthodox Bayesian policy for accommodating uncertain learning experiences. Two of the great insights of holist epistemology are that (i) the effects of experience ought to be mediated by one’s background beliefs, and (ii) the support provided by one’s learning experience can and often is undercut by subsequent learning. Jeffrey conditioning fails to vindicate either of these insights. My aim is to describe and defend a new updating policy that does better. In addition to showing that this new policy is more holism-friendly than Jeffrey conditioning, I will also show that it has an accuracy-centered justification.

How confident are you that you will make it through the afternoon without a heart attack? 99%? 99.999%? Whatever your answer, you probably agree on this much: whether you spend an hour reading the news and leaning lazily on your left elbow is *irrelevant* to your prospects. It provides no evidence one way or another about whether you will have a heart attack (unless an oracle told you, “Lean on your elbow and meet your doom!” or something of the sort). Finding out that you have high levels of “bad cholesterol,” for example, is bad news. But finding out that you will spend the next hour reading and leaning is neither bad news nor good news. It is *no* news (no *relevant* news anyway). It does not tell you much one way or another about the risk of cardiovascular catastrophe.

But now imagine that you feel your left arm tingling. As a result, your credence that you are about to have a heart attack shoots up. As you start to panic, the kind stranger next to you asks you whether they can help. You explain your situation. They respond calmly, “You do realise that you have been leaning on your left elbow while reading for the last hour, don’t you?” Prior to feeling your left arm tingle and spiraling into a panic, you thought this bit of information was no news. It did not, in your view, give you any evidence about whether you will have a heart attack. But now the situation is different. You are well aware that the ulnar nerve runs along the back of your elbow; a nerve that might get pinched if you lean on it for too long and cause your arm to tingle. So *now* the kind stranger’s information is *highly relevant news*. Learning that you have been leaning on your left elbow changes your opinion dramatically. It causes you to become pretty sure that you are fine; no heart attack in sight. Irrelevant before, but irrelevant no more!

Your learning experience in this case has multiple effects. It raises your credence that you will have a heart attack. But it also introduces new *undercutting defeaters*. Undercutting defeaters cause you to re-evaluate the evidential import of some prior learning experience. More carefully, if a proposition X undercuts the support that a learning experience \mathcal{E} provides for a hypothesis Y , then learning X makes you think that \mathcal{E} provides *no support whatsoever*

for Y .¹ However we cash this out, the following seems true: X undercuts the support that \mathcal{E} provides for Y only to the extent that your post- \mathcal{E} credence for Y conditional on X (immediately after \mathcal{E} , before acquiring additional information) is close to your pre- \mathcal{E} credence for Y . The more X undercuts \mathcal{E} 's support for Y , the more learning X would push your confidence in Y back near its pre- \mathcal{E} levels. This means that X must be **relevant to Y** to count as an undercutting defeater for Y .

In the case at hand, your learning experience \mathcal{E}^* turns the proposition L that you lean on your elbow for an hour into an undercutting defeater for the proposition H that you will have a heart attack. Prior to the experience, L is irrelevant to H . And irrelevant propositions do not undercut anything. But after the experience you see L as an undercutting defeater for H . Experience \mathcal{E}^* causes you to think that L undercuts the support that \mathcal{E}^* itself provides for H . After \mathcal{E}^* you think: Learning L would be good reason to drop my high posterior (post- \mathcal{E}^*) credence in H back down to something like its pre-arm-tingling level.

According to Christensen (1992) and Weisberg (2015), cases like this are both utterly ubiquitous, and cause big problems for *Jeffrey conditioning*. Jeffrey conditioning says that you ought to update your credences in light of new learning experiences as follows:

J-Con. If you have prior credences *old* and undergo a learning experience \mathcal{E} that shifts your credences on some partition $\{E_1, \dots, E_n\}$ to x_1, \dots, x_n , and nothing more, then your new credence for any proposition X ought to be:

$$new(X) = \sum_{i \leq n} x_i \cdot old(X | E_i).$$

According to Weisberg, J-Con bungles the introduction of new undercutting defeaters. A rational agent's confidence in most propositions (maybe all propositions) can be undermined by theoretical considerations. This is one of the basic insights of holist epistemology. Moreover, rational agents change their mind about what undercuts what. Learning experiences introduce *new* undercutting defeaters. In our heart attack case, \mathcal{E}^* causes you to think that L undercuts the support that \mathcal{E}^* itself provides for H , despite the fact that you previously thought L was irrelevant to H . But Jeffrey conditioning simply does not allow for this. J-Con is a "rigid" updating rule (§1). And rigid updating rules preserve opinions about irrelevance. So if you thought L was irrelevant to H before \mathcal{E}^* , then you must think it is irrelevant after \mathcal{E}^* . But this means that if you update by Jeffrey conditioning in our heart attack case, then \mathcal{E}^* will *not* push you to treat L as an undercutting defeater for H . Hence J-Con makes bad predictions about rational learning.

Christensen's problem is a bit different. Christensen focuses on whether J-Con has the resources to explain how inputs to updating depend on background beliefs. According to Christensen, learning experiences never push your credences up or down *directly*. Rather, the doxastic effects of experience are always mediated by your background beliefs. Whether the tingling in your arm pushes your credence

¹Rebutting defeaters, in contrast, undermine your confidence in Y not by undercutting the support for Y provided by some learning experience \mathcal{E} , but rather by "telling directly against" the truth of Y .

that you are going to have a heart attack up, or down, or leaves it unchanged, ought to depend on your prior opinions on a whole host of matters: your current state of cardiovascular health, what exactly the warning signs of a heart attack are, whether you have been leaning on your left elbow, etc. This is another basic insight of holist epistemology. But Jeffrey conditioning treats the initial effect of experience—a shifting of credences over some partition (sometimes called a “Jeffrey shift”)—as an exogenous factor; an output of some “black box” process that serves as an input to Jeffrey conditioning. Hence J-Con fails to explain an important feature of rational learning: how or why the initial effects of experience (the Jeffrey shift) depend on prior opinion.

Jeffrey was comfortable with this. Training the messy network of neurons in your skull to translate perceptive and proprioceptive inputs into sensible Jeffrey shifts is simply not something that formal epistemology *should* speak to. That training happens in one’s PhD programme, in the lab, etc. It yields a domain-specific skill; a skill that goes far beyond the skills one might be said to have simply in virtue of being epistemically rational. Christensen, on the other hand, sees Jeffrey’s ‘concession’ as placing “an important cognitive or structural aspect of justification outside the area our theory purports to describe” (Christensen, 1992, p. 547).

My aim in this paper is to describe and defend a new updating policy that answers both Christensen and Weisberg’s concerns. Because it shares structural features with J-Con, I call it **J-Kon**. Unlike J-Con, J-Kon makes explicit the way in which inputs to updating depend on background beliefs. J-Kon also naturally accounts for the introduction of new undercutting defeaters. In addition to being more holism-friendly than J-Con, it also has an accuracy-centered justification.

Here is the plan in a bit more detail. In §1, I will briefly explain why “rigid” updating rules like J-Con have trouble introducing new undercutting defeaters. In §2, I will describe J-Kon and show how it works in a few simple cases. In §3, I will show that J-Kon is in fact equivalent to what Jeffrey called “superconditioning” (with a few extra bells and whistles). In §4, I will extol the epistemic virtues of J-Kon. More carefully, I will provide a *chance-dominance* argument for J-Kon. If you fail to update by J-Kon, then your epistemic life—the sequence of credal states that you adopt in response to your learning experiences—is chance-dominated by some other J-Kon satisfying life. What it means to be chance-dominated is this: *every possible chance function* expects your life to accrue strictly less total epistemic value than the J-Kon satisfying life. If you update by J-Kon, in contrast, your epistemic life is never chance-dominated in this way. This provides a purely alethic (rather than evidential or pragmatic) rationale for updating by J-Kon. In §5, I will explore the ways in which J-Kon makes inputs to updating depend on background beliefs. I will also run through a few cases that illustrate how J-Kon introduces new undercutting defeaters. In §6, I will respond to some pressing objections. Finally, in §7, I will wrap up.

1 Rigidity and Undercutting Defeat

Consider a histopathologist looking at stained cells under a microscope. Careful examination might make her fairly confident that the relevant sample is, *e.g.*, Ductal cell carcinoma. Given her training and expertise, she may well have *good reason* for her confidence. Nevertheless, she might be wholly unaware of what

features of the cells she is picking up on. Despite the fact that there *are* definite features that she is tracking, there is no reason to expect that she will have any kind of *privileged access* to information about those features. There is no reason to expect she will store such information as a “passive data structure,” or in a more action-guiding fashion (directly in her sensorimotor system), or in any other way that makes that info available as an input to action, decision-making, or reasoning. As Weiskrantz puts it, “a large amount of bodily processing. . . in sensory channels proceeds quite detached from any awareness. . . Awareness of an event is a form of privileged access that allows further perceptual and cognitive manipulations to occur; as far as neural processes are concerned, it is probably a minority privilege” ((Jeffrey, 1992, p. 197), (Weiskrantz, 1986, pp. 168-9)).

None of this is unique to histopathologists, of course. It is “typical of our most familiar sorts of updating, as when we recognize friends’ faces or voices or handwritings pretty surely, and when we recognize familiar foods pretty surely by their look, smell, taste, feel, and heft.” (Jeffrey, 1992, p. 79). In these run-of-the-mill learning situations, there are definite features that you are picking up on. And in virtue of your sensitivity to these features, you may well have *good reason* for being fairly confident that the person in the distance is your friend, that Abbi Jacobson is narrating your new favourite podcast, etc. But there typically is no proposition describing those features that you even have an opinion about, let alone become *certain* of. The reason is familiar. For you to count as having an opinion about those features, propositions describing them must be available (in some way or other) as inputs to action, decision-making, or reasoning. But such information typically does *not* enjoy this ‘minority privilege’. It is processed and has various downstream effects on your doxastic, affective and conative state, but is not *itself* made available as an input to action, decision-making, or reasoning.

How should you update your credences, then, when you have learning experiences like our histopathologist’s? Jeffrey conditioning offers a partial answer. Suppose that you have prior credences $old : \mathcal{F} \rightarrow \mathbb{R}$ for propositions in a σ -algebra \mathcal{F} on a set of worlds \mathcal{W} .² Suppose also that you undergo a learning experience \mathcal{E} that (i) does not make you *certain* of anything, but nevertheless (ii) shifts your credences on some partition $\{E_1, \dots, E_n\}$ of \mathcal{W} to x_1, \dots, x_n (and nothing more). That is, \mathcal{E} induces a “Jeffrey shift” over $\{E_1, \dots, E_n\}$. Then Jeffrey conditioning (J-Con) says that your new credence for any proposition X ought to be:

$$new(X) = \sum_{i \leq n} x_i \cdot old(X | E_i).$$

When is Jeffrey conditioning appropriate? According to Jeffrey, you ought to update by J-Con just in case two conditions hold.

1. $new(E_i) = x_i$.
2. $new(X | E_i) = old(X | E_i)$.

The first condition says that you take the credences produced directly by \mathcal{E} at face value. For each E_i , you stick with x_i as your new credence (rather than treating this shift in your credences as the aberrant result of momentary mania).

² \mathcal{W} contains the finest-grained possibilities you can distinguish between. Saying that \mathcal{F} is a σ -algebra means that (i) \mathcal{F} contains the tautology \mathcal{W} , and (ii) \mathcal{F} is closed under complement (negation), countable union (disjunction), and countable intersection (conjunction). For simplicity, we will assume that \mathcal{W} is finite.

The second condition says that you treat \mathcal{E} as the sort of learning experience that provides information about *which* of the E_i is true, but *not* about what else is true *given* that E_i is true. So you hold your credence in X conditional on E_i fixed for each X and each E_i . When this second condition holds, we say that your update is “rigid.”

For example, our histopathologist’s learning experience (examining stained cells under a microscope) might make her 80% confident that the sample is Ductal cell carcinoma and 20% confident that it is not. And she might take these “direct effects” of experience at face value (stick with them as her new credences). Moreover, she might think that while \mathcal{E} provides information about whether the sample is Ductal cell carcinoma or not, it does *not* provide information about, for example, the patient’s survival prospects S *given* that it is (or is not) Ductal cell carcinoma. If the same holds not just for S , but for any proposition X , and she holds her conditional credences fixed as a result (*i.e.*, $new(X | E_i) = old(X | E_i)$), then both conditions (1) and (2) hold.

And in *that* case, J-Con is simply mandated by probabilistic coherence. Given that $new : \mathcal{F} \rightarrow \mathbb{R}$ satisfies the probability axioms, (1) and (2) are jointly equivalent to:

$$new(X) = \sum_{i \leq n} x_i \cdot old(X | E_i).$$

But—and here’s the rub—both (1) and (2) are quite strong conditions. Indeed, the rigidity condition (2) is *so* strong that it plausibly *never* holds.

Our heart attack example illustrates the point. Recall, in that example, your learning experience \mathcal{E}^* (tingling arm) raises your credence for the proposition H (that you will have a heart attack).

$$new(H) > old(H).$$

It also turns the proposition L (that you lean on your elbow for an hour) into an undercutting defeater for H . Prior to the experience, L is irrelevant to H .

$$old(H | L) = old(H).$$

But *after* the experience you come to see L as an undercutting defeater for H . After \mathcal{E}^* you think: learning L would be good reason to drop my high posterior (post- \mathcal{E}^*) credence in H back down to something like its pre-arm-tingling level.

$$new(H) > new(H | L).$$

The problem is this. Suppose that you accommodate \mathcal{E}^* by Jeffrey conditioning, using your new distribution

$$new(H) = high, \quad new(\neg H) = low$$

over $\{H, \neg H\}$ as an input. Then we must have:

$$new(H) = new(H | L)$$

contra our assumption that $new(H) > new(H | L)$. So J-Con simply fails to introduce L as a new undercutting defeater for H . The proof is dead simple. Firstly, note that $old(H) = old(H | L)$ implies

$$old(L) = old(L | H) = old(L | \neg H).$$

Next, note that since J-Con implies the rigidity condition (2), we have

$$\text{new}(L | H) = \text{old}(L | H) \text{ and } \text{new}(L | \neg H) = \text{old}(L | \neg H).$$

This straightaway gives us

$$\begin{aligned} \text{new}(H | L) &= \frac{\text{new}(L | H)\text{new}(H)}{\text{new}(L | H)\text{new}(H) + \text{new}(L | \neg H)\text{new}(\neg H)} \\ &= \frac{\text{old}(L | H)\text{new}(H)}{\text{old}(L | H)\text{new}(H) + \text{old}(L | \neg H)\text{new}(\neg H)} \\ &= \frac{\text{old}(L)\text{new}(H)}{\text{old}(L)\text{new}(H) + \text{old}(L)\text{new}(\neg H)} \\ &= \text{new}(H). \end{aligned}$$

The moral is this. Rigidity forces you to preserve your old opinions about irrelevance. So if your updating rule is rigid with respect to $\{E_1, \dots, E_n\}$ (it leaves your credences conditional on the E_i unchanged), then that rule cannot take you from thinking that X was irrelevant to E_i before \mathcal{E} to thinking that X is relevant to E_i after \mathcal{E} . But turning an irrelevant X into a relevant X is *precisely* what is involved in introducing a new undercutting defeater. So whenever \mathcal{E} *should* introduce a new undercutting defeater for one of the E_i , you should *not* update according to J-Con, or any other rigid updating rule.

This is *really* bad for J-Con. Learning experiences should *almost always* introduce new undercutting defeaters. For any learning experience \mathcal{E} that pushes your credences over some partition $\{E_1, \dots, E_n\}$ around, there will be *some* proposition X that you previously took to be irrelevant to E_i , but now should take to be relevant; now you should take X to undercut the support that \mathcal{E} provides for E_i . Any X that describes conditions that compromise the reliability of your credence-formation process—the process that pushed your credences for the E_i around—will do the trick. This, again, is one of the basic insights of holist epistemology. And if this insight is right, then according to Jeffrey’s own criteria, **it is almost never appropriate to update by Jeffrey conditioning.**

To recap, J-Con bungles the introduction of new undercutting defeaters. Since nearly all learning experiences introduce new undercutting defeaters, J-Con almost never applies. This is Weisberg’s main concern. The question now is: Can we do better? The answer: yes.

2 J-Kon

We seem to be back at square one. When our histopathologist examines her cells, she becomes *fairly confident* that the sample is one type of carcinoma rather than another. But she does not become *certain* of anything new. The orthodox Bayesian method for updating in light of such “uncertain learning experiences” is Jeffrey conditioning. But J-Con is almost never appropriate. How then *should* we update in light of uncertain learning experiences?

My aim now is to describe and defend a new updating policy for uncertain learning: **J-Kon**. Firstly, I will describe J-Kon briefly and show how it works in a few cases (§2.1-2.2). Then I will show that J-Kon is in fact equivalent to a restricted form of what Jeffrey called *superconditioning*. I will also provide

a second characterisation of J-Kon to help illuminate what it demands of us (§3). These characterisations will be important for showing that J-Kon has an accuracy-centered justification (§4). Finally, with our main argument for J-Kon in place, I will return to the issues surrounding confirmational holism that we began with. In particular, I will show that J-Kon naturally explains how inputs to updating depend on background beliefs, and naturally accounts for the introduction of new undercutting defeaters (§5).

2.1 Formal Description

Before outlining J-Kon, we need to introduce one key concept: the *conditional chance estimate*. Estimates are familiar enough. For example, the Met office’s best estimate of the amount of rainfall that London will receive next year might be 597mm. Your best estimate of the number of children that your brother will have might be 2.7. And so on. Estimates are numbers; numbers that fall between the possible values of the quantity that you are estimating; numbers that are evaluated principally on the basis of their *accuracy* (the closer they are to the true value of the quantity, the better).³ In the Bayesian tradition, rational choice is a matter of choosing the option that you estimate to produce the most utility. And your best estimates of utility (and other quantities) are determined by your best estimates of truth-values, which are captured by your credences, or degrees of belief.

Conditional estimates are also familiar. For example, the Met office’s best estimate of the amount of rainfall that London will receive next year might be 597mm *on the supposition that a catastrophic climate event does not derail the Gulf stream*. Your best estimate of the number of children that your brother will have might be 2.7 *on the supposition that he and his partner work through their issues*. And so on. Like plain old unconditional estimates, conditional estimates are numbers; numbers that fall between the possible values of the quantity that you are estimating. But where unconditional estimates are something like epistemic bets—bets that pay out in an epistemic currency, *viz.*, accuracy—conditional estimates are more like *called-off epistemic bets*. They are evaluable for accuracy only in worlds where their condition holds. In worlds where the condition does not hold, the epistemic bet that they represent is called off.

Conditional *chance* estimates, then, are just conditional estimates of a particular quantity: *chance*. In particular, a conditional chance estimate of the form:

$$est[ch(X | Y) | Y]$$

is an estimate of the conditional chance $ch(X | Y)$ *on the supposition that Y is true*. Put differently, it is an estimate of the chance of X when both you and chance take the supposition Y on board.

Following de Finetti, Jeffrey thought of estimation as basic, and credences as capturing a particular type of estimate, *viz.*, an agent’s best estimates of truth-values.⁴ More carefully, Jeffrey thought of propositions X as “indicator variables” that take the value 1 at worlds where X is true, and 0 where X is false. (In what follows, we will slip between talking of propositions as sets of worlds

³Indeed, one might think that it is *constitutive* of estimates that they are evaluable on the basis of their accuracy. See Konek (2019).

⁴See (Jeffrey, 1986, p.51).

and indicator variables.) Truth-value estimates, then, are simply estimates of the value, 0 or 1, that the proposition takes at the actual world. And an agent's credence for a proposition is just her best estimate of its truth-value.

Treating estimation as basic allows us to see the laws of (finitely) additive probability, *i.e.*,

Normalisation. $p(\mathcal{W}) = 1$.

Non-negativity. $p(X) \geq 0$.

Finite Additivity. $p(X \cup Y) + p(X \cap Y) = p(X) + p(Y)$.

as straightforward consequences of de Finetti's *laws of estimation* (or what are sometimes called *the axioms of linear prevision*). De Finetti's laws say that whenever the estimates of variables $\mathcal{V} : \mathcal{W} \rightarrow \mathbb{R}$ and $\mathcal{Q} : \mathcal{W} \rightarrow \mathbb{R}$ are well-defined, they ought to satisfy the following conditions (de Finetti, 1974, §3.1.5):

Boundedness. $\inf_{w \in \mathcal{W}} \mathcal{V}(w) \leq \text{est}(\mathcal{V}) \leq \sup_{w \in \mathcal{W}} \mathcal{V}(w)$.

Linearity. $\text{est}(\mathcal{V} + \mathcal{Q}) = \text{est}(\mathcal{V}) + \text{est}(\mathcal{Q})$.

The boundedness condition says roughly that your estimate of \mathcal{V} should fall somewhere between the minimum and maximum possible values of \mathcal{V} . The linearity condition says that your estimate of the sum of \mathcal{V} and \mathcal{Q} , $\mathcal{V} + \mathcal{Q}$ —*i.e.*, the variable whose value at a world is the sum of \mathcal{V} 's value and \mathcal{Q} 's value, respectively—should equal the sum of your individual estimates for \mathcal{V} and \mathcal{Q} . If the unconditional estimates captured by *est* satisfy de Finetti's laws, then we say that *est* is *coherent*. Similarly, we can see the Rényi-Popper axioms for conditional probability, *i.e.*,

Conditional Probability. $c(\cdot | X)$ is a probability function with $c(X | X) = 1$.

Generalised Ratio Constraint.

$c(Y \& Z | X) = c(Z | X \& Y)c(Y | X)$ if $X \cap Y \neq \emptyset$.

as straightforward consequences of the *laws of conditional estimation*. The laws of conditional estimation say that whenever the involved conditional estimates are well-defined, they ought to satisfy:

Conditional Boundedness. $\inf_{w \in X} \mathcal{V}(w) \leq \text{est}(\mathcal{V} | X) \leq \sup_{w \in X} \mathcal{V}(w)$.

Conditional Linearity. $\text{est}(\mathcal{V} + \mathcal{Q} | X) = \text{est}(\mathcal{V} | X) + \text{est}(\mathcal{Q} | X)$.

Multiplicativity. $\text{est}(Y\mathcal{V} | X) = \text{est}(\mathcal{V} | X \& Y)\text{est}(Y | X)$ if $X \cap Y \neq \emptyset$ (where $Y\mathcal{V}$ is the product of Y and \mathcal{V} , *i.e.* $(Y\mathcal{V})(w) = Y(w)\mathcal{V}(w)$).

If the conditional estimates captured by *est* satisfy the laws of conditional estimation, then we say that *est* is *coherent* (or a *conditional linear prevision*).

With these preliminary remarks in place, we can now outline J-Kon. Note though that I will initially present J-Kon as a very deliberate, cognitively demanding updating procedure. This will help get the basic picture across without too much fuss. But the official version of J-Kon only involves updating *as if* you are deliberately following this procedure. It is much less cognitively demanding than it appears at first blush. We return to this issue in §3 and §6.

J-Kon proceeds in two stages. The first stage is the *expansion stage*. In response to any learning experience, \mathcal{E} , you ought to settle on a new conditional credence, $new(\mathcal{E} | w)$, for each world w . These new conditional credences reflect how likely you think you are to have that very experience *conditional on being in this, that, or the other world*. Typically, you will not have any *prior* (pre- \mathcal{E}) credences about \mathcal{E} , conditional or not. \mathcal{E} itself puts you in a position to have opinions about \mathcal{E} . So settling on these new conditional credences involves *expanding* the range of propositions that you have opinions about.

J-Kon updaters, however, do not expand willy nilly. Rather, they settle on new conditional credences by settling on *new conditional chance estimates*. They estimate the *chance* of having the experience \mathcal{E} conditional on being in this, that, or the other world. Then they use these new conditional chance estimates as their new conditional credences.

The second stage is the *update stage*. Once you have expanded, you should input your *old unconditional credences* for atoms (worlds) and your *new conditional chance estimates* into Bayes' theorem. Then conditionalize on \mathcal{E} . This specifies new credences for each atom of your algebra. Together with the probability axioms, this fixes your entire new credal state.

In a little bit more detail, J-Kon says that you ought to accommodate uncertain learning experiences as follows:

J-Kon (Unofficial Version). Suppose you have prior credences *old* and have accommodated past learning experiences \mathcal{P} via J-Kon using conditional chance estimates est_{old} . You now undergo learning experience \mathcal{E} . Then you ought to update *old* as follows.

1. **Expansion stage.** Expand est_{old} to include new conditional chance estimates of the form:

$$est[ch(\mathcal{E} | w \& \mathcal{P}) | w \& \mathcal{P}].$$

The only constraints on the newly expanded assignment of conditional estimates est are that (i) est should be coherent, *i.e.*, satisfy the laws of conditional estimation, and (ii) est should be consistent with the Principal Principle.⁵

2. **Update stage.** Use these new conditional chance estimates to update your credence for each atom w of \mathcal{F} as follows:

$$new(w) = \frac{est[ch(\mathcal{E} | w \& \mathcal{P}) | w \& \mathcal{P}] \cdot old(w)}{\sum_{w'} est[ch(\mathcal{E} | w' \& \mathcal{P}) | w' \& \mathcal{P}] \cdot old(w')}.$$

The best way to get a feel for J-Kon is to see it in action. But before we do, we should set aside a few potential concerns. Firstly, a word about learning experiences and these heretofore nebulous propositions \mathcal{E} that we have been using to describe them. Learning experiences are rich, complex events with a range of properties. An agent's learning experience might have a certain phenomenological character, for example. In addition, it might "directly affect" her doxastic, affective or conative state in various ways. For example, it might

⁵ Est is consistent with the Principal Principle iff expanding it to include conditional estimates of the form $est[X | \phi_{ch} \& Y] = ch(X | Y)$ preserves coherence. See the appendix for more detail.

shift her credences over some partition, as Jeffrey imagined. Or it could push her conditional credences around, or her expectations, or other properties of her credal state. And of course learning experiences do not happen in a vacuum. Some learning experiences are shaped by years of training. Others are not. And on, and on, and on.

Given how rich and varied learning experiences are, it is a fool’s errand to try to pin down any single property, or even a cluster of properties that we can always use to characterise them. Simply describing the phenomenological character of an agent’s learning experience, or its direct effects on her credal state, etc., will almost certainly leave out epistemologically important information; information that ought to have some impact on how she accommodates that learning experience. To avoid this sort of concern, we will pick out learning experiences using demonstrative propositions. That is, we will describe learning experiences using propositions \mathcal{E} of the form *agent A had **that** learning experience*. Of course, this is the sort of proposition that you are only in a position to have opinions about *once you have already had the learning experience*. This would be problematic if J-Kon required you to have *prior* credences about \mathcal{E} . But it does not.

Secondly, J-Kon may seem to just dress up an old Bayesian story in new garb. After all, the unofficial version says that if you undergo a non-dogmatic learning experience \mathcal{E} —one that does not make you *certain* of anything, but nevertheless shifts your credences around in some way or other—then even though you cannot conditionalize straightaway (because there is no proposition learned with certainty to conditionalize on), you should nonetheless (i) put yourself in a position to conditionalize (on \mathcal{E} in particular) and then (ii) conditionalize. But this concern misses the mark in two ways. Firstly, the unofficial version of J-Kon recommends responding to non-dogmatic learning experiences by *expanding*. Moreover, it constrains *how* you should expand (at least if there are interesting constraints on the space of possible chance functions). None of this is standard Bayesian fare. Secondly, the official version of J-Kon (§3) ditches the recommendation to explicitly expand and conditionalize, and rather recommends updating *as if* you were expanding and conditionalizing. Toward the end of §3, I will characterise when exactly you count as updating in this fashion. As will be abundantly clear, what we end up with is a genuinely new story about rational updating; more than just an old tune in a new key.

So much for the preliminary remarks. Onto the applications!

2.2 Applications

1. Dogmatic learning. Vanji is having a routine sexual health check-up. Though she does not suspect that she is at risk for HIV, she decides to have an HIV test anyway. She has prior credences over the power set \mathcal{F} of \mathcal{W} , which contains:

$$\begin{array}{ll} w_1 = HIV \ \& \ Test \ + & w_2 = HIV \ \& \ Test \ - \\ w_3 = No \ HIV \ \& \ Test \ + & w_4 = No \ HIV \ \& \ Test \ - \end{array}$$

Her prior credences are as follows:

$$\begin{array}{ll} old(w_1) = 0.00095 & old(w_2) = 0.00005 \\ old(w_3) = 0.04995 & old(w_4) = 0.94905 \end{array}$$

So Vanji's prior credence that she has HIV is 0.001. Her prior credence that the test will come back positive given that she has HIV is 0.95. Her prior credence that the test will come back negative given that she has does not have HIV is also 0.95. Finally, her prior credence that she has HIV given that the test comes back positive is about 0.019 (19 times higher than her unconditional prior for HIV, but still rather low).

The doctor hands her the results of the test. This causes her to have a learning experience \mathcal{E} that pushes her credence that the test is positive up to 1 (*i.e.*, a dogmatic learning experience).

Question: What should Vanji's new credences be after \mathcal{E} ?

Answer: According to J-Kon, she ought to update as follows. Firstly, adopt conditional chance estimates of the form:

$$est[ch(\mathcal{E} | w) | w]$$

(To simplify the problem, we will imagine that this learning experience is first of Vanji's epistemic life. So the proposition \mathcal{P} describing her past learning experiences is just the tautology. We will make the same assumption in subsequent examples. But the results in the appendix do not make this assumption.) For example, Vanji might adopt the following conditional chance estimates:

$$est[ch(\mathcal{E} | w_1) | w_1] = est[ch(\mathcal{E} | w_3) | w_3] = 1$$

and

$$est[ch(\mathcal{E} | w_2) | w_2] = est[ch(\mathcal{E} | w_4) | w_4] = 0.$$

These estimates reflect the opinion that Vanji is an *infallible learner*, at least on this occasion. There is *no* chance of failing to have this sort of learning experience—one which sends her credence in a positive test result up to 1—if the test does in fact come back positive. Likewise, there is *no* chance of mistakenly having this sort of learning experience if the test comes back negative.

Next, she should treat these new conditional chance estimates, $est[ch(\mathcal{E} | w_i) | w_i]$, as her old conditional credences, $old(\mathcal{E} | w_i)$, and conditionalize on \mathcal{E} . This yields a new credal state given by:

$$new(w_i) = \frac{est[ch(\mathcal{E} | w_i) | w_i] \cdot old(w_i)}{\sum_j est[ch(\mathcal{E} | w_j) | w_j] \cdot old(w_j)}.$$

So, for example, Vanji's new credence for w_1 is:

$$\begin{aligned} new(w_1) &= \frac{est[ch(\mathcal{E} | w_1) | w_1] \cdot old(w_1)}{\sum_j est[ch(\mathcal{E} | w_j) | w_j] \cdot old(w_j)} \\ &= \frac{1 \cdot 0.00095}{1 \cdot 0.00095 + 0 \cdot 0.00005 + 1 \cdot 0.04995 + 0 \cdot 0.94905} \\ &= 0.018664. \end{aligned}$$

More generally, her new credal state is given by:

$$\begin{aligned} new(w_1) &= 0.018664 & new(w_2) &= 0 \\ new(w_3) &= 0.981336 & new(w_4) &= 0 \end{aligned}$$

This is precisely the same posterior credal state that Vanji would end up with if she updated by conditioning on the proposition that the test came back positive. This is no coincidence. **Orthodox conditionalization is equivalent to J-Kon for infallible dogmatic learners.** For agents whose learning experiences always push their credences up to 1 in some proposition E (and nothing more), and who are infallible in the sense that (i) there *no* chance of this occurring when E is false, and (ii) *no* chance of it *not* occurring when E is true, Orthodox Bayesian conditionalization is the way to go.

But even dogmatic learners are not always infallible. For example, Vanji might be extremely nervous about a positive test result. She might think: there is a marginal chance that I will mistakenly have \mathcal{E} even if the test comes back negative. Perhaps her anxiety will cause her to not properly register the words printed on the report. In that case, the following conditional chance estimates estimates might seem appropriate:

$$est[ch(\mathcal{E} | w_1) | w_1] = est[ch(\mathcal{E} | w_3) | w_3] = 1$$

and

$$est[ch(\mathcal{E} | w_2) | w_2] = est[ch(\mathcal{E} | w_4) | w_4] = .01.$$

And if she uses *those* estimates to accommodate \mathcal{E} via J-Kon, she will end up with a different posterior credal state:

$$\begin{aligned} new^*(w_1) &= 0.0157308 & new^*(w_2) &= 8.27938 \cdot 10^{-6} \\ new^*(w_3) &= 0.82711 & new^*(w_4) &= 0.157151 \end{aligned}$$

This is the same posterior credal state that Vanji would end up with if \mathcal{E} “directly” affected her credences for $Test +$ and $Test -$, pushing them to 0.842841 and 0.157159, respectively, and she then updated by J-Con. The moral: her conditional estimates reflect the opinion that there is a marginal chance of mistakenly having \mathcal{E} . J-Kon tells Vanji that, in light of this, she should hedge her epistemic bets that \mathcal{E} is on the money by not quite conditioning, but Jeffrey conditioning instead. As is clear from this example, such hedging can have a *big* impact on your posterior credences. *Marginal* chances of error are not necessarily *negligible*.

2. Uncertain learning with Jeffrey shifts.⁶ Nahdika is a histopathologist. She recently received a section of tissue surgically removed from a pancreatic tumor. She hopes to settle on a diagnosis by examining the tissue under a microscope. (To simplify matters, suppose that exactly one of the three diagnoses is correct.) Nahdika has prior credences over the power set \mathcal{F} of \mathcal{W} , which contains:

$$\begin{aligned} w_1 &= A \& \neg B \& \neg C \& L & w_2 &= A \& \neg B \& \neg C \& \neg L \\ w_3 &= \neg A \& B \& \neg C \& L & w_4 &= \neg A \& B \& \neg C \& \neg L \\ w_5 &= \neg A \& \neg B \& C \& L & w_6 &= \neg A \& \neg B \& C \& \neg L \end{aligned}$$

where

$$\begin{aligned} A &= \text{Islet cell carcinoma} & B &= \text{Ductal cell carcinoma} \\ C &= \text{Benign tumor} & L &= \text{Patient lives} \end{aligned}$$

Her prior credences are as follows:

⁶This case is adapted from (Jeffrey, 1992, pp. 7-9).

$$\begin{aligned}
old(w_1) &= 0.2 & old(w_2) &= 0.3 \\
old(w_3) &= 0.15 & old(w_4) &= 0.1 \\
old(w_5) &= 0.225 & old(w_6) &= 0.025
\end{aligned}$$

So Nahdika's priors for A , B and C are as follows:

$$old(A) = 0.5 \text{ and } old(B) = old(C) = 0.25.$$

Likewise, her priors for L conditional on A , B and C , respectively, are as follows:

$$old(L | A) = 0.4, \text{ } old(L | B) = 0.6 \text{ and } old(L | C) = 0.9.$$

Nahdika looks in the microscope. This causes her to have a learning experience \mathcal{E} that pushes her credence for A , B and C to $1/3$, $1/6$ and $1/2$, respectively.

Question: What should Nahdika's new credences be after \mathcal{E} ?

Answer: According to J-Kon, she ought to update as follows. Firstly, adopt conditional chance estimates of the form $est[ch(\mathcal{E} | w) | w]$. For example, Nahdika might adopt the following conditional chance estimates:

$$est[ch(\mathcal{E} | w_1) | w_1] = est[ch(\mathcal{E} | w_2) | w_2] = est[ch(\mathcal{E} | w_3) | w_3] = est[ch(\mathcal{E} | w_4) | w_4] = 0.3$$

as well as:

$$est[ch(\mathcal{E} | w_5) | w_5] = est[ch(\mathcal{E} | w_6) | w_6] = 0.9.$$

(We will explore why these might be sensible conditional chance estimates shortly.)

Next, she should use these new conditional chance estimates to update her old credences as follows:

$$new(w_i) = \frac{est[ch(\mathcal{E} | w_i) | w_i] \cdot old(w_i)}{\sum_j est[ch(\mathcal{E} | w_j) | w_j] \cdot old(w_j)}.$$

This yields:

$$\begin{aligned}
new(w_1) &= 0.133333 & new(w_2) &= 0.2 \\
new(w_3) &= 0.1 & new(w_4) &= 0.0666667 \\
new(w_5) &= 0.45 & new(w_6) &= 0.05
\end{aligned}$$

So Nahdika's posteriors for A , B and C are:

$$new(A) = 1/3, \text{ } new(B) = 1/6 \text{ and } new(C) = 1/2.$$

Her posteriors for L conditional on A , B and C remain unchanged:

$$new(L | A) = 0.4, \text{ } new(L | B) = 0.6 \text{ and } new(L | C) = 0.9.$$

This is precisely the same posterior credal state that Nahdika would end up with if she accommodated the Jeffrey shift induced by \mathcal{E} via J-Con. But what is it about Nahdika's conditional chance estimates that forces J-Kon to agree with J-Con? And when might it make sense to adopt conditional chance estimates like Nahdika's?

The answer to the first question is given by (★). Choose any learning experience \mathcal{E} that induces a Jeffrey shift over a partition $\{E_1, \dots, E_n\}$. To keep the new credences recommended by J-Kon separate from the "direct" effects of \mathcal{E} , let dir capture the latter. So in Nahdika's case we have:

$$\text{dir}(A) = 1/3, \text{dir}(B) = 1/6 \text{ and } \text{dir}(C) = 1/2.$$

Then J-Kon agrees with J-Con if and only if:

(★) For all $X \in \{E_1, \dots, E_n\}$ and $w_i \in X$, if $\text{old}(w_i) = 0$ then $\text{new}(w_i) = 0$, and if $\text{old}(w_i) > 0$ then

$$\frac{\text{dir}(X)}{\text{old}(X)} = \frac{\text{est}[\text{ch}(\mathcal{E} \mid w_i) \mid w_i]}{\sum_j \text{est}[\text{ch}(\mathcal{E} \mid w_j) \mid w_j] \cdot \text{old}(w_j)}.$$

The second question, then, amounts to the following: When might it make sense to adopt conditional chance estimates that satisfy (★), as Nahdika's do? The answer is not obvious. (★) is a strong condition. Spelling out its consequences requires work. Here is one such consequence. (★) implies (♥):

(♥) For all $X \in \{E_1, \dots, E_n\}$ and $w_i, w_j \in X$ with $\text{old}(w_i) > 0$ and $\text{old}(w_j) > 0$

$$\text{est}[\text{ch}(\mathcal{E} \mid w_i) \mid w_i] = \text{est}[\text{ch}(\mathcal{E} \mid w_j) \mid w_j].$$

Having conditional chance estimates that satisfy (♥) is only appropriate if you take a particular view of the credence formation process \mathbb{P} that induces the Jeffrey shift; the one that translates perceptive and proprioceptive inputs into new credences for elements of $\{E_1, \dots, E_n\}$. In particular, you must think that \mathbb{P} is transparent about what it is causally sensitive to. You must think that if \mathbb{P} is causally sensitive to the differences between w_i and w_j , in the sense that the chance that \mathbb{P} produces \mathcal{E} if you are in w_i is *different*, in your view, than the chance if you are in w_j —*i.e.*, $\text{est}[\text{ch}(\mathcal{E} \mid w_i) \mid w_i] \neq \text{est}[\text{ch}(\mathcal{E} \mid w_j) \mid w_j]$ —then \mathbb{P} will be transparent about this fact. It will announce it to the world by inducing a shift over a sufficiently fine partition; one that slots w_i and w_j into different cells.

But when is it appropriate to adopt conditional chance estimates that reflect this sort of opinion about \mathbb{P} ? Partial answer: not if \mathcal{E} introduces a new undercutting defeater D for one of the E_i . Consider, for example, the learning experience \mathcal{E}^* in our heart attack example (§1). Recall, \mathcal{E}^* involved a tingling arm which caused your credence that you are about to have a heart attack to shoot up toward 1. \mathcal{E}^* ought to cause the proposition L (that you lean on your elbow for an hour) to become an undercutting defeater for the proposition H (that you will have a heart attack) only if leaning on your elbow makes a difference to the chance of having a tingling arm and subsequent credence-in- H boost. If leaning on your elbow makes no difference to the chance of all that (*i.e.*, to \mathcal{E}^*), then learning L ought to leave your post- \mathcal{E}^* credence in H intact.

The upshot: whenever \mathcal{E} introduces a new undercutting defeater D for one of the E_i , (♥) will fail. And whenever (♥) fails, J-Kon and J-Con come apart. This helps to explain how J-Kon might naturally introduce new undercutting defeaters where J-Con stumbles.

(★) has other consequences too. For example, it implies (♦):

(♦) For all $X, Y \in \{E_1, \dots, E_n\}$, $w_i \in X$ and $w_j \in Y$ with $\text{old}(w_i) > 0$ and $\text{old}(w_j) > 0$:

$$\frac{\text{est}[\text{ch}(\mathcal{E} \mid w_i) \mid w_i]}{\text{est}[\text{ch}(\mathcal{E} \mid w_j) \mid w_j]} = \frac{\text{dir}(X)/\text{dir}(Y)}{\text{old}(X)/\text{old}(Y)}.$$

Having conditional chance estimates that satisfy (\blacklozenge) is only appropriate if you think that \mathbb{P} “tracks the chances” in a certain sense. For example, if you used to think that X and Y were equally likely (*i.e.*, $old(X)/old(Y) = 1$), then you must think that \mathbb{P} will produce an experience \mathcal{E} that makes you think that X is twice as likely as Y just in case the *chance* of having \mathcal{E} in any X -world is, in your view, twice as great as the chance of having \mathcal{E} in any Y -world. In this sense, \mathbb{P} must “track the chances.” More generally, you must think that \mathbb{P} will produce an \mathcal{E} that sets your new odds for X and Y (*i.e.*, $dir(X)/dir(Y)$) to be k times your old odds (*i.e.*, $old(X)/old(Y)$) just in case the chance of having \mathcal{E} in any X -world is, in your view, k -times as great as the chance of having \mathcal{E} in any Y -world (*i.e.*, $est[ch(\mathcal{E} | w_i) | w_i]/est[ch(\mathcal{E} | w_j) | w_j] = k$).

Back to our second question then. When might it make sense to adopt conditional chance estimates that satisfy (\blackstar)? We now have an informative partial answer: only if you think that the credence formation process \mathbb{P} in play is “transparent about what it is causally sensitive to” in the sense of (\blackheartsuit), and “tracks the chances” in the sense of (\blacklozenge).

Now consider Nahdika again. In particular, consider her view about the chance of having a learning experience \mathcal{E} that impacts her like so:

$$dir(A) = 1/3, dir(B) = 1/6 \text{ and } dir(C) = 1/2.$$

She might well think that the chance of \mathcal{E} is greater in A -worlds (the patient has Islet cell carcinoma) than in B -worlds (the patient has Ductal cell carcinoma). In that case, Nahdika’s original chance estimates are inappropriate. They do not reflect this opinion. After all, w_1 is an A -world, w_3 a B -world, but nevertheless:

$$\frac{est[ch(\mathcal{E} | w_1) | w_1]}{est[ch(\mathcal{E} | w_3) | w_3]} = \frac{0.3}{0.3} = 1.$$

Instead, Nahdika might opt for the following:

$$\begin{aligned} est^*[ch(\mathcal{E} | w_1) | w_1] &= est^*[ch(\mathcal{E} | w_2) | w_2] = 0.6 \\ est^*[ch(\mathcal{E} | w_3) | w_3] &= est^*[ch(\mathcal{E} | w_4) | w_4] = 0.3 \\ est^*[ch(\mathcal{E} | w_5) | w_5] &= est^*[ch(\mathcal{E} | w_6) | w_6] = 0.9 \end{aligned}$$

And if she uses *those* estimates to accommodate \mathcal{E} via J-Kon, she will end up with a different posterior credal state:

$$\begin{aligned} new^*(w_1) &= 0.2 & new^*(w_2) &= 0.3 \\ new^*(w_3) &= 0.075 & new^*(w_4) &= 0.05 \\ new^*(w_5) &= 0.3375 & new^*(w_6) &= 0.0375 \end{aligned}$$

So Nahdika’s posteriors for A , B and C are:

$$new^*(A) = 1/2, new^*(B) = 1/8 \text{ and } new^*(C) = 3/8.$$

Indeed, new^* is the same posterior credal state that Nahdika would end up with if she used this distribution over $\{A, B, C\}$, rather than the “direct effect” of \mathcal{E} , as an input to J-Con. The moral seems to be this: by taking the “direct effect” of \mathcal{E} at face value—using it as an input to updating—J-Con tacitly presupposes that the credence formation process \mathbb{P} that produced \mathcal{E} satisfies (\blackstar). This means presupposing that \mathbb{P} is “transparent about what it is causally

sensitive to” in the sense of (♥), and “tracks the chances” in the sense of (♦). If (★) holds, and consequently (♥) and (♦) hold, then J-Kon and J-Con agree. If not, not.

When (♦) fails, but (♥) holds, J-Kon says: use your views about the *way* in which \mathbb{P} fails to track the chances to adjust the *input* to J-Con. (More specifically, keep the input *partition* fixed, but adjust the *distribution* over that partition.) When (♥) fails, J-Kon recommends more radical departures from J-Con; departures which are inconsistent with any naïve application of J-Con.

3. Uncertain learning without Jeffrey shifts. Aamilah is walking past the abandoned coin factory at night. She picks up an old coin. She plans to flip it three times. She has prior credences over the power set \mathcal{F} of \mathcal{W} , which contains:

$$\begin{array}{ll} w_1 = H_1 \& H_2 \& H_3 & w_2 = H_1 \& H_2 \& T_3 \\ w_3 = H_1 \& T_2 \& H_3 & w_4 = H_1 \& T_2 \& T_3 \\ w_5 = T_1 \& H_2 \& H_3 & w_6 = T_1 \& H_2 \& T_3 \\ w_7 = T_1 \& T_2 \& H_3 & w_8 = T_1 \& T_2 \& T_3 \end{array}$$

Her prior credences are as follows:

$$\begin{array}{ll} \text{old}(w_1) = .25 & \text{old}(w_2) = .08\bar{3} \\ \text{old}(w_3) = .08\bar{3} & \text{old}(w_4) = .08\bar{3} \\ \text{old}(w_5) = .08\bar{3} & \text{old}(w_6) = .08\bar{3} \\ \text{old}(w_7) = .08\bar{3} & \text{old}(w_8) = .25 \end{array}$$

(There are just the credences she would have if she had a uniform distribution over hypotheses about the bias of the coin.) Aamilah flips the coin three times. In the black of night, she glimpses the outcome of each flip. The first flip causes her to have a fairly ambiguous heads-ish visual experience \mathcal{E}_1 . The second and third flips cause her to have fairly ambiguous tails-ish visual experiences, \mathcal{E}_2 and \mathcal{E}_3 .



Question: What should Aamilah’s new credences be after \mathcal{E}_1 ? \mathcal{E}_2 ? \mathcal{E}_3 ?

Answer: According to J-Kon, she ought to update as follows. Firstly, adopt conditional chance estimates of the form:

$$\text{est}[ch(\mathcal{E}_1 | w) | w].$$

For example, Aamilah might adopt the following conditional chance estimates:

$$\text{est}[ch(\mathcal{E}_1 | w_1) | w_1] = \text{est}[ch(\mathcal{E}_1 | w_2) | w_2] = \text{est}[ch(\mathcal{E}_1 | w_3) | w_3] = \text{est}[ch(\mathcal{E}_1 | w_4) | w_4] = 0.8$$

as well as:

$$\text{est}[ch(\mathcal{E}_1 | w_5) | w_5] = \text{est}[ch(\mathcal{E}_1 | w_6) | w_6] = \text{est}[ch(\mathcal{E}_1 | w_7) | w_7] = \text{est}[ch(\mathcal{E}_1 | w_8) | w_8] = 0.2.$$

Such estimates seem appropriate if you think (i) that there is a high (low) chance of having an ambiguously heads-ish visual experience in a dark environment given that the first flip comes up heads (tails), and (ii) the chance of having that experience is independent of the outcome of the second and third toss conditional on the outcome of the first.

Next, she should use these conditional chance estimates to update her old credences as follows:

$$new_1(w_i) = \frac{est[ch(\mathcal{E} | w_i) | w_i] \cdot old(w_i)}{\sum_j est[ch(\mathcal{E} | w_j) | w_j] \cdot old(w_j)}.$$

This yields:

$$\begin{aligned} new_1(w_1) &= .4 & new_1(w_2) &= .1\bar{3} \\ new_1(w_3) &= .1\bar{3} & new_1(w_4) &= .1\bar{3} \\ new_1(w_5) &= .0\bar{3} & new_1(w_6) &= .0\bar{3} \\ new_1(w_7) &= .0\bar{3} & new_1(w_8) &= .1 \end{aligned}$$

So Aamilah's ambiguously heads-ish learning experience makes her fairly confident (credence 0.8) that the coin came up heads on the first flip. This, in turn, is evidence that the coin is biased towards heads, and hence increases her confidence that it will come up heads on flips two and three.

$$new_1(H_1) = .8 \text{ and } new_1(H_2) = new_1(H_3) = .6.$$

This is precisely the same posterior credal state that Aamilah would end up with if \mathcal{E}_1 "directly" affected her credences for H_1 and T_1 , pushing them to 0.8 and 0.2, respectively, and she then updated by J-Con.

What credences should Aamilah have after \mathcal{E}_2 ? According to J-Kon, she ought to adopt new conditional chance estimates of the form:

$$est[ch(\mathcal{E}_2 | w \& \mathcal{E}_1) | w \& \mathcal{E}_1].$$

And she ought to do so in a way that preserves coherence, so that her new expanded set of conditional chance estimates satisfies the laws of conditional estimation. For example, Aamilah might adopt the following:

$$\begin{aligned} & est[ch(\mathcal{E}_2 | w_1 \& \mathcal{E}_1) | w_1 \& \mathcal{E}_1] & (1) \\ & = est[ch(\mathcal{E}_2 | w_2 \& \mathcal{E}_1) | w_2 \& \mathcal{E}_1] & (2) \\ & = est[ch(\mathcal{E}_2 | w_7 \& \mathcal{E}_1) | w_7 \& \mathcal{E}_1] & (3) \\ & = est[ch(\mathcal{E}_2 | w_8 \& \mathcal{E}_1) | w_8 \& \mathcal{E}_1] & (4) \\ & = .001 & (5) \end{aligned}$$

as well as:

$$\begin{aligned} & est[ch(\mathcal{E}_2 | w_3 \& \mathcal{E}_1) | w_3 \& \mathcal{E}_1] & (6) \\ & = est[ch(\mathcal{E}_2 | w_4 \& \mathcal{E}_1) | w_4 \& \mathcal{E}_1] & (7) \\ & = .8 & (8) \end{aligned}$$

and:

$$\begin{aligned} & est[ch(\mathcal{E}_2 | w_5 \& \mathcal{E}_1) | w_5 \& \mathcal{E}_1] & (9) \\ & = est[ch(\mathcal{E}_2 | w_6 \& \mathcal{E}_1) | w_6 \& \mathcal{E}_1] & (10) \\ & = .2 & (11) \end{aligned}$$

(1)-(5) seem appropriate if you think your visual system is stable. In what sense? The following: if you had an ambiguously heads-ish visual experience in response to a heads on the first flip, then there is almost no chance that you will have a different (tails-ish) experience in response to the same outcome (heads) on the second. Ditto for tails. If you had a heads-ish experience in response to a *tails* on the first flip, there is almost no chance that you will have a different (tails-ish) experience in response to the same outcome (tails) on the second.

(6)-(8) seem appropriate if you think that having a heads-ish visual experience in response to heads on the first flip is good evidence that the lighting conditions are just good enough for your visual system to be sensitive to the outcome. In that case, you think: there is a reasonably high chance of having a tails-ish experience if you get tails on the second.

(9)-(11) seem appropriate if you think that having a heads-ish visual experience in response to *tails* on the first flip is good evidence that the lighting conditions are just *bad* enough for your visual system to *not* be sensitive to the outcome. In that case, you think: there is a high chance of having the same heads-ish experience even if you get tails on the second. So there is a *low* chance of having the tails-ish experience described by \mathcal{E}_2 .

With these estimates in hand, Aamilah ought to update as follows:

$$new_2(w_i) = \frac{est[ch(\mathcal{E}_2 | w_i \& \mathcal{E}_1) | w_i \& \mathcal{E}_1] \cdot new_1(w_i)}{\sum_j est[ch(\mathcal{E}_2 | w_j \& \mathcal{E}_1) | w_j \& \mathcal{E}_1] \cdot new_1(w_j)}.$$

This yields:

$$\begin{array}{ll} new_2(w_1) = 0.00175953 & new_2(w_2) = 0.00058651 \\ new_2(w_3) = 0.469208 & new_2(w_4) = 0.469208 \\ new_2(w_5) = 0.0293255 & new_2(w_6) = 0.0293255 \\ new_2(w_7) = 0.000146628 & new_2(w_8) = 0.000439883 \end{array}$$

So Aamilah's credence that the first and second flips resulted in heads and tails, respectively, are now roughly 0.94. Her credence that the third flip will come up heads still hovers around 0.5.

$$new_2(H_1) = 0.940762, \quad new_2(T_2) = 0.939003, \quad new_2(H_3) = 0.50044.$$

But why would Aamilah's credence that the *first* flip came up heads shoot up (from 0.8 to 0.94) in response to an ambiguously tails-ish visual experience following the *second* flip? The reason: her conditional chance estimates reflect the opinion that having *different* visual experiences in response to the first two flips is good news about the reliability of her visual system. It is evidence that the lighting conditions are just good enough for her visual system to track the outcome. Wiping away this uncertainty about the reliability of her visual system pushes her credence in H_1 up from 0.8 to 0.94.

Finally, what credences should Aamilah have after \mathcal{E}_3 ? According to J-Kon, she ought to adopt new conditional chance estimates of the form:

$$est[ch(\mathcal{E}_3 | w \& \mathcal{E}_1 \& \mathcal{E}_2) | w \& \mathcal{E}_1 \& \mathcal{E}_2]$$

while again ensuring to preserve coherence. She might, for example, adopt the following:

$$est[ch(\mathcal{E}_3 | w_1 \& \mathcal{E}_1 \& \mathcal{E}_2) | w_1 \& \mathcal{E}_1 \& \mathcal{E}_2] = .4 \quad (12)$$

$$est[ch(\mathcal{E}_3 | w_2 \& \mathcal{E}_1 \& \mathcal{E}_2) | w_2 \& \mathcal{E}_1 \& \mathcal{E}_2] = .6 \quad (13)$$

$$est[ch(\mathcal{E}_3 | w_3 \& \mathcal{E}_1 \& \mathcal{E}_2) | w_3 \& \mathcal{E}_1 \& \mathcal{E}_2] = .0001 \quad (14)$$

$$est[ch(\mathcal{E}_3 | w_4 \& \mathcal{E}_1 \& \mathcal{E}_2) | w_4 \& \mathcal{E}_1 \& \mathcal{E}_2] = .9999 \quad (15)$$

$$est[ch(\mathcal{E}_3 | w_5 \& \mathcal{E}_1 \& \mathcal{E}_2) | w_5 \& \mathcal{E}_1 \& \mathcal{E}_2] = .8 \quad (16)$$

$$est[ch(\mathcal{E}_3 | w_6 \& \mathcal{E}_1 \& \mathcal{E}_2) | w_6 \& \mathcal{E}_1 \& \mathcal{E}_2] = .2 \quad (17)$$

$$est[ch(\mathcal{E}_3 | w_7 \& \mathcal{E}_1 \& \mathcal{E}_2) | w_7 \& \mathcal{E}_1 \& \mathcal{E}_2] = .4 \quad (18)$$

$$est[ch(\mathcal{E}_3 | w_8 \& \mathcal{E}_1 \& \mathcal{E}_2) | w_8 \& \mathcal{E}_1 \& \mathcal{E}_2] = .6 \quad (19)$$

(12)-(13) and (18)-(19) seem appropriate if having *different* visual experiences (a heads-ish and tails-ish experience, respectively) in response to the *same* outcome on flips 1 and 2—which is precisely what happens in w_1/w_2 and w_7/w_8 —makes you think that, in the current lighting conditions, a small tilt of the head this way, or a half step that way might well make you have an entirely different experience regardless of how the coin lands. The output of your visual system in current conditions is rather fragile. In that case, one “success” (*e.g.*, head-ish experience in response to heads on trial 1) and one “failure” (*e.g.*, tails-ish experience in response to heads on trial 2) might cause you to think that your chance of success is not much better than 0.5, *e.g.*, 0.6 (down from $est[ch(\mathcal{E}_1 | w_1) | w_1] = 0.8$.)

(14)-(15) seem appropriate if you think that having a heads-ish visual experience in response to heads on the first flip, and a tails-ish experience in response to tails on the second flip provide overwhelming evidence that the lighting conditions are good enough for your visual system to reliably track the outcome.

(16)-(17) seem appropriate if you think that having a heads-ish visual experience in response to *tails* on the first flip, and a tails-ish experience in response to *heads* on the second flip provide good evidence that (i) there is enough light for your visual system to be sensitive to *something*, but (ii) it is sensitive to an oddly misleading set of features; a set of features that renders your visual system *anti-reliable*.

Then yet again Aamilah ought to update as follows:

$$new_3(w_i) = \frac{est[ch(\mathcal{E}_3 | w_i \& \mathcal{E}_1 \& \mathcal{E}_2) | w_i \& \mathcal{E}_1 \& \mathcal{E}_2] \cdot new_2(w_i)}{\sum_j est[ch(\mathcal{E}_3 | w_j \& \mathcal{E}_1 \& \mathcal{E}_2) | w_j \& \mathcal{E}_1 \& \mathcal{E}_2] \cdot new_2(w_j)}.$$

This yields:

$$\begin{aligned} new_3(w_1) &= 0.00140787 & new_3(w_2) &= 0.000703936 \\ new_3(w_3) &= 0.0000938582 & new_3(w_4) &= 0.938488 \\ new_3(w_5) &= 0.0469291 & new_3(w_6) &= 0.0117323 \\ new_3(w_7) &= 0.000117323 & new_3(w_8) &= 0.000527953 \end{aligned}$$

So Aamilah’s post- \mathcal{E}_3 credences for heads on the first flip, tails on the second flip, and tails on the third flip are as follows:

$$new_3(H_1) = 0.940693, new_3(T_2) = 0.939227, new_3(T_3) = 0.951452.$$

Of course, the conditional chance estimates that we imagined Aamilah adopting at each stage of our little inference problem are merely illustrative. Nothing in J-Kon *forces* you to adopt such estimates. J-Kon permits using *any* conditional chance estimates as inputs to updating so long as (i) each time you expand that set of chance estimates you preserve coherence, and (ii) your estimates are consistent with the Principal Principle, in the sense that expanding to include estimates of the form $est[X | \phi_{ch} \& Y] = ch(X | Y)$ —where ϕ_{ch} is the proposition that ch is the true chance function—also preserves coherence. So the examples in §2 describe only *one* way you might apply J-Kon, not *the* way.

3 Equivalence to Gilded Superconditioning

We now have a new policy for rational learning on the table. To recap, J-Kon in its unofficial form says: when you have a learning experience, \mathcal{E} , you ought to estimate the chance of having that very experience conditional on being in this, that, or the other world (*expansion stage*). Then you ought to treat these new conditional chance estimates, $est[ch(\mathcal{E} | w) | w]$, as your old conditional credences, $old(\mathcal{E} | w)$ (conditional credences which you in fact lack, since you have no pre- \mathcal{E} opinions about \mathcal{E}). Input your old unconditional credences, $old(w)$, and your new conditional chance estimates, $est[ch(\mathcal{E} | w) | w]$, into Bayes' theorem. Then conditionalize on \mathcal{E} . In its official form (detailed shortly), J-Kon says that you ought to update *as if* you were expanding and conditioning in this way.

This updating policy might seem as if it was plucked out of thin air. But in fact it characterises a brand of what Jeffrey called “superconditioning.” Your new credences $new : \mathcal{F} \rightarrow \mathbb{R}$ come from your old credences $old : \mathcal{F} \rightarrow \mathbb{R}$ by superconditioning when you can:

- (i) Expand \mathcal{F} to a larger σ -algebra \mathcal{F}^+
- (ii) Extend $old : \mathcal{F} \rightarrow \mathbb{R}$ to $old^+ : \mathcal{F}^+ \rightarrow \mathbb{R}$
- (iii) Obtain new^+ from old^+ by conditioning on some proposition E in the larger algebra \mathcal{F}^+ , so that $new^+(\cdot) = old^+(\cdot | E)$
- (iv) Recover new by restricting new^+ to \mathcal{F} .

Pictorially, superconditioning looks like this:

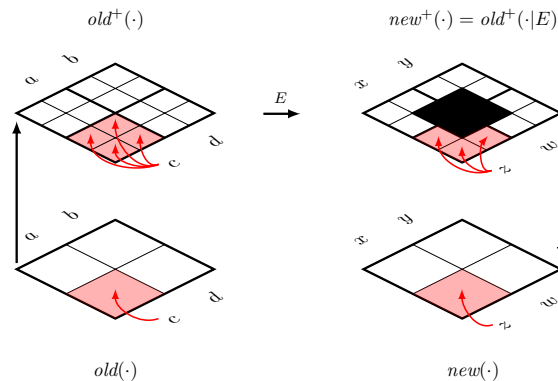


Figure 1: $old(\cdot)$ can be extended to $old^+(\cdot)$ defined on \mathcal{F}^+ and $new(\cdot)$ can be recovered by cutting $old^+(\cdot | E)$ back down to \mathcal{F} .

Diaconis and Zabell (1982, p. 824) provide necessary and sufficient conditions for your new credences to come from your old ones by superconditioning. They prove that $new : \mathcal{F} \rightarrow \mathbb{R}$ comes from $old : \mathcal{F} \rightarrow \mathbb{R}$ by superconditioning just in case there is some upper bound $b \geq 1$ on your probability ratios, so that for every $X \in \mathcal{F}$ we have

$$b \geq \frac{new(X)}{old(X)}.$$

Diaconis and Zabell are concerned with what you might call *austere superconditioning*. Austere superconditioning places no constraints at all on what the larger algebra \mathcal{F}^+ looks like. Neither does it place any constraints on the extended prior old^+ , save for probabilistic coherence.

But we might hope that old^+ satisfies norms of epistemic rationality beyond probabilism. For example, we might require old^+ to satisfy Lewis' Principal Principle (PP), so that once chance is brought up to speed on your past learning experiences \mathcal{P} , old^+ treats it as an expert worthy of full deference:⁷

$$old^+(X | Y \& \phi_{ch}) = ch(X | Y \& \mathcal{P}).$$

We might also hope that we can recover new not simply by conditioning old^+ on *some* proposition or other in \mathcal{F}^+ and then cutting back down to \mathcal{F} , but rather by conditioning old^+ on the proposition \mathcal{E} describing *your learning experience*.

Finally, we might hope that not only can we obtain new —your post- \mathcal{E} credences—from old —your pre- \mathcal{E} credences—in this way, but moreover that your *entire epistemic life* hangs together in the right way. Let $c_0 : \mathcal{F} \rightarrow \mathbb{R}$ be your initial credence function, $c_1 : \mathcal{F} \rightarrow \mathbb{R}$ be your credence function after learning experience \mathcal{E}_1 , $c_2 : \mathcal{F} \rightarrow \mathbb{R}$ be your credence function after learning experience \mathcal{E}_2 , and so on. Let $\mathbf{c} = \langle c_0, \dots, c_n \rangle$ be the sequence of credal states that you adopt over the course of your life. Call \mathbf{c} your *epistemic life*. Then we might hope that the various stages of your epistemic life hang together in the following sense: we can extend c_0 to some Principal Principle satisfying c_0^+ defined on \mathcal{F}^+ and recover each c_i by conditioning c_0^+ on the proposition $\mathcal{E}_1 \& \dots \& \mathcal{E}_i$ describing your learning experiences to that point in your life and cutting back down to \mathcal{F} .

When \mathbf{c} hangs together in this way, call it a *gilded superconditioning life*. More carefully, say that \mathbf{c} is a gilded superconditioning life when you can:

- (v) Expand \mathcal{F} to the σ -algebra \mathcal{F}^+ generated by $\mathcal{F} \cup \{\mathcal{E}_1, \dots, \mathcal{E}_n\} \cup \{\phi_{ch} | ch \in \mathcal{C}\}$, where $\mathcal{E}_1, \dots, \mathcal{E}_n$ are demonstrative propositions describing your learning experiences over the course of your life, ϕ_{ch} is the proposition that ch is the true chance function, and \mathcal{C} is the set of possible chance functions.
- (vi) Extend $c_0 : \mathcal{F} \rightarrow \mathbb{R}$ to $c_0^+ : \mathcal{F}^+ \rightarrow \mathbb{R}$, where c_0^+ is not only probabilistically coherent, but also satisfies the Principal Principle.
- (vii) Obtain c_i^+ from c_0^+ by conditioning on $\mathcal{E}_1 \& \dots \& \mathcal{E}_i$, for all $i \leq n$, so that $c_i^+(\cdot) = c_0^+(\cdot | \mathcal{E}_1 \& \dots \& \mathcal{E}_i)$.
- (viii) Recover c_i by restricting c_i^+ to \mathcal{F} .

A natural question, then: when exactly is your epistemic life a gilded superconditioning life? The answer: exactly when you update by the official version of J-Kon.

⁷This is a variant of what (Pettigrew, 2016, p. 135) calls the “Extended Principal Principle.” See chapters 9 and 10 of Pettigrew (2016) for a careful discussion of the strengths and weaknesses of various formulations of the Principal Principle.

J-Kon (Official Version). Suppose $\mathbf{c} = \langle c_0, \dots, c_n \rangle$ is your epistemic life. So $c_0 : \mathcal{F} \rightarrow \mathbb{R}$ is your initial credence function, $c_1 : \mathcal{F} \rightarrow \mathbb{R}$ is the credence function you adopt in response to learning experience \mathcal{E}_1 , and so on. Let \mathcal{A} be the set of atoms of \mathcal{F} , \mathcal{W} be the set of atoms of \mathcal{F}^+ , \mathcal{L}_i be shorthand for $\mathcal{E}_1 \& \dots \& \mathcal{E}_i$, and \mathbb{V} be the set of variables $\mathcal{V} : \mathcal{W} \rightarrow \mathbb{R}$. Then \mathbf{c} ought to be such that there is a coherent partial assignment of conditional estimates, $est : \mathbb{V} \times \mathcal{F}^+ \rightarrow \mathbb{R}$, which satisfies two conditions: (i) est is consistent with the Principal Principle, in the sense that extending it to include estimates of the form

$$est[X | \phi_{ch} \& Y] = ch(X | Y) \quad (\text{for all } X, Y \in \mathcal{F}^+ \text{ and } ch \in \mathcal{C})$$

preserves coherence. Moreover, (ii) for all $w \in \mathcal{A}$ and $0 < i \leq n$

$$c_0(w) = est[ch(w | \mathcal{W}) | \mathcal{W}]$$

and

$$c_i(w) = \frac{est[ch(\mathcal{E}_i | w \& \mathcal{L}_{i-1}) | w \& \mathcal{L}_{i-1}] \cdot c_{i-1}(w)}{\sum_{w' \in \mathcal{A}} est[ch(\mathcal{E}_i | w' \& \mathcal{L}_{i-1}) | w' \& \mathcal{L}_{i-1}] \cdot c_{i-1}(w')}$$

Your epistemic life satisfies J-Kon just in case it is a gilded superconditioning life (see appendix for proof).

What this means is that even when you are not in a position to condition on the proposition \mathcal{E} describing your learning experience—perhaps because you have no prior opinions about \mathcal{E} (it is not in your algebra)—you should nonetheless, according to J-Kon, update your “small space credences,” *i.e.*, your credences over “first order” propositions in \mathcal{F} , just as a more opinionated Bayesian agent would. Such an agent would update her Principal Principle satisfying “big space credences”—her credences over “first order” propositions as well as “higher order” credences in \mathcal{E} and relevant chance hypotheses ϕ_{ch} —by conditioning on \mathcal{E} . You should mimic her in your small space. Your update should be something like a 2D cross-section of her 3D update.




Figure 2: 3D dynamics (top) vs. 2D cross-section dynamics (bottom).

Of course, there are *many* ways of extending your old small space credences to PP-satisfying big space credences. Accordingly, there will be *many* ways

of mimicking a big space Bayesian agent. This is why there is slack in J-Kon that there is not in J-Con. There is a one-to-one correspondence between PP-satisfying extensions and the sets of coherent conditional chance estimates that J-Kon takes as inputs. J-Kon simply recommends using *some* such set as an input. It does not specify *which* set to use.

You might wonder, though, what J-Kon *really* demands of us. It is fine to say that you should update *as if* you were explicitly settling on new conditional chance estimates (expanding), treating them as your old conditional credences, and conditionalizing on \mathcal{E} . Equivalently, it is fine to say that you should update in such a way that *there exists* some conditional chance estimates that you *could in principle* treat as your old conditional credences and use to obtain your new credences by conditionalizing on \mathcal{E} . But what sorts of updates have this property? What we want here is something like Diaconis and Zabell's characterisation of superconditioning. We want a characterisation of J-Kon—or equivalently, gilded superconditioning—that only references intrinsic features of your epistemic life.

Luckily, such a characterisation is not too hard to find. Let $\mathbf{c} = \langle c_0, \dots, c_n \rangle$ be your epistemic life. Note that each stage of your life, c_i , determines an expected value for any variable $\mathcal{V} : \mathcal{A} \rightarrow \mathbb{R}$. Then \mathbf{c} satisfies J-Kon just in case the following holds. Firstly, your initial expectation for \mathcal{V} —the expectation determined by c_0 —must always be at least as great as the smallest possible unconditional *objective* expected value for \mathcal{V} —*i.e.*, the smallest possible expectation determined by some chance function $ch(\cdot | \mathcal{W})$. Secondly, the expectations determined by c_1 must always be at least as great as the smallest possible objective expectations *conditional on* \mathcal{E}_1 —*i.e.*, the smallest possible expectations determined by some chance function $ch(\cdot | \mathcal{E}_1)$. But they may need to be even larger than that. How much larger depends on c_0 . Consider all of the convex combinations of expectation operators (linear prevision) determined by some chance function $ch(\cdot | \mathcal{W})$. They are themselves expectation operators (linear prevision). Now consider the subset of those expectation operators that agree with c_0 on the probability of all $X \in \mathcal{F}$. Finally, condition each of those on \mathcal{E}_1 . c_1 's expectation of any \mathcal{V} must always be at least as great as the smallest possible estimate determined by *that* set of expectation operators. And that value will always be at least as great as the smallest possible expectation determined by some chance function $ch(\cdot | \mathcal{E}_1)$. Similarly, the expectations determined by c_2 must always be at least as great as the smallest possible objective expectations conditional on \mathcal{E}_2 . But they may need to be even larger. How much larger depends in a similar way on c_1 . And so on, and so on.⁸

This tells us precisely what J-Kon demands of us. It tells us precisely what it takes to mimic a big space Bayesian agent in the way that J-Kon (or gilded superconditioning) recommends. But even with a better understanding of *what* J-Kon demands of us, you might wonder *why* any of this is a good idea. There are good epistemic reasons to condition when you have infallible dogmatic learning experiences. For example, Greaves and Wallace (2006) provide an expected accuracy argument for conditioning, and Briggs and Pettigrew (2018) provide an accuracy-dominance argument.⁹ Likewise, there are good epistemic reasons

⁸See proposition 2 in the appendix for formal details.

⁹Both Greaves and Wallace (2006) and Briggs and Pettigrew (2018) tacitly assume that rational agents are infallible dogmatic learners. They measure the epistemic value of an updating plan at a world by the accuracy of the credal state that the plan recommends at the

to satisfy the Principal Principle. Pettigrew (2016), for example, provides an objective expected accuracy (*i.e.*, chance-dominance) argument for PP. But these are all reasons for *actually being* a PP-satisfying conditioner. Is there any reason to *mimic* such an agent? Does small-space updating that mimics PP-satisfying big-space conditioning enjoy any of the epistemic benefits that the latter is privy to?

The answer is *yes*. In the next section I will provide a *chance-dominance* argument for J-Kon, or equivalently, gilded superconditioning. Put roughly: failing to update by J-Kon/gilded superconditioning reduces your chances of living the epistemic good life, or accruing accuracy over the course of your life.

4 Chance-Dominance Argument for J-Kon

Here is how we will proceed. Firstly, we will attempt to pin down what it is that makes one’s credences epistemically valuable at a world, and specify “epistemic utility functions” that measure this sort of value. Secondly, we will use the machinery of decision theory to show that J-Kon can be given an epistemic-value-based rationale. The story in a nutshell is this: if you fail to update by J-Kon, your epistemic life is *chance-dominated* by some other J-Kon satisfying life, *i.e.*, every possible chance function expects your life to accrue strictly less total epistemic value than the J-Kon satisfying life. In this sense, you are *guaranteed* to have a worse chance of living the epistemic good life than you could have had by satisfying J-Kon. If you update by J-Kon, in contrast, you are never chance-dominated in this way. This reveals an *evaluative* defect—chance-dominance—that mars the epistemic lives of J-Kon violators. And facts about such defects ought to inform your preferences over epistemic lives. They give you reason to prefer updating via J-Kon to not.

Here is our main argument:

1. **Veritism:** The ultimate source of epistemic value is accuracy. So the epistemic value of an unconditional credence function $c : \mathcal{F} \rightarrow \mathbb{R}$ at a world w is given by $-\mathcal{I}(c, w)$, where \mathcal{I} is some reasonable measure of inaccuracy. We assume that \mathcal{I} is an additive inaccuracy score generated by a continuous, bounded, strictly proper component function. Moreover, the epistemic value of an epistemic life, or sequence of credal states $\mathbf{c} = \langle c_0, \dots, c_n \rangle$ at w is given by

$$\mathcal{I}(\mathbf{c}, w) = \sum_{i:w \in \mathcal{L}_i} \mathcal{I}(c_i, w)$$

where $\mathcal{L}_i = \mathcal{E}_1 \& \dots \& \mathcal{E}_i$. \mathcal{I} reflects the view that life stages c_i ought to be evaluated as conditional credence functions. More carefully, $c_i(X)$ should be evaluated as your credence for X conditional on the learning experiences that produced c_i . As such, c_i is evaluable for accuracy only at worlds w in which the learning experiences that produced that stage take place, *viz.*, $w \in \mathcal{L}_i$. The total inaccuracy of your epistemic life \mathbf{c} at w is the sum of the degrees of inaccuracy of the life stages c_i that are evaluable for accuracy at w .

world. But this only makes sense if there is no chance that the agent will mis-execute the plan. In this way, they bake infallibility into their measure of epistemic value.

2. **Chance-dominance:** If one option O **strictly chance-dominates** another option O' , in the sense that every possible chance function ch expects O to be strictly more valuable than O' , *i.e.*

$$\sum_{w \in \mathcal{W}} ch(w | \mathcal{W}) \mathcal{U}(O, w) > \sum_{w \in \mathcal{W}} ch(w | \mathcal{W}) \mathcal{U}(O', w)$$

then any rational agent ought to strictly prefer O to O' .

3. **Theorem**¹⁰: If \mathbf{b} is not a J-Kon sequence (or gilded superconditioning sequence), then \mathbf{b} is **strictly chance-dominated** by a J-Kon sequence \mathbf{c} , *i.e.*,

$$\sum_{w \in \mathcal{W}} ch(w | \mathcal{W}) \mathcal{I}(\mathbf{b}, w) > \sum_{w \in \mathcal{W}} ch(w | \mathcal{W}) \mathcal{I}(\mathbf{c}, w)$$

for every possible chance function ch . If \mathbf{b} is a J-Kon sequence, on the other hand, then it is not even **weakly chance-dominated**, *i.e.*, there is no $\mathbf{c} \neq \mathbf{b}$ such that

$$\sum_{w \in \mathcal{W}} ch(w | \mathcal{W}) \mathcal{I}(\mathbf{b}, w) \geq \sum_{w \in \mathcal{W}} ch(w | \mathcal{W}) \mathcal{I}(\mathbf{c}, w)$$

for all possible chance functions ch .

- C. **J-Kon:** If a rational agent fails to update via J-Kon, then she ought to strictly prefer some other J-Kon satisfying epistemic life to her own.

Premise 1 identifies a theory of epistemic value: *veritism*. Veritists claim that accuracy is the principal epistemic good-making feature of credal states. Your credences may well be epistemically laudable for a range of reasons. They may be specific, informative, verisimilar, encode simple, unified explanations, be justified, etc. But these are all *instrumental* epistemic good-making features, on the veritist view. They are good, roughly speaking, as a means to the end of accuracy. Accuracy—how close your credences are to the truth—is the fundamental epistemic good. It is the primary source of epistemic value. The higher your credence in truths and the lower your credence in falsehoods, the more valuable your credal state is from the epistemic perspective.

And this applies not just to your credal state at a single time—to an epistemic time slice—but to your *whole epistemic life*. The principal epistemic good-making feature of your epistemic life—the *sequence* of credal states that you adopt over time—is the total accuracy that your life accrues. This really captures two distinct veritist thoughts. Firstly, the epistemic value of your life supervenes on accuracy of the credal states you adopt over the course of your life. Secondly, the *shape* of your epistemic life does not contribute to its distinctively epistemic value. A life of gradual gains in accuracy is no better or worse, from the epistemic perspective, than a life of alethic ups and downs. Any two lives that accrue the same total accuracy are equally epistemically valuable. The shape of one's

¹⁰We should note that the proof of this theorem makes a number of simplifying assumptions for ease of presentation. For example, we assume that both “small space” and “big space” credence functions are defined on finite algebras. We also assume that all small space credence function are regular, in the sense that they assign positive probability to each world, and that chance functions are quasi-regular, in the sense that they assign positive probability to the proposition $\mathcal{E}_1 \& \dots \& \mathcal{E}_n$ describing the totality of learning experiences that you have over the course of your life. The results in the appendix will be generalised in future work.

epistemic life is not the “right kind of alethic fact” to determine its epistemic value.

Premise 1 also delineates a class of reasonable measures of accuracy, or what’s better for technical purposes, *inaccuracy*. We will assume that inaccuracy is measured by an additive, continuous, bounded, strictly proper inaccuracy score \mathcal{I} . **Additivity** says that \mathcal{I} takes the following form:

$$\mathcal{I}(c, w) = \sum_{X \in \mathcal{F}} s(c(X), w(X))$$

where $s : \mathbb{R} \times \{0, 1\} \rightarrow \mathbb{R}$ is what Joyce (2009) calls a *component function*; a function which measures the inaccuracy of an individual credence $c(X)$ when X ’s truth value is $w(X)$ (0 or 1). **Strict Propriety** says that every probabilistically coherent credence function expects itself to be strictly more accurate than any other credence function. See appendix for technical details. For a philosophical rich discussion of these properties, see Joyce (2009); Pettigrew (2016).

Now extend \mathcal{I} to measure not only the accuracy of individual credence functions c , but also epistemic lives $\mathbf{c} = \langle c_0, \dots, c_n \rangle$ as follows:

$$\mathcal{I}(\mathbf{c}, w) = \sum_{i:w \in \mathcal{L}_i} \mathcal{I}(c_i, w)$$

\mathcal{I} reflects the view that epistemic time slices c_i ought to be evaluated as *conditional credence functions*. More carefully, $c_i(X)$ should be evaluated as an epistemic bet on the truth-value X , but only a conditional one; one which “takes for granted” the learning experiences \mathcal{L}_i that produced it, and so is “called-off”—not evaluable for accuracy—at worlds w in which those learning experiences do not take place. According to \mathcal{I} , the total inaccuracy of your epistemic life \mathbf{c} at w is the sum of the degrees of inaccuracy of the epistemic time slices or life stages c_i that are evaluable for accuracy at w . When \mathcal{I} takes this form, call it *temporally additive*.

Premise 2 identifies a constraint on rational preference. It specifies how rational agents structure their preferences over options in light of facts about their value, together with facts about the chances. In particular, it says that if every possible chance function expects one option to be better than another, then you ought to prefer the one to the other.

Imagine, for example, that a friend tells you that she is going to pick a ball from one of two urns, A or B . Whichever urn she draws from, she will pick the ball at random. Both urns contain red and blue balls in the exact same proportion. But you have no idea what that proportion is. Finally, your friend drops one additional red ball into urn B . Now she gives you a choice. You can select the urn, A or B , and she will give you £100 if she draws a red. (What a friend!) Which should you prefer? Urn B , of course. Whatever proportion you started with, urn B now has a higher objective expected payout than urn A . You have a better chance of making off with the goods if you go with B . Chance dominance says: in light of these facts, you ought to prefer B to A .

Premise 3 is the spine of the argument. It shows that if you violate J-Kon, then your epistemic life is chance-dominated. There is some other J-Kon satisfying life that has higher objective expected value relative to every possible chance function. (See proposition 7 in the appendix.) In conjunction with

premises 1 and 2, this establishes that J-Kon violators should prefer some other J-Kon satisfying life to their own.

This argument provides a purely alethic (rather than evidential or pragmatic) rationale for abiding by J-Kon. The epistemic lives of J-Kon violators are marred by a particular sort of evaluative defect. They are chance-dominated. This is a “wide scope” defect; a defect of epistemic lives which may not be reflected in “lower level defects.” Each stage of an agent’s epistemic life—the credal state that she adopts at any particular time—may be epistemically unimpugnable considered individually. But nevertheless, those stages may together constitute an epistemically defective life.

Upon reflection, this is perhaps not very surprising. Wide scope evaluative defects are precisely the sorts of defects you would *expect* to underpin a genuinely diachronic updating policy like J-Kon.

5 Confirmational Holism

To recap, we began our journey in §1 by outlining Weisberg’s concern. J-Con bungles the introduction of new undercutting defeaters. Since nearly all learning experiences introduce new undercutting defeaters, J-Con almost never applies. We then introduced a new updating policy, J-Kon, designed to apply where J-Con does not. After seeing how J-Kon works in practice, we explored its equivalence to gilded superconditioning. We also provided a purely alethic (accuracy-centred) rationale for abiding by J-Kon. Now we will return to the issues surrounding confirmational holism that set us on this path.

Consider our heart attack case again. For the last hour you have been sitting on the train, reading the news, and, unbeknownst to you, leaning lazily on your left elbow. You have prior credences over the power set \mathcal{F} of \mathcal{W} , which contains:

$$\begin{aligned} w_1 &= L \& H & & w_2 &= L \& \neg H \\ w_3 &= \neg L \& H & & w_4 &= \neg L \& \neg H \end{aligned}$$

where

L = You have been leaning on your elbow for an hour

H = You will have a heart attack.

Your prior credence that you will have a heart attack is low—maybe 0.001. You have no idea whether you have been leaning on your elbow for an hour. So we will say that you are 50-50. Your prior credence for L is 0.5. Finally, you think that leaning lazily on your left elbow is *irrelevant* to your prospects of having a heart attack. Taking on the supposition that you have been leaning on your elbow does nothing to raise or lower your credence that you will have a heart attack:

$$\text{old}(H \mid L) = \text{old}(H) = 0.001.$$

This fixes your prior credal state:

$$\begin{aligned} \text{old}(w_1) &= 0.0005 & \text{old}(w_2) &= 0.4995 \\ \text{old}(w_3) &= 0.0005 & \text{old}(w_4) &= 0.4995 \end{aligned}$$

Then you feel your left arm tingling. This learning experience \mathcal{E} “directly affects” your credence for H , pushing it up to 0.99. But what *should* your new credences

be after \mathcal{E} ? Should you take this “direct product” of \mathcal{E} at face value? Should you stick with 0.99 as your new credence for H , rather than sloughing it off as you would if it were the aberrant result of momentary mania? Or should you correct for biases in your credence-producing processes in some other way? And what about the rest of your new credences? What should they be?

According to Christensen, the doxastic effects of experience should almost always be mediated by your background beliefs. This is one of confirmational holism’s great insights. Whether the tingling in your arm pushes your credence that you are going to have a heart attack up, or down, or leaves it unchanged, ought to depend for example on your background beliefs about your cardiovascular health, what exactly the symptoms of a heart attack are, whether you have been leaning on your left elbow, etc. But J-Con treats the doxastic effects of experience as an exogenous factor; an output of some “black box” process that serves as an input to J-Con. So it fails to vindicate this insight.

The inputs to J-Kon, on the other hand, are not the *direct effects* of experience. Rather, they are quantities that might plausibly be mediated by your background beliefs. Recall, two types of quantities serve as inputs to J-Kon. The first quantity is your old credence function. The second quantity is a set of conditional chance estimates of the form:

$$est[ch(\mathcal{E} \mid w \& \mathcal{P}) \mid w \& \mathcal{P}].$$

The conditional chance estimates that you are disposed to adopt (or behave *as if* you were adopting) reflect your background beliefs. In particular, they reflect your beliefs about how various factors causally influence your learning experiences. For example, you might estimate the chance of having an arm-tingling experience conditional on being in a heart-attack world to be roughly 0.999. This reflects your background belief that heart attacks cause arm tingling. Of course, conditional chance estimates will often reflect background beliefs in varied and complex ways. But the basic point stands: background beliefs influence which conditional chance estimates you are disposed to adopt (or behave as if you were adopting). (Maybe such dispositions even *constitute* background beliefs.) Conditional chance estimates, in turn, serve as an input to J-Kon. As such, they provide a vehicle for background beliefs to systematically influence the doxastic effects of experience.

Of course, J-Kon does not specify precisely *how* an agent’s background beliefs ought to determine how she expands; which conditional chance estimates she ought to adopt. But we side with Jeffrey here. Learning how to change one’s mind involves training the messy network of neurons in your skull to translate perceptive and proprioceptive inputs into sensible inputs to updating—in our case, sensible conditional chance estimates. This will undoubtedly involve domain-specific skill; skill that goes far beyond the skills one might be said to have simply in virtue of being epistemically rational. As such, it will not be something that constraints on *rational updating* will speak to.

Importantly, though, J-Kon makes sensible updating rather less mysterious than J-Con does. It leaves less to domain-specific skill than J-Con. According to Jeffrey, settling on a sensible input to J-Con (a sensible input partition and distribution) is a matter of skill; outside the purview of epistemic rationality. But according to J-Kon, this is only half right. Settling on sensible conditional chance estimates may be a matter of domain-specific skill. From there, however, J-Kon determines whether you should update by J-Con, and if so, what the

input partition and distribution ought to be. So we have a story about where the inputs to J-Con come from (at least in those situations where J-Kon agrees that J-Con applies). We have lightened how much of the explanatory burden we offload onto domain-specific skill, just as Christensen requested.

Back to our heart attack case. What *should* your new credences be after \mathcal{E} ? According to J-Kon, you ought to update as follows. Firstly, adopt conditional chance estimates of the form $est[ch(\mathcal{E} | w_i) | w_i]$. For example, you might adopt the following conditional chance estimates:

$$est[ch(\mathcal{E} | w_1) | w_1] = est[ch(\mathcal{E} | w_3) | w_3] = 0.99$$

as well as:

$$est[ch(\mathcal{E} | w_2) | w_2] = 0.1$$

and:

$$est[ch(\mathcal{E} | w_4) | w_4] = 0.01.$$

The first two estimates seem appropriate if you think that there is a high ($\approx 99\%$) chance of having an arm-tingling, credence-pushing learning experience if you are indeed having a heart attack. The third estimate seems appropriate if you think that there is a significant, but nevertheless fairly low ($\approx 10\%$) chance of having an arm-tingling, credence-pushing experience if you have been leaning on your elbow, potentially pinching your ulnar nerve, even if you are *not* having a heart attack. The final estimate seems appropriate if you think there is a rather low chance ($\approx 1\%$) of having such an experience otherwise.

Next, you should use these conditional chance estimates to update as follows:

$$new(w_i) = \frac{est[ch(\mathcal{E} | w_i) | w_i] \cdot old(w_i)}{\sum_j est[ch(\mathcal{E} | w_j) | w_j] \cdot old(w_j)}.$$

This yields:

$$\begin{aligned} new(w_1) &= 0.00884956 & new(w_2) &= 0.893001 \\ new(w_3) &= 0.00884956 & new(w_4) &= 0.0893001 \end{aligned}$$

So your posteriors for H and L are:

$$new(H) = 0.0176991 \text{ and } new(L) = 0.90185.$$

So despite the fact your learning experience \mathcal{E} “directly affects” your credence for H , pushing it up to 0.99, J-Kon recommends adopting a significantly lower (≈ 0.02) credence for H . Why? Because your conditional chance estimates reflect the opinion that there is a significant ($\approx 10\%$) chance of having \mathcal{E} even if H is false—in particular, if L is true. In light of this and your prior near-certainty in $\neg H$ (before \mathcal{E} , you were 99.9% confident that you are *not* going to have a heart attack), J-Kon recommends hedging your epistemic bets that \mathcal{E} is on the money by only bumping up your credence in H to roughly 0.02.

More importantly, your posterior for H conditional on L is:

$$new(H | L) = 0.00981267.$$

So J-Kon naturally predicts that \mathcal{E} turns L into an undercutting defeater for H . Prior to \mathcal{E} , L is irrelevant to H . And irrelevant propositions do not undercut anything. But *after* \mathcal{E} , you *do* see L as an undercutting defeater for

H . After \mathcal{E} you think: Learning L would be good reason to cut your posterior confidence in H in half, dropping it from roughly 0.02 to 0.01.

J-Kon does not do this by hand, in some ad hoc fashion. Take yourself back to the moment when you feel your arm tingling and your credence in H shoots up to 0.99. Ask yourself about the chance of having that experience if you are in w_2 (no heart attack, but leaning on elbow). What's your best estimate of that chance? Now ask yourself about the chance of having that experience if you are in w_4 (no heart attack, no leaning). What's your best estimate? Is it different from your earlier estimate? If so—and I am betting it is—then J-Kon *demand*s that you accommodate your experience by applying J-Con to a sufficiently fine partition; one that slots w_2 and w_4 into different cells.¹¹ (See our discussion of (♥) in §2.2.) But then *new* cannot come from *old* by applying J-Con naïvely to any Jeffrey shift over $\{H, \neg H\}$. And it is precisely this fact—that we were applying J-Con to some Jeffrey shift over $\{H, \neg H\}$ —that prevented \mathcal{E} from turning L into an undercutting defeater for H .

The moral: J-Kon handles the introduction of new undercutting defeaters in a natural and principled fashion, just as Weisberg requested. J-Kon also specifies the sense in which the doxastic effects of experience are mediated by one's background beliefs, just as Christensen required. So J-Kon is pretty holism-friendly!¹² Even more importantly, there is a purely alethic (accuracy-centred) justification for updating via J-Kon.

6 Objections

Before wrapping up, let's address a few pressing concerns.

Objection. If you start your epistemic life with an initial credence function and an extremely rich set of conditional chance estimates, as J-Kon requires, then you never really *learn to learn*. How you update over the course of your epistemic life is fixed by (i) those estimates, which do not change with domain-specific training, and (ii) one's experiences, which do not obviously change with training either. So it seems as though we have lost Jeffrey's insight that learning to translate perceptive and proprioceptive inputs into the sort of information that feeds into your updating policy is a matter of acquiring a domain-specific skill.

Reply. J-Kon does *not* require you to start your epistemic life with an extremely rich set of conditional chance estimates. It just requires you to expand that set in a coherent manner after each learning experience (or rather to update as if you were doing so). Becoming skilled at settling on sensible inputs is a matter of becoming skilled at estimating the conditional chance of having various types of experience. And that *is* something you can plausibly improve on with domain-specific training.

Objection. If J-Kon is equivalent to gilded superconditioning, then it in effect requires having prior credences over an extremely rich algebra \mathcal{F}^+ . \mathcal{F}^+ must include demonstrative propositions describing all possible learning experiences.

¹¹ Any shift from *old* to *new* can be accommodated by applying J-Con to a sufficiently fine partition. The question is simply: How fine? And is there a principled story about why to choose *that* input distribution over *that* input partition?

¹² Gallow's (2013) HCondi is also holist-friendly. An extended comparison of J-Kon and HCondi is beyond the scope of this paper.

But that is massively implausible! You are only in a position to have opinions about those demonstrative propositions (I had *that* learning experience) *after* you have the experience. Secondly, *even if* you somehow have such prior credences, then J-Kon just recommends conditioning on the proposition describing your learning experience. But this is just Bayesian orthodoxy!

Reply. J-Kon/gilded superconditioning is not conditioning. It does *not* require you to have prior credences over propositions describing all possible learning experiences. It merely requires that in response to any learning experience, you “settle on” conditional chance estimates that cohere with the conditional chance estimates that you used earlier. And it only requires you to “settle on” and “use” conditional chance estimates in an extremely thin sense. J-Kon updaters only need to be causally sensitive to learning experiences in such a way that they transition from old to new credences *as if* they are explicitly estimating conditional chances. But those conditional chance estimates, which are in this thin sense available for the purposes of updating, *need not be available for other purposes*. You need not be able to announce those estimates, for example. You need not be able to use them to guide action (*e.g.*, to decide whether to accept or reject a bet). So while these estimates reflect your opinions *in some weak sense*, they need not play the theoretical roles characteristic of *credences* or *expectations* in the Bayesian tradition. This more subtle story is quite clearly not standard Bayesian orthodoxy.

Objection. J-Kon recommends updating as if you were (i) generating new conditional chance estimates $est[ch(\mathcal{E} | w) | w]$ and then (ii) conditionalizing on \mathcal{E} . In this second step, $est[ch(\mathcal{E} | w) | w]$ functions as your old credence for \mathcal{E} conditional on w (a conditional credence which you lack). Why take this detour through conditional chance estimates? Why not simply update by expanding your old credal state to include new credences for \mathcal{E} conditional on w , and then updating as if you were conditionalizing on \mathcal{E} ?

Reply. J-Kon is equivalent to gilded superconditioning. Gilded superconditioning requires that your new credences for \mathcal{E} conditional on w take a particular form, *viz.*, the form of conditional chance estimates. This makes J-Kon more restrictive than austere superconditioning (at least if there are interesting constraints on the space of possible chance functions). And gilded superconditioning, not austere conditioning, is what our chance dominance argument yields. So this restriction is forced on us by accuracy considerations.

Objection. Garber (1980) objected to Field’s (1978) reformulation of J-Con—Field conditioning—on (roughly) the following grounds. Suppose that you have phenomenologically identical uncertain learning experiences (glimpsing a coin in poor lighting conditions, hearing a voice that sounds familiar but you cannot quite place, etc.) back to back to back—9 times instead of once, let’s say. Then you get no more information from the last 8 than you do from the first. Nevertheless, Field conditioning predicts that your credence in some cell of the input partition will keep going up and up and up, in the end rendering you “virtually certain” that the given cell is the true one (that the coin came up heads, that the voice belongs to your friend Rashad, etc.). But J-Kon is just a species of Field conditioning! It is what you get when you set the e^{α_i} ’s in (5’) of (Field, 1978, p. 367) as follows:

$$e^{\alpha_i} = est[ch(\mathcal{E} | w_i \& \mathcal{P}) | w_i \& \mathcal{P}]$$

But then surely J-Kon falls to the same objection.

Reply. It does not. Garber and Field both assume that the α_i 's are "given by experience." Since you have phenomenologically identical uncertain learning experiences in Garber's example, he uses the same α_i 's as inputs to Field conditioning after each experience. J-Kon does not use the same inputs after each experience. And rightly so. Garber's intuition that you get no more information from the last 8 experiences than you do from the first is undergirded by the thought that *if you had a certain type of experience the first time around, and nothing relevant has changed, then you are virtually certain—chance of 1—to have the same type of experience the next 8 times.* This information is reflected in the conditional chance estimates that serve as inputs to J-Kon, and prevent the sort of compounding that plagues naïve applications of Field conditioning. (In fact, we already saw J-Kon handle this sort of case in §2.2, example 3.)

7 Summary and Epilogue

Confirmational holism causes fits for J-Con. According to holism, the effects of experience ought to be mediated by one's background beliefs. J-Con fails to adequately explain how the latter influences the former. Holism also maintains that learning experiences ought nearly always to introduce new undercutting defeaters. But J-Con bungles the introduction of new undercutting defeaters.

Our question was this: Can we provide a more holism-friendly alternative to J-Con? And if so, can we provide a purely epistemic rationale for employing this alternative updating rule?

To that end, we detailed and defended J-Kon. J-Kon says, roughly, that when you have a learning experience, \mathcal{E} , you should *expand* by settling on new conditional chance estimates, $est[ch(\mathcal{E} | w) | w]$; estimates of the chance of having that very experience conditional on being in this, or that, or the other world. Then you should to treat these new conditional chance estimates, $est[ch(\mathcal{E} | w) | w]$, as your old conditional credences, $old(\mathcal{E} | w)$ (conditional credences which you in fact lack, since you have no pre- \mathcal{E} opinions about \mathcal{E}). Input your old unconditional credences, $old(w)$, and your new conditional chance estimates, $est[ch(\mathcal{E} | w) | w]$, into Bayes' theorem. Then conditionalize on \mathcal{E} . This specifies new credences for each atom of your algebra. Together with the probability axioms, this fixes your entire new credal state.

After putting J-Kon to work in a range of applications, we showed that J-Kon was in fact equivalent to what we called *gilded superconditioning*. We also provided a second characterisation of J-Kon similar to Diaconis and Zabell's characterisation of (austere) superconditioning. Then we used this equivalence to provide a purely alethic (accuracy-centred) rationale for updating via J-Kon (rather than an evidential or pragmatic rationale). Finally, we showed that J-Kon is indeed more holism-friendly than J-Con. It pins down a precise sense in which the effects of experience are mediated by one's background beliefs. It also elegantly handles the introduction of new undercutting defeaters.

At the end of his life, Jeffrey rejected the claim that epistemic rationality requires you to update in accordance with any particular updating policy. As he himself put it, "it no longer goes without saying that you will change your mind by conditioning or generalized conditioning (probability kinematics)," *i.e.*,

J-Con, “any more than it goes without saying that your changes of mind will be quite spontaneous or unconsidered... these are questions about which you may make up your mind about changing your mind in specific cases.” (Jeffrey, 1992, p. 6). This was part of Jeffrey’s “radical probabilism.”

Our project here shows us that accuracy considerations seem to push us toward something of a middle ground. Gone are the old days of Bayesian orthodoxy. In the old days, there was conditioning and Jeffrey conditioning (and perhaps minimisation of Kullback-Leibler divergence). And according to each of these updating policies, one’s old credences and learning experience *uniquely determines* one’s new credences. Rational agents have no wiggle room in updating.

In Jeffrey’s radical paradise, quite the opposite is true. In this paradise, rational agents have nothing *but* wiggle room. They may choose which properties of their old credal state to preserve, choose which new properties seem appropriate in light of their learning experience, and adopt any new coherent credal state that has the whole lot of them. Sometimes these properties, together with the probability axioms, will be sufficient to single out a unique new credal state. Sometimes not. Of course, that does not mean that just anything goes. They may make these choices well or poorly. And their skill at making these choices will often be the product of hard-earned domain-specific skill. But as we emphasised earlier, such skill will far outstrip the skill one might be said to have in virtue of being epistemically rational. As such, it will simply not be reflected in constraints of rationality.

The accuracy considerations that motivate J-Kon seem to land us somewhere in the middle. Unlike the old days of Bayesian orthodoxy, J-Kon permits significant wiggle room in updating. Rational agents can accommodate new learning experiences using *any* set of coherent conditional chance estimates that are consistent with the Principal Principle and their past updates. And there are *many* such sets. But unlike Jeffrey’s radical paradise, the wiggle room does not stretch all the way to the horizon. If there are interesting constraints on the space of possible chance functions, these will also constrain how you can update. So while you have quite a bit of wiggle room in updating, according to J-Kon, you do not have all the room in the world.

J-Kon also clears up a mystifying bit of Jeffrey’s radical probabilism. Learning to change one’s mind, on Jeffrey’s picture, involves acquiring skill at deciding which properties of one’s old credal state to preserve, and which new properties seem appropriate in light of their learning experience. But without some insight into what makes properties of credal states good or bad quite generally—without a theory of value for properties of credal states (not for credal states *themselves*, but for *properties* of credal states)—it is slightly perplexing what it might mean to make such choices well or poorly.

J-Kon swaps out this perplexing skill for a more understandable one. Learning to change one’s mind does not involve acquiring skill at choosing properties of credal states. Instead, it involves acquiring skill at estimating conditional chances. But we already have a theory of value for estimates! According to veritism, estimates are good exactly to the extent that they are accurate; exactly to the extent that they are close to the true value of the estimated quantity. The upshot: it is much easier to wrap one’s head around what learning to change one’s mind might actually amount to.

To conclude, it is worth pointing to a number of exciting new questions that

J-Kon raises. Firstly, J-Kon makes room for a healthy subjectivism/objectivism debate about rational learning. Subjectivists might be swayed only by accuracy-dominance considerations, rather than chance-dominance considerations. Such considerations seem to support superconditioning, rather than gilded superconditioning, or J-Kon. And plain old superconditioning *really does* permit epistemic lives comprising nearly any sequence of coherent credal states. On the other hand, objectivists might be swayed by considerations of worst-case epistemic loss avoidance (*cf.* Pettigrew (2014)). Such considerations seem to support an even more restrictive form of superconditioning than gilded superconditioning; one that requires *old*⁺ to maximise entropy, perhaps. As a result, objectivists will tighten up what slack J-Kon permits.

Secondly, J-Kon opens up new questions about *diachronic deference principles*. Synchronic deference principles answer the following question: When I learn that another agent's credence for X is x (maybe she is an expert, maybe an epistemic peer), how should I adjust my credence in X ? Diachronic deference principles, in contrast, answer this question: When I learn that another agent accommodated learning experiences $\mathcal{E}_1, \dots, \mathcal{E}_n$ by adopting a particular sequence of credal states, how should that influence how I update in the future? Perhaps the tools used to provide an alethic justification for J-Kon can help answer this question.

Finally, J-Kon raises questions about *diachronic aggregation principles*. Synchronic aggregation principles answer the following question: Given the credal states of some group members, how can we arrive at a single credal state that captures not the members' individual opinions, but the group's opinions as a whole? Diachronic aggregation principles answer a slightly different question: Given how different group members accommodated past learning experiences, how can we arrive at a single updating rule that captures not the members' individual inductive dispositions, but the group's inductive dispositions as a whole? For a group of J-Kon updaters, you might aggregate the conditional chance estimates that the members used to update in the past, and then accommodate group learning experiences by expanding that aggregated set going forward. Or perhaps the tools used to provide an alethic justification for J-Kon will offer up a different solution.

References

- A.V. Arkhangel'skii and V.V. Fedorchuk. *General topology I*. Springer, 1990.
- R.A. Briggs and Richard Pettigrew. An accuracy-dominance argument for conditionalization. *Nous*, 2018.
- David Christensen. Confirmational holism and bayesian epistemology. *Philosophy of Science*, 59(4):540–557, 1992.
- Bruno de Finetti. Sull'impostazione assiomatica del calcolo delle probabilit. *Annali Triestini dell'Universita di Trieste XIX*, 2:29–81, 1949.
- Bruno de Finetti. *Theory of Probability. A Critical Introductory Treatment*. John Wiley & Sons, 1974.

- Persi Diaconis and Sandy Zabell. Updating subjective probability. *J. Am. Stat. Assn.*, 77:822–830, 1982.
- Hartry Field. A note on jeffrey conditionalization. *Philosophy of Science*, 45(3): 361–367, 1978.
- Dmitri Gallow. How to learn from theory-dependent evidence; or commutativity and holism: A solution for conditionalizers. *British Journal for the Philosophy of Science*, 65(3):493–519, 2013.
- Daniel Garber. Field and jeffrey conditionalization. *Philosophy of Science*, 47 (1):142–145, 1980.
- Hilary Greaves and David Wallace. Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459):607–632, 2006.
- P.J. Huber. *Robust Statistics*. Wiley, New York, 1981.
- Richard Jeffrey. Probabilism and induction. *Topoi*, 5(1):51–58, 1986.
- Richard Jeffrey. *The Art of Judgment*. Cambridge University Press, 1992.
- James M Joyce. Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In Franz Huber and Christoph Schmidt-Petri, editors, *Degrees of Belief*, volume 342. Springer, Dordrecht, 2009.
- Jason Konek. Comparative probabilities. In Jonathan Weisberg and Richard Pettigrew, editors, *The Open Handbook of Formal Epistemology*. PhilPapers, 2019.
- Richard Pettigrew. Accuracy, risk, and the principle of indifference. *Philosophy and Phenomenological Research*, pages 1–24, 2014.
- Richard Pettigrew. *Accuracy and the Laws of Credence*. Oxford University Press, Oxford, 2016.
- Joel Predd, Robert Seiringer, Elliott Lieb, Daniel Osherson, H. Vincent Poor, , and Sanjeev R. Kulkarni. Probabilistic coherence and proper scoring rules. *IEEE Transaction on Information Theory*, 55(10):4786–4792, 2009.
- Peter Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, New York, 1991.
- Jonathan Weisberg. Updating, undermining, and independence. *British Journal for the Philosophy of Science*, 66(1):121–159, 2015.
- L. Weiskrantz. *Blindsight, a Case Study and Implication*. Clarendon Press, Oxford, 1986.
- P. M. Williams. Notes on conditional previsions. *International Journal of Approximate Reasoning*, 44(3):366–383, 2007.

A Appendix

A.1 The Framework

Let Ω be a finite set of worlds.

Let \mathcal{F}^+ be a σ -algebra on Ω . Let \mathcal{F} be a sub σ -algebra of \mathcal{F}^+ . Let \mathcal{A} be the atoms of \mathcal{F} and \mathcal{W} be the atoms of \mathcal{F}^+ . (We will abuse notation by using ‘ w ’ to refer both to elements of \mathcal{A}/\mathcal{W} , as well as indicator functions with $w(X) = 1$ iff $w \subseteq X$.)

Intuitively, \mathcal{F} contains all of the propositions that you have opinions about. \mathcal{F}^+ contains additional propositions that you do not have opinions about: (i) propositions $\mathcal{E}_1, \dots, \mathcal{E}_n$ describing your learning experiences at times t_1, \dots, t_n , and (ii) and the proposition ϕ_{ch} that ch is the true chance function, for every possible chance function ch .

For convenience, let $\mathcal{L}_0 = \mathcal{W}$ and $\mathcal{L}_i = \mathcal{E}_1 \& \dots \& \mathcal{E}_i$ for $i > 0$.

Let \mathbb{S} be the set of all sequences $\mathbf{b} = \langle b_0, \dots, b_n \rangle$ of “small space” credence functions $b_i : \mathcal{F} \rightarrow [0, 1]$. We will assume that every $\mathbf{b} \in \mathbb{S}$ is **regular** in the sense that $b_i(w_j) > 0$ for all $i \leq n$ and $w_j \in \mathcal{A} = \{w_1, \dots, w_s\}$.

Let \mathbb{B} be the set of all sequences $\mathbf{p} = \langle p_0, \dots, p_n \rangle$ of “big space” credence functions $p_i : \mathcal{F}^+ \rightarrow [0, 1]$. Choose any $\mathbf{p} \in \mathbb{B}$.

Definition 1. \mathbf{p} is **coherent** iff p_i satisfies the laws of probability for all $i \leq n$.

Definition 2. \mathbf{p} is **quasi-regular** iff $p_0(\mathcal{L}_n) > 0$.

Definition 3. \mathbf{p} is a **condi sequence** iff

$$p_i(X)p_0(\mathcal{L}_i) = p_0(X \& \mathcal{L}_i)$$

for all $X \in \mathcal{F}^+$ and $i \leq n$.

Let \mathcal{C} be the set of possible conditional chance functions $ch : \mathcal{F}^+ \times \mathcal{F}^+ \rightarrow \mathbb{R}$ with $ch(\mathcal{L}_n | \mathcal{W}) > 0$. We assume that every $ch \in \mathcal{C}$ is coherent (*i.e.*, satisfies the Rényi-Popper axioms) and *non-self-undermining*, in the sense that $ch(\phi_{ch'} | \mathcal{W})$ is 1 if $ch = ch'$ and 0 otherwise. We also assume that \mathcal{C} is finite.

For any $ch \in \mathcal{C}$, let

$$\mathbf{ch} = \langle ch_0(\cdot) = ch(\cdot | \mathcal{L}_0), \dots, ch_n(\cdot) = ch(\cdot | \mathcal{L}_n) \rangle$$

Let \mathfrak{C} be the set of all \mathbf{ch} with $ch \in \mathcal{C}$. Then every $\mathbf{ch} \in \mathfrak{C}$ is a coherent, quasi-regular, condi sequence.

Definition 4. \mathbf{p} is a **PP sequence** iff

$$p_i(Y \& \phi_{ch})ch_i(X | Y) = p_i(X \& Y \& \phi_{ch})$$

for all $i \leq n$, $ch \in \mathcal{C}$ and $X, Y \in \mathcal{F}^+$.

Definition 5. Let $\mathbb{Q} \subseteq \mathbb{B}$ be the set of coherent, quasi-regular, condi sequences.

Definition 6. Let $\mathbb{P} \subseteq \mathbb{Q} \subseteq \mathbb{B}$ be the set of coherent, quasi-regular, condi, PP sequences.

Now choose any “small space” sequence $\mathbf{b} = \langle b_0, \dots, b_n \rangle \in \mathbb{S}$.

Definition 7. \mathbf{b} is **coherent** iff b_i satisfies the laws of probability for all $i \leq n$.

Definition 8. $\mathbf{b} = \langle b_0, \dots, b_n \rangle$ is a **superconditioning sequence** iff there is some coherent condi sequence $\mathbf{p} = \langle p_0, \dots, p_n \rangle$ that extends \mathbf{b} to \mathcal{F}^+ , i.e., p_i extends b_i for all $i \leq n$.

Definition 9. \mathbf{b} is a **quasi-regular superconditioning sequence** iff there is some coherent, quasi-regular, condi sequence \mathbf{p} that extends \mathbf{b} to \mathcal{F}^+ .

Definition 10. \mathbf{b} is a **gilded superconditioning sequence** iff there is some coherent, quasi-regular, condi, PP sequence \mathbf{p} that extends \mathbf{b} to \mathcal{F}^+ .

Let $est: \mathbb{V} \times \mathcal{F}^+ \rightarrow \mathbb{R}$ be a partial assignment of conditional chance estimates of the form

$$est[ch(\mathcal{E}_i | w\&\mathcal{L}_{i-1}) | w\&\mathcal{L}_{i-1}]$$

where $i \leq n$, $w \in \mathcal{A}$ and \mathbb{V} is the set of all (bounded) variables $\mathcal{V}: \mathcal{F}^+ \rightarrow \mathbb{R}$.

Definition 11. est is **coherent** (or a **conditional linear prevision**) iff the following holds whenever the involved conditional estimates (previsions) are well-defined:

1. $est(\mathcal{V} | X) \geq \inf_{w \in X} \mathcal{V}(w)$
2. $est(\lambda \mathcal{V} | X) = \lambda est(\mathcal{V} | X)$
3. $est(\mathcal{V} + \mathcal{Q} | X) = est(\mathcal{V} | X) + est(\mathcal{Q} | X)$
4. $est(Y\mathcal{V} | X) = est(\mathcal{V} | X\&Y)est(Y | X)$ if $X \cap Y \neq \emptyset$ (where $Y\mathcal{V}$ is the product of Y and \mathcal{V} , i.e., $(Y\mathcal{V})(w) = Y(w)\mathcal{V}(w)$).

Definition 12. est is **PP-consistent** iff extending est to include conditional estimates of the form

$$est[X | \phi_{ch}\&Y] = ch(X | Y) \quad (\text{for all } X, Y \in \mathcal{F}^+)$$

preserves coherence.

Definition 13. $\mathbf{b} = \langle b_0, \dots, b_n \rangle$ is a **J-Kon sequence** iff there is coherent, PP-consistent partial assignment of conditional chance estimates, est , such that for all $w \in \mathcal{A}$

$$b_0(w) = est[ch(w | \mathcal{L}_0) | \mathcal{L}_0]$$

and for all $0 < i \leq n$

$$b_i(w) = \frac{est[ch(\mathcal{E}_i | w\&\mathcal{L}_{i-1}) | w\&\mathcal{L}_{i-1}] \cdot b_{i-1}(w)}{\sum_{w' \in \mathcal{A}} est[ch(\mathcal{E}_i | w'\&\mathcal{L}_{i-1}) | w'\&\mathcal{L}_{i-1}] \cdot b_{i-1}(w')}$$

A.2 Characterizing J-Kon Sequences

Proposition 1. *For any sequence $\mathbf{b} \in \mathbb{S}$, \mathbf{b} is a J-Kon sequence iff \mathbf{b} is a gilded superconditioning sequence.*

Proof. Suppose that $\mathbf{b} = \langle b_0, \dots, b_n \rangle$ is a J-Kon sequence. So there is coherent, PP-consistent partial assignment of conditional chance estimates, est , such that for all $w \in \mathcal{A}$

$$b_0(w) = est[ch(w \mid \mathcal{L}_0) \mid \mathcal{L}_0]$$

and for all $0 < i \leq n$

$$b_i(w) = \frac{est[ch(\mathcal{E}_i \mid w \& \mathcal{L}_{i-1}) \mid w \& \mathcal{L}_{i-1}] \cdot b_{i-1}(w)}{\sum_{w' \in \mathcal{A}} est[ch(\mathcal{E}_i \mid w' \& \mathcal{L}_{i-1}) \mid w' \& \mathcal{L}_{i-1}] \cdot b_{i-1}(w')}$$

Since est is PP-consistent, we can assume WLOG that

$$est[X \mid \phi_{ch} \& Y] = ch(X \mid Y) \quad (\text{for all } X, Y \in \mathcal{F}^+)$$

Theorem 3 of Williams (2007) establishes that any conditional linear prevision defined on a domain \mathcal{D} is extendable to a conditional linear prevision defined on a superset \mathcal{D}^+ of \mathcal{D} . So in particular there is some coherent extension est^+ of est to $\mathbb{V} \times \mathcal{F}^+$.

For any $i \leq n$ and $X \in \mathcal{F}^+$, let

$$p_i(X) = est[X \mid \mathcal{L}_i]$$

Since est is a conditional linear prevision, it follows trivially that p_i is a finitely additive probability function. So \mathbf{p} is coherent. And

$$\begin{aligned} p_0(\mathcal{L}_n) &= est[\mathcal{L}_n \mid \mathcal{W}] \\ &= \sum_{ch \in \mathcal{C}} est[\mathcal{L}_n \mid \phi_{ch} \& \mathcal{W}] est[\phi_{ch} \mid \mathcal{W}] \\ &= \sum_{ch \in \mathcal{C}} ch(\mathcal{L}_n \mid \mathcal{W}) est[\phi_{ch} \mid \mathcal{W}] > 0 \end{aligned}$$

So \mathbf{p} is quasi-regular. To see that \mathbf{p} is a condi sequence, note that for any $X \in \mathcal{F}^+$ and $i \leq n$

$$p_i(X)p_0(\mathcal{L}_i) = p_0(X \& \mathcal{L}_i)$$

iff

$$est[X \mid \mathcal{L}_i \& \mathcal{W}] est[\mathcal{L}_i \mid \mathcal{W}] = est[X \& \mathcal{L}_i \mid \mathcal{W}]$$

which follows immediately from axiom 4 for conditional linear previsions. To see that \mathbf{p} is a PP sequence, note that for any $i \leq n$, $ch \in \mathcal{C}$ and $X, Y \in \mathcal{F}^+$

$$p_i(Y \& \phi_{ch}) ch_i(X \mid Y) = p_i(X \& Y \& \phi_{ch})$$

iff

$$est[Y \& \phi_{ch} \mid \mathcal{L}_i] ch(X \mid Y \& \mathcal{L}_i) = est[X \& Y \& \phi_{ch} \mid \mathcal{L}_i]$$

Since $ch(X \mid Y \& \mathcal{L}_i) = est[X \mid \phi_{ch} \& Y \& \mathcal{L}_i]$, this follows immediately from axiom 4 for conditional linear previsions.

To show that \mathfrak{b} is a gilded superconditioning sequence, it only remains to show that \mathfrak{p} extends \mathfrak{b} to \mathcal{F}^+ . Choose $w \in \mathcal{A}$. First note that

$$\begin{aligned}
p_0(w) &= est[w \mid \mathcal{L}_0] \\
&= \sum_{ch \in \mathcal{C}} est[w \mid \phi_{ch} \& \mathcal{L}_0] est[\phi_{ch} \mid \mathcal{L}_0] \\
&= \sum_{ch \in \mathcal{C}} ch(w \mid \mathcal{L}_0) est[\phi_{ch} \mid \mathcal{L}_0] \\
&= est[ch(w \mid \mathcal{L}_0) \mid \mathcal{L}_0] \\
&= b_0(w)
\end{aligned}$$

Now suppose that p_{i-1} extends b_{i-1} to \mathcal{F}^+ . We will show that p_i must extend b_i to \mathcal{F}^+ as well.

$$\begin{aligned}
b_i(w) &= \frac{est[ch(\mathcal{E}_i \mid w \& \mathcal{L}_{i-1}) \mid w \& \mathcal{L}_{i-1}] \cdot b_{i-1}(w)}{\sum_{w' \in \mathcal{A}} est[ch(\mathcal{E}_i \mid w' \& \mathcal{L}_{i-1}) \mid w' \& \mathcal{L}_{i-1}] \cdot b_{i-1}(w')} \\
&= \frac{\left[\sum_{ch \in \mathcal{C}} ch(\mathcal{E}_i \mid w \& \mathcal{L}_{i-1}) est[\phi_{ch} \mid w \& \mathcal{L}_{i-1}] \right] \cdot p_{i-1}(w)}{\sum_{w' \in \mathcal{A}} \left[\sum_{ch \in \mathcal{C}} ch(\mathcal{E}_i \mid w' \& \mathcal{L}_{i-1}) est[\phi_{ch} \mid w' \& \mathcal{L}_{i-1}] \right] \cdot p_{i-1}(w')} \\
&= \frac{\left[\sum_{ch \in \mathcal{C}} est[\mathcal{E}_i \mid \phi_{ch} \& w \& \mathcal{L}_{i-1}] est[\phi_{ch} \mid w \& \mathcal{L}_{i-1}] \right] \cdot est[w \mid \mathcal{L}_i]}{\sum_{w' \in \mathcal{A}} \left[\sum_{ch \in \mathcal{C}} est[\mathcal{E}_i \mid \phi_{ch} \& w' \& \mathcal{L}_{i-1}] est[\phi_{ch} \mid w' \& \mathcal{L}_{i-1}] \right] \cdot est[w' \mid \mathcal{L}_i]} \\
&= \frac{est[\mathcal{E}_i \& w \mid \mathcal{L}_{i-1}]}{est[\mathcal{E}_i \mid \mathcal{L}_{i-1}]} = est[w \mid \mathcal{E}_i \& \mathcal{L}_{i-1}] = est[w \mid \mathcal{L}_i] = p_i(w)
\end{aligned}$$

Therefore \mathfrak{b} is a gilded superconditioning sequence.

Now suppose that \mathfrak{b} is a gilded superconditioning sequence. So there is some coherent, quasi-regular, condi, PP sequence \mathfrak{p} that extends \mathfrak{b} to \mathcal{F}^+ . Let $est : \mathbb{V} \times \mathcal{F}^+ \rightarrow \mathbb{R}$ be the partial assignment of conditional estimates defined by

$$est[\mathcal{V} \mid X] = \sum_{w \in \mathcal{W}} \frac{p_0(w \& X)}{p_0(X)} \mathcal{V}(w)$$

for any $X \in \mathcal{F}^+$ such that $p_0(X) > 0$. It is straightforward to check that est is a conditional linear prevision. To see that est is PP-consistent, note that if $p_0(\phi_{ch} \& Y) > 0$, then

$$\begin{aligned}
est[X \mid \phi_{ch} \& Y] &= \sum_{w \in \mathcal{W}} \frac{p_0(w \& \phi_{ch} \& Y)}{p_0(\phi_{ch} \& Y)} \mathbb{1}_X(w) \\
&= \sum_{w \in \mathcal{W}} ch(w \mid Y) \mathbb{1}_X(w) = ch(X \mid Y)
\end{aligned}$$

where $\mathbb{1}_X : \mathcal{W} \rightarrow \{0, 1\}$ is the indicator function for X . And since every $ch \in \mathcal{C}$ satisfies the Rényi-Popper axioms, extending est to include conditional estimates of the form

$$est[X \mid \phi_{ch} \& Y] = ch(X \mid Y)$$

even when $p_0(\phi_{ch} \& Y) = 0$ obviously preserves coherence. It only remains to show that for all $w \in \mathcal{A}$

$$b_0(w) = est[ch(w \mid \mathcal{L}_0) \mid \mathcal{L}_0]$$

and for all $0 < i \leq n$

$$b_i(w) = \frac{est[ch(\mathcal{E}_i \mid w \& \mathcal{L}_{i-1}) \mid w \& \mathcal{L}_{i-1}] \cdot b_{i-1}(w)}{\sum_{w' \in \mathcal{A}} est[ch(\mathcal{E}_i \mid w' \& \mathcal{L}_{i-1}) \mid w' \& \mathcal{L}_{i-1}] \cdot b_{i-1}(w')}$$

Choose $w \in \mathcal{A}$.

$$\begin{aligned} b_0(w) &= p_0(w) \\ &= \sum_{ch \in \mathcal{C}: p_0(\phi_{ch}) > 0} \frac{p_0(w \& \phi_{ch})}{p_0(\phi_{ch})} p_0(\phi_{ch}) \\ &= \sum_{ch \in \mathcal{C}: p_0(\phi_{ch}) > 0} ch(w \mid \mathcal{L}_0) \frac{p_0(\phi_{ch} \& \mathcal{L}_0)}{p_0(\mathcal{L}_0)} \\ &= est[ch(w \mid \mathcal{L}_0) \mid \mathcal{L}_0] \end{aligned}$$

Now choose $i \leq n$.

$$\begin{aligned} b_i(w) &= p_i(w) = \frac{p_{i-1}(\mathcal{E}_i \& w)}{p_{i-1}(\mathcal{E}_i)} = \frac{\frac{p_{i-1}(\mathcal{E}_i \& w)}{p_{i-1}(w)} \cdot b_{i-1}(w)}{\sum_{w' \in \mathcal{A}} \frac{p_{i-1}(\mathcal{E}_i \& w')}{p_{i-1}(w')} \cdot b_{i-1}(w')} \\ &= \frac{\left[\sum_{ch \in \mathcal{C}: p_{i-1}(\phi_{ch} \& w) > 0} \frac{p_{i-1}(\phi_{ch} \& w)}{p_{i-1}(w)} \frac{p_{i-1}(\mathcal{E}_i \& \phi_{ch} \& w)}{p_{i-1}(\phi_{ch} \& w)} \right] \cdot b_{i-1}(w)}{\sum_{w' \in \mathcal{A}} \left[\sum_{ch \in \mathcal{C}: p_{i-1}(\phi_{ch} \& w') > 0} \frac{p_{i-1}(\phi_{ch} \& w')}{p_{i-1}(w')} \frac{p_{i-1}(\mathcal{E}_i \& \phi_{ch} \& w')}{p_{i-1}(\phi_{ch} \& w')} \right] \cdot b_{i-1}(w')} \\ &= \frac{\left[\sum_{ch \in \mathcal{C}: p_{i-1}(\phi_{ch} \& w) > 0} \frac{p_{i-1}(\phi_{ch} \& w)}{p_{i-1}(w)} ch(\mathcal{E}_i \mid w \& \mathcal{L}_{i-1}) \right] \cdot b_{i-1}(w)}{\sum_{w' \in \mathcal{A}} \left[\sum_{ch \in \mathcal{C}: p_{i-1}(\phi_{ch} \& w') > 0} \frac{p_{i-1}(\phi_{ch} \& w')}{p_{i-1}(w')} ch(\mathcal{E}_i \mid w' \& \mathcal{L}_{i-1}) \right] \cdot b_{i-1}(w')} \\ &= \frac{est[ch(\mathcal{E}_i \mid w \& \mathcal{L}_{i-1}) \mid w \& \mathcal{L}_{i-1}] \cdot b_{i-1}(w)}{\sum_{w' \in \mathcal{A}} est[ch(\mathcal{E}_i \mid w' \& \mathcal{L}_{i-1}) \mid w' \& \mathcal{L}_{i-1}] \cdot b_{i-1}(w')} \end{aligned}$$

This establishes that \mathbf{b} is J-Kon sequence. □

Choose $\mathbf{b} = \langle b_0, \dots, b_n \rangle \in \mathbb{S}$. Let

$$\Psi_0 = \left\{ E \mid E = \sum_{ch \in \mathcal{C}} \lambda_{ch} E_{ch_{\mathcal{L}_0}} \text{ for some } \lambda_{ch} \in [0, 1] \text{ s.t. } \sum_{ch \in \mathcal{C}} \lambda_{ch} = 1 \right\}$$

where

$$E_{ch_{\mathcal{L}_0}}(\mathcal{V}) = \sum_{w \in \mathcal{W}} ch(w \mid \mathcal{L}_0) \mathcal{V}(w)$$

So Ψ_0 is the set of convex combinations of coherent linear previsions generated by chance functions $ch_{\mathcal{L}_0}$. Let

$$\underline{P}_0(\mathcal{V}) = \inf \{ E(\mathcal{V}) \mid E \in \Psi_0 \}$$

For any $i > 0$, let

$$\begin{aligned} \Psi_i &= \{ E^{\mathcal{L}_i} \mid E \in \Psi_{i-1} \text{ and } E(w) = b_{i-1}(w) \text{ for all } w \in \mathcal{A} \} \\ &= \{ E^{\mathcal{L}_i} \mid E \in \Psi_0 \text{ and } E^{\mathcal{L}_j}(w) = b_j(w) \text{ for all } j < i \text{ and } w \in \mathcal{A} \} \end{aligned}$$

where

$$E^{\mathcal{L}_i}(\mathcal{V}) = \frac{E(\mathcal{V} \cdot \mathcal{L}_i)}{E(\mathcal{L}_i)}$$

and let

$$\underline{P}_i(\mathcal{V}) = \inf \{ E(\mathcal{V}) \mid E \in \Psi_i \}$$

Finally let $\mathbb{A} = \{ \mathcal{V} \in \mathbb{V} \mid \mathcal{V}(w) = \mathcal{V}(w') \text{ for any } \omega \in \mathcal{A} \text{ and } w, w' \in \omega \}$.

Proposition 2. *For any sequence $\mathbf{b} \in \mathbb{S}$, \mathbf{b} is a gilded superconditioning sequence iff \mathbf{b} is coherent and $E_{b_i}(\mathcal{V}) \geq \underline{P}_i(\mathcal{V})$ for all $i \leq n$ and $\mathcal{V} \in \mathbb{A}$.*

Proof. Suppose that \mathbf{b} is a gilded superconditioning sequence. So there is some coherent, quasi-regular, condi, PP sequence \mathbf{p} that extends \mathbf{b} to \mathcal{F}^+ . Trivially \mathbf{b} is coherent. Let $E : \mathbb{V} \rightarrow \mathbb{R}$ be defined by

$$E(\mathcal{V}) = \sum_{w \in \mathcal{W}} p_0(w) \mathcal{V}(w)$$

for any $\mathcal{V} \in \mathbb{V}$. Note that

$$\begin{aligned} E(\mathcal{V}) &= \sum_{w \in \mathcal{W}} p_0(w) \mathcal{V}(w) \\ &= \sum_{ch \in \mathcal{C}} \sum_{w \in \mathcal{W}} p_0(w \& \phi_{ch}) \mathcal{V}(w) \\ &= \sum_{ch \in \mathcal{C}} \sum_{w \in \mathcal{W}} p_0(\phi_{ch}) ch(w \mid \mathcal{L}_0) \mathcal{V}(w) \\ &= \sum_{ch \in \mathcal{C}} p_0(\phi_{ch}) E_{ch_{\mathcal{L}_0}}(w) \end{aligned}$$

So $E \in \Psi_0$. Now choose $\mathcal{Q} \in \mathbb{A}$. Since \mathfrak{p} extends \mathfrak{b} to \mathcal{F}^+ , we have

$$\begin{aligned} E_{b_0}(\mathcal{Q}) &= \sum_{\omega \in \mathcal{A}} \sum_{w \in \omega} b_0(\omega) \mathcal{Q}(w) \\ &= \sum_{\omega \in \mathcal{A}} \sum_{w \in \omega} p_0(\omega) \mathcal{Q}(w) \\ &= \sum_{w \in \mathcal{W}} p_0(w) \mathcal{Q}(w) \\ &= E(\mathcal{Q}) \end{aligned}$$

So $E_{b_0}(\mathcal{Q}) = E(\mathcal{Q}) \geq \underline{P}_0(\mathcal{Q})$ for all $\mathcal{Q} \in \mathbb{A}$.

Now choose $0 < i, j \leq i$ and $\omega \in \mathcal{A}$. Then

$$E^{\mathcal{L}_j}(\omega) = \frac{E(\omega \& \mathcal{L}_j)}{E(\mathcal{L}_j)} = \frac{p_0(\omega \& \mathcal{L}_j)}{p_0(\mathcal{L}_j)} = p_j(\omega) = b_j(\omega)$$

So $E^{\mathcal{L}_i} \in \Psi_i$. Choose $\mathcal{Q} \in \mathbb{A}$. Again, since \mathfrak{p} extends \mathfrak{b} to \mathcal{F}^+ , we have

$$\begin{aligned} E_{b_i}(\mathcal{Q}) &= \sum_{\omega \in \mathcal{A}} \sum_{w \in \omega} b_i(\omega) \mathcal{Q}(w) \\ &= \sum_{\omega \in \mathcal{A}} \sum_{w \in \omega} p_i(\omega) \mathcal{Q}(w) \\ &= \sum_{w \in \mathcal{W}} p_i(w) \mathcal{Q}(w) \\ &= E^{\mathcal{L}_i}(\mathcal{Q}) \end{aligned}$$

So $E_{b_i}(\mathcal{Q}) = E^{\mathcal{L}_i}(\mathcal{Q}) \geq \underline{P}_i(\mathcal{Q})$ for all $\mathcal{Q} \in \mathbb{A}$.

Now suppose that \mathfrak{b} is coherent and $E_{b_i}(\mathcal{Q}) \geq \underline{P}_i(\mathcal{Q})$ for all $i \leq n$ and $\mathcal{Q} \in \mathbb{A}$. Then in particular $E_{b_n}(\mathcal{Q}) \geq \underline{P}_n(\mathcal{Q})$ for all $\mathcal{Q} \in \mathbb{A}$. Since \mathbb{A} is a linear subspace of \mathbb{V} , \underline{P}_n is a coherent lower prevision on \mathbb{V} , and E_{b_n} is a coherent linear prevision on \mathbb{A} , theorem 3.4.5 of Walley (1991) guarantees that there is a coherent linear prevision E^* that extends E_{b_n} to \mathbb{V} which dominates \underline{P}_n , *i.e.*, $E^*(\mathcal{V}) \geq \underline{P}_n(\mathcal{V})$ for all $\mathcal{V} \in \mathbb{V}$.

Now observe that Ψ_n is a closed, convex set of linear previsions on \mathbb{V} —the set of all (bounded) gambles on the finite sample space \mathcal{W} . To see this, note that Ψ_0 is the convex hull of $\{E_{ch_{\mathcal{L}_0}} \mid ch \in \mathcal{C}\}$. Since \mathcal{C} is finite, Ψ_0 is closed. Now suppose that Ψ_j is closed and convex for $j < i \leq n$. Then it is straightforward to show that Ψ_i is closed and convex. Firstly, note that $\Psi_i = \{E^{\mathcal{E}_i} \mid E \in \Phi\}$ where

$$\Phi = \Psi_{i-1} \cap \{E \mid E(w) = b_{i-1}(w) \text{ for all } w \in \mathcal{A}\}$$

Since Φ is the intersection of closed convex sets, Φ is closed and convex. Moreover, note that $\Psi_i = \{E^{\mathcal{E}_i} \mid E \in \Phi\}$ and $E(\mathcal{E}_i) > 0$ for all $E \in \Phi$. So Ψ_i is the convex hull of its extreme points, and its extreme points are a subset of the extreme points of Φ conditioned on \mathcal{E}_i (Walley, 1991, §6.4.2). Since Φ is closed, the extreme points of Φ are elements of Φ . Hence the extreme points of Φ conditioned on \mathcal{E}_i are elements of Ψ_i . Ψ_i therefore contains its extreme points, and moreover is the convex hull of those points. Hence Ψ_i contains its limit points and therefore is closed.

Ψ_n in particular then is closed and convex. So by proposition 2.1 of Huber (1981)

$$\Psi_n = \{E \mid E(\mathcal{V}) \geq \underline{P}_n(\mathcal{V}) \text{ for all } \mathcal{V} \in \mathbb{V}\}$$

So $E^* \in \Psi_n$. As a result, there is some $F \in \Psi_0$ such that $E^* = F^{\mathcal{L}_n}$, and moreover $F^{\mathcal{L}_i}(\omega) = b_i(\omega)$ for all $i < n$ and $\omega \in \mathcal{A}$. Since $F \in \Psi_0$, $F = \sum_{ch \in \mathcal{C}} \lambda_{ch} E_{ch\mathcal{L}_0}$ for some $\lambda_{ch} \in [0, 1]$ with $\sum_{ch \in \mathcal{C}} \lambda_{ch} = 1$.

Let $p_i(X) = F^{\mathcal{L}_i}(X)$ for all $X \in \mathcal{F}^+$. Then $\mathbf{p} = \langle p_0, \dots, p_n \rangle$ is coherent, since $F^{\mathcal{L}_i}$ is a coherent linear prevision for all $i \leq n$. \mathbf{p} is quasi-regular, since $p_0(\mathcal{L}_n) = F(\mathcal{L}_n) = \sum_{ch \in \mathcal{C}} \lambda_{ch} E_{ch\mathcal{L}_0}(\mathcal{L}_n)$ and $E_{ch\mathcal{L}_0}(\mathcal{L}_n) = ch(\mathcal{L}_n \mid \mathcal{L}_0) > 0$ for all $ch \in \mathcal{C}$. \mathbf{p} is a condi sequence, since for any $i \leq n$ and $X \in \mathcal{F}^+$

$$p_i(X) = F^{\mathcal{L}_i}(X) = \frac{F(X \& \mathcal{L}_i)}{F(\mathcal{L}_i)} = \frac{p_0(X \& \mathcal{L}_i)}{p_0(\mathcal{L}_i)}$$

And \mathbf{p} is a PP sequence, since for any $i \leq n$, $ch \in \mathcal{C}$ and $X, Y \in \mathcal{F}^+$

$$p_i(Y \& \phi_{ch}) ch_i(X \mid Y) = p_i(X \& Y \& \phi_{ch})$$

iff

$$p_0(Y \& \phi_{ch} \& \mathcal{L}_i) ch_i(X \mid Y) = p_0(X \& Y \& \phi_{ch} \& \mathcal{L}_i)$$

iff

$$F(Y \& \phi_{ch} \& \mathcal{L}_i) ch_i(X \mid Y) = F(X \& Y \& \phi_{ch} \& \mathcal{L}_i)$$

Since $F = \sum_{ch \in \mathcal{C}} \lambda_{ch} E_{ch\mathcal{L}_0}$ and $ch(\phi_{ch} \mid \mathcal{L}_0) = 1$ for all $ch \in \mathcal{C}$, we have $F(Y \& \phi_{ch} \& \mathcal{L}_i) = ch(Y \& \mathcal{L}_i \mid \mathcal{L}_0)$ and $F(X \& Y \& \phi_{ch} \& \mathcal{L}_i) = ch(X \& Y \& \mathcal{L}_i \mid \mathcal{L}_0)$. So the above equality holds iff

$$ch(Y \& \mathcal{L}_i \mid \mathcal{L}_0) ch(X \mid Y \& \mathcal{L}_i) = ch(X \& Y \& \mathcal{L}_i \mid \mathcal{L}_0)$$

Since every $ch \in \mathcal{C}$ satisfies the Rényi-Popper axioms, this final equality holds (by the generalised ratio constraint). Hence \mathbf{p} is a coherent, quasi-regular, condi, PP sequence that extends \mathbf{b} to \mathcal{F}^+ . So \mathbf{b} is a gilded superconditioning sequence. \square

A.3 Quasi Bregman Divergences

Let $\mathcal{I} : \mathbb{S} \times \mathcal{W} \rightarrow \mathbb{R}_{\geq 0}$ be an **inaccuracy measure**

$$\mathcal{I}(b, w) = \sum_{X \in \mathcal{F}} s(b(X), w(X))$$

defined by a continuous, bounded, strictly proper **component function**

$$s : [0, 1] \times \{0, 1\} \rightarrow \mathbb{R}_{\geq 0}$$

s is **strictly proper** iff

$$x \cdot s(x, 1) + (1 - x) \cdot s(x, 0) < x \cdot s(y, 1) + (1 - x) \cdot s(y, 0)$$

for any $x, y \in [0, 1]$ with $x \neq y$. A component function s measures the inaccuracy of the credence $b(X)$ when X 's truth-value is $w(X)$. When \mathcal{I} is defined by a component function in this way, we call it *additive*.

Now extend \mathcal{I} to measure not only the accuracy of individual credence functions b , but also epistemic lives $\mathbf{b} = \langle b_0, \dots, b_n \rangle \in \mathbb{S}$ as follows:

$$\mathcal{I}(\mathbf{b}, w) = \sum_{i:w \in \mathcal{L}_i} \mathcal{I}(b_i, w)$$

\mathcal{I} reflects the view that life stages b_i ought to be evaluated as conditional credence functions. More carefully, $b_i(X)$ should be evaluated as your credence for X conditional on the learning experiences that produced b_i . As such, b_i is evaluable for accuracy only at worlds w in which the learning experiences that produced that stage take place, *viz.*, $w \in \mathcal{L}_i$.

According to \mathcal{I} , the total inaccuracy of your epistemic life \mathbf{b} at w is the sum of the degrees of inaccuracy of the life stages b_i that are evaluable for accuracy at w . When \mathcal{I} takes this form, call it *temporally additive*.

For any temporally additive \mathcal{I} defined by a continuous, bounded, strictly proper component function, we can generate a divergence between big space lives, $\mathbf{p} = \langle p_0, \dots, p_n \rangle$ and $\mathbf{q} = \langle q_0, \dots, q_n \rangle$, as follows.

Let $\mathbf{b} = \langle b_0, \dots, b_n \rangle$ be the restriction of \mathbf{p} to \mathcal{F} , *i.e.*, b_i is the restriction of p_i to \mathcal{F} for all $i \leq n$. Similarly, let \mathbf{c} be the restriction of \mathbf{q} to \mathcal{F} . Now let

$$H(\mathbf{p}, \mathbf{q}) = \sum_{w \in \mathcal{W}} p_0(w) \mathcal{I}(\mathbf{c}, w)$$

Then let

$$\mathcal{D}(\mathbf{p}, \mathbf{q}) = H(\mathbf{p}, \mathbf{q}) - H(\mathbf{p}, \mathbf{p})$$

So the divergence from \mathbf{p} to \mathbf{q} is given by the difference between the expected inaccuracy of their respective restrictions, from p_0 's perspective.

Say that \mathbf{p} and \mathbf{q} are **small space equivalent** iff $p_i(X) = q_i(X)$ for all $i \geq 0$ and all $X \in \mathcal{F}$. When \mathbf{p} and \mathbf{q} are small space equivalent, we write $\mathbf{p} \approx \mathbf{q}$.

Say that \mathcal{D} is a **Quasi Bregman divergence** iff

(I) For any \mathbf{p} and \mathbf{q} in the set $\mathbb{Q} \subseteq \mathbb{B}$ of coherent, quasi-regular, conditioning sequences

- $\mathcal{D}(\mathbf{p}, \mathbf{q}) = 0$ if $\mathbf{p} \approx \mathbf{q}$
- $\mathcal{D}(\mathbf{p}, \mathbf{q}) > 0$ if $\mathbf{p} \not\approx \mathbf{q}$

(II) $\mathcal{D}(\cdot, \mathbf{r})$ is **quasi convex** on the set \mathbb{Q} of coherent, quasi-regular, conditioning sequences in the following sense. Choose $\mathbf{r} \in \mathbb{B}$ and $\mathbf{p}, \mathbf{q}^1, \dots, \mathbf{q}^k \in \mathbb{Q}$. If (i) $p_0 = \sum_{i \leq k} \mu^i q_0^i$ for some $\mu^i \in [0, 1]$ with $\sum_{i \leq k} \mu^i = 1$ and (ii) $\mathbf{q}^a \not\approx \mathbf{q}^b$ for some $a, b \leq k$ with $\mu^a, \mu^b > 0$, then

$$\mathcal{D}(\mathbf{p}, \mathbf{r}) < \sum_{i \leq k} \mu^i \mathcal{D}(\mathbf{q}^i, \mathbf{r})$$

(III) There is some $\Phi : \mathbb{B} \rightarrow [0, \infty]$ that is bounded and continuously differentiable on \mathbb{Q} , and moreover

$$\mathcal{D}(\mathbf{p}, \mathbf{q}) = \Phi(\mathbf{p}) - \Phi(\mathbf{q}) - \nabla \Phi(\mathbf{q}) \cdot (\mathbf{p} - \mathbf{q})$$

for any $\mathbf{p}, \mathbf{q} \in \mathbb{Q}$.

Proposition 3. For any temporally additive inaccuracy measure \mathcal{I} defined by a continuous, bounded, strictly proper component function

$$\mathcal{D}(\mathbf{p}, \mathbf{q}) = H(\mathbf{p}, \mathbf{q}) - H(\mathbf{p}, \mathbf{p})$$

is a Quasi Bregman divergence.

Proof. Let \mathcal{I} be a temporally additive inaccuracy measure defined by a continuous, bounded, strictly proper component function. Let

$$\mathcal{D}(\mathbf{p}, \mathbf{q}) = H(\mathbf{p}, \mathbf{q}) - H(\mathbf{p}, \mathbf{p})$$

Proof of (I). Choose $\mathbf{p}, \mathbf{q} \in \mathbb{Q}$. Let \mathbf{b} and \mathbf{c} be the restrictions of \mathbf{p} and \mathbf{q} to \mathcal{F} , respectively.

Case 1. $\mathbf{p} \approx \mathbf{q}$. Then $\mathbf{b} = \mathbf{c}$. Hence

$$H(\mathbf{p}, \mathbf{q}) = \sum_{w \in \mathcal{W}} p_0(w) \mathcal{I}(\mathbf{c}, w) = \sum_{w \in \mathcal{W}} p_0(w) \mathcal{I}(\mathbf{b}, w) = H(\mathbf{p}, \mathbf{p})$$

So

$$\mathcal{D}(\mathbf{p}, \mathbf{q}) = H(\mathbf{p}, \mathbf{q}) - H(\mathbf{p}, \mathbf{p}) = 0$$

Case 2. $\mathbf{p} \not\approx \mathbf{q}$. Then $\mathbf{b} \neq \mathbf{c}$.

Since s is strictly proper, for any $i \leq n$ and any $X \in \mathcal{F}$ we have

$$\begin{aligned} & p_i(X) s(q_i(X), 1) + (1 - p_i(X)) s(q_i(X), 0) \\ \geq & p_i(X) s(p_i(X), 1) + (1 - p_i(X)) s(p_i(X), 0) \end{aligned}$$

with equality iff $p_i(X) = q_i(X)$. Moreover, since \mathbf{p} is a coherent, quasi-regular, condi sequence we have

$$p_i(X) = \frac{p_0(X \& \mathcal{L}_i)}{p_0(\mathcal{L}_i)}$$

So the above inequality holds iff

$$\begin{aligned} & p_0(X \& \mathcal{L}_i) s(q_i(X), 1) + p_0(\neg X \& \mathcal{L}_i) \cdot s(q_i(X), 0) \\ \geq & p_0(X \& \mathcal{L}_i) s(p_i(X), 1) + p_0(\neg X \& \mathcal{L}_i) s(p_i(X), 0) \end{aligned}$$

In addition, since $\mathbf{p} \not\approx \mathbf{q}$, $q_j(Y) \neq p_j(Y)$ for some $j \leq n$ and $Y \in \mathcal{F}$. Hence

$$\begin{aligned} & p_0(Y \& \mathcal{L}_j) s(q_j(Y), 1) + p_0(\neg Y \& \mathcal{L}_j) \cdot s(q_j(Y), 0) \\ > & p_0(Y \& \mathcal{L}_j) s(p_j(Y), 1) + p_0(\neg Y \& \mathcal{L}_j) s(p_j(Y), 0) \end{aligned}$$

Finally we have

$$\begin{aligned} H(\mathbf{p}, \mathbf{q}) &= \sum_{w \in \mathcal{W}} p_0(w) \mathcal{I}(\mathbf{c}, w) \\ &= \sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(c_i, w) \\ &= \sum_{i \leq n} \sum_{w \in \mathcal{L}_i} p_0(w) \mathcal{I}(c_i, w) \\ &= \sum_{i \leq n} \sum_{w \in \mathcal{L}_i} p_0(w) \sum_{X \in \mathcal{F}} s(c_i(X), w(X)) \end{aligned}$$

$$\begin{aligned}
&= \sum_{i \leq n} \sum_{w \in \mathcal{L}_i} p_0(w) \sum_{X \in \mathcal{F}} s(q_i(X), w(X)) \\
&= \sum_{i \leq n} \sum_{X \in \mathcal{F}} p_0(X \& \mathcal{L}_i) s(q_i(X), 1) + p_0(\neg X \& \mathcal{L}_i) s(q_i(X), 0) \\
&> \sum_{i \leq n} \sum_{X \in \mathcal{F}} p_0(X \& \mathcal{L}_i) s(p_i(X), 1) + p_0(\neg X \& \mathcal{L}_i) s(p_i(X), 0) \\
&= H(\mathbf{p}, \mathbf{p})
\end{aligned}$$

Therefore $\mathcal{D}(\mathbf{p}, \mathbf{q}) = H(\mathbf{p}, \mathbf{q}) - H(\mathbf{p}, \mathbf{p}) > 0$.

Proof of (II). Choose $\mathbf{r} \in \mathbb{B}$ and $\mathbf{p}, \mathbf{q}^1, \dots, \mathbf{q}^k \in \mathbb{Q}$. Suppose that (i) $p_0 = \sum_{i \leq k} \mu^i q_0^i$ for some $\mu^i \in [0, 1]$ with $\sum_{i \leq k} \mu^i = 1$ and (ii) $\mathbf{q}^a \not\approx \mathbf{q}^b$ for some $a, b \leq k$ with $\mu^a, \mu^b > 0$. We must show that

$$\mathcal{D}(\mathbf{p}, \mathbf{r}) < \sum_{i \leq k} \mu^i \mathcal{D}(\mathbf{q}^i, \mathbf{r})$$

Note that

$$\begin{aligned}
\mathcal{D}(\mathbf{p}, \mathbf{r}) &= H(\mathbf{p}, \mathbf{r}) - H(\mathbf{p}, \mathbf{p}) \\
&= \sum_{i \leq n} \sum_{X \in \mathcal{F}} p_0(X \& \mathcal{L}_i) s(r_i(X), 1) + p_0(\neg X \& \mathcal{L}_i) s(r_i(X), 0) \\
&\quad - \sum_{i \leq n} \sum_{X \in \mathcal{F}} p_0(X \& \mathcal{L}_i) s(p_i(X), 1) + p_0(\neg X \& \mathcal{L}_i) s(p_i(X), 0) \\
&= \sum_{i \leq n} \sum_{X \in \mathcal{F}} \left[\sum_{j \leq k} \mu^j q_0^j(X \& \mathcal{L}_i) \right] s(r_i(X), 1) + \left[\sum_{j \leq k} \mu^j q_0^j(\neg X \& \mathcal{L}_i) \right] s(r_i(X), 0) \\
&\quad - \sum_{i \leq n} \sum_{X \in \mathcal{F}} \left[\sum_{j \leq k} \mu^j q_0^j(X \& \mathcal{L}_i) \right] s(p_i(X), 1) + \left[\sum_{j \leq k} \mu^j q_0^j(\neg X \& \mathcal{L}_i) \right] s(p_i(X), 0) \\
&= \sum_{j \leq k} \mu^j \sum_{i \leq n} \sum_{X \in \mathcal{F}} q_0^j(X \& \mathcal{L}_i) s(r_i(X), 1) + q_0^j(\neg X \& \mathcal{L}_i) s(r_i(X), 0) \\
&\quad - \sum_{j \leq k} \mu^j \sum_{i \leq n} \sum_{X \in \mathcal{F}} q_0^j(X \& \mathcal{L}_i) s(p_i(X), 1) + q_0^j(\neg X \& \mathcal{L}_i) s(p_i(X), 0)
\end{aligned}$$

Similarly

$$\begin{aligned}
\mathcal{D}(\mathbf{q}^j, \mathbf{r}) &= H(\mathbf{q}^j, \mathbf{r}) - H(\mathbf{q}^j, \mathbf{q}^j) \\
&= \sum_{i \leq n} \sum_{X \in \mathcal{F}} q_0^j(X \& \mathcal{L}_i) s(r_i(X), 1) + q_0^j(\neg X \& \mathcal{L}_i) s(r_i(X), 0) \\
&\quad - \sum_{i \leq n} \sum_{X \in \mathcal{F}} q_0^j(X \& \mathcal{L}_i) s(q_i^j(X), 1) + q_0^j(\neg X \& \mathcal{L}_i) s(q_i^j(X), 0)
\end{aligned}$$

So

$$\mathcal{D}(\mathbf{p}, \mathbf{r}) < \sum_{j \leq k} \mu^j \mathcal{D}(\mathbf{q}^j, \mathbf{r})$$

iff

$$\begin{aligned} & \sum_{j \leq k} \mu^j \sum_{i \leq n} \sum_{X \in \mathcal{F}} q_0^j(X \& \mathcal{L}_i) s(p_i(X), 1) + q_0^j(\neg X \& \mathcal{L}_i) s(p_i(X), 0) \\ > & \sum_{j \leq k} \mu^j \sum_{i \leq n} \sum_{X \in \mathcal{F}} q_0^j(X \& \mathcal{L}_i) s(q_i^j(X), 1) + q_0^j(\neg X \& \mathcal{L}_i) s(q_i^j(X), 0) \end{aligned}$$

Since s is strictly proper, and \mathfrak{q}^j is a coherent, quasi-regular, condi sequence

$$\begin{aligned} & q_0^j(X \& \mathcal{L}_i) s(p_i(X), 1) + q_0^j(\neg X \& \mathcal{L}_i) s(p_i(X), 0) \\ \geq & q_0^j(X \& \mathcal{L}_i) s(q_i^j(X), 1) + q_0^j(\neg X \& \mathcal{L}_i) s(q_i^j(X), 0) \end{aligned}$$

with equality iff $p_i(X) = q_i^j(X)$. Now recall that $\mathfrak{q}^a \not\approx \mathfrak{q}^b$ for some $a, b \leq k$ with $\mu^a, \mu^b > 0$. So $q_i^a(X) \neq q_i^b(X)$ for some $i \leq n$ and $X \in \mathcal{F}$. Hence either $p_i(X) \neq q_i^a(X)$ or $p_i(X) \neq q_i^b(X)$. This ensures that

$$\begin{aligned} & \sum_{j \leq k} \mu^j \sum_{i \leq n} \sum_{X \in \mathcal{F}} q_0^j(X \& \mathcal{L}_i) s(p_i(X), 1) + q_0^j(\neg X \& \mathcal{L}_i) s(p_i(X), 0) \\ > & \sum_{j \leq k} \mu^j \sum_{i \leq n} \sum_{X \in \mathcal{F}} q_0^j(X \& \mathcal{L}_i) s(q_i^j(X), 1) + q_0^j(\neg X \& \mathcal{L}_i) s(q_i^j(X), 0) \end{aligned}$$

Therefore

$$\mathcal{D}(\mathfrak{p}, \mathfrak{r}) < \sum_{j \leq k} \mu^j \mathcal{D}(\mathfrak{q}^j, \mathfrak{r})$$

Proof of (III). Represent any $\mathfrak{p} \in \mathbb{B}$ as a vector in $[0, 1]^{t(n+1)}$

$$\langle p_0(X_1), \dots, p_0(X_t), \dots, p_n(X_1), \dots, p_n(X_t) \rangle \in [0, 1]^{t(n+1)}$$

where $\mathcal{F}^+ = \{X_1, \dots, X_t\}$. Assume WLOG that $\mathcal{W} = \{w_1, \dots, w_s\}$ and that $X_1 = w_1, \dots, X_s = w_s$.

Now let

$$\Phi(\mathfrak{p}) = -H(\mathfrak{p}, \mathfrak{p})$$

Let \mathfrak{b} be the restriction of \mathfrak{p} to \mathcal{F} . And let

$$\Psi(\mathfrak{p}) = \left\langle - \sum_{i: w_1 \in \mathcal{L}_i} \mathcal{I}(b_i, w_1), \dots, - \sum_{i: w_s \in \mathcal{L}_i} \mathcal{I}(b_i, w_s), 0, \dots, 0 \right\rangle$$

We will now show that (i) Φ is bounded and continuously differentiable on \mathbb{Q} , (ii) $\nabla \Phi = \Psi$, and (iii) for any $\mathfrak{p}, \mathfrak{q} \in \mathbb{Q}$

$$\mathcal{D}(\mathfrak{p}, \mathfrak{q}) = \Phi(\mathfrak{p}) - \Phi(\mathfrak{q}) - \Psi(\mathfrak{q}) \cdot (\mathfrak{p} - \mathfrak{q})$$

The continuity and boundedness of Φ follows trivially from the fact that \mathcal{I} is defined by a continuous, bounded component score.

Choose $\mathfrak{p} \in \mathbb{Q}$. We will show that Φ is partially differentiable with respect to x_j at \mathfrak{p} , for each $j \leq t(n+1)$, and moreover that these partial derivatives ensure that $\nabla \Phi(\mathfrak{p}) = \Psi(\mathfrak{p})$.

Case 1: $1 \leq j \leq s$. Choose $\epsilon > 0$. Let \mathbf{q} be

$$\langle p_0(w_1), \dots, p_0(w_j) + \epsilon, \dots, p_0(w_s), \dots, p_0(X_t), \dots, p_n(X_1), \dots, p_n(X_t) \rangle$$

Let \mathbf{b} and \mathbf{c} be the restrictions of \mathbf{p} and \mathbf{q} to \mathcal{F} . Then

$$\begin{aligned} & \frac{1}{\epsilon} [\Phi(\mathbf{q}) - \Phi(\mathbf{p})] \\ &= \frac{1}{\epsilon} [H(\mathbf{p}, \mathbf{p}) - H(\mathbf{q}, \mathbf{q})] \\ &= \frac{1}{\epsilon} \left[\sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(b_i, w) - \sum_{w \in \mathcal{W}} q_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(c_i, w) \right] \\ &= \frac{1}{\epsilon} \left[\sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(b_i, w) - \sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(c_i, w) - \epsilon \sum_{i: w_j \in \mathcal{L}_i} \mathcal{I}(c_i, w_j) \right] \\ &= - \sum_{i: w_j \in \mathcal{L}_i} \mathcal{I}(c_i, w_j) + \frac{1}{\epsilon} \left[\sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(b_i, w) - \sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(c_i, w) \right] \end{aligned}$$

Since \mathcal{I} is defined by a strictly proper component score, and \mathbf{p} is a coherent, quasi-regular, condi sequence, this second term is less than or equal to zero. Hence

$$\begin{aligned} \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} [\Phi(\mathbf{q}) - \Phi(\mathbf{p})] &\leq \lim_{\epsilon \rightarrow 0} - \sum_{i: w_j \in \mathcal{L}_i} \mathcal{I}(c_i, w_j) \\ &= - \sum_{i: w_j \in \mathcal{L}_i} \mathcal{I}(b_i, w_j) \end{aligned}$$

Now let \mathbf{r} be

$$\langle p_0(w_1), \dots, p_0(w_j) - \epsilon, \dots, p_0(w_s), \dots, p_0(X_t), \dots, p_n(X_1), \dots, p_n(X_t) \rangle$$

A similar argument shows that

$$\begin{aligned} \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} [\Phi(\mathbf{p}) - \Phi(\mathbf{r})] &\geq \lim_{\epsilon \rightarrow 0} - \sum_{i: w_j \in \mathcal{L}_i} \mathcal{I}(c_i, w_j) \\ &= - \sum_{i: w_j \in \mathcal{L}_i} \mathcal{I}(b_i, w_j) \end{aligned}$$

Since Φ is continuous, this shows that Φ is partially differentiable with respect to x_j at \mathbf{p} , and that

$$\frac{\partial \Phi}{\partial x_j}(\mathbf{p}) = - \sum_{i: w_j \in \mathcal{L}_i} \mathcal{I}(b_i, w_j)$$

Case 2: $s < j$. Choose $\epsilon > 0$. Let \mathbf{q} be

$$\langle p_0(X_1), \dots, p_0(X_t), \dots, p_a(X_b) + \epsilon, \dots, p_n(X_1), \dots, p_n(X_t) \rangle$$

where $p_a(X_b)$ is the j^{th} entry of \mathbf{p} . Let \mathbf{b} and \mathbf{c} be the restrictions of \mathbf{p} and \mathbf{q} to \mathcal{F} . Then

$$\begin{aligned} \frac{1}{\epsilon}[\Phi(\mathbf{q}) - \Phi(\mathbf{p})] &= \frac{1}{\epsilon}[H(\mathbf{p}, \mathbf{p}) - H(\mathbf{q}, \mathbf{q})] \\ &= \frac{1}{\epsilon} \left[\sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(b_i, w) - \sum_{w \in \mathcal{W}} q_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(c_i, w) \right] \\ &= \frac{1}{\epsilon} \left[\sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(b_i, w) - \sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(c_i, w) \right] \end{aligned}$$

Since \mathcal{I} is defined by a strictly proper component score, and \mathbf{p} is a coherent, quasi-regular, condi sequence, the term in brackets is less than or equal to zero. Hence

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon}[\Phi(\mathbf{q}) - \Phi(\mathbf{p})] \leq 0$$

Let \mathbf{r} be

$$\langle p_0(X_1), \dots, p_0(X_t), \dots, p_a(X_b) - \epsilon, \dots, p_n(X_1), \dots, p_n(X_t) \rangle$$

where again $p_a(X_b)$ is the j^{th} entry of \mathbf{p} . A similar argument shows that

$$\lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon}[\Phi(\mathbf{p}) - \Phi(\mathbf{r})] \geq 0$$

Since Φ is continuous, this shows that Φ is partially differentiable with respect to x_j at \mathbf{p} , and that

$$\frac{\partial \Phi}{\partial x_j}(\mathbf{p}) = 0$$

Together cases 1 and 2 establish that

$$\nabla \Phi(\mathbf{p}) = \Psi(\mathbf{p}) = \left\langle - \sum_{i: w_1 \in \mathcal{L}_i} \mathcal{I}(b_i, w_1), \dots, - \sum_{i: w_s \in \mathcal{L}_i} \mathcal{I}(b_i, w_s), 0, \dots, 0 \right\rangle$$

Finally we must show that for any $\mathbf{p}, \mathbf{q} \in \mathbb{Q}$

$$\mathcal{D}(\mathbf{p}, \mathbf{q}) = \Phi(\mathbf{p}) - \Phi(\mathbf{q}) - \Psi(\mathbf{q}) \cdot (\mathbf{p} - \mathbf{q})$$

Choose $\mathbf{p}, \mathbf{q} \in \mathbb{Q}$.

$$\begin{aligned} \mathcal{D}(\mathbf{p}, \mathbf{q}) &= H(\mathbf{p}, \mathbf{q}) - H(\mathbf{p}, \mathbf{p}) \\ &= \sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(c_i, w) - \sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(b_i, w) \\ &= \sum_{w \in \mathcal{W}} [p_0(w) - q_0(w)] \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(c_i, w) + \sum_{w \in \mathcal{W}} q_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(c_i, w) \\ &\quad - \sum_{w \in \mathcal{W}} p_0(w) \sum_{i: w \in \mathcal{L}_i} \mathcal{I}(b_i, w) \\ &= -\Psi(\mathbf{q}) \cdot (\mathbf{p} - \mathbf{q}) - \Phi(\mathbf{q}) + \Phi(\mathbf{p}) \end{aligned}$$

Hence \mathcal{D} is a Quasi Bregman divergence. □

A.4 Chance-Dominance Argument for J-Kon

Proposition 4. *Let \mathcal{D} be a Quasi Bregman divergence. Then for any $\mathbf{q} \in \mathbb{Q} - \mathbb{P}$ there is a point $\pi^{\mathbf{q}} \in \mathbb{P}$, called the projection of \mathbf{q} onto \mathbb{P} , such that*

$$\mathcal{D}(\pi^{\mathbf{q}}, \mathbf{q}) \leq \mathcal{D}(\mathbf{p}, \mathbf{q})$$

for any $\mathbf{p} \in \mathbb{P}$. Moreover, $\pi^{\mathbf{q}}$ is unique up to small space equivalence.

Proof. It is trivial to show that \mathbb{P} is the set of $\mathbf{p} \in \mathbb{Q}$ such that $p_0 = \sum_{ch} \mu_{ch} ch_0$ for some sequence of $\mu_{ch} \in [0, 1]$ with $\sum_{ch \in \mathcal{C}} \mu_{ch} = 1$. Let Γ be the function mapping any such sequence of μ_{ch} 's to the $\mathbf{p} \in \mathbb{P}$ with $p_0 = \sum_{ch} \mu_{ch} ch_0$. The domain of Γ is compact. Γ is continuous. And the image of Γ is \mathbb{P} . Theorem 5.2.2 of Arkhangel'skii and Fedorchuk (1990) shows that the image of a compact set under a continuous function is compact. So \mathbb{P} is compact, and hence closed and bounded.

Now choose $\mathbf{q} \in \mathbb{Q} - \mathbb{P}$. Note that $\mathcal{D}(\cdot, \mathbf{q})$ is continuous for fixed \mathbf{q} . Hence $\mathcal{D}(\cdot, \mathbf{q})$ takes a minimum $\pi^{\mathbf{q}}$ on the closed bounded set \mathbb{P} .

To see that this minimum is unique up to small space equivalence, suppose that $\mathcal{D}(\cdot, \mathbf{q})$ takes a minimum at both $\pi^{\mathbf{q}}$ and \mathbf{r} , where the restrictions of $\pi^{\mathbf{q}}$ and \mathbf{r} to \mathcal{F} are distinct. Since $\pi^{\mathbf{q}} \in \mathbb{P}$, there are $\mu_{ch} \in [0, 1]$ with $\sum_{ch \in \mathcal{C}} \mu_{ch} = 1$ such that $\pi_0^{\mathbf{q}} = \sum_{ch} \mu_{ch} ch_0$. Likewise, there are $\theta_{ch} \in [0, 1]$ with $\sum_{ch \in \mathcal{C}} \theta_{ch} = 1$ such that $r_0 = \sum_{ch} \theta_{ch} ch_0$. Let \mathbf{p} be the coherent, quasi-regular, condi sequence with $p_0 = \frac{1}{2}\pi_0^{\mathbf{q}} + \frac{1}{2}r_0$. Then $p_0 = \sum_{ch} (\frac{1}{2}\mu_{ch} + \frac{1}{2}\theta_{ch}) ch_0$. So $\mathbf{p} \in \mathbb{P}$. But since \mathcal{D} is quasi convex, we have

$$\mathcal{D}(\mathbf{p}, \mathbf{q}) < \frac{1}{2}\mathcal{D}(\pi^{\mathbf{q}}, \mathbf{q}) + \frac{1}{2}\mathcal{D}(\mathbf{r}, \mathbf{q}) = \mathcal{D}(\pi^{\mathbf{q}}, \mathbf{q})$$

which contradicts the assumption that $\mathcal{D}(\cdot, \mathbf{q})$ takes a minimum on \mathbb{P} at $\pi^{\mathbf{q}}$. \square

Proposition 5. *Let \mathcal{D} be a Quasi Bregman divergence. Then for any $\mathbf{q} \in \mathbb{Q} - \mathbb{P}$, the projection of \mathbf{q} onto \mathbb{P} , $\pi^{\mathbf{q}}$, is such that*

$$\mathcal{D}(\mathbf{p}, \pi^{\mathbf{q}}) \leq \mathcal{D}(\mathbf{p}, \mathbf{q}) - \mathcal{D}(\pi^{\mathbf{q}}, \mathbf{q})$$

for all $\mathbf{p} \in \mathbb{P}$.

Proof. Choose $\mathbf{q} \in \mathbb{Q} - \mathbb{P}$ and $\mathbf{p} \in \mathbb{P}$. Let $\pi^{\mathbf{q}}$ be the projection of \mathbf{q} onto \mathbb{P} . We will show that

$$\mathcal{D}(\mathbf{p}, \pi^{\mathbf{q}}) \leq \mathcal{D}(\mathbf{p}, \mathbf{q}) - \mathcal{D}(\pi^{\mathbf{q}}, \mathbf{q})$$

Let $\mathbf{m} \in \mathbb{P}$ be the coherent, quasi-regular, condi, PP sequence with

$$m_0 = (1 - \epsilon)\pi_0^{\mathbf{q}} + \epsilon p_0$$

for some $0 < \epsilon \leq 1$. Then since $\mathcal{D}(\cdot, \mathbf{q})$ takes a minimum on \mathbb{P} at $\pi^{\mathbf{q}}$ (unique up to small space equivalence) we have

$$0 \leq \mathcal{D}(\mathbf{m}, \mathbf{q}) - \mathcal{D}(\pi^{\mathbf{q}}, \mathbf{q})$$

This gives us

$$\begin{aligned}
0 &\leq \mathcal{D}(\mathbf{m}, \mathbf{q}) - \mathcal{D}(\pi^{\mathbf{q}}, \mathbf{q}) \\
&= [\Phi(\mathbf{m}) - \Phi(\mathbf{q}) - \nabla\Phi(\mathbf{q}) \cdot (\mathbf{m} - \mathbf{q})] - [\Phi(\pi^{\mathbf{q}}) - \Phi(\mathbf{q}) - \nabla\Phi(\mathbf{q}) \cdot (\pi^{\mathbf{q}} - \mathbf{q})] \\
&= -\nabla\Phi(\mathbf{q}) \cdot (\mathbf{m} - \pi^{\mathbf{q}}) + [\Phi(\mathbf{m}) - \Phi(\pi^{\mathbf{q}})] \\
&= -\epsilon\nabla\Phi(\mathbf{q}) \cdot (\mathbf{p} - \pi^{\mathbf{q}}) + [\Phi(\mathbf{m}) - \Phi(\pi^{\mathbf{q}})]
\end{aligned}$$

We also have

$$\begin{aligned}
\lim_{\epsilon \rightarrow 0} \frac{\Phi(\mathbf{m}) - \Phi(\pi^{\mathbf{q}})}{\epsilon} &= \lim_{\epsilon \rightarrow 0} \frac{\nabla\Phi(\pi^{\mathbf{q}}) \cdot (\mathbf{m} - \pi^{\mathbf{q}})}{\epsilon} \\
&= \lim_{\epsilon \rightarrow 0} \frac{\nabla\Phi(\pi^{\mathbf{q}}) \cdot (\epsilon\mathbf{p} - \epsilon\pi^{\mathbf{q}})}{\epsilon} \\
&= \nabla\Phi(\pi^{\mathbf{q}}) \cdot (\mathbf{p} - \pi^{\mathbf{q}})
\end{aligned}$$

Hence

$$\begin{aligned}
0 &\leq \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} [-\epsilon\nabla\Phi(\mathbf{q}) \cdot (\mathbf{p} - \pi^{\mathbf{q}}) + [\Phi(\mathbf{m}) - \Phi(\pi^{\mathbf{q}})]] \\
&= -\nabla\Phi(\mathbf{q}) \cdot (\mathbf{p} - \pi^{\mathbf{q}}) + \lim_{\epsilon \rightarrow 0} \frac{\Phi(\mathbf{m}) - \Phi(\pi^{\mathbf{q}})}{\epsilon} \\
&= -\nabla\Phi(\mathbf{q}) \cdot (\mathbf{p} - \pi^{\mathbf{q}}) + \nabla\Phi(\pi^{\mathbf{q}}) \cdot (\mathbf{p} - \pi^{\mathbf{q}}) \\
&= [\nabla\Phi(\pi^{\mathbf{q}}) - \nabla\Phi(\mathbf{q})] \cdot (\mathbf{p} - \pi^{\mathbf{q}})
\end{aligned}$$

Finally note that

$$\mathcal{D}(\mathbf{p}, \mathbf{q}) - \mathcal{D}(\pi^{\mathbf{q}}, \mathbf{q}) - \mathcal{D}(\mathbf{p}, \pi^{\mathbf{q}}) = [\nabla\Phi(\pi^{\mathbf{q}}) - \nabla\Phi(\mathbf{q})] \cdot (\mathbf{p} - \pi^{\mathbf{q}})$$

So

$$0 \leq \mathcal{D}(\mathbf{p}, \mathbf{q}) - \mathcal{D}(\pi^{\mathbf{q}}, \mathbf{q}) - \mathcal{D}(\mathbf{p}, \pi^{\mathbf{q}})$$

and therefore

$$\mathcal{D}(\mathbf{p}, \pi^{\mathbf{q}}) \leq \mathcal{D}(\mathbf{p}, \mathbf{q}) - \mathcal{D}(\pi^{\mathbf{q}}, \mathbf{q})$$

□

Before moving on, it will be useful to prove the following lemma.

Lemma 6. *For any $\mathbf{b} \in \mathbb{S}$, \mathbf{b} is quasi-regular superconditioning sequence iff \mathbf{b} is coherent and certainty-preserving, i.e., $b_i(X) = 0$ implies $b_j(X) = 0$ for all $X \in \mathcal{F}$, $i \leq n$ and $j > i$.*

Proof. Choose $\mathbf{b} \in \mathbb{S}$. Suppose that \mathbf{b} is a quasi-regular superconditioning life. So there is some coherent, quasi-regular, condi sequence \mathbf{p} that extends \mathbf{b} to \mathcal{F}^+ .

\mathbf{b} 's coherence follows trivially from \mathbf{p} 's. To see that \mathbf{b} is certainty preserving, choose $X \in \mathcal{F}$, $i \leq n$ and $j > i$. Suppose that $b_i(X) = 0$. Since p_i is an extension of b_i , and \mathbf{p} is a condi sequence with $p_0(\mathcal{L}_i) > 0$, we have

$$p_i(X) = \frac{p_0(X \& \mathcal{L}_i)}{p_0(\mathcal{L}_i)} = b_i(X) = 0$$

Hence $p_0(X \& \mathcal{L}_i) = 0$. And since $\mathcal{L}_j \subseteq \mathcal{L}_i$, $p_0(X \& \mathcal{L}_j) = 0$. As a result

$$b_j(X) = p_j(X) = \frac{p_0(X \& \mathcal{L}_j)}{p_0(\mathcal{L}_j)} = 0$$

Therefore \mathfrak{b} is certainty preserving.

Now suppose that \mathfrak{b} is coherent and certainty preserving. We will first recursively define a partial function $q : \mathcal{F}^+ \rightarrow \mathbb{R}$. Then we will use q to construct a coherent, quasi-regular, condi extension \mathfrak{p} of \mathfrak{b} .

Base Case.

- $\delta_1 = \frac{1}{2} \min \{b_0(a) \mid b_0(a) > 0, a \in \mathcal{A}\}$
- $q(a \& \mathcal{E}_1) = \begin{cases} \delta_1 b_1(a) & \text{if } b_0(a) > 0 \\ 0 & \text{otherwise} \end{cases}$
- $q(a \& \neg \mathcal{E}_1) = b_0(a) - q(a \& \mathcal{E}_1)$

Inductive Clause.

- $\delta_i = \frac{1}{2} \min \{q(a \& \mathcal{L}_{i-1}) \mid q(a \& \mathcal{L}_{i-1}) > 0, a \in \mathcal{A}\}$
- $q(a \& \mathcal{L}_i) = \begin{cases} \delta_i b_i(a) & \text{if } q(a \& \mathcal{L}_{i-1}) > 0 \\ 0 & \text{otherwise} \end{cases}$
- $q(a \& \mathcal{L}_{i-1} \& \neg \mathcal{E}_i) = q(a \& \mathcal{L}_{i-1}) - q(a \& \mathcal{L}_i)$
- $q(a \& \mathcal{L}_{i-1} \& \neg \mathcal{E}_i \& \pm \mathcal{E}_{i+1} \& \dots \& \pm \mathcal{E}_n) = \frac{q(a \& \mathcal{L}_{i-1} \& \neg \mathcal{E}_i)}{2^{n-(i-1)}}$

Finally let

$$q(a \& \phi_{ch} \& \pm \mathcal{E}_1 \& \dots \& \pm \mathcal{E}_n) = \frac{q(a \& \pm \mathcal{E}_1 \& \dots \& \pm \mathcal{E}_n)}{|\mathcal{C}|}$$

Now construct $p_0 : \mathcal{F}^+ \rightarrow \mathbb{R}$ as follows:

$$p_0(X) = \sum_{w \in X} q(w)$$

p_0 is non-negative and additive by construction. Hence to check that p_0 is

coherent, we only need to check that $p_0(\mathcal{W}) = 1$.

$$\begin{aligned}
p_0(\mathcal{W}) &= \sum_{w \in \mathcal{W}} q(w) \\
&= \sum_{a \in \mathcal{A}} \sum_{ch \in \mathcal{C}} q(a \& \phi_{ch} \& \pm \mathcal{E}_1 \& \dots \& \pm \mathcal{E}_n) \\
&= \sum_{a \in \mathcal{A}} q(a \& \pm \mathcal{E}_1 \& \dots \& \pm \mathcal{E}_n) \\
&= \sum_{a \in \mathcal{A}} q(a \& \neg \mathcal{E}_1 \& \pm \mathcal{E}_2 \& \dots \& \pm \mathcal{E}_n) + q(a \& \mathcal{E}_1 \& \neg \mathcal{E}_2 \& \pm \mathcal{E}_3 \& \dots \& \pm \mathcal{E}_n) \\
&+ q(a \& \mathcal{E}_1 \& \mathcal{E}_2 \& \neg \mathcal{E}_3 \& \pm \mathcal{E}_4 \& \dots \& \pm \mathcal{E}_n) + \dots + q(a \& \mathcal{E}_1 \& \dots \& \mathcal{E}_{n-1} \& \neg \mathcal{E}_n) \\
&+ q(a \& \mathcal{E}_1 \& \dots \& \mathcal{E}_n) \\
&= \sum_{a \in \mathcal{A}} \frac{q(a \& \neg \mathcal{E}_1)}{2^{n-1}} \cdot 2^{n-1} + \frac{q(a \& \mathcal{E}_1 \& \neg \mathcal{E}_2)}{2^{n-2}} \cdot 2^{n-2} \\
&+ \frac{q(a \& \mathcal{E}_1 \& \mathcal{E}_2 \& \neg \mathcal{E}_3)}{2^{n-3}} \cdot 2^{n-3} + \dots + q(a \& \mathcal{E}_1 \& \dots \& \mathcal{E}_{n-1} \& \neg \mathcal{E}_n) + q(a \& \mathcal{E}_1 \& \dots \& \mathcal{E}_n) \\
&= \sum_{a \in \mathcal{A}} [b_0(a) - q(a \& \mathcal{E}_1)] + [q(a \& \mathcal{E}_1) - q(a \& \mathcal{E}_1 \& \mathcal{E}_2)] + [q(a \& \mathcal{E}_1 \& \mathcal{E}_2) - q(a \& \mathcal{E}_1 \& \mathcal{E}_2 \& \mathcal{E}_3)] \\
&+ \dots + [q(a \& \mathcal{E}_1 \& \dots \& \mathcal{E}_{n-1}) - q(a \& \mathcal{E}_1 \& \dots \& \mathcal{E}_n)] + q(a \& \mathcal{E}_1 \& \dots \& \mathcal{E}_n) \\
&= \sum_{a \in \mathcal{A}} b_0(a) = 1
\end{aligned}$$

This establishes that p_0 is coherent. Now choose $i \leq n$. Then

$$p_0(\mathcal{L}_i) = \sum_{a \in \mathcal{A}} q(a \& \mathcal{L}_i) = \sum_{a: b_i(a) > 0} \delta_i b_i(a) = \delta_i > 0$$

Let $p_i(X) = p_0(X \& \mathcal{L}_i) / p_0(\mathcal{L}_i)$ for all $X \in \mathcal{F}^+$ and

$$\mathbf{p} = \langle p_0, \dots, p_n \rangle$$

Then \mathbf{p} is a coherent, quasi-regular, condi sequence. It only remains to show that \mathbf{p} extends \mathbf{b} to \mathcal{F}^+ . Choose $i \leq n$ and $X \in \mathcal{F}$.

$$p_i(X) = \frac{p_0(X \& \mathcal{L}_i)}{p_0(\mathcal{L}_i)} = \frac{\sum_{a \in X} q(a \& \mathcal{L}_i)}{\sum_{a \in \mathcal{A}} q(a \& \mathcal{L}_i)} = \frac{\sum_{a \in X: b_i(a) > 0} \delta_i b_i(a)}{\sum_{a \in \mathcal{A}: b_i(a) > 0} \delta_i b_i(a)} = b_i(X)$$

This establishes that \mathbf{b} is a quasi-regular superconditioning life. \square

Proposition 7. *Let \mathcal{I} be a temporally additive inaccuracy measure defined by a continuous, bounded, strictly proper component function. Then for any $\mathbf{b} \in \mathbb{S}$ we have the following:*

- (I) *If \mathbf{b} is not a J-Kon sequence, then \mathbf{b} is **strictly chance-dominated** by a J-Kon sequence \mathbf{c} , i.e.,*

$$\sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathbf{b}, w) > \sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathbf{c}, w)$$

for all $ch \in \mathcal{C}$.

(II) If \mathbf{b} is a J-Kon sequence, then it is not even **weakly chance-dominated**, i.e., there is no $\mathbf{c} \neq \mathbf{b}$ such that

$$\sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathbf{b}, w) \geq \sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathbf{c}, w)$$

for all $ch \in \mathcal{C}$.

Proof. Choose $\mathbf{b} \in \mathbb{S}$. Suppose that \mathbf{b} is not a J-Kon sequence. Then by proposition 1, \mathbf{b} is not a gilded superconditioning. So no extension of \mathbf{b} to \mathcal{F}^+ is a coherent, quasi-regular, condi, PP sequence.

Case 1. No extension of \mathbf{b} to \mathcal{F}^+ is coherent. de Finetti (1949) shows that any coherent b_i is coherently extendable to \mathcal{F}^+ . So $\mathbf{b} = \langle b_0, \dots, b_n \rangle$ must be incoherent. Suppose WLOG that b_i is incoherent, but b_j is coherent for all $j \neq i$. Then theorem 1 of Predd et al. (2009) implies that there is some coherent c such that

$$\mathcal{I}(b_i, w) > \mathcal{I}(c, w)$$

for all $w \in \mathcal{W}$. Let \mathbf{c} be

$$\mathbf{c} = \langle b_0, \dots, b_{i-1}, c, b_{i+1}, \dots, b_n \rangle$$

Choose $w \in \mathcal{L}_i$. Suppose WLOG that $w \in \mathcal{L}_j$ for some $j \geq i$, but $w \notin \mathcal{L}_k$ for any $k > j$. Then

$$\begin{aligned} \mathcal{I}(\mathbf{c}, w) &= \sum_{k \leq j} \mathcal{I}(c_k, w) \\ &= \mathcal{I}(c, w) + \sum_{k \leq j: k \neq i} \mathcal{I}(b_k, w) \\ &< \mathcal{I}(b_i, w) + \sum_{k \leq j: k \neq i} \mathcal{I}(b_k, w) \\ &= \mathcal{I}(\mathbf{b}, w) \end{aligned}$$

Now choose $w \notin \mathcal{L}_i$. Suppose WLOG that $w \in \mathcal{L}_j$ for some $j < i$, but $w \notin \mathcal{L}_k$ for any $k > j$. Then

$$\mathcal{I}(\mathbf{c}, w) = \sum_{k \leq j} \mathcal{I}(c_k, w) = \sum_{k \leq j} \mathcal{I}(b_k, w) = \mathcal{I}(\mathbf{b}, w)$$

Since $ch_0(\mathcal{L}_i) > 0$ for all $ch \in \mathcal{C}$, this implies

$$\begin{aligned} \sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathbf{b}, w) &= \sum_{w \in \mathcal{L}_i} ch_0(w) \mathcal{I}(\mathbf{b}, w) + \sum_{w \notin \mathcal{L}_i} ch_0(w) \mathcal{I}(\mathbf{b}, w) \\ &> \sum_{w \in \mathcal{L}_i} ch_0(w) \mathcal{I}(\mathbf{c}, w) + \sum_{w \notin \mathcal{L}_i} ch_0(w) \mathcal{I}(\mathbf{c}, w) \\ &= \sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathbf{c}, w) \end{aligned}$$

Case 2. Some extension \mathbf{q} of \mathbf{b} to \mathcal{F}^+ is coherent, but no coherent extension is a quasi-regular, condi, PP sequence.

Since \mathfrak{q} is coherent, \mathfrak{b} must be coherent. Since \mathfrak{b} is regular, it is trivially certainty-preserving. So by lemma 5, \mathfrak{b} is a quasi-regular superconditioning sequence. So we can assume WLOG that \mathfrak{q} is a coherent, quasi-regular, condi sequence, but not a PP sequence. That is, $\mathfrak{q} \in \mathbb{Q} - \mathbb{P}$.

By theorem 5, the projection of \mathfrak{q} onto \mathbb{P} , $\pi^{\mathfrak{q}}$, is such that

$$\mathcal{D}(\mathfrak{p}, \pi^{\mathfrak{q}}) \leq \mathcal{D}(\mathfrak{p}, \mathfrak{q}) - \mathcal{D}(\pi^{\mathfrak{q}}, \mathfrak{q})$$

for all $\mathfrak{p} \in \mathbb{P}$. So in particular

$$\mathcal{D}(\mathfrak{ch}, \pi^{\mathfrak{q}}) \leq \mathcal{D}(\mathfrak{ch}, \mathfrak{q}) - \mathcal{D}(\pi^{\mathfrak{q}}, \mathfrak{q})$$

for all $\mathfrak{ch} \in \mathfrak{C} \subseteq \mathbb{P}$.

Let \mathfrak{c} be the restriction of $\pi^{\mathfrak{q}}$ to \mathcal{F} .

Since \mathfrak{b} is not extendable to any $\mathfrak{p} \in \mathbb{P}$, $\pi^{\mathfrak{q}} \not\approx \mathfrak{q}$. So $\mathcal{D}(\pi^{\mathfrak{q}}, \mathfrak{q}) > 0$ and hence

$$\mathcal{D}(\mathfrak{ch}, \pi^{\mathfrak{q}}) < \mathcal{D}(\mathfrak{ch}, \mathfrak{q})$$

for all $\mathfrak{ch} \in \mathfrak{C}$. But this inequality holds iff

$$H(\mathfrak{ch}, \pi^{\mathfrak{q}}) < H(\mathfrak{ch}, \mathfrak{q})$$

And this is the case iff

$$\sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathfrak{c}, w) < \sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathfrak{b}, w)$$

for all $ch \in \mathfrak{C}$. This suffices to establish (I).

We will now prove (II). Choose $\mathfrak{b} \in \mathbb{S}$. Suppose that \mathfrak{b} is a J-Kon sequence. Suppose for reductio that there is some $\mathfrak{c} \neq \mathfrak{b}$ such that

$$(\star) \quad \sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathfrak{b}, w) \geq \sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathfrak{c}, w)$$

for all $ch \in \mathfrak{C}$.

Since \mathfrak{b} is a J-Kon sequence, \mathfrak{b} is extendable to some $\mathfrak{p} \in \mathbb{P}$, by proposition 1.

Let \mathfrak{q} be any extension of \mathfrak{c} to \mathcal{F}^+ . (\star) implies

$$\sum_{ch \in \mathfrak{C}} p_0(\phi_{ch}) \sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathfrak{b}, w) \geq \sum_{ch \in \mathfrak{C}} p_0(\phi_{ch}) \sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathfrak{c}, w)$$

Since \mathfrak{p} is a PP sequence, this holds iff

$$\sum_{w \in \mathcal{W}} \sum_{ch \in \mathfrak{C}} p_0(w \& \phi_{ch}) \mathcal{I}(\mathfrak{b}, w) \geq \sum_{w \in \mathcal{W}} \sum_{ch \in \mathfrak{C}} p_0(w \& \phi_{ch}) \mathcal{I}(\mathfrak{c}, w)$$

iff

$$\sum_{w \in \mathcal{W}} p_0(w) \mathcal{I}(\mathfrak{b}, w) \geq \sum_{w \in \mathcal{W}} p_0(w) \mathcal{I}(\mathfrak{c}, w)$$

And this implies

$$0 = \mathcal{D}(\mathfrak{p}, \mathfrak{p}) \geq \mathcal{D}(\mathfrak{p}, \mathfrak{q}) > 0$$

Hence there is no $\mathfrak{c} \neq \mathfrak{b}$ such that

$$\sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathfrak{b}, w) \geq \sum_{w \in \mathcal{W}} ch_0(w) \mathcal{I}(\mathfrak{c}, w)$$

for all $ch \in \mathfrak{C}$.

□