Chancy Accuracy and Imprecise Credence

Accuracy-first epistemology is the view that norms of epistemic rationality can be explained in terms of the aim of accurate belief. Some epistemologists hold that in the face of unspecific or ambiguous evidence, rationality requires having imprecise credences. Can this norm be justified in terms of the aim of accuracy? Will our belief states tend to be more accurate if we respond to ambiguous or unspecific evidence with imprecise credences?

There are different possible ways of measuring the accuracy of imprecise credences. And there are different decision rules that might apply when epistemic options are imprecise. Only some accuracy measures, paired with only some decision rules, will recommend imprecision. But vindicating imprecision as such is not the aim. The question is: can we vindicate the *kind* of imprecise credences that are thought to be rationally required?

One common epistemic motivation for imprecise credences is to accommodate the so-called "Chance Grounding Thesis", which says that a rational agent's spread of credence should cover the range of chance hypotheses left open by the agent's evidence. I offer an epistemic utility theoretic argument for chance-grounded imprecise credences. My proposal uses relatively intuitive accuracy measures and decision rules, but reinterprets vindication (ideal accuracy) as chancy instead of alethic.[1]

1 Background

Epistemic utility theory is a form of epistemologized decision theory, wherein the rationality of credal states is determined by their tendency to promote a specifically epistemic form of utility. Epistemic utility is typically claimed to be characterized in terms of gradational accuracy. Accuracy first-epistemology aims to provide a vindication of various epistemic norms as means to the end of accuracy. These norms include probabilism (de Finetti, 1974, Joyce, 1998, 2009, Leitgeb & Pettigrew, 2010), conditionalization (Greaves & Wallace, 2006; Leitgeb & Pettigrew, 2010), versions of the principal principle (Pettigrew, 2012), the principle of indifference (Pettigrew, 2014), and others.

Epistemic utility theory uses the following tools:

[1] Caveat: the author of this paper is skeptical of the claim that imprecise credences are ever epistemically required, and also has reservations about the idea that norms of epistemic rationality derive from their conduciveness to accurate mental states.

- *Credence function:* a function from propositions to real numbers in the unit interval. Call the space of possible credence functions over an algebra $\mathcal{A}$ of propositions $\mathcal{C}$.

- *Vindication function*: a function from propositions to real numbers, relative to which credal states' accuracy is measured. An *alethic* vindication function is a characteristic function of a possible world: it specifies, for every proposition, whether that proposition is true (1) or false (0) at the world.

- For an algebra $\mathcal{A}$ of possible worlds propositions, each w has an alethic vindication function, $v_w : \mathcal{A} \rightarrow \{0, 1\}$. Call the space of alethic vindication functions V.

- A *precise local inaccuracy measure* is a function from credence functions and propositions and vindication functions to real numbers, representing the degree of inaccuracy the credence a credence function lends to a proposition at a vindicator (here, a world):

$$\mathcal{I} : \mathcal{C} \times \mathcal{A} \times V \rightarrow \mathbb{R}_{[0,\infty]}$$

- A *precise global inaccuracy measure* is a function from credence functions and vindication functions to real numbers, representing the total degree of inaccuracy of the credence at a vindicator:

$$\mathcal{I}_g : \mathcal{C} \times V \rightarrow \mathbb{R}_{[0,\infty]}$$

In order to count as an inaccuracy measure, a utility function should meet certain constraints. Here are two candidate weak, intuitive constraints on alethic inaccuracy measures:

Vindication-directedness: For credence functions c and d, if c's credences are at at least as close to v_w as d's and sometimes closer,[2] then $\mathcal{I}_g(c, v_w) < \mathcal{I}_g(d, v_w)$.)

Nontriviality: for all $c, d \in \mathcal{C}$, if $c(A) = v_w(A)$ for all A that c is defined over, and there is some A that c is defined over s.t. $d(A) \neq v_w(A)$, then $\mathcal{I}_g(c, v_w) < \mathcal{I}_g(d, v_w)$.[3]

---

[2] I.e. for all $A \in \mathcal{A}$, $|v_w(A) - c(A)| \leq |v_w(A) - d(A)|$ with inequality between $c(A)$ and $d(A)$ for at least one A. (This is a generalization of Joyce's (2009) truth-directedness. Joyce characterized "at least as close to the truth" as: either $d(A) \geq c(A) \geq v(A)$ or $d(A) \leq c(A) \leq v(A)$. I characterize it in terms of difference because this generalization extends to cases where credences can fall on either side of their vindicator.

[3] This is a generalization of Pettigrew's (2012) Nontriviality constraint. Pettigrew's constraint presupposes that vindicator functions are defined over the same algebras of propositions as the relevant credence functions.

On to imprecise credences.

Many have argued, against traditional Bayesianism, that rational subjects typically have imprecise credences. Instead of sharp, real number values like .4, rational agents have credences that are spread out over multiple real numbers, e.g. intervals like [.2, .8]. This claim is intended to be normative: even idealized agents without our cognitive limitations sometimes shouldn't have sharp credences. Epistemic rationality sometimes demands imprecise credences: they are the appropriate response to evidence that's unspecific or ambiguous.[4]

I'll represent imprecise credences with sets of precise credences.

Some notation:

- I'll use c for (precise) credence functions and C for (precise or imprecise) credal states.
- $C(A) = \{x : c(A) = x \text{ for some } c \in C\}$.[5]

2 Inaccuracy of imprecise credence

2.1 False start

Is there an intuitive notion of the inaccuracy of an imprecise credence? In epistemic utility theory, the notion of gradational accuracy is intuitively glossed as "distance from truth." But how can this extend to imprecise credences?

Consider the following analogy: what's the distance between Candle Cafe (which is situated at a particular street corner) and Central Park (which is extended over many blocks)? If you ask a local, they are likely to tell you: four blocks. Similarly, one might ask: what's the distance between the Metropolitan Museum and Central Park? The natural answer: none. The Metropolitan Museum is in Central Park.

If we take seriously the distance metaphor (and it's a good question whether we should), and interpret it according to ordinary language (also questionable), then the distance between an individual and a set is the minimum distance between the individual and the closest member of the set.

In the case of imprecise credences, then, the inaccuracy of an imprecise credence in a proposition A at a world w will be determined by the upper credence in $C(A)$ if $v_w(A) = 1$ and the lower credence in $C(A)$ if $v_w(A) = 0$.

---

[4] See e.g. Levi (1974, 1980, 1985), Walley (1991), Joyce (2005, 2010), Weatherson (2008), Sturgeon (2008), Hájek & Smithson (2012), and Moss (2014).

[5] Following Joyce, I use C ambiguously as a set of probability functions and a function from propositions to a set of real numbers.

On this interpretation of accuracy, how can an agent avoid possessing accuracy dominated credences? By adopting maximally imprecise credences. The only credence function that is accuracy nondominated will be the credence function such that, for all $A \in \mathcal{A}$ except $\top$ and $\bot$, $C(A) = [0, 1]$.

This generates an unattractive epistemic norm:

Imprecise skepticism. Rationality requires maximally imprecise credence in all contingent propositions.

Obvious objection: shouldn't an epistemic utility theoretic defense of imprecise credences make them appropriately responsive to *evidence*?

Let's suppose an agent in this state sometimes updates on (factive) information, such that (somehow)[6] her credence in the information goes to 1 and its negation to 0. Relative to the space of possibilities compatible with the agent's evidence, assigning a credence other than $[0, 1]$ to some contingent propositions—namely, those the agent learns—will no longer be weakly dominated. And so it'll be rationally permissible to have credence 1 or 0 in some contingent propositions.

Still: there'll be no rational uncertainty state other than a state with 0 as its lower credence and 1 as its upper credence. It will never be permissible to have credence .5, or $[.2, .6]$, or anything other than 0, 1, or $[0, 1]$.

More generally, we might question: does this conception of inaccuracy motivate the right kind of imprecise credences? Can it accommodate the motivations for the view that imprecise credences can be rationally required?

2.2 Epistemically required imprecise credences

A common motivation for rational imprecise credences comes from examples like Mystery Coin:

> **Mystery Coin**
>
> You have a coin that was made at a factory where they can make coins of pretty much any bias. You have no idea whatsoever what bias your coin has. What should your credence be that when you toss the coin, it'll land heads?

The proponent of imprecise credences holds that it's rationally inappropriate to have credence .5 in heads, relative to this background information. Credence .5 is

[6] Note that if she updates in this way, she violates conditionalization, since for every contingent A in $\mathcal{A}$, her prior credal state contained a credence function that assigned A credence 1 and another that assigned it credence 0. So if she conditionalizes all $c \in C$ on incoming evidence, where such an update is defined, her imprecise credal state will still have 0 and 1 as its lower and upper bounds.

appropriate for agents with more evidence about the coin, i.e. agents who know the coin is fair. But with a mystery coin, the proponent of imprecise credences says, the only rational response is to spread your credence in heads over the possible objective chances of heads.

Most proponents of the evidential argument endorse or presuppose something like the Chance Grounding Thesis.[7]

Chance Grounding Thesis. Rational imprecise credences cover the range of evidentially possible chance hypotheses.[8]

My interpretation: for each proposition A, evidence determines some set of possible chances of A at an appropriate time t. A rational agent's upper credence in A will be equal to the upper evidentially possible chance of A at t, and her lower credence in A will be equal to the lower evidentially possible chance of A at t.

Suppose you accept the Chance Grounding Thesis. You might hope for something more than an accuracy-based argument for why imprecise credences are in some circumstances or other rationally required. You might want there to be an accuracy-based argument in favor of the right kind of imprecise credences—namely, those that satisfy the Chance Grounding Thesis.

For example, Konek (forthcoming) argues that, with a certain form of numerical inaccuracy measure for imprecise credences and a certain form of decision rule, imprecise credences are epistemically required. On Konek's view, the inaccuracy of an imprecise credal state is determined by some weighted average of the maximal and minimal inaccuracy of credence functions that are elements of the credal state. (Which weighting? On Konek's view, the weighting captures an agent's degree of epistemic "conservativity": the degree to which they value seeking truth weighed against shunning error.) Konek defends a family of decision rules for imprecise credences, the Hurwicz criteria, which assign options a number on the basis of the weighted average of each option's best and worst possible outcomes. These decision rules differ only in the weightings they use. As with Konek's inaccuracy measure,

7 Most explicitly perhaps in Walley (1991) Joyce (2005, 2010), Hájek & Smithson (2012). Joyce (2010) disputes an aspect of the thesis, but I take him to accept a subtly modified version of the thesis. His caveat: suppose in the mystery coin case, you have additional information that the coin was chosen randomly, and that for every coin with bias β, there was an equal chance of choosing a coin with bias $1 - \beta$. In this case, he accepts that it's rational to have precise credence .5, even though, in once sense, you have no idea what the chances are. It seems to me that you do have chance information: in particular, information that before the coin was selected, the chances of heads was .5. So you have chance information about a different time, when the information was more informative. This will be important later in the paper.

8 Named by (White, 2009).

the weighting in the decision rule is meant to capture the agent's degree of epistemic conservativity.

The predictions of his view make the degree of imprecision that rational credences have a function of how much an agent values proximity to the truth and disvalues distance from error. The view does not take into account the character of an agent's evidence and the extent to which it is informative about objective chances.[9]

What kind of epistemic utility theory can justify imprecise credences that are chance-grounded?

3 First pass proposal

First, we introduce the notion of chancy vindication. Full belief is vindicated by truth; precise credence, I suggest, is vindicated by matching chance.[10] Here are some possible justifications:

First, the notion of vindication is perhaps best understood in terms of objective correctness. Full belief in A is objectively correct iff A is true. Can there be a similar notion of objective correctness for credences? As Hájek (manuscript) argues, objective chance might be the best candidate. Hájek offers a thought experiment: suppose some event E has chance .3 at t. And at t, Avital has credence .3 in E; Ben has credence 1. Then at t' E occurs. Whose credence was more accurate at t? There is a sense, Hájek argues, in which Ben's credence in E was objectively incorrect and Avital's was objectively correct. [11]

Second, on some views of future contingents, future contingent propositions have indeterminate truth values.[12] And on some views of chance, only future contingent propositions are chancy. If these two views are correct, then when a proposition A has nontrivial chance, its truth value is indeterminate. When a proposition has no determinate truth value, it has no determinate alethic inaccuracy; arguably, it has no alethic inaccuracy at all. So the inaccuracy of one's credence in a chance proposition can only be assessed relative to its chance. Its chance is as close as it gets to a truth value.

---

[9] Another possible challenge to Konek's view: it uses different decision rules for priors and subsequent credal states. It's unclear what justifies this distinction—why there should be separate epistemic norms for superbabies. The view I propose does not have this feature, though it may avoid it only by dodging the question of how the space of epistemic possibilities contracts.

[10] This hypothesis is endorsed, for different purposes, by Pettigrew (2012).

[11] Joyce suggests that this view is prima facie plausible in his (1998 and 2009): "Having a credence far from a proposition's objective chance seems like a defect even if that credence is close to the proposition's truth-value" (2009, 274).

[12] E.g. Belnap & Green 1994.

Finally, there are possible cases where an agent's precise credence in A affects the chance of A such that they match. (See Greaves 2014, Carr manuscript). Truth-directedness dictates that the agent, informed of this situation, must have credence 1 or 0 in the proposition to avoid accuracy domination. But many find it intuitive that in these circumstances, any credence in A is rationally permissible. We can accommodate this intuition in accuracy first epistemology by allowing that all are equally accurate—but in that case, we are measuring accuracy as distance from chance, not distance from truth.

My suggestion: instead of characteristic functions of worlds, our vindication functions should be chance functions. That is: we should evaluate credences' accuracy relative to chances.

3.1 Chancy inaccuracy

I assume that $\mathcal{I}$ satisfies chancy versions of Truth-Directedness and Nontriviality, where the vindication functions are chances.

Local imprecise inaccuracy will be set-valued:

$$\mathcal{I}(C, A, ch) = \{\mathcal{I}(c, A, ch) : c \in C\}$$

From here, we can define the upper and lower inaccuracy of an imprecise credence:

$$\mathcal{I}^-(C, A, ch) = \min_{c \in C} \mathcal{I}(c, A, ch)$$
$$\mathcal{I}^+(C, A, ch) = \max_{c \in C} \mathcal{I}(c, A, ch)$$

Global imprecise inaccuracy: the set of global precise inaccuracies.

$$\mathcal{I}_g(c, ch) = \sum_{A \in \mathcal{A}} \mathcal{I}(c, A, ch)$$
$$\mathcal{I}_g^-(C, ch) = \min_{c \in C} \mathcal{I}_g(c, ch)$$
$$\mathcal{I}_g^+(C, ch) = \max_{c \in C} \mathcal{I}_g(c, ch)$$

3.2 Evidence-relativity

I'm not going to provide an update rule. I'll just assume that evidence eliminates chance functions from the state space.[13]

[13] One might wonder how evidence could do this: how one could ever learn with certainty any facts about non-trivial chances. The possibility (and perhaps typicality) of doing so seems to be a commitment of the imprecise credence view that accepts any version of the Chance Grounding Thesis.

I stipulate for the decision rules that the state space relevant for our epistemic decision rules is the space of evidentially possible chance functions E, rather than the total space of possible chance functions.[14]

3.3 Decision rule

Now, we can formulate decision rules for pursuing accuracy with imprecise credences. First candidate:

Lower-Dominance Avoidance: it is irrational to adopt a credal state that is weakly lower-dominated.

C weakly lower-dominates D iff at each $ch \in E$, $\mathcal{I}_g^-(C, ch) \leq \mathcal{I}_g^-(D, ch)$ and for some $ch \in E$, $\mathcal{I}_g^-(C, ch) < \mathcal{I}_g^-(D, ch)$

Any state that contains all evidentially possible chance functions will have minimal lower inaccuracy at every $ch \in E$. Lower dominance avoidance ensures that rational credal states contain all evidentially possible chance hypotheses.

But some lower-inaccuracy nondominated credal states exhibit more imprecision than is evidentially appropriate. So lower dominance avoidance permits ignoring information by, e.g., having maximally imprecise credences even where evidence entails chance information. When the evidence is informative about chance, a rational agent shouldn't wallow in uncertainty. In short, lower dominance avoidance is too permissive.

To reflect this intuition (such as it is), we add a second (sub-)rule:

Upper-Dominance Avoidance: it is irrational to adopt a credal state that is weakly upper-dominated.

C weakly upper-dominates D iff at each $ch \in E$, $\mathcal{I}_g^+(C, ch) \leq \mathcal{I}_g^+(D, ch)$ and for some $ch \in E$, $\mathcal{I}_g^+(C, ch) < \mathcal{I}_g^+(D, ch)$

Adding upper dominance avoidance ensures that credal states are no more imprecise than the Chance Grounding Thesis requires.

And so we arrive at a lexicographic decision rule:

Lower-Then-Upper Dominance Avoidance: it is irrational to adopt D if either

(i) D is lower-dominated, or

(ii) D is upper-dominated by any credal state that is not lower-dominated.

[14] This sort of stipulation is standard in practical decision theory but not epistemic decision theory.

3.4 Toy case

Let's walk through how this proposal handles a simplified version of the mystery coin case. Suppose the coin could be fair, and could have a bias .2 or .8 toward heads. All other chance hypotheses are incompatible with total evidence.

	ch_1	ch_2	ch_3
H	.2	.5	.8
T	.8	.5	.2

Now, let's define an agent's upper and lower credence in a proposition:

$$\overline{c}(A) = \max_{c \in C} c(A)$$
$$\underline{c}(A) = \min_{c \in C} c(A)$$

The Chance Grounding Thesis entails that the rational credal state contains credence functions equivalent to ch_1, ch_2, and ch_3[15] and is such that $\overline{c}(H) = \overline{c}(T) = .8$ and $\underline{c}(H) = \underline{c}(T) = .2$. Call one such state C.[16]

$$\mathcal{I}_g^-(C, ch_1) = \mathcal{I}_g^-(C, ch_2) = \mathcal{I}_g^-(C, ch_3) = 0$$

C will have minimal lower inaccuracy, and hence will be lower-non-dominated. Any credal state D that doesn't include all epistemically possible chance functions will have positive inaccuracy at some possible chance function. So according to Nontriviality, its inaccuracy will not be minimal.

Now consider the credal states that are lower-nondominated. Suppose some credal state D contains ch_1, ch_2, and ch_3, but is such that $\underline{d}(H) < \underline{c}(H)$. Then its local upper inaccuracy for H is greater than C's at ch_3 and ch_2. What about ch_1?

- Suppose $\mathcal{I}^+(\underline{d}, H, ch_1) \leq \mathcal{I}^+(ch_3, H, ch_1)$. Because D is lower nondominated, it contains ch_3, and so its upper inaccuracy at ch_1 is the same as C's.
- Suppose $\mathcal{I}^+(\underline{d}, H, ch_1) > \mathcal{I}^+(ch_3, H, ch_1)$. Then D's upper local inaccuracy is greater than C's upper local inaccuracy.

15 Henceforth I'll assume for simplicity that all ch are only defined over propositions that the relevant credence functions are also defined over.

16 I don't take a stand on the structure of imprecise credences—whether they're determined by upper and lower credences in each proposition or whether they're more structured. Depending on how we model imprecise credences, the Chance Grounding Thesis may be interpreted as permissive or impermissive.

So relative to all three chance functions, $\mathcal{I}^+(D, H, ch_i) \geq \mathcal{I}^+(C, H, ch_i)$, with inequality in at least two cases.

For the same reasons, D's local upper inaccuracy for T can't be lower than C's, hence can't offset D's local upper inaccuracy for H in the computation of global upper inaccuracy. So D is weakly upper-dominated, and therefore inadmissible according to the decision rule under discussion. (Mutatis mutandis for $\overline{c}(H), \underline{c}(T), \overline{c}(T)$.) Of the lower-nondominated credal states, C's upper inaccuracy is minimal.

In general, a credal state satisfies lower-then-upper dominance avoidance iff it also satisfies the Chance Grounding Thesis.

Note: with this toy case, we can also illustrate why, in order to predict the Chance Grounding Thesis, we need a lexicographic rule (lower-then-upper dominance, in that order) rather than simply the conjunction of lower and upper dominance. Suppose we apply upper dominance first. Then a precise credence function equal to ch_2 would upper-dominate C. Indeed, all imprecise credences would be upper-dominated by the precise credence function determined by the midpoints between the upper and lower evidentially possible chances for each proposition.

4 Complications

4.1 Evidence about different times

> **Problem #1**
>
> Suppose at t_1 you know that a fair coin will be tossed out of your view. Later, at t_2, you know that $ch(H) = 0$ or 1, but you don't know which.

Lower-then-upper dominance avoidance apparently requires adopting a credal state at t_2 such that $\underline{c}(H) = 0$ and $\overline{c}(H) = 1$. But intuitively, credence .5 in H is rational.

Fans of imprecise credence treat imprecision as constituting, in some sense, *more uncertainty* than middling precise credence. But just by learning that time has passed, you haven't lost any information. So you shouldn't be any more uncertain.

> **Problem #2**
>
> Suppose an oracle has told you that a fair coin will land heads at t_2. At t_1 (now), the chance of heads is .5, but you have "inadmissible" evidence—that is, evidence that allows you to violate the Principal Principle.

Lower-then-upper dominance avoidance apparently require having $c(H) = .5$, ignoring your oracular evidence. One might worry that this challenges the very idea of using chance functions as vindication functions, rather than using characteristic functions of worlds. (Or one might worry that this challenges the Chance Grounding Thesis as a view about why or how credences should be imprecise.)

4.2 Final proposal

Problems #1 and #2 suggest that chances change over time—unlike characteristic functions of worlds.[17] Our vindication functions, therefore, should instead be sequences of chances $\mathbf{ch} = \langle ch_{t_1}, ch_{t_2}, \ldots \rangle$, indexed to times. Moreover, we cannot avoid the problem of changing chances merely by confining our assessments of credal states' accuracy to times. Problems #1 and #2 show that rational credences are sometimes sensitive to chances at other times: the past in problem #1, the future in problem #2.

$\mathcal{I}$ should be time-relativized: $\mathcal{I}(c, A, \mathbf{ch}, t)$. We confine lower- and upper-dominance avoidance to shared times:

C *weakly lower-dominates* D relative to t_i iff at each $\mathbf{ch} \in E$, $\mathcal{I}_g^-(C, \mathbf{ch}, t_i) \leq \mathcal{I}_g^-(D, \mathbf{ch}, t_i)$ and for some $\mathbf{ch} \in E$, $\mathcal{I}_g^-(C, \mathbf{ch}, t_i) < \mathcal{I}_g^-(D, \mathbf{ch}, t_i)$

C *weakly upper-dominates* D relative to t_i iff at each $\mathbf{ch} \in E$, $\mathcal{I}_g^+(C, \mathbf{ch}, t_i) \leq \mathcal{I}_g^+(D, \mathbf{ch}, t_i)$ and for some $\mathbf{ch} \in E$, $\mathcal{I}_g^+(C, \mathbf{ch}, t_i) < \mathcal{I}_g^+(D, \mathbf{ch}, t_i)$

These two rules determine a set of credal-state–time pairs: the credal states that are lower-then-upper nondominated relative to particular times. (Note that the relevant times are not the times when the credal states are held, but the times when the chances obtain relative to which the credal states' inaccuracy is assessed.)

We introduce a third (sub-)rule:

Cross-Temporal-Upper-Dominance Avoidance: it is irrational to adopt a credal state that is weakly cross-temporally upper-dominated.

$\langle C, t_i \rangle$ *weakly crosstemporally upper-dominates* $\langle D, t_j \rangle$ iff at each $\mathbf{ch} \in E$, $\mathcal{I}_g^+(C, \mathbf{ch}, t_i) \leq \mathcal{I}_g^+(D, \mathbf{ch}, t_j)$ and for some $\mathbf{ch} \in E$, $\mathcal{I}_g^+(C, \mathbf{ch}, t_i) < \mathcal{I}_g^+(D, \mathbf{ch}, t_j)$

Our new sequential rule—lower-then-upper-then-CT-upper dominance avoidance(!)—in effect says: your credences should be regulated by the chances that you're *most informed* about. If your evidence is most informative about the chances at t_n, then your credence should be spread over the chance hypotheses that your evidence treats as possible chances at t_n.

This generates the intended result for **problem #1**: credence $[0, 1]$, relative to t_1, has a higher upper inaccuracy than credence .5 has relative to t_2 at all $\mathbf{ch} \in E$.

[17] I aim to be as neutral as possible about the metaphysics of chances. If chances don't change over time, then neither of the problems in the previous section poses a challenge to lower-then-upper dominance avoidance as a satisfactory decision rule for imprecise credences.

We can represent chance (and credence) functions as $\langle ch(H), ch(T)\rangle$ sequences and points of evaluation as $\langle ch_{t_1}, ch_{t_2}\rangle$ sequences.

$\mathbf{ch}_a = \langle\langle .5, .5\rangle, \langle 1, 0\rangle\rangle =$ the chance sequence where the coin lands heads
$\mathbf{ch}_b = \langle\langle .5, .5\rangle, \langle 0, 1\rangle\rangle =$ the chance sequence where the coin lands tails

Then it's easy to see that, given Nontriviality, $\langle\langle .5, .5\rangle, t_1\rangle$ dominates $\langle\langle [0,1], [0,1]\rangle, t_2\rangle$:

	$\mathbf{ch}_a$	$\mathbf{ch}_b$
$\langle\langle [0,1], [0,1]\rangle, t_2\rangle$	not minimal	not minimal
$\langle\langle .5, .5\rangle, t_1\rangle$	minimal	minimal

But our decision rule doesn't help with **problem #2**: there, two different credal states satisfy our decision rule. But intuitively, one is rationally required and the other impermissible.

The evidence provided by the oracle entails that $ch_{t_2}(H) = 1$. But the fairness of the coin entails that $ch_{t_1}(H) = .5$. $c_1 = ch_{t_1}$ and $c_2 = ch_{t_2}$ are both minimally upper inaccurate (relative to different times). Intuitively, c_2 is rationally required.

So a final tweak is required: we should modify cross-temporal upper-dominance avoidance.

Cross-Temporal Upper Dominance Avoidance*: choose the credal state that is cross-temporally upper-nondominated relative to the latest time.[18]

Our resulting lexicographic decision rule:

Lower-then-upper-then-CT*-upper-dominance avoidance. Apply lower-dominance avoidance, then upper dominance avoidance, then Cross-Temporal Upper Dominance Avoidance*.

What is the accuracy-based justification for choosing accuracy relative to later chances? Here's a suggestion: the direction of time is such that later chances have higher alethic accuracy (or at least expected alethic accuracy) than earlier chances. At the hypothetical end of time, the final chance function will be a characteristic function of a world.

[18] That is, latest of all the times relative to which any credence function is cross-temporally upper non-dominated within the relevant space of evidentially possible chance sequences.

5 Further challenges

Objection 1. Schoenfield (2015) proves, with minimal and plausible assumptions, that however we characterize the inaccuracy of imprecise credences—numerical or otherwise—for every imprecise credal state C, there is a precise credence c such that, at every possible world, the inaccuracy of c is not less than the inaccuracy of C. And so, Schoenfield concludes, there can be no accuracy-based argument for rationally required imprecise credences. Schoenfield's argument does not depend on the claim that the vindicators of credal states are alethic. So why does her proof not discredit your argument for rationally required imprecise credences?

Reply 1. I think Schoenfield is entirely correct that we cannot justify a requirement for imprecise credences in terms of "the inaccuracy of" a credal state. But I resist the conclusion that no accuracy-based argument for imprecise credences can be given. In place of the quantity that Schoenfield and others call "the inaccuracy of" an imprecise credal state—with its uniqueness presupposition intact—I suggest that imprecise credences have multiple accuracy properties—in particular upper and lower inaccuracy. And if we can adequately motivate decision rules that take these accuracy properties into consideration and that sometimes require imprecise credences, then this argument seems to me properly accuracy-based.

An analogy: suppose you've moved to a new city to work at a local university, and you're deciding where to rent an apartment. Suppose all that matters to you in an apartment is distance from campus. But the campus is large and you don't know which buildings you'll be teaching in over the years. You might care only about an apartment's distance from your office. Or you might care about the apartment's average distance from all the buildings on campus. Or you might care only about the distance between an apartment and the nearest point to it on campus. (Maybe there's an underground tunnel system on the campus that will shield you from the winter weather.) Or maybe you care about the distance between an apartment and the farthest point on campus. (Maybe the risk of ever having a 45-minute commute is a dealbreaker for you.)

Clearly, there are a number of distance properties that are relevant for characterizing "the distance between" an apartment and the large university, and that are potentially relevant for your choice. As with extended objects, so with extended credences. The notion of gradational accuracy is most intuitive when it is characterized as "distance" from the truth or from vindication. And so, I suggest, when credal states are imprecise, a number of accuracy properties are potentially relevant for epistemic utility theory.

Objection 2. Why treat credences as vindicated by chance? It's (arguably) part of the concept of belief that it aims for truth. The idea that credence aims for chance doesn't have any similar intuitive support.

Rather, the claim that credences are vindicated by chance seems tantamount to treating credences as full beliefs about propositions concerning chance—commitments which are *alethically* vindicated when the right chances obtain.

Reply 2. This is a problem for my proposal only inasmuch as it's a problem for anything like the Chance Grounding Thesis. The Chance Grounding Thesis has the same commitments.

Objection 3. Suppose you receive testimony from someone you believe to be 90% reliable that a mystery coin landed heads. What credence should your credence be that the coin landed heads? You can't rule out worlds where the coin landed tails; so according to the decision rules defended above, your lower credence in heads should be zero. But then your credence is not more informative than before you received the testimonial evidence, when the coin's outcome is still a mystery. So your credence fails to reflect the testimonial evidence you've received.

Reply 3. Same as reply 2.

Objection 4. The great thing about Joyce's (1998; 2009) accuracy-based argument for probabilism: by contrast to (de Finetti, 1974), the only restrictions Joyce placed on epistemic utility functions were (arguably) intuitively justified—as capturing something about the concept of belief or any measure that could plausibly be characterized as "distance from truth." And the only decision rule he needs is accuracy dominance avoidance.

Your proposal, by contrast, uses various notions the inaccuracy of imprecise credence. And your decision rule is more complex. One can make different stipulations—with equally plausible arguments—for inaccuracy measures and decision rules that generate very different epistemic norms.

Reply 4. Yes! For example: suppose we accept certain of Joyce's constraints on inaccuracy measures for precise credences (Truth-Directedness, Extensionality, Convexity.) Then why not let $\mathcal{I}(C, A, w)$ be the mean of $\mathcal{I}^{-}(C, A, w)$ and $\mathcal{I}^{+}(C, A, w)$? Then every imprecise credence function will be dominated by a precise credence function, e.g. the function determined by the midpoints of C. So all admissible credences are precise. And so we have an easy argument to the conclusion that, in fact, imprecise credences are never rationally permissible, let alone mandatory.

For my part, I'm skeptical that any intuitive notion of numerical inaccuracy of imprecise credences can be given. As for imprecise inaccuracy, I'm skeptical that any intuitive decision rules can be given for rational epistemic choice among options with imprecise epistemic utility. Stipulations are evidently inevitable.

If we stipulate that our vindication functions are chance functions, then at least we can get by with very weak stipulations about precise inaccuracy (Chance-Directedness and Nontriviality), imprecise inaccuracy (namely, that lower (upper) global inaccuracy of a credal state is determined by the lower (upper) global inaccuracy of precise credences within the credal state), and decision rules for imprecise inaccuracy (generalization of dominance: lower-then-upper dominance secures the Chance Grounding Thesis; further decision rules refine it).

So, this objection is correct. But it's only a problem for the present account to the extent that it's a problem for the very idea of rational imprecise credences.[19]

References

de Finetti, Bruno (1974). *Theory of probability*. New York: John Wiley.

Greaves, Hilary & David Wallace (2006). "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115(459), 607–632.

Hájek, Alan (manuscript). "A Puzzle about Degree of Belief."

Hájek, Alan & Michael Smithson (2012). "Rationality and Indeterminate Probabilities." *Synthèse* 187(1), 33–48.

Joyce, James (1998). "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65(4), 575–603.

Joyce, James (2009). "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In *Degrees of belief*, vol. 342, 263–297. Synthèse.

Joyce, James M. (2005). "How Probabilities Reflect Evidence." *Philosophical Perspectives* 19(1), 153–8211.

Joyce, James M. (2010). "A Defense of Imprecise Credences in Inference and Decision Making." *Philosophical Perspectives* 24(1), 281–323.

Konek, Jason (forthcoming). "Epistemic Conservativity and Imprecise Credence." *Philosophy and Phenomenological Research* .

Leitgeb, Hannes & Richard Pettigrew (2010). "An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy." *Philosophy of Science* 77(2), 236–272.

Levi, Isaac (1974). "On Indeterminate Probabilities." *Journal of Philosophy* 71(13), 391–418.

[19] Thanks to Luke Fenton-Glynn, Pavel Janda, Jim Joyce, Jason Konek, Miriam Schoenfield, Jonathan Weisberg, Greg Wheeler, Robbie Williams, and audiences at the Accuracy and Imprecise Credence workshop at Bristol and Jonathan Weisberg's graduate seminar at the University of Toronto. The research for this paper has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 312938.

Levi, Isaac (1980). *The enterprise of knowledge: An essay on knowledge, credal probability, and chance*. Cambridge, MA: MIT Press.

Levi, Isaac (1985). "Imprecision and Indeterminacy in Probability Judgment." *Philosophy of Science* 52(3), 390–409.

Moss, Sarah (2014). "Credal Dilemmas." *Noûs* .

Pettigrew, Richard (2012). "Accuracy, Chance, and the Principal Principle." *Philosophical Review* 121(2), 241–275.

Pettigrew, Richard (2014). "Accuracy, Risk, and the Principle of Indifference." *Philosophy and Phenomenological Research* 91(3),

Schoenfield, Miriam (2015). "The Accuracy and Rationality of Imprecise Credences." *Noûs* 49(4), n/a–n/a.

Sturgeon, Scott (2008). "Reason and the Grain of Belief." *Noûs* 42(1), 139–165.

Walley, Peter (1991). *Statistical reasoning with imprecise probabilities*. Chapman & Hall.

Weatherson, Brian (2008). "Decision Making with Imprecise Probabilities." (Unpublished).

White, Roger (2009). "Evidential Symmetry and Mushy Credence." In T. Szabo Gendler & J. Hawthorne (eds.), *Oxford studies in epistemology*, vol. 3, 161–186. Oxford University Press.