Epistemic Conservativity and Imprecise Credence

Jason Konek
University of Bristol
jason.konek@bristol.ac.uk

Abstract

Unspecific evidence calls for imprecise credence. My aim is to vindicate this thought. First, I will pin down what it is that makes one's imprecise credences more or less epistemically valuable. Then I will use this account of epistemic value to delineate a class of reasonable *epistemic scoring rules* for imprecise credences. Finally, I will show that if we plump for one of these scoring rules as our measure of epistemic value or utility, then a popular family of decision rules recommend imprecise credences. In particular, a range of *Hurwicz criteria*, which generalise the *Maximin* decision rule, recommend imprecise credences. If correct, the moral is this: an agent who adopts precise credences, rather than imprecise ones, in the face of unspecific and incomplete evidence, goes wrong by gambling with the epistemic utility of her doxastic state in too risky a fashion. Precise credences represent an overly risky epistemic bet, according to the Hurwicz criteria.

Keywords. Imprecise Credence, Epistemic Utility, Risk Aversion, Hurwicz Criterion

Hillary Clinton will win the 2016 U.S. presidential election. Tesla Motors' stock price will be over \$310 on January 1, 2016. Dinosaurs were wiped out by a giant asteroid, rather than gradual climatic change.

If you have a *perfectly precise* credence on any of these matters, you *might* be a little off your rocker. Having a precise credence for a proposition X means having opinions that are are so rich and specific that they pin down a *single* estimate $c(X)$ of the truth-value of X. (I follow de Finetti and Jeffrey in thinking of propositions as quantities that take the value 1 at worlds where they are true, and 0 where false. Truth-value estimates are estimates of the value, 0 or 1, that the proposition takes at the actual world.) If your evidence is anything like mine, it is too incomplete and ambiguous to justify such a rich and specific range of opinions.

Or maybe you are not off your rocker. Maybe precise credences are an acceptable characterisation of your uncertainty even when your evidence is unspecific and incomplete. Jeffrey (1983, 1987) thought so anyway. In any case, surely *I* am not off my rocker for *not* having a precise credence on these matters. So we can ask: How can I rationalise my doxastic response? What reasons do I have to adopt imprecise credences when my evidence is incomplete and unspecific? In particular, what *epistemic* reasons do I have?

The aim of this paper is to provide a new answer to this question. In particular, I hope to elucidate how adopting imprecise credences is a good means to the end of epistemic *value* or *utility*.

In §1, I walk through the most prominent extant rationale for adopting imprecise credences. In §2, I explain what it is missing, from the perspective of *epistemic utility theory*. In §3, I try to sort out what it is that makes one's imprecise credences more or less epistemically valuable. Then I use this account of epistemic value to delineate a class of reasonable *epistemic scoring rules* for imprecise credences. In §4, I show that if we plump for one of these scoring rules as our measure of epistemic value, then a popular family of decision rules recommend imprecise credences. In particular, a range of *Hurwicz criteria*, which generalise the *Maximin* decision rule, recommend imprecise credences. Finally, in §5, I summarise the preceding discussion, and respond briefly to a few objections.

1 Motivations for Imprecision

Suppose that you walk into a factory that makes green, black, and yellow balls. Before you: a seemingly bottomless opaque urn. You have no evidence whatsoever about the proportion of green, black, and yellow balls in the urn. Your tour guide then walks over to a towering machine, presses a button, and a giant mechanical arm reaches into the urn. It draws ten balls at random with replacement. The result: 2 green, 1 black, 7 yellow. Question: How confident should you be that the next ball drawn will be green?

Maybe you should turn to Laplace's Principle of Indifference.

POI. If an agent has incomplete and unspecific evidence about hypotheses $H_1, ..., H_n$, and so no reason to think that any one is more or less probable than any other, then she ought to be equally confident in each: $c(H_i) = c(H_j)$ for all H_i and H_j.

Since you have no information about the proportion of green/black/yellow balls in the urn, prior to observing the mechanical arm's ten draws, the POI says: You ought to spread your credence evenly over the various hypotheses about the relevant proportions. You ought to adopt uniform prior credences u over these hypotheses. Then, to figure out how confident should you be that the next ball drawn will be green, just update those prior credences u by conditionalizing them on your new data D, *viz.*, 2 green, 1 black, and 7 yellow in 10 trials. If you follow these instructions, your credence for green on the next draw will be precisely 3/13.

The problem, according to imprecise Bayesians like Levi, Joyce, and Kaplan is this: the POI, and any rule like it, "requires you to effect a precision in your state of opinion that your evidence often does not warrant" (Kaplan, 2010, p. 47). This, in turn, results in you making a wide range of unwarranted judgments. To drive the point home, consider a passage from Joyce (2010, pp. 284-5), amended to fit our example:

> The uniform distribution u commits you to thinking that in 100 independent draws (with replacement) the probability of seeing a green ball fewer than 15 times is exactly 0.274137, just a smidgen (≈ 0.02) more probable than drawing a diamond from a normal deck of cards. Do you really think that your evidence justifies such a specific probability assignment? Or, to take another example, are you comfortable with the idea that seeing s greens and $N - s$ black-or-yellows (in any proportion) should lead you to expect a green on the next draw with a credence of exactly $s + 1/N + 3$? This is W. E. Johnson's generalisation of the rule of succession. If you adopt the uniform distribution, you are stuck with it. The prior evidence you have about the proportion of green/black/yellow balls in the urn (viz., nada!) comes nowhere close to warranting such definite beliefs about repeated events and such specific inductive policies.

We might reframe the issue as follows: having precise credences requires having opinions that are are so rich and specific that they pin down a single estimate $c(X)$ of the truth-value of every proposition X that you are aware of. And this really is *incredibly* rich and specific. It means, *inter alia*, that your *comparative beliefs* — judgments of the form *X is at least as probable as Y* — must be *total*. That is, you must either think X is at least as probable as Y, or vice versa, for any propositions X and Y that you are aware of. No abstaining from judgment. Same goes for your conditional comparative beliefs. You must either think that X given D is at least as probable as Y given D', or vice versa, for any propositions X and Y, and any potential new data, D and D'. Likewise for your preferences. You must either think that bet A is at least as choiceworthy as bet B, or vice versa, for any A and B.

But incomplete and unspecific evidence requires non-totality. It requires you *not* to have total comparative beliefs, conditional comparative beliefs, preferences, and so on. Your prior evidence in the urn example comes nowhere near justifying the judgment that seeing a green ball fewer than 15 times is more probable/equally probable/less probable than than drawing a diamond from a normal deck of cards. Remember, you have *no prior information whatsoever* about the proportion of green, black, and yellow balls in the urn. So intuitively you should abstain from judgment on the matter. Your comparative beliefs should *not* be total. At the very least, you are rationally *permitted* to abstain from judgment, given the unspecific nature of your evidence.

Imprecise credences allow for non-totality. Suppose that you have opinions about some Boolean algebra of propositions $\mathcal{F}$ (closed under negation and disjunction). An imprecise credal state is a set $\mathcal{C}$ of credence functions c defined on $\mathcal{F}$. Adopting $\mathcal{C}$ is just a matter of making exactly the judgments that $\mathcal{C}$ is univocal about, *i.e.*, that all the credence functions in $\mathcal{C}$ agree about. For example:

- Judge that X is at least as probable as Y if and only if $c(X) \geq c(Y)$ for all $c \in \mathcal{C}$.
- Judge that X is probable to at least degree 0.6 and at most 0.9 if and only if $0.6 \leq c(X) \leq 0.9$ for all $c \in \mathcal{C}$.
- Judge that X is evidentially independent of Y if and only if $c(X|Y) = c(X)$ for all $c \in \mathcal{C}$.

So if an imprecise credal state $\mathcal{C}$ is not univocal about the comparative probability of propositions X and Y — *i.e.*, $b(X) \geq b(Y)$ and $c(X) < c(Y)$ for some $b, c \in \mathcal{C}$ — then an agent who adopts $\mathcal{C}$ will simply abstain from judgment on the matter. Her comparative beliefs will not be total. *Mutatis mutandis* for conditional comparative beliefs, preferences, etc.

The moral: If your evidence is incomplete and unspecific, then your comparative beliefs, conditional comparative beliefs, preferences, etc., should be correspondingly non-total, to reflect the unspecific nature of that evidence. This is the response that is most *justified*, or *warranted*, or *appropriate* in light of such evidence. (No other response is *more* justified anyway.) And only imprecise credences allow for non-totality. Having precise credences requires having total comparative beliefs, conditional comparative beliefs, preferences, etc. So, as Joyce puts it, "imprecise credences have a clear epistemological motivation: they are the proper response to unspecific evidence" (Joyce, 2005, p. 171).

Though this argument is forceful, it is also incomplete, according to epistemic utility theory. It does not tell the *full* story about what reasons we have to adopt imprecise credences. In §2, I explore why. In §3-4, I try to fill out this story, by illuminating how adopting imprecise credences is a good means to the end of epistemic *value* or *utility*.

2 Epistemic Utility Theory

According to epistemic utility theory, the principal tasks of any epistemology are two-fold: pinning down what it is that makes doxastic states more or less epistemically valuable, and illuminating why epistemic norms have the force that they do by showing that they are a good means to the end of such value.

The first task is *evaluative*. It involves specifying the epistemic good-making features of doxastic states, and saying how they conspire to make those states epistemically valuable.

The second task is *normative*. It involves first constructing a system of epistemic norms: *coherence norms*, which specify properties that our doxastic states ought to have at any given time, and how we ought to update those states, regardless of the peculiarities of our evidential circumstances; *deference principles*, such as the Principal Principle, which specify the sort of epistemic deference that we ought to afford certain experts (*e.g.*, chance); and *evidence norms*, which specify how we ought to respond to certain kinds of evidence, *e.g.*, the Principle of Indifference, or what we might dub the *Principle of Imprecision*:

PI. If an agent has incomplete and unspecific evidence about some hypotheses, then rather than assigning probabilities symmetrically to those hypotheses, she ought to refrain from assigning precise probabilities altogether. She ought to adopt some *imprecise* credal state.

The second part of this normative task is to show that our preferred norms *cohere with the rational pursuit of epistemic value*. We must: (i) identify some reasonable measures of epistemic value or utility, which place the appropriate emphasis on the appropriate good-making features (per our theory of epistemic value); (ii) identify some set of plausible choice rules; and (iii) show that these choice rules, together with our epistemic utility functions, deliver the same prescriptions as the system of epistemic norms.

There are two reasons you might want epistemic norms to cohere with the pursuit of epistemic value, in this sense. The first is that it provides a sort of *symbiotic support* for the resulting epistemology. If they cohere, then abiding by the aforementioned independently plausible epistemic norms *guarantees* that you will behave just as if you were explicitly pursuing epistemic value, according to independently plausible choice rules. Likewise, pursuing epistemic value in the right way *guarantees* that you will abide by the relevant norms. And *that* is good reason to think that each component of your epistemology was on the right track.

Alternatively, you might want epistemic norms to cohere with the pursuit of epistemic value in order to tell a *reductive* story about epistemic normativity. You might hope to provide a self-standing theory of epistemic value — one that can be justified without appeal to any particular epistemic norms — and then *fully explain* why various norms have the force that they do by elucidating the ways in which they are a good means to the end of epistemic value. Joyce (2013) and Fitelson and McCarthy (2014) are proponents of the non-hierarchical view. Pettigrew (2013, 2014) is a proponent of the reductive view.

Either way, with just the "justification rationale" in hand — imprecise credences are the proper response to unspecific evidence — we lack a *full* story about our reasons to adopt imprecise credences. We lack a full story until we explain the sense in which they are a good means to the end of epistemic value.

The aim now is to fill this gap. To accomplish this, we will first try to sort out what it is that makes imprecise credences more or less epistemically valuable. Then we will use this account of epistemic value to delineate a class of reasonable *epistemic scoring rules* for imprecise credences. Finally, we will show that relative to any such scoring rule, a popular family of choice rules — the *Hurwicz criteria* — recommend imprecise credences.

3 Scoring Rules for Imprecise Credences

William James tells us that we have two great commandments as would-be knowers: Avoid error! Seek truth! What's more, these are...

> ...two materially different laws; and by choosing between them we may end by coloring differently our whole intellectual life... We may regard the chase for truth as paramount, and the avoidance of error as secondary; or we may, on the other hand, treat the avoidance of error as more imperative, and let truth take its chance. (James, 1897, §VII)

Inspired by this, Isaac Levi (1967, 1984, 2004) proposed measuring the epistemic value or utility of a state of full belief K at a world w by the following quantity:[1]

$$\mathcal{V}(K,w) = \alpha \cdot \mathbb{E}(K,w) + (1-\alpha) \cdot \mathbb{T}(K,w).$$

The first component, $\mathbb{E}(K,w)$, is the truth-value of belief state K at w (0 if false, 1 if true), *i.e.*, the truth-value of the conjunction of all propositions believed in K. It measures the extent to which K promotes the first of our two principal epistemic aims: Avoid error! If K is true, it promotes it fully, *i.e.*, to degree 1 ($\mathbb{E}(K,w) = 1$). If false, it promotes it minimally, *i.e.*, to degree 0 ($\mathbb{E}(K,w) = 0$).

The second component, $\mathbb{T}(K,w)$ (a number between 0 and 1), is K's degree of informativeness, which is meant to reflect how virtuous K is at w, in terms of "simplicity, explanatory and predictive power, and other allegedly scientific or epistemic" desiderata (Levi (1991), p. 83; Levi (2012), p. 179). It measures the extent to which K promotes the second of our two principal aims: Seek truth! or as Levi reframes it: Seek Valuable Information! The more informative K is, the closer it comes to promoting it fully, *i.e.*, to degree 1 ($\mathbb{T}(K,w) \approx 1$). The less informative K is, the closer it comes to promoting it minimally, *i.e.*, to degree 0 ($\mathbb{T}(K,w) \approx 0$).

The third component, α (again, a number between 0 and 1), reflects the extent to which you let one or the other of our two principal epistemic aims

[1] See (Levi, 2004, §3.1) and Rott (2006).

— Avoid error! Seek Truth! — "color your intellectual life." It measures the respective degree of priority that you give to them.

Putting these components together as follows:

$$\mathcal{V}(K, w) = \alpha \cdot \mathbb{E}(K, w) + (1 - \alpha) \cdot \mathbb{T}(K, w).$$

gives us a "summary statistic" that measures the extent to which K succeeds at striking the optimal balance between promoting our two principal epistemic aims — Avoid error! Seek Truth! — at w.

We might call Levi's underlying theory of epistemic value the *bipartite theory*:

BTEV. A doxastic state D is more or less epistemically valuable at a world w to the extent that it succeeds at striking the optimal balance between promoting our two principal epistemic aims at w: Avoid error! Seek Truth!

We will now use the bipartite theory of epistemic value to delineate a class of reasonable *epistemic scoring rules* for imprecise credences. Much like Levi's measure, a scoring rule for imprecise credences is a function $\mathcal{I}$ which maps credal states $\mathcal{C}$ (sets of credence functions) and worlds w to non-negative real numbers, $\mathcal{I}(\mathcal{C}, w)$. And like before, $\mathcal{I}(\mathcal{C}, w)$ is a summary statistic that measures the extent to which $\mathcal{C}$ succeeds at striking the optimal balance between promoting our two principal epistemic aims at w: Avoid error! Seek Truth! The only difference: Levi's quantity, $\mathcal{V}(K, w)$, captures K's positive epistemic value at w. Our quantity, $\mathcal{I}(\mathcal{C}, w)$, on the other hand, captures $\mathcal{C}$'s epistemic *dis*value at w. If $\mathcal{I}(\mathcal{C}, w)$ equals 0, then $\mathcal{C}$ is minimally *dis*valuable at w (maximally epistemically valuable). It strikes the *best* possible balance between avoiding error and seeking truth. The larger $\mathcal{I}(\mathcal{C}, w)$ is, the less epistemically valuable.

Following Levi still, we will assume that $\mathcal{I}$ takes the following form:

$$\mathcal{I}(\mathcal{C}, w) = \alpha \cdot \mathcal{E}(\mathcal{C}, w) + (1 - \alpha) \cdot \mathcal{T}(\mathcal{C}, w).$$

Just as before, the first component, $\mathcal{E}(\mathcal{C}, w)$, is meant to measure the extent to which $\mathcal{C}$ promotes the first of our two principal epistemic aims: Avoid error! If it promotes it fully — if $\mathcal{C}$ avoids error to the greatest possible degree — then this first component contributes nothing to overall epistemic disutility: $\mathcal{E}(\mathcal{C}, w) = 0$. The worse $\mathcal{C}$ does at avoiding error at w, the larger $\mathcal{E}(\mathcal{C}, w)$ grows.

Similarly, the second component, $\mathcal{T}(\mathcal{C}, w)$, is meant to measure the extent to which $\mathcal{C}$ promotes the second of our two principal epistemic aims: Seek truth! If it promotes it fully — if $\mathcal{C}$ positively pins down the truth to the greatest possible degree — then this second component contributes nothing to overall epistemic disutility: $\mathcal{T}(\mathcal{C}, w) = 0$. The worse $\mathcal{C}$ does at pinning down the truth at w, the larger $\mathcal{T}(\mathcal{C}, w)$ grows.

Finally, α measures the degree to which you prioritise avoiding error over seeking truth, or vice versa.

To nail down $\mathcal{E}(\mathcal{C}, w)$ and $\mathcal{T}(\mathcal{C}, w)$, we we need some sense of *what it is* for an imprecise credal state $\mathcal{C}$ to avoid error, and to positively pin down the truth. Here's one plausible account, in a nutshell: credal states leave open and rule out various more or less accurate credence functions. They *avoid error* to the extent that they *leave open* accurate credences functions. They *pin down the truth* to the extent that they *rule out* inaccurate credence functions.

A bit more methodical: a credal state $\mathcal{C}$ leaves open all credence functions $b \in \mathcal{C}$ and rules out all credence functions $c \notin \mathcal{C}$. And these credence functions are more or less *accurate*. A credence function $c : \mathcal{F} \to \mathbb{R}$ is accurate at a world w to the extent that its credences $c(X)$ for propositions $X \in \mathcal{F}$ are close to the actual truth-values $w(X)$ of those propositions at w. Now, there is a wide range of reasonable measures of inaccuracy, each of which captures a slightly different way of valuing closeness to the truth (*cf.* §5-6). But for now, we will focus focus on the most popular one, which shares all of its most important properties with these other measures, *viz.*, the *Brier score*:

$$\mathcal{B}(c, w) = \frac{1}{|\mathcal{F}|} \cdot \sum_{X \in \mathcal{F}} (c(X) - w(X))^2$$

The Brier score proceeds in two stages. First, it looks at how far each individual credence, $c(X)$, is from X's actual truth-value, $w(X)$, and then penalises it the further away it is. In particular, it penalises that credence $c(X)$ by its squared Euclidean distance from X's actual truth-value, $(c(X) - w(X))^2$. Then it takes a weighted average of these individual penalties, to arrive at a measure of c's overall inaccuracy at w, $\mathcal{B}(c,w)$.

What can we say, then, about the extent to which a credal state $\mathcal{C}$ leaves open accurate credence functions, and rules out *in*accurate credence functions? In particular, can we find statistics that measure how well $\mathcal{C}$ performs in these respects?

There are two seemingly obvious things to say about our desired statistics. Suppose that we have two credal states, $\mathcal{C}$ and $\mathcal{C}'$. First obvious fact: if avoiding error is a matter of leaving open accurate credences, and $\mathcal{C}$ leaves open (contains) an unequivocally more accurate credence function than $\mathcal{C}'$ — a credence function that is more accurate (at w) than *every single credence function in* $\mathcal{C}'$ — then $\mathcal{C}$ does a better job avoiding error than $\mathcal{C}'$: $\mathcal{E}(\mathcal{C},w) < \mathcal{E}(\mathcal{C}',w)$. This is true just in case:

$$(\heartsuit)\ \text{If } \min_{c\in\mathcal{C}} \mathcal{B}(c,w) < \min_{c'\in\mathcal{C}'} \mathcal{B}(c',w), \text{ then}$$
$$\mathcal{E}(\mathcal{C},w) < \mathcal{E}(\mathcal{C}',w).$$

This means, for example, that if $\mathcal{C}$ is maximally inclusive — it contains every credence function, and so suspends judgment on every issue — then it does better at avoiding error at w (for any w) than any $\mathcal{C}'$ that determinately makes even the smallest mistake (*i.e.*, any $\mathcal{C}'$ such that $\mathcal{B}(c',w) > 0$ for all $c' \in \mathcal{C}'$).

Second obvious fact: if pinning down the truth is a matter of *not* leaving open (ruling out) *in*accurate credences, and $\mathcal{C}$ leaves open an unequivocally *less* accurate credence function than $\mathcal{C}'$ — a credence function that is less accurate (at w) than *every single credence function in* $\mathcal{C}'$ — then $\mathcal{C}$ does a *worse* job pinning down the truth than $\mathcal{C}'$: $\mathcal{T}(\mathcal{C},w) > \mathcal{T}(\mathcal{C}',w)$. And this is true just in case:

$$(\diamondsuit)\ \text{If } \max_{c\in\mathcal{C}} \mathcal{B}(c,w) > \max_{c'\in\mathcal{C}'} \mathcal{B}(c',w), \text{ then}$$
$$\mathcal{T}(\mathcal{C},w) > \mathcal{T}(\mathcal{C}',w).$$

($\heartsuit$) and ($\diamondsuit$) seem non-negotiable. The converse of ($\heartsuit$) and ($\diamondsuit$) are less obvious. For example, it is less than obvious that $\mathcal{C}$ does a determinately better job avoiding error than $\mathcal{C}'$ *only if* $\mathcal{C}$ leaves open an unequivocally more accurate credence function: a credence function that is more accurate than *every single credence function left open by* $\mathcal{C}'$. Imagine, for example, that $\mathcal{C}$ and $\mathcal{C}'$ contain the same most-accurate credence function, b. But every other c in $\mathcal{C}$ is more accurate than every other c' in $\mathcal{C}'$. Perhaps under these circumstances we want to say that $\mathcal{C}$ does a better job avoiding error than $\mathcal{C}'$. In that case, the converse of ($\heartsuit$):

$$(\heartsuit\heartsuit)\ \text{If } \mathcal{E}(\mathcal{C},w) < \mathcal{E}(\mathcal{C}',w), \text{ then}$$
$$\min_{c\in\mathcal{C}} \mathcal{B}(c,w) < \min_{c'\in\mathcal{C}'} \mathcal{B}(c',w).$$

would be false. You might doubt the converse of ($\diamondsuit$) for similar reasons.

$$(\diamondsuit\diamondsuit)\ \text{If } \mathcal{T}(\mathcal{C},w) > \mathcal{T}(\mathcal{C}',w), \text{ then}$$
$$\max_{c\in\mathcal{C}} \mathcal{B}(c,w) > \max_{c'\in\mathcal{C}'} \mathcal{B}(c',w).$$

Nevertheless, we will run with ($\heartsuit\heartsuit$) and ($\diamondsuit\diamondsuit$). This means adopting a rather exacting conception of what it is to avoid error and seek truth. On this view, the *only* way to do a better job avoiding error is by moving from an old credal state $\mathcal{C}$ to a new one $\mathcal{C}'$ that leaves open an unequivocally more accurate credence function: a credence function that is more accurate than *every single credence function in* $\mathcal{C}$. And the *only* way to do a better job pinning down the truth is by moving to a new credal state $\mathcal{C}'$ from an old one $\mathcal{C}$ that left open an unequivocally *less* accurate credence function: a credence function that is less accurate than *every single credence function in* $\mathcal{C}'$.

Of course, this is not the *only* way to think about what it is to avoid error and seek truth. But that's no problem. Our task is modest: to elucidate one plausible way in which adopting imprecise credences can be seen as a good means to the end of epistemic value. So we do not need our way to be the only way.

Given ($\heartsuit$) and ($\diamondsuit$), and their converses ($\heartsuit\heartsuit$) and ($\diamondsuit\diamondsuit$), we have:

$$(\blacktriangledown)\ \mathcal{E}(\mathcal{C},w) < \mathcal{E}(\mathcal{C}',w) \textit{ iff } \min_{c\in\mathcal{C}} \mathcal{B}(c,w) < \min_{c'\in\mathcal{C}'} \mathcal{B}(c',w)$$

$$(\blacklozenge)\ \mathcal{T}(\mathcal{C},w) < \mathcal{T}(\mathcal{C}',w) \textit{ iff } \max_{c\in\mathcal{C}} \mathcal{B}(c,w) < \max_{c'\in\mathcal{C}'} \mathcal{B}(c',w)$$

The most straightforward to guarantee that ($\blacktriangledown$) and ($\blacklozenge$) are satisfied is by assuming:

$$(\clubsuit)\ \mathcal{E}(\mathcal{C},w) = \min_{c\in\mathcal{C}} \mathcal{B}(c,w)$$

$$(\spadesuit)\ \mathcal{T}(\mathcal{C},w) = \max_{c\in\mathcal{C}} \mathcal{B}(c,w)$$

Again, this is not the *only* way to guarantee that (♥) and (♦) are satisfied. (Setting $\mathcal{E}(\mathcal{C}, w)$ and $\mathcal{T}(\mathcal{C}, w)$ equal to any monotonic transformation of $\min_{c \in \mathcal{C}} \mathcal{B}(c, w)$ and $\max_{c \in \mathcal{C}} \mathcal{B}(c, w)$ will do.) But it does not need to be.

With (♣) and (♠) in hand, our Levi-style scoring rule for imprecise credences looks like this:

$$\begin{aligned} \mathcal{I}(\mathcal{C}, w) &= \alpha \cdot \mathcal{E}(\mathcal{C}, w) + (1-\alpha) \cdot \mathcal{T}(\mathcal{C}, w) \\ &= \alpha \cdot \min_{c \in \mathcal{C}} \mathcal{B}(c, w) + (1-\alpha) \cdot \max_{c \in \mathcal{C}} \mathcal{B}(c, w) \end{aligned}$$

This is our official scoring rule moving forward. The first component, $\min_{c \in \mathcal{C}} \mathcal{B}(c, w)$, measures the extent to which $\mathcal{C}$ avoids error at w. The second component, $\max_{c \in \mathcal{C}} \mathcal{B}(c, w)$, measures the extent to which $\mathcal{C}$ positively pins down the truth at w. Finally, α measures the relative importance of avoiding error and seeking truth, respectively. The closer α is to 1, the more important it is to avoid error; the more *epistemically conservative* $\mathcal{I}$ is. The closer α is to 0, the more important it is to positively pin down the truth; the more *epistemically liberal* $\mathcal{I}$ is.

The end result: a summary statistic, $\mathcal{I}(\mathcal{C}, w)$, which captures the extent to which $\mathcal{C}$ succeeds at striking the optimal balance between promoting our two principal epistemic aims at w: Avoid error! Seek Truth! On the bipartite theory, this is *precisely* the statistic that captures $\mathcal{C}$'s overall epistemic value.

Now I will show that if we plump for $\mathcal{I}$ as our measure of epistemic value or utility, and if $\mathcal{I}$ is sufficiently *conservative*, then a popular family of decision rules — the *Hurwicz criteria* — recommends imprecise credences.

4 A Hurwicz-Style Argument for the Principle of Imprecision

Pettigrew (2014) offers the following argument for the Principle of Indifference:

1. **Measure of epistemic value**: The epistemic value or utility of a precise credal state $c : \mathcal{F} \to \mathbb{R}$ at a world w is measured by its accuracy, which is given by the negative of its Brier score.
2. **Minimax decision rule**: If an agent has incomplete and unspecific evidence about the value of options in $\mathcal{O}$, then she ought to choose an o in $\mathcal{O}$ that maximises worst-case utility.
3. **Theorem**: the uniform distribution $u : \mathcal{F} \to \mathbb{R}$, defined by $u(X) = 1/|\mathcal{F}|$ for all $X \in \mathcal{F}$, uniquely minimises worst-case Brier score, and hence uniquely maximises worst-case epistemic utility:

$$\max_{w \in \mathcal{W}} \mathcal{B}(u, w) < \max_{w \in \mathcal{W}} \mathcal{B}(c, w)$$

for all other credence functions $c : \mathcal{F} \to \mathbb{R}$.

C. **Principle of Indifference**: If an agent has incomplete and unspecific evidence about the propositions in $\mathcal{F}$, and hence incomplete and unspecific evidence about the epistemic value of possible credence functions $c : \mathcal{F} \to \mathbb{R}$, then she ought to adopt the uniform distribution $u : \mathcal{F} \to \mathbb{R}$.

This argument reveals something important about the Principle of Indifference. It shows that the POI coheres with the rational pursuit of epistemic value: there is some reasonable measure of epistemic value for precise credences — the Brier score — and some *prima facie* plausible choice rule — Minimax — which together deliver the same prescriptions as the POI.

To the extent, then, that you see the Brier score and Minimax as reasonable expressions of your epistemic values — to the extent that they evaluate doxastic states in just the right way, to your mind, on the basis of just the right good-making features — you can justify the POI as follows. Any agent who adopts non-uniform precise credences, rather than uniform ones, in the face of unspecific evidence, goes wrong by gambling with the epistemic utility of her credences in too risky a fashion. She risks greater inaccuracy, as measured by the Brier score, and hence lower epistemic value, than she needs to risk — a no-no, according to Minimax.

Non-reductivists will take this to provide *symbiotic support* for the POI, the Brier Score, and Minimax. Reductivists will take it to *fully explain* why the POI has the normative force that it does. Either way, it helps to provide a fuller story about our purely epistemic reasons for distributing credence uniformly in the face of unspecific evidence.

But Pettigrew's argument also comes with a caveat: Only apply if you are restricting your attention to precise credal states! Once we countenance the possibility that we might respond to unspecific evidence by adopting *imprecise* credences, all bets are off. Pettigrew's theorem shows that Minimax recommends uniform credences over any other *precise* credal state. But it does *not* show that it recommends uniform credence over any other credal state *full stop* — precise or imprecise.

Our argument for the Principle of Imprecision generalises Pettigrew's in a way that allows us to do away with this caveat:

1. **Measure of epistemic value**: The epistemic value or utility of a credal state $\mathcal{C}$ — precise or imprecise — at a world w is measured by

$$\begin{aligned}\mathcal{I}(\mathcal{C},w) &= \alpha\cdot\mathcal{E}(\mathcal{C},w)+(1-\alpha)\cdot\mathcal{T}(\mathcal{C},w)\\ &= \alpha\cdot\min_{c\in\mathcal{C}}\mathcal{B}(c,w)+(1-\alpha)\cdot\max_{c\in\mathcal{C}}\mathcal{B}(c,w)\end{aligned}$$

   As a special case, when $\mathcal{C}$ is precise, *i.e.*, $\mathcal{C}=\{c\}$ for some credence function $c:\mathcal{F}\to\mathbb{R}$, its epistemic value is just given by the negative of its Brier score: $\mathcal{I}(\mathcal{C},w)=\mathcal{B}(c,w)$.

2. **Hurwicz decision rule**: If an agent has incomplete and unspecific evidence about the value of options in $\mathcal{O}$, then she ought to choose an o in $\mathcal{O}$ that strikes the best balance between worst-case utility $\mathcal{U}^-$ and best-case utility $\mathcal{U}^+$, *i.e.*, she ought to choose an o that maximises $\beta\cdot\mathcal{U}^-+(1-\beta)\cdot\mathcal{U}^+$, where $0\leq\beta\leq 1$ measures the relative importance of avoiding worst-case catastrophe and maximising best-case spoils.

3. **Theorem**: If our epistemic utility measure and our Hurwicz rule are both *conservative*, in the sense that $\alpha>1/2$ and $\beta>1/2$ — *i.e.*, they place more of an emphasis on avoiding error/worst-case catastrophe than on positively pinning down the truth/maximising best-case spoils — then there is some imprecise credal state $\mathcal{C}$ that our Hurwicz rule recommends over all precise credal states:

$$\begin{aligned}\mathcal{H}(\mathcal{C}) &= \beta\cdot\max_{w\in\mathcal{W}}\mathcal{I}(\mathcal{C},w)+(1-\beta)\cdot\min_{w\in\mathcal{W}}\mathcal{I}(\mathcal{C},w)\\ &< \beta\cdot\max_{w\in\mathcal{W}}\mathcal{I}(\{c\},w)+(1-\beta)\cdot\min_{w\in\mathcal{W}}\mathcal{I}(\{c\},w)\\ &= \beta\cdot\max_{w\in\mathcal{W}}\mathcal{B}(c,w)+(1-\beta)\cdot\min_{w\in\mathcal{W}}\mathcal{B}(c,w)\\ &= \mathcal{H}(c)\end{aligned}$$

   for any precise credence function $c:\mathcal{F}\to\mathbb{R}$. (See appendix for proof.)

C. **Principle of Imprecision**: If an agent has incomplete and unspecific evidence about the propositions in $\mathcal{F}$, and hence incomplete and unspecific evidence about the epistemic value of possible credal states (precise or imprecise), then she ought to adopt some *imprecise* credal state.

This argument helps to illuminate our most basic epistemic reasons for adopting imprecise credences. It reveals one way in which adopting imprecise credences is a good means to the end of epistemic value or utility. If you measure epistemic utility by a conservative Levi-style scoring rule, then any conservative Hurwicz decision rule recommends adopting imprecise credences. It judges imprecise credences to be a better means to the end of such value. Any agent who adopts precise credences, rather than imprecise ones, in the face of unspecific evidence, goes wrong by gambling with the epistemic utility of her doxastic state in too risky a fashion. Precise credences represent an overly risky epistemic bet, according to the Hurwicz criteria.

Non-reductivists will take this to provide *symbiotic support* for the Principle of Imprecision, conservative Levi-style scoring rules, and conservative Hurwicz decision rules. Reductivists will take it to *fully explain* why the Principle of Imprecision has the normative force that it does. Either way, it helps to provide a fuller story about our purely epistemic reasons for adopting imprecise credences in the face of unspecific evidence.

Our argument tells us something about the Principle of Indifference too. Minimax does *not* recommend the POI — at least not if we countenance the possibility of responding to unspecific evidence by adopting imprecise credences, and we measure their value by conservative Levi-style scoring rules. In that case, Minimax — the most extreme example of a conservative Hurwicz rule — rather than recommending uniform credences, recommends imprecise credences.

We *can*, it turns out, rehabilitate Pettigrew's argument for the POI, even once countenancing imprecise credences. But it requires plumping for an extremely *liberal* Levi-style scoring rule — one that prioritises positively pinning down the truth over avoiding error — together with an extremely *conservative* Hurwicz rule — one that prioritises avoiding catastrophic worst-case epistemic disutility over maximising best-case epistemic spoils. Perhaps the fact that the POI requires this sort of mismatch between your epistemic utility function and choice rule, and hence requires seemingly discordant or incoherent epistemic values, is bad news for the POI. Or perhaps it is bad news for our Levi-style scoring rules, or Hurwicz decision rules.

5 Summary, Generalisations, and Objections

Evidence is often incomplete and unspecific. And when it is, it seems, our credences ought to be imprecise, to reflect that character of the evidence. This is the *proper* response to incomplete and unspecific evidence.

Our question was this: How can we rationalise this response? What are our most basic epistemic reasons for adopting imprecise credences?

The most prominent extant rationale — *viz.*, that unspecific evidence *justifies* or *warrants* decidedly *non*-total comparative beliefs, preferences, etc., which require imprecise credences — fails to tell the *full* story about why we ought to "go imprecise," we argued. Adopting imprecise credences is also a good means to the end of epistemic *value* or *utility*. Any fully adequate epistemology ought to vindicate this thought.

So we set out to vindicate it. To that end, we tried to sort out what it is that makes imprecise credences more or less epistemically valuable. Our account in a nutshell: imprecise credences are epistemically valuable to the extent that they succeed at striking the optimal balance between promoting our two principal epistemic aims: Avoid error! Seek Truth! They avoid error to the extent that they leave open accurate credences functions. They pin down the truth to the extent that they rule out inaccurate credence functions.

Then we used this account of epistemic value to delineate a class of *prima facie* reasonable Levi-style scoring rules for imprecise credences. Finally, we showed that that if we plump for one of these scoring rules as our measure of epistemic value or utility, then a popular family of decision rules recommend imprecise credences. In particular, a range of conservative Hurwicz decision rules recommend imprecise credences.

The upshot: an agent who adopts precise credences, rather than imprecise ones, in the face of unspecific evidence, goes wrong by gambling with the epistemic utility of her doxastic state in too risky a fashion. Precise credences represent an overly risky epistemic bet, according to the Hurwicz criteria.

Before concluding, we ought to briefly address a few potential concerns. Firstly, our preferred measures of epistemic value or utility — conservative Levi-style scoring rules — extend one particular score for precise credal states, *viz.*, the Brier score, to the space of imprecise credal states. But surely there are other reasonable scoring rules for precise credal states, *e.g.*, the logarithmic score, or power scores. We could just as well have used *those* to generate Levi-style scoring rules for imprecise credal states. If we had, would our argument for the Principle of Imprecision still have gone through?

In the appendix, I prove that any conservative Hurwicz rule, together with any conservative Levi-style scoring rule that extends a *truth-directed*, *extensional*, *convex*, *strictly proper* precise score, recommends imprecise credences. See Joyce (2009) for discussion of these properties. Our argument for the Principle of Imprecision goes through on a wide range of epistemic utility measures.

Another potential concern: the conservative Hurwicz choice rules are not the only conservative choice rules in town. Even if you find epistemic conservatism plausible — you think that, when our evidence is unspecific, we ought to choose between doxastic states in a way that prioritises avoiding worst-case epistemic catastrophe over maximising best-case epistemic spoils — there are other choice rules open to you. For example, rather than using a conservative Hurwicz rule, you might have used *Gamma-Minimax*, which recommends choosing the doxastic state that maximises worst-case *expected* epistemic disutility, relative to the set of credence functions left open by your evidence. Why a conservative Hurwicz rule rather than Gamma-Minimax?

I do not insist that conservative Hurwicz rules are the *only* plausible conservative choice rules. Rather, I follow Levi in taking your choice rule to be an expression of your epistemic values. Your epistemic values are reflected in both the epistemic utility measures and the package of choice rules that rationalise your doxastic behaviour, *e.g.*, adopting imprecise credences. And there may well be many *prima facie* plausible epistemic utility measures and choice rules that do so. This paper only aims to outline one pair. Outlining others is a task for future research.

I conclude by raising a few additional questions to be addressed in future research.

- Seidenfeld et al. (2012), Mayo-Wilson and Wheeler (2014a) and Schoenfield (2014) all provide related impossibility results, which show that real-valued scoring rules for imprecise credences must fail to have certain *prima facie* desirable properties, *e.g.*, they cannot be strictly proper, in a certain sense. What does this mean for our real-valued, Levi-style scoring rules? Are they unreasonable measures of epistemic value? Or are those *prima facie* desirable properties less desirable than they seem on their face?

- Mayo-Wilson and Wheeler (2014b) question the motivation for assuming that epistemic value or utility can be precisely quantified, in the way required for *real-valued* scoring rules to serve as measures of such value. Even on the supervaluationist approach, according to which determinate facts about epistemic value are just the facts that are invariant across *all* reasonable real-valued scoring rules, there is a genuine question about what hidden assumptions we are sneaking in via the underlying numerical structure. Can we instead run our argument for the Principle of Imprecision with *comparative* epistemic utility orderings, which simply rank credal states in terms of their comparative utility?

* Thanks especially to Richard Pettigrew, who provided key insight into the main theorem of this paper. Thanks also to a number of other people for helpful feedback: Jim Joyce, Sarah Moss, Allan Gibbard, Ben Levinstein, Gregory Wheeler, Seamus Bradley, Catrin Campbell-Moore, and others at the MCMP. I was supported by ERC Starting Grant "Epistemic Utility Theory: Foundations and Applications" during my work on this paper.

6 Appendix

Let $\mathcal{F} = \{X_1, ..., X_n\}$ be a finite partition.

Let $\mathbb{P}$ be the set of probability distributions on $\mathcal{F}$.

Let w_i be the world such that $w_i(X_i) = 1$ and $w_i(X_j) = 0$ for all $j \neq i$.

Let $\mathcal{W} = \{w_1, ..., w_n\}$.

Let $\mathcal{I}$ be an inaccuracy measure

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

that satisfies Truth-Directedness, Extensionality, Convexity, and Strict-Propriety, where $s : [0,1] \times \{0,1\} \to [0, \infty]$ is what Joyce (2009) calls a *component function*, which measures the inaccuracy of the credence $c(X)$ when X's truth-value is $w(X)$.

- *Truth-Directedness*: For any b and c, and any world w, if b's truth-value estimates are uniformly closer to the truth than c's at w, so that:

  $|b(X) - w(X)| \leq |c(X) - w(X)|$ for all X

  and:

  $|b(Y) - w(Y)| < |c(Y) - w(Y)|$ for some Y

  then $\mathcal{I}(b, w) < \mathcal{I}(c, w)$.

- *Extensionality*: The inaccuracy of c at w, $\mathcal{I}(c, w)$, is solely a function of the c's credences for propositions in $\mathcal{F}$ and their truth-values at w.

- *Convexity*. For any b and c, and any $\lambda \in (0,1)$,

  $\mathcal{I}(\lambda \cdot b + (1-\lambda) \cdot c, w) < \lambda \cdot \mathcal{I}(b, w) + (1-\lambda) \cdot \mathcal{I}(c, w)$

- *Strict Propriety*. For any probabilistically coherent c and $b \neq c$, $\mathbf{Exp}_c(\mathcal{I}(c)) < \mathbf{Exp}_c(\mathcal{I}(b))$.

Extend $\mathcal{I}$ to the space $\mathbb{C} = 2^{\mathbb{P}}$ of imprecise credal states $\mathcal{C}$ as follows:

$$\mathcal{I}(\mathcal{C}, w) = \alpha \cdot \min_{c \in \mathcal{C}} \mathcal{I}(c, w) + (1 - \alpha) \cdot \max_{c \in \mathcal{C}} \mathcal{I}(c, w)$$

Assume that $\mathcal{I}$ is conservative, in the sense that $1/2 < \alpha \leq 1$.

Now choose a Hurwicz penalty:

$$\mathcal{H}(\mathcal{C}) = \beta \cdot \max_{w \in \mathcal{W}} \mathcal{I}(\mathcal{C}, w) + (1 - \beta) \cdot \min_{w \in \mathcal{W}} \mathcal{I}(\mathcal{C}, w)$$

Assume that $\mathcal{H}$ is conservative as well, in the sense that $1/2 < \beta \leq 1$.

<u>To Show</u>: There is some imprecise (non-singleton) credal state $\mathcal{C}$ that $\mathcal{H}$ recommends over every precise credal state: $\mathcal{H}(\mathcal{C}) < \mathcal{H}(c)$ for every $c \in \mathbb{P}$.

Let

$$c_i(X_j) := \begin{cases} \alpha & \text{if } i = j \\ \frac{1-\alpha}{n-1} & \text{if } i \neq j \end{cases}$$

So,

- $c_1 = \langle \alpha, \frac{1-\alpha}{n-1}, \ldots, \frac{1-\alpha}{n-1} \rangle$
- $c_n = \langle \frac{1-\alpha}{n-1}, \ldots, \frac{1-\alpha}{n-1}, \alpha \rangle$

Similarly, let

$$b_i(X_j) := \begin{cases} \frac{1}{2} & \text{if } i = j \\ \frac{1}{2(n-1)} & \text{if } i \neq j \end{cases}$$

Note that since $1/2 < \alpha \leq 1$,

$$b_i = \lambda_i \cdot c_i + \sum_{j \neq i} \lambda_j w_j,$$

for some $0 < \lambda_1, ..., \lambda_n < 1$ with $\sum_k \lambda_k = 1$. We will return to b_i toward the end of the proof.

Let $\mathcal{C}$ be the convex hull of $\{c_1, ..., c_n\}$.

Choose $c \in \mathbb{P}$.

<u>To Show</u>: $H(\mathcal{C}) < H(c)$.

To prove this, we will prove four auxiliary claims:

1. $\mathcal{I}(\mathcal{C}, w_i) = \mathcal{I}(\mathcal{C}, w_j)$ for all $1 \leq i, j \leq n$

In that case, $H(\mathcal{C}) = \mathcal{I}(\mathcal{C}, w_i)$. Then we will prove

2. $\mathcal{I}(\mathcal{C}, w_i) = \text{Exp}_{c_i}(\mathcal{I}(c_i))$

We will go on to show that

3. $\text{Exp}_{c_i}(\mathcal{I}(c_i)) < \text{Exp}_{b_j}(\mathcal{I}(c))$

Finally, we will show that

4. $\text{Exp}_{b_j}(\mathcal{I}(c)) \leq \mathcal{H}(c)$

Together, 1-4 establish that $\mathcal{H}(\mathcal{C}) < \mathcal{H}(c)$.

<u>Proof of 1 and 2</u>: Choose $1 \leq i \leq n$.

First we will prove

1a. $\min_{c \in \mathcal{C}} \mathcal{I}(c, w_i) = \mathcal{I}(c_i, w_i)$

1b. $\max_{c \in \mathcal{C}} \mathcal{I}(c, w_i) = \mathcal{I}(c_j, w_i)$ for any $j \neq i$

<u>Proof of 1a</u>: For any $c \in \mathcal{C}$ with $c \neq c_i$,

$$c(X_i) < c_i(X_i) = \alpha \leq w_i(X_i) = 1$$

and

$$c(X_j) \geq c_i(X_j) = \tfrac{1-\alpha}{n-1} \geq w_i(X_j) = 0$$

for all $j \neq i$. So $\mathcal{I}(c_i, w_i) < \mathcal{I}(c, w_i)$ by Truth-Directedness. Hence $\min_{c \in \mathcal{C}} \mathcal{I}(c, w_i) = \mathcal{I}(c_i, w_i)$.

<u>Proof of 1b</u>: Choose $j \neq i$.

Suppose $\max_{c \in \mathcal{C}} \mathcal{I}(c, w_i) \neq \mathcal{I}(c_j, w_i)$ for reductio.

Then there must be some $b \in \mathcal{C}$ with $\mathcal{I}(c_j, w_i) < \mathcal{I}(b, w_i)$.

Note: $b \neq c_i$, since we can't have $\mathcal{I}(c_j, w_i) < \mathcal{I}(c_i, w_i)$ (that would violate Truth-Directedness).

And $b \neq c_k$ for any $k \neq i$, since we can't have $\mathcal{I}(c_j, w_i) \neq \mathcal{I}(c_k, w_i)$ (that would violate Extensionality).

So $b \neq c_k$ for any $1 \leq k \leq n$.

Since $\mathcal{C}$ is the convex hull of $\{c_1, ..., c_n\}$, $b = \sum_k \lambda_k c_k$ for some $\sum_k \lambda_k = 1$. And since $\mathcal{I}$ is convex, $\mathcal{I}(b, w_i) < \sum_k \lambda_k \mathcal{I}(c_k, w_i) \leq \mathcal{I}(c_j, w_i)$. $\Rightarrow\Leftarrow$.

Putting $(1a)$ and $(1b)$ together gives us:

$$\begin{aligned} \mathcal{I}(\mathcal{C}, w_i) &= \alpha \cdot \min_{c \in \mathcal{C}} \mathcal{I}(c, w_i) + (1-\alpha) \cdot \max_{c \in \mathcal{C}} \mathcal{I}(c, w_i) \\ &= \alpha \cdot \mathcal{I}(c_i, w_i) + (1-\alpha) \cdot \mathcal{I}(c_j, w_i) \end{aligned}$$

Now note that $\mathcal{I}(c_j, w_i) = \mathcal{I}(c_i, w_j)$, by Extensionality. And $\mathcal{I}(c_i, w_j) = \mathcal{I}(c_i, w_k)$ for any $k \neq i$, by Extensionality. Hence:

$$\begin{aligned} \mathcal{I}(\mathcal{C}, w_i) &= \alpha \cdot \min_{c \in \mathcal{C}} \mathcal{I}(c, w_i) + (1-\alpha) \cdot \max_{c \in \mathcal{C}} \mathcal{I}(c, w_i) \\ &= \alpha \cdot \mathcal{I}(c_i, w_i) + (1-\alpha) \cdot \mathcal{I}(c_j, w_i) \\ &= \alpha \cdot \mathcal{I}(c_i, w_i) + (1-\alpha) \cdot \mathcal{I}(c_i, w_j) \\ &= \alpha \cdot \mathcal{I}(c_i, w_i) + \frac{(1-\alpha)}{n-1} \cdot \sum_{k \neq i} \mathcal{I}(c_i, w_k) \\ &= \text{Exp}_{c_i}(\mathcal{I}(c_i)) \end{aligned}$$

A similar argument shows $\mathcal{I}(\mathcal{C}, w_j) = \text{Exp}_{c_j}(\mathcal{I}(c_j))$.

But of course

$$\text{Exp}_{c_i}(\mathcal{I}(c_i)) = \text{Exp}_{c_j}(\mathcal{I}(c_j))$$

for all $1 \leq i, j \leq n$. So $\mathcal{I}(\mathcal{C}, w_i) = \mathcal{I}(\mathcal{C}, w_j)$ as well.

This means that

$$\begin{aligned}
\mathcal{H}(\mathcal{C}) &= \beta \cdot \max_{w \in \mathcal{W}} \mathcal{I}(\mathcal{C}, w) + (1-\beta) \cdot \min_{w \in \mathcal{W}} \mathcal{I}(\mathcal{C}, w) \\
&= \mathcal{I}(\mathcal{C}, w_i) \\
&= \text{Exp}_{c_i}(\mathcal{I}(c_i))
\end{aligned}$$

Proof of 3:

Assume WLOG that $w_i \in \arg\min_{w \in \mathcal{W}} \mathcal{I}(c, w)$.

Then since $1/2 < \beta \leq 1$, we have

$$\begin{aligned}
\mathcal{H}(c) &= \beta \cdot \max_{w \in \mathcal{W}} \mathcal{I}(c, w) + (1-\beta) \cdot \min_{w \in \mathcal{W}} \mathcal{I}(c, w) \\
&\geq \frac{1}{2} \cdot \max_{w \in \mathcal{W}} \mathcal{I}(c, w) + \frac{1}{2} \cdot \min_{w \in \mathcal{W}} \mathcal{I}(c, w) \\
&= \frac{1}{2} \cdot \max_{w \in \mathcal{W}} \mathcal{I}(c, w) + \frac{1}{2} \cdot \mathcal{I}(c, w_i) \\
&\geq \frac{1}{2(n-1)} \cdot [\sum_{j \neq i} \mathcal{I}(c, w_j)] + \frac{1}{2} \cdot \mathcal{I}(c, w_i) \\
&= \text{Exp}_{b_i}(\mathcal{I}(c))
\end{aligned}$$

Proof of 4: Since $\mathcal{I}$ is strictly proper,

$$\text{Exp}_{b_i}(\mathcal{I}(c)) > \text{Exp}_{b_i}(\mathcal{I}(b_i)).$$

In addition, for any strictly proper $\mathcal{I}$, the entropy function $Ent(b) = \text{Exp}_b(\mathcal{I}(b))$ is concave.

Now recall that

$$b_i = \lambda_i \cdot c_i + \sum_{j \neq i} \lambda_j w_j,$$

for some $0 < \lambda_1, \ldots, \lambda_n < 1$. So

$$\begin{aligned}
&\text{Exp}_{b_i}(\mathcal{I}(b_i)) \\
> \;& \lambda_i \cdot \text{Exp}_{c_i}(\mathcal{I}(c_i)) + \sum_{j \neq i} \lambda_j \cdot \text{Exp}_{w_j}(\mathcal{I}(w_j)) \\
= \;& \lambda_i \cdot \text{Exp}_{c_i}(\mathcal{I}(c_i)) \\
> \;& \text{Exp}_{c_i}(\mathcal{I}(c_i))
\end{aligned}$$

Together, (3)-(4) show that $\mathcal{H}(c) > \text{Exp}_{c_i}(\mathcal{I}(c_i))$.

Therefore, by (1)-(4),

$$\begin{aligned}
\mathcal{H}(\mathcal{C}) &= \mathcal{I}(\mathcal{C}, w_i) \\
&= \text{Exp}_{c_i}(\mathcal{I}(c_i)) \\
&< \mathcal{H}(c)
\end{aligned}$$

Q.E.D.

References

Fitelson, B. and D. McCarthy (2014). Toward an epistemic foundation for comparative confidence. *Ms.*.

James, W. (1897). The will to believe. In *The Will to Believe and Other Essays in Popular Philosophy*. Longmans Green and Co.

Jeffrey, R. (1983). Bayesianism with a human face. *Testing Scientific Theories, Minnesota Studies in the Philosophy of Science 10*, 133–156.

Jeffrey, R. (1987). Indefinite probability judgment: A reply to levi. *Philosophy of Science 54*(4), 586–591.

Joyce, J. (2005). How probabilities reflect evidence. *Philosophical Perspectives 19*, 153–178.

Joyce, J. (2010). A defense of imprecise credences in inference and decision making. *Philosophical Perspectives 24*, 281–323.

Joyce, J. (2013). Why evidentialists need not worry about the accuracy argument for probabilism. *Ms.*, 1–29.

Joyce, J. M. (2009). Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In F. Huber and C. Schmidt-Petri (Eds.), *Degrees of Belief*, Volume 342. Dordrecht: Springer.

Kaplan, M. (2010). In defense of modest probabilism. *Synthese 176*(1), 41–55.

Levi, I. (1967). *Gambling with Truth*. New York: Knopf.

Levi, I. (1984). *Decisions and Revisions*. Cambridge: Cambridge University Press.

Levi, I. (1991). *The Fixation of Belief and its Undoing: Changing Beliefs through Inquiry*. Cambridge: Cambridge University Press.

Levi, I. (2004). *Mild Contraction: Evaluating Loss of Information due to Loss of Belief*. Oxford: Oxford University Press.

Levi, I. (2012). *Pragmatism and Inquiry: Selected Essays*. Oxford: Oxford University Press.

Mayo-Wilson, C. and G. Wheeler (2014a). Accuracy and imprecision: A mildly immodest proposal. *Ms.*.

Mayo-Wilson, C. and G. Wheeler (2014b). Accuracy without purity. *Ms.*.

Pettigrew, R. (2013). A new epistemic utility argument for the Principal Principle. *Episteme 10*(01), 19–35.

Pettigrew, R. (2014). Accuracy, risk, and the principle of indifference. *Philosophy and Phenomenological Research*, 1–24.

Rott, H. (2006). The Value of Truth and the Value of Information. *Knowledge and inquiry: Essays on the pragmatism of Isaac Levi*, 179.

Schoenfield, M. (2014). The accuracy and rationality of imprecise credences. *Ms.*.

Seidenfeld, T., M. J. Schervish, and J. B. Kadane (2012). Forecasting with imprecise probabilities. *International Journal of Approximate Reasoning 53*, 1248–1261.