Deference Principles for Imprecise Credences

Giacomo Molinari

April 19, 2024

Abstract

This essay gives an account of epistemic deference for agents with imprecise credences. I look at the two main imprecise deference principles in the literature, known as Identity Reflection and Pointwise Reflection (Moss, 2021). I show that Pointwise Reflection is strictly weaker than Identity Reflection, and argue that, if you are certain you will update by conditionalisation, you should defer to your future self according to Identity Reflection. Then I give a more general justification for Pointwise and Identity Reflection from the assumption that you defer to someone whenever you consider their doxastic state to be better than yours, in the sense of leading to better decision-making.

1 Overview

Deference plays a central role in our epistemic life. As members of an epistemic community we acquire many of our beliefs through the testimony of trusted informers and experts. And even as individuals, our beliefs are constrained by deference relations: we ought to defer to the objective chances, or to the opinions best supported by our total evidence. Any epistemological theory is incomplete without an account of deference.

Orthodox Bayesian epistemology has been very successful in providing formal characterisations of deference relations, called *deference principles*, which have been used to articulate and defend various epistemic norms (Van Fraassen, 1984; Lewis, 1980; Elga, 2013). But the orthodox Bayesian approach assumes that a rational agent's credences are precise, i.e. that they can be adequately captured by a precise probability function. Imprecise Bayesians deny this assumption. They seek to develop an epistemological theory that allows (and for some, even demands) rational agents to have imprecise credences.[1]

Can Imprecise Bayesians give a convincing treatment of deference? Some have argued that they cannot (White, 2010; Topey, 2012). These arguments show that once we allow imprecise credences, we must give up some very plausible intuitions about deference to our future selves. Responses have tended to deny some of the arguments' premises, but

[1]Early versions of this view have been articulated by Keynes (1921), Smith (1961), Levi (1974), and Williams (1976). See Bradley (2019) and Mahtani (2019) for recent overviews of imprecise credences in epistemology.

they have not provided a systematic treatment of deference analogous to the orthodox Bayesian one. That will be the aim of this essay: to lay the foundations of an account of epistemic deference that allows for imprecise credences.

Here is the essay plan. I start (Section 2) by giving an overview of the various contexts where deference principles have been fruitfully applied in Orthodox Bayesian epistemology. In particular, a popular principle called Global Reflection has been applied in all those contexts. Section 3 looks at the existing literature on deference principles for imprecise probabilities. Here I outline the argument that Imprecise Bayesians cannot give a convincing account of deference to one's future self, and look at two deference principles that have been introduced in response to this argument, called Identity Reflection and Pointwise Reflection. I also propose a small tweak of Pointwise Reflection. In Section 4 I highlight an important assumption that is needed to justify all deference principles discussed in this essay, both precise and imprecise. It is often called "Immodesty" in the literature, but I will call it "Clarity" here. While there are criticisms of Clarity, I will assume it throughout this essay. Section 5 explores whether Identity and Pointwise Reflection vindicate our intuitions about deference to one's future self, and how they are related to one another. Here I show that Pointwise Reflection is equivalent to one part of Identity Reflection for a certain class of imprecise credences. I argue that we should think of it as a local consequence of the latter principle, rather than as an alternative to it. In Section 6 I argue that Identity and Pointwise Reflection can characterise deference more generally than just in cases where you defer to your future self. This argument is based on the assumption that you defer to someone when you consider their doxastic state to be better than yours, in the sense of leading to better decision-making. Section 7 sums up the main conclusions.

2 The Role of Deference

At the most general level, a *deference principle* is a relationship between two doxastic states.[2] Deference principles can be used to characterise the relationship between two doxastic states when one considers the other to be worthy of deference, in the sense of being a trusted source of opinions and judgements. I will refer to an agent/state that is considered worthy of deference by some other agent/state as an *epistemic superior*, where the term is intended to be inclusive of epistemic peers.

Within precise Bayesian epistemology, there are a number of different contexts where some notion of deference is relevant:

1. (Deference to future self). If you are certain that you'll remain rational between now and tomorrow, you should regard your future self as epistemically superior to you, since they have access to at least as much evidence as you do, and possibly more. Deference

[2] I will use the terms "doxastic state" and "credences" very loosely here to refer to anything that encodes probabilities, estimates, comparative likelihood judgements, and the like. In particular, a doxastic state need not depend on some agent's mental state.

principles have been used to detail how a rational agent's credences at some time t relate to that agent's credences at some future time t' (Van Fraassen, 1984; Briggs, 2009).

2. (Characterise genuine learning). Intuitively, genuine learning is something that leads you to have epistemically superior opinions to the ones you had before. So deference principles have been used to provide necessary conditions for a credal update to be considered a genuine learning experience (Huttegger, 2014; Skyrms, 1990).
3. (Characterise external expertise). There are many agents beyond ourselves that we regard as experts, relying on their opinions to form our own. Deference principles have been used to detail the relationship that holds between your credences and those of someone you regard as an expert (Pettigrew and Titelbaum, 2014; Dorst et al., 2021).
4. (Rationality principles). Deference principles have been used to express the requirement that you regard the objective chances as an epistemic superior (Lewis, 1980; Hall, 1994). Many epistemologists also think that, given some facts about you (e.g. your total evidence, your epistemic values, etc.) there is a set of opinions you are rationally required to have. Intuitively, you are rationally required to regard these opinions as worthy of deference. Deference principles have thus been used to specify the rationality requirement that you should defer to the opinions that are rationally required of you (Elga, 2013; Dorst, 2020).

Although we could use a different deference principle (i.e. a different notion of deference) in each of these contexts, it is an interesting *result* of precise Bayesian epistemology that certain deference principles can be fruitfully applied to play all of these roles. Perhaps the most famous such principle is known as the (Global) Reflection principle. Before we can define it, it's useful to introduce some notation.

Denote by $\mathcal{W}$ a finite set of worlds, and $\mathcal{F}_{\mathcal{W}}$ the set of propositions induced by $\mathcal{W}$ (i.e. the set of all subsets of $\mathcal{W}$). We represent a doxastic state with a *credence function*: a function $\pi : \mathcal{F}_{\mathcal{W}} \to [0, 1]$ which assigns to each proposition a real value, interpreted as the agent's degree of belief in that proposition. In particular, I will focus on coherent doxastic states, whose corresponding credence functions are probability functions. Denote by $\mathcal{P}_{\mathcal{W}}$ the set of all probability functions over $\mathcal{W}$. When $X : \mathcal{W} \to \mathbb{R}$ is a random variable, I will write $\pi(X)$ as shorthand for the expectation of X under π, i.e. $\sum_{w_i \in \mathcal{W}} \pi(\{w_i\})X(w_i)$, and I will write $\pi(w_i)$ as shorthand for $\pi(\{w_i\})$.

A *definite description* of a credence function is a function $p : \mathcal{W} \to \mathcal{P}_{\mathcal{W}}$ that assigns a credence function to each possible world. If p is a definite description of a credence function and w_i is a world, I write p_i as shorthand for the credence function $p(w_i)$. For any definite description of a credence function p and any property ϕ of a credence function, I write $[\phi(p)]$ for the proposition that p has property ϕ (more precisely: that the credence function denoted by p has property ϕ). This proposition corresponds to the set of worlds $\{w_i : \phi(p_i) \text{ is true}\}$. For example, for any proposition a and $s \in \mathbb{R}$, $[p(a) = s]$ is the proposition $\{w_i : p_i(a) = s\}$. To help distinguish between definite descriptions of credence functions and rigidly designated credence functions, I will use latin letters (p, q) to denote the former, and greek letters (π, γ) to denote the latter.

We are now ready to give a definition of the Global Reflection principle:

- **Global Reflection**:

$$p(\cdot|[q = \gamma]) = \gamma \tag{1}$$

  for every credence function γ such that the above conditional credence function is defined.

The letters p and q are definite descriptions of credence functions, whereas γ is a rigid designator of a credence function, i.e. it denotes the same credence function at every possible world. $[q = \gamma]$ is the proposition that q denotes the function γ. This is just the set of all worlds w_i such that $q_i = \gamma$. The principle says that the credence function denoted by p, conditional on the proposition that $q = \gamma$, is equal to γ. So the principle is respected at a world w_i iff $p_i(\cdot|[q = \gamma]) = \gamma$ for every γ such that $\pi([q = \gamma]) > 0$. When the principle is respected at every possible world, we say that p *Global Reflects* q. We will normally assume that p is constant across the possible worlds, in which case the principle becomes:

$$\pi(\cdot|[q = \gamma]) = \gamma \tag{2}$$

where π is a rigid designator of the deferring credence function.

Throughout the essay I will assume that π is *regular*: for every $w \in \mathcal{W}$, $\pi(w) > 0$. I think of this as a technical assumption that considerably simplifies many of the proofs in this chapter. In particular, regularity is helpful when comparing the relative strength of different deference principles, since it ensures that all conditional credences involved in these principles are defined whenever the conditioning proposition is not empty.

Global Reflection has been used to characterise the relevant notion of deference in all four contexts listed above. Let π denote your credence function at present time. For context (1), let q denote your credences at some future time, about which you are currently uncertain. If q is obtained by conditioning your current credence function on the true element of a partition, then π Global Reflects q (Briggs, 2009; Weisberg, 2007).

For context (2), let q denote the credences you would have after undergoing some update whose outcome you are currently uncertain about. Under an assumption which I call Clarity, and which I discuss at greater length in Section 4, Huttegger (2014) has shown that you expect this update to be pragmatically valuable, in the sense of leading to better decision-making, if and only if π Global Reflects q. Huttegger argues that this sort of pragmatic value is necessary for an update to count as genuine learning, and thus that Global Reflection is necessary for genuine learning.

For (3), let q denote the credence function of some other agent whose opinions you are uncertain about. The fact that conditional on this agent having a certain set of opinions, you would adopt the same opinions, is a natural way to formalise the fact that you treat this agent as an epistemic expert (although as we shall see, this view has been subject to some criticism (Hall, 1994; Pettigrew and Titelbaum, 2014; Dorst et al., 2021)).

For (4), let q denote the credence function which is rational for you to adopt in light of your total evidence. Then we can use Global Reflection to express a rationality requirement: conditional on your total evidence warranting some credence function γ, your credence function should be γ. This constraint has interesting consequences for how your credences towards any ("first-order") proposition a should relate to your credences towards "higher order" propositions about the rationality of your credences (Elga, 2013; Titelbaum, 2015; Dorst, 2019).

By applying the same principle in all four roles we can make interesting connections between the relevant notions of deference in each context. Because (3) Global Reflection is a natural way to cash out the idea that you treat some agent as an expert, (2) shows that your update is valuable in terms of decision-making if and only if you should treat your updated self as you would treat an external expert, and (1) shows that conditioning on the true element of a partition is valuable this way. And the claim (4) that you should Global Reflect the credence function warranted by your total evidence is similarly motivated by the idea that you should treat your evidence as an epistemic expert.

3 Deference and Imprecise Credences

A number of theorists have argued that credence functions are inadequate as representations of a (rational) agent's doxastic state. In particular, they argue that sometimes rational agents have *imprecise credences*, which cannot be satisfactorily represented by a single probability function.[3] For this reason, they introduce a class of more complex mathematical objects to do the job. Examples include, but are not limited to: lower previsions (Walley, 1991), sets of probability functions (Levi, 1980; Joyce, 2010), sets of random variables (Quaeghebeur, 2014), and sets of sets of probability functions (Moss, 2018; Campbell-Moore, 2021). I refer to these theorists as Imprecise Bayesians.

In this essay I will focus on sets of credence functions as representations of imprecise doxastic states. Following (Levi, 1980) I call these *credal sets* (although note that, unlike Levi and other authors, I do not assume that credal sets are convex). So a credal set on $\mathcal{W}$ is just a set Π probability functions on $\mathcal{W}$. If Π is a credal set and $a \subseteq \mathcal{W}$ is a proposition, I write $\Pi(a)$ for the *value set* that Π assigns to a, defined by $\Pi(a) = \{\pi(a) : \pi \in \Pi\}$. Just like for credence functions, we can define definite descriptions of credal sets as functions $P : \mathcal{W} \to 2^{\mathcal{P}_{\mathcal{W}}}$ mapping worlds to credal sets. As in the precise case, I write P_i as shorthand for the credal set $P(w_i)$, and denote by $[\phi(P)]$ the proposition $\{w_i \in \mathcal{W} : \phi(P_i) \text{ is true}\}$.

I will interpret credal set *supervaluationally*. For example, if a and b are propositions, you judge a to be more probable than b iff for every credence function π in your credal set, $\pi(a) > \pi(b)$. More generally, we can think of any *probabilistic judgement* as a set of probability functions. For example, the judgement that a is between 50% and 80% likely corresponds to the set $\{\pi \in \mathcal{P}_{\mathcal{W}} : .5 \leq \pi(a) \leq .8\}$, and the judgement that a is just as

[3] See Bradley (2019) and Mahtani (2019) for recent overviews of these arguments.

likely as b correponds to the set $\{\pi \in \mathcal{P}_{\mathcal{W}} : \pi(a) = \pi(b)\}$. Then an agent with credal set Π makes a probabilistic judgement iff Π is contained in the corresponding set of probability functions. That is, iff every probability function in Π makes that judgement.

Despite the fact that epistemic deference is as relevant for precise Bayesians as for their Imprecise counterparts, deference principles have received fairly limited attention in the Imprecise Bayesian literature. The only exception I am aware of involves two deference principles that have been introduced in response to an argument by Roger White (2010), who criticises imprecise credences on the grounds that they violate some plausible intuitions about (1) deference to one's future self.

This section will briefly outline White's argument and define the two imprecise deference principles introduced in response to it: Identity Reflection and Pointwise Reflection. I will also argue for a slight modification of Pointwise Reflection. The rest of the essay will clarify the relationship between these two principles, and show whether they can characterise the relevant notion of epistemic superiority in contexts (1-4).

3.1 Identity and Pointwise Reflection

Here is a quick overview of White's (2010) argument against imprecise credences.[4] Denote your credal set at present time t by Π, and denote your credal set at some future time t' by Q. Intuitively, if you knew that you will rationally come to have attitude ϕ towards some proposition at some future time t', after you learn some more (non-misleading) evidence, then you should have attitude ϕ towards that proposition *now*. In the precise case, this is captured by the following local version of Reflection:

- **Local Reflection**:

$$\pi(a|[q(a) = r]) = r \tag{3}$$

  for every proposition a and real number $r \in \mathbb{R}$ such that the above conditional credence is defined.

Briggs (2009) has shown that if you are currently certain that your future credence q is obtained by conditionalising your current credence π on the true element of a partition, then you Local Reflect q. In fact, we shall see that you also Global Reflect q. This shows that precise credences satisfy our starting intuition: conditional on your more-informed future self having credence r in a, you now should have credence r in a.

White argues that a popular way of updating imprecise credences, known as *pointwise conditionalisation*, violates this intuition about how your current and future attitudes ought to be related.[5] In particular, suppose that Q is obtained by (pointwise) conditionalising your initial credence Π on the true element of a partition. It may be the case that

[4]Note that this just one of many arguments against imprecise credences given by White (2010).

[5]Pointwise conditionalisation is at least the standard update rule for closed and convex credal sets, which suffices to run White's argument. See Walley (1991) for a defense of pointwise conditionalisation for closed and convex credal sets, and Bradley and Steele (2014) for a discussion of alternative update rules.

for some proposition a, you are initially certain that your future credal set will assign the set of probability values S to a, and yet you currently assign a different set of probability values to a. That is:

$$\Pi(a|[Q(a) = S]) \neq S \tag{4}$$

White takes this to show that imprecise credences violate our starting intuition. Sometimes, you can know your future self will rationally come to have an imprecise attitude towards a and yet fail to have the same attitude towards a.

It has been pointed out that White's argument goes wrong by assuming that the set of probability values assigned by your credal set to a proposition is a good representation of your attitude towards that proposition (Joyce, 2010; Bradley, 2019). Learning that $Q(a)$ equals S does not by itself tell you what your attitude towards a will be, since this depends on other features of your future credal set beyond the set of probability values assigned to S.

While this suffices to defuse White's argument, I think the argument still raises a challenge for Imprecise Bayesians. White's starting intuition remains appealing: there should be at least *some sense* in which, conditional on your better-informed attitudes being so-and-so, your current attitudes should be so-and-so. And more generally, insofar as they are interested in characterising deference relations in various contexts (1-4), Imprecise Bayesians should be interested in what kind of principle, if any, can do so within their theory.

To address these questions, two imprecise deference principles have been proposed. They are usually discussed with a focus on cases like those described by White (2010), where Q is your future credal set. The first principle is briefly mentioned by Topey (2012) and Schoenfield (2012), and it mirrors the Global Reflection principle given earlier. The name and formulation given here are due to Moss (2021).

- **Identity Reflection**

  $$\Pi(\cdot|[Q = \Gamma]) = \Gamma \tag{5}$$

  for every credal set Γ such that the above conditional credence is defined.

Here Π and Γ are rigid designators of credal sets, and Q is a definite description of a credal set. When Π respects this condition with regards to Q, we say that Π *Identity Reflects* Q. Intuitively Π Identity Reflects Q when, conditional on the entire credal set Q being so-and-so, Π is so-and-so. As in the precise case I will assume throughout that Π is regular, in the sense that every $\pi \in \Pi$ is regular.

To get a better grasp of what Identity Reflection involves, it helps to think of it as the conjunction of two independently interesting deference principles:

- **Subset Reflection**

  $$\Pi(\cdot|[Q = \Gamma]) \subseteq \Gamma \tag{6}$$

for every credal set Γ such that the above conditional credence is defined.

- **Superset Reflection**

$$\Pi(\cdot|[Q = \Gamma]) \supseteq \Gamma \tag{7}$$

for every credal set Γ such that the above conditional credence is defined.

These two principles have a natural supervaluational reading. Subset Reflection says that, conditional on Q being the credal set Γ, you make every judgement that Γ makes (and possibly more). Superset Reflection says that, conditional on Q being the credal set Γ, you make no more (and possibly less) judgements than those Γ makes. Note that, in the precise case, both principles above are equivalent to Global Reflection. In the imprecise case, Identity Reflection is clearly equivalent to their conjunction.

In Section 5 I will prove a result analogous to the one in Briggs (2009): when Q is the credal set obtained by (pointwise) conditioning your current credal set on the true element of a partition, you Identity Reflect Q. This shows that there is at least some sense in which, conditional on your better-informed attitudes being so-and-so, your current attitudes should be so-and-so.

But let's turn now to the second deference principle discussed in the Imprecise Bayesian literature. Moss (2021) has raised the complaint that Identity Reflection is hard to operationalise. This is because it constrains your opinions conditional on extremely fine-grained propositions —propositions which fully specify your entire future credal set Q. Moss makes it clear that her worry is not that Identity Reflection *only* constrains you conditional on such propositions. As she points out, by constraining your conditional credences even on very strong conditions, the principle constrains your unconditional credences too. So even if you are quite unsure about your future credal set, Identity Reflection can still impose substantive constraints on your current one. As I understand it, her worry is more about the guidance value of Identity Reflection: because we normally lack, or find it difficult to access, opinions conditional on such detailed propositions, we might find it difficult to *use* Identity Reflection when evaluating our own rationality or the rationality of others.

For this reason, Moss (2021) proposes the following alternative imprecise deference principle:

- **Pointwise Reflection**

$$\Pi(a|[Q(a) = S]) \subseteq S \tag{8}$$

for every proposition a and real number set $S \subseteq \mathbb{R}$ such that the above conditional credal set is defined.

When Π respects this condition with regards to Q, we say that Π *Pointwise Reflects* Q. As we shall see in Section 5, Moss (2021) has proven a Briggs-style for Pointwise Reflection: if Q is obtained by (pointwise) conditioning Π on the true element of a partition, *and* Π *is convex*, then

Π Pointwise Reflects Q.[6] So once again, we can show imprecise credences satisfy White's motivating intuition, at least for convex credal sets: conditional on your future imprecise doxastic state having a certain feature, your current imprecise doxastic state should have that feature. But this time, unlike with Identity Reflection, the feature in question is localised to a single proposition, making the principle easier to operationalise.

Now that we have two imprecise deference principles on the table, the question is whether they characterise the various notions of deference that Global Reflection characterises within precise Bayesian epistemology. I will tackle this question in the next two sections. Before doing so, however, I want to slightly amend Pointwise Reflection in a way that makes it a stronger and, I think, more plausible deference principle. The discussion gets a bit techical: those less interested in the formal details can skip to Section 4 by just noting that Pointwise Reflection will be defined as in (22) on all random variables, rather than being restricted to propositions as in (8).

3.2 Amending Pointwise Reflection

My problem with the above version of Pointwise Reflection is that it treats propositions and random variables differently. But it's hard to think of a good reason for doing so. To see the problem consider the following example, adapted from Walley (1991)[pp.82-83].

Example 3.1. Let $\mathcal{W} = \{w_1, w_2, w_3\}$ and define three probability functions as follows:

$$p_1 = (2/3, 1/3, 0) \tag{10}$$
$$p_2 = (1/3, 0, 2/3) \tag{11}$$
$$p_3 = (2/3, 0, 1/3) \tag{12}$$

we can use these probabilities to define two credal sets:

$$\Gamma = \text{ch}(\{p_1, p_2\}) \tag{13}$$
$$\Gamma' = \text{ch}(\{p_1, p_2, p_3\}) \tag{14}$$

where ch(P) denotes the *convex hull* of the set P.[7]

We can show that for every proposition a, $\Gamma(a) = \Gamma'(a)$. This can be proven formally, but the easiest way to see it is by looking at a barycentric plot of the two credal sets (Figure 2).

[6] A credal set Γ is convex iff for every $\lambda \in (0, 1)$, and every $\gamma_1, \gamma_2 \in \Gamma$, the credence function γ_λ is also a member of Γ, where:

$$\gamma_\lambda(a) = \lambda\gamma_1(a) + (1 - \lambda)\gamma_2(a) \text{ for every } a \in \mathcal{F}_{\mathcal{W}} \tag{9}$$

[7] If P is a credal set then its convex hull is the set of all convex combinations of elements in P. That is:

$$ch(P) = \{\sum_{i=1}^{n} \lambda_i p_i : p_i \in P, \lambda_i > 0, \text{ and } \sum_{i=1}^{n} \lambda_i = 1\} \tag{15}$$

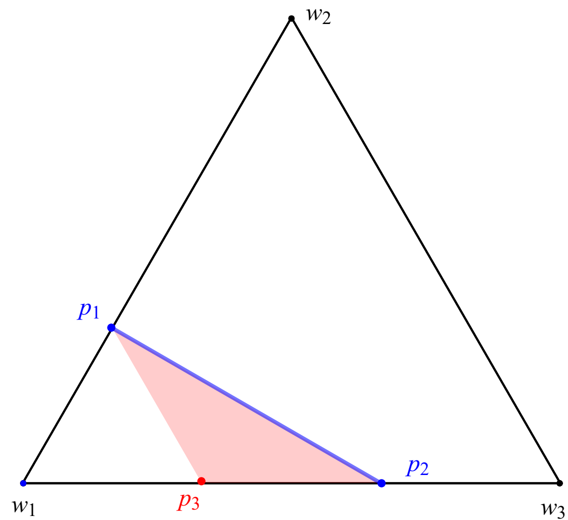

Figure 1: Barycentric plot of the credal sets in Example 3.1. Γ corresponds to the blue line between p_1 and p_2, whereas Γ' corresponds to the red highlighted triangle with vertices p_1, p_2, and p_3.

In a barycentric plot of this kind, each world is represented by a vertex of an equilateral triangle with height 1 unit. Each point p within the triangle corresponds to a probability function: the probability assigned to world w_i is determined by the distance between p and the edge opposite w_i. For example, p_1 is $2/3$ units away from the edge $\overline{w_2 w_3}$, $1/3$ units away from the edge $\overline{w_1 w_3}$, and 0 units away from the edge $\overline{w_1 w_2}$.

A credal set P is represented by a set of points in the plot. In this case, Γ is represented by the blue line between p_1 and p_2, and Γ' is represented by the red shaded area with vertices p_1, p_2, and p_3. Given a credal set P and a world $\{w_i\}$, the value set $P(w_i)$ is just the set of distances between points in P and the vertex opposite to w_i.

We can now easily show that $\Gamma(a) = \Gamma'(a)$ for every proposition a. First, it's easy to check from the figure that $\Gamma(\{w_i\}) = \Gamma'(\{w_i\})$ for every w_i. Furthermore, note that:

$$\begin{aligned}
\Gamma'(\{w_1, w_2\}) &= \{\pi(\{w_1, w_2\}) : \pi \in \Pi\} && (16)\\
&= \{1 - \pi(\{w_3\}) : \pi \in \Pi\} && (17)\\
&= \{(1 - r) : r \in \Pi(\{w_3\})\} && (18)\\
&= \{(1 - r) : r \in \Gamma(\{w_3\})\} && (19)\\
&= \Gamma(\{w_1, w_2\}) && (20)
\end{aligned}$$

and the same holds for every 2-world proposition. Finally, we have $\Gamma(\mathcal{W}) = \{1\} = \Gamma'(\mathcal{W})$.

Despite this, we can find a random variable X such that $\Gamma'(X) \neq \Gamma(X)$.

To see this, consider the random variable X defined by:

$$X = \begin{cases} 1 \text{ if } w_1 \\ -1 \text{ if } w_2 \\ 0 \text{ if } w_3 \end{cases} \tag{21}$$

A random variable divides the unit triangle into two (possibly empty) half-spaces: one corresponds to the set of all probability functions with non-negative expectation for the random variable, the other to the set of all probability functions with negative expectation for the random variable. So we can represent X in our figure by a dashed line, with an arrow indicating which side of the triangle contains the probability functions that assign positive expectation to X. The points on the dashed line assign expectation 0 to X. Among the points assigning positive expectation

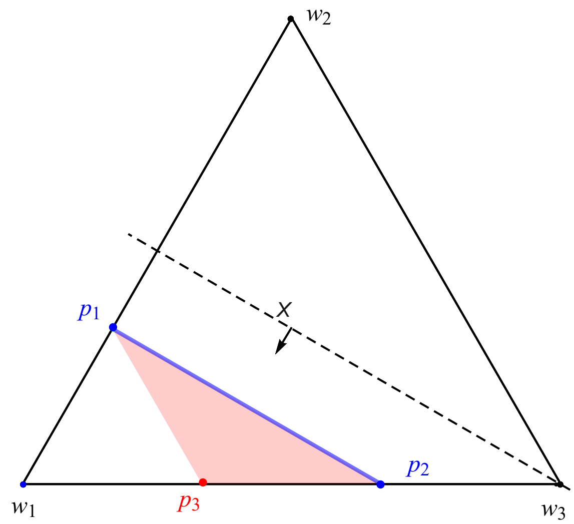

Figure 2: Barycentric plot of the credal sets in Example 3.1, with dashed line representing random variable $X = (1, -1, 0)$.

to X, those further from the dashed line assign larger expectation to X. So we can tell from the figure that $\Gamma(X) \neq \Gamma'(X)$. In particular, $\Gamma(X) = \{p_1(X)\} = \{1/3\}$, whereas $\Gamma'(X) = [p_1(X), p_3(X)] = [1/3, 2/3]$.

This example shows that two credal sets can assign the exact same set of probablity values to every proposition, while representing importantly different doxastic states, with different sets of expectation values for some random variables. It would be odd if our deference principle made it *impossible* to defer to either credal set without also deferring to the other.

A simple way to avoid this is to extend the Pointwise Reflection constraint to all random variables:

- **Pointwise Reflection***

$$\Pi(A|[Q(X) = S]) \subseteq S \tag{22}$$

for every random variable X and real number set $S \subseteq \mathbb{R}$ such that the above conditional credal set is defined.

From now on, I will use the term Pointwise Reflection to refer to this amended version of the principle.

4 Clarity and Deference

After introducing a few deference principles for imprecise credences, the remainder of this essay will look at whether and how we can argue that these principles characterise the notions of deference relevant to contexts (1-4), which are characterised by Global Reflection in the precise case. My arguments here will rely on a seemingly innocuous, yet surprisingly consequential assumption, which I call *Clarity* (borrowing the term from Dorst (2023)). So I want to start by stating this assumption and briefly outlining its current status within contemporary (precise) Bayesian epistemology.

Here is a definition of Clarity:

- **Clarity**: The definite description Q of a doxastic state is *Clear at world* w iff the state Q_w it refers to at world w is certain that Q refers to Q_w. That is, Q is Clear at w iff $Q_w([Q = Q_w]) = \{1\}$. I say that Q is *Clear* when it is Clear at every possible world.[8]

Albeit under different names, and sometimes in slightly different shapes, Clarity has been widely discussed in the precise Bayesian literature. This is because it is deeply tied to Global Reflection, which I introduced in Section 2 as one of the most popular precise deference principles. In fact if a regular credence function π Global Reflects a definite description q, then q is Clear. The proof is very simple:

Proposition 1 (Global Reflection $\Rightarrow$ Clarity)**.** *Let π be a credence function and q a definite description of a credence function, both defined on the same domain $\mathcal{W}$. Then π Global Reflects q only if q is Clear at w_i.*

Proof. Assume by way of contradiction that π Global Reflects q, and yet q is not Clear at some $w_i \in \mathcal{W}$. The conditional credence function $\pi(\cdot|[q = q_i])$ is well-defined by regularity. So we can write:

$$\begin{aligned} \pi([q = q_i]|[q = q_i]) &= q_i([q = q_i]) \text{ by Global Reflection} && (23) \\ &< 1 \text{ because } q \text{ not Clear at } w_i && (24) \end{aligned}$$

But since π is a probability function, $\pi([q = q_i]|[q = q_i]) = 1$, contradiction. $\square$

[8] This is commonly known as "Immodesty", particularly in the literature about deference to external experts. While the term works well in that context (being certain that you are the expert sounds quite "immodest") it is a bit strange if we apply it e.g. to your current/future credal set. Your current/future credal set knowing it is your current/future credal set does not intuitively make it "immodest".

The relationship between Local Reflection and Clarity is somewhat more complicated. Gaifman (1988) has shown that Local Reflection does not entail Clarity, by providing examples where a regular credence function π Local Reflects q even though q is not Clear. However, a recent paper by Gallow (2023) has shown that examples of this sort are exceedingly rare, only occurring when q has a very specific cyclic structure. Gallow convincingly argues that, even though logically distinct, Local Reflection and Global Reflection are philosophically equivalent. So Local Reflection commits us to Clarity in the same way as Global Reflection does.

Clarity often appears in the precise Bayesian literature as a (more or less explicit) assumption used to motivate Global and Local Reflection. But it has also been criticised by several authors as excessively restrictive for many of the contexts we are interested in (Hall, 1994; Elga, 2013; Pettigrew and Titelbaum, 2014; Dorst, 2020). While I find many of these criticisms convincing, I will assume Clarity throughout the rest of this essay, either directly or as a consequence of other assumptions. This is for three reasons. First, it's usually easier to study deference principles when the expert is assumed to be Clear, so it makes sense to start our treatment of imprecise deference principles by focusing on these simpler cases. Secondly, Clarity still has its philosophical supporters in some contexts (Titelbaum, 2015). And finally, some authors have argued that once we give up Clarity, we have no more reason to introduce imprecise doxastic states. For example, Carr (2020) argues that if we allow that the rational credence function in light of your evidence may not be Clear (i.e. if we allow rational uncertainty about the requirements of rationality) precise credences can do all the work that imprecise credences have been introduced to do. If she is right, it might be pointless to worry about deference principles for imprecise credences without the Clarity assumption. Before abandoning Clarity, we need to argue that it still makes sense to care about imprecise credences once we do so, and that goes beyond the scope of this essay. So I will leave it for future work to determine whether and how imprecise deference principles can be adapted to accommodate epistemic superiors that are not Clear.

Let's take stock. We have outlined a number of contexts (1-4) where the epistemic deference relation plays an important role, and presented two imprecise deference principles, called Identity Reflection and Pointwise Reflection. These have been introduced in response to White's objection that imprecise credences are inconsistent with our intuitions about (1) deference to one's future self. The next section will show how these two principles are related to one another, and prove that under assumptions suitable for context (1), rational agents defer to their future selves according to both principles. Although Clarity is not assumed explicitly, we will see that it follows from the other assumptions. Then Section 6 will look at a more general way to motivate Identity and Pointwise Reflection for the other contexts we are interested in.

5 Certain Conditionalisation and Deference to Future Self

I have mentioned in Section 3 that we can show, under some conditions, that a rational agent should defer to their future doxastic state under both Identity and Pointwise Reflection. Now it's time to look at these arguments in some detail.

Let me start by outlining the general idea of the argument. Intuitively, we tend to think of our future selves as being in a privileged epistemic position with respect to us. This is attested by a number of sayings and expressions ("Hindsight is 20/20", "Wise after the event", "Monday morning quarterback"...), all capturing the idea that future-you is generally more informed than current-you. So we would expect a rational agent to defer to their future attitudes: conditional on your future self having attitude ϕ, surely you should have attitude ϕ right now. This is the same sort of intuition appealed to by White's (2010) argument, discussed in Section 3.

This intuition needs a bit of sharpening before it becomes philosophically viable. It's easy to think about cases where you *definitely should not* defer to your future self: sometimes we are misled, we become delusional, we forget things, or our opinions change in strange ways due to forces outside our control (e.g. drugs, brainwashing). Even if you are rational now, and rational in the future, things like this could happen to you between now and then, making current-you suspicious of the epistemic worth of future-you's opinions.

To save our starting intuition we must assume away all of these cases. More precisely, we must assume that the only way your doxastic state can be different between now and the future time is that future-you has updated their beliefs rationally on some non-misleading evidence. And for many Bayesians, the most uncontroversial case of rational updating on non-misleading evidence involves conditioning on the true element of some partition. At least when you are sure that your future doxastic state is obtained in this way, you should reflect your future attitudes. Brigg's (2009) result starts by assuming precisely this. Here is one way to formulate this assumption in the imprecise case:

- **Certain Conditionalisation**:[9] Let Π be the agent's credal set at some initial time t_0, and Q the definite description of the agent's credal set at some future time t_1. There is a finite partition $\mathcal{E}$ =

[9]Moss (2021)[Claim 4] makes an analogous assumption, although in a different form. Expressed in my notation, for some partition $\mathcal{E} = \{e_1, ..., e_m\}$ she requires that the following three conditions hold for every $a \subseteq \mathcal{W}, e_j \in \mathcal{E}, S \subseteq \mathbb{R}$:

i $\Pi([\Pi(a|e_j) = Q(a|e_j)]|e_j) = \{1\}$.

ii $\Pi(e_j \Leftrightarrow [Q(e_j) = \{1\}]) = \{1\}$.

iii $\Pi(\Pi(a|e_j) = S)$ iff $\Pi(a|e_j) = S$.

Note that (iii) is unnecessary in my treatment, since I'm assuming that Π denotes the same credence function at every possible world. See Remark 1 in the Appendix for a proof that, once we extend them to all random variables, (i) and (ii) are equivalent to Certain Conditionalisation.

$\{e_1, ..., e_m\}$ such that for every $e_j \in \mathcal{E}$:

$$\Pi([Q = \Pi(\cdot|e_j)]|e_j) = \{1\} \tag{25}$$

whenever this conditional credal set is defined. Informally: Π is certain, conditional on e_j being true, that $Q = \Pi(\cdot|e_j)$.

Since we are assuming that Π is regular, Certain Conditionalisation entails Clarity. To see this it helps to prove the following technical lemma:

Lemma 1. *Let Π be a regular credal set, Q a definite description of a credal set, and $\mathcal{E} = \{e_1, ..., e_m\}$ a finite partition. Assume Π respects Certain Conditionalisation w.r.t. Q and $\mathcal{E}$. Then for every $e_j \in \mathcal{E}$:*

$$[Q = \Pi(\cdot|e_j)] = e_j \tag{26}$$

Proof. See Appendix. □

Using this lemma we can see that, for every w_i, letting e_j be the element of $\mathcal{E}$ such that $w_i \in e_j$, we have:

$$Q_i([Q = Q_i]) = \Pi([Q = \Pi(\cdot|e_j)]|e_j) = \Pi(e_j|e_j) = \{1\} \tag{27}$$

So even though we won't be assuming Clarity explicitly, the arguments in this section will rely on an assumption which entails Clarity.

We can now show that Certain Conditionalisation entails Identity Reflection.

Proposition 2 (Certain Conditionalisation $\Rightarrow$ Identity Reflection)**.** *Let Π be a regular credal set, Q a definite description of a credal set, and $\mathcal{E} = \{e_1, ..., e_m\}$ a finite partition. If Π respects Certain Conditionalisation w.r.t. Q and $\mathcal{E}$, then Π Identity Reflects Q.*

Proof. See Appendix □

This shows that, when you are certain that your future doxastic state is obtained by conditionalising your current doxastic state on the true element of a partition, you should Identity Reflect your future self. In turn this means that conditional on your future self having a certain doxastic state, you should adopt that doxastic state right now. As a special case of this result, when all credence functions are precise and the agent obeys Certain Conditionalisation, the agent should Global Reflect their future self.

Taking a step back: how does this result fit into the broader debate around deference principles for imprecise credences? Recall White's (2010) starting intuition, that conditional on your future self having a certain attitude (as a result of rationally updating on factive evidence), you should have that attitude. By taking your attitude towards a proposition to be encoded by the set of probability values you assign to that proposition, White was able to show that agents with imprecise credences violate this intuitive requirement.

We have seen that a good response to his argument is the (independently motivated) point that the set of probability values you assign to a proposition does not suffice to encode your attitude towards that proposition. So conditional on your future self assigning a certain value set to some proposition, you need not assign that proposition the same value set now. However this response leaves open what, if anything, *does* encode your future attitudes in a way that, supposing you had such attitudes in the future, you should adopt the same attitudes now. A natural suggestion is that this is the entirety of your future credal set. After all, if anything encodes your future attitudes, surely your *entire credal set* does. Proposition 2 shows that this is in fact the case. If you know you'll update by conditioning on the true element of a partition, then conditional on your future credal set being so-and-so, your current credal set should be so-and-so.

As mentioned earlier, even though Identity Reflection is expressed in terms of extremely strong conditioning propositions, it can constrain your attitudes conditional on weaker propositions, and even your unconditional attitudes. Having said that, it would be interesting, if only for pragmatic reasons, to find out whether there are any more localised attitudes of your future self which you ought to reflect.

We can find such localised constraints in the precise case. Here, Certain Conditionalisation entails that you should Local Reflect your future self. Indeed, this is the form of the result that was originally proven by Briggs (2009). We can show that this follows from the relationship between Local and Global Reflection under Clarity:

Lemma 3. *Let π be a regular credence function and q a definite description of a credence function. Assume q is Clear. Then π Local Reflects q iff π Global Reflects q.*

Proof. See Appendix. □

Since Certain Conditionalisation entails Clarity (Proposition 2), Lemma 3 shows that Certain Conditionalisation entails Local Reflection in the precise case.

What about localised constraints in the imprecise case? White's (2010) argument shows that the most natural imprecise analogue of Local Reflection is not entailed by Certain Conditionalisation (recall Equation (4)). For this reason, Moss (2021) has proposed Pointwise Reflection as an alternative local principle. But does Pointwise Reflection follow from Certain Conditionalisation? And what is the relationship between this principle and Identity Reflection?

I will now answer both questions. First I will show that Pointwise Reflection does not follow from Certain Conditionalisation in general. However, Pointwise Reflection follows from Certain Conditionalisation when Q is a *convex* credal set. Then I will show that, under this assumption, Pointwise Reflection is equivalent to Subset Reflection, and strictly weaker than Identity Reflection.

To see that Certain Conditionalisation does not generally entail Pointwise Reflection, consider the following variant of the Coin Toss example in White (2010):

Example 5.1 (Coin Game)**.** Jack has a coin which you know is fair, and you also know that he knows whether some proposition a is true. Jack paints over the two sides of the coin so you can't tell which one is heads and which one is tails. Then depending on the truth of a, he writes something on each side: if a is true, he writes 'a' on the heads side, and '$\neg a$'on the tails side; and if a is false, he writes '$\neg a$' on the heads side and a on the tails side.

Assume your starting credal set Π contains just two probability functions, π_1 and π_2, such that $\pi_1(a) = .4$ and $\pi_2(a) = .6$ (for example, you might know that the objective chance of a is either .4 or .6). Letting h be the proposition that the coin lands with the heads face up, you have $\pi_1(h) = \pi_2(h) = 1/2$ since you know the coin is fair. Furthermore, you judge the coin toss to be independent of a, in the sense that:

$$\pi_k(h|a) = \pi_k(h) = 0.5 \tag{28}$$

for $k = 1, 2$. Denote by e_a the proposition that the coin lands with the 'a' face up. We can identify each possible world in this example with the propositions that are true at that world. For example the possible world where a is true, e_a is false, and h is false will be denoted by $a\bar{e}_a\bar{h}$. Then we can write the two probability functions in your credal set as follows:

	ahe_a	$a\bar{h}\bar{e}_a$	$\bar{a}\bar{h}e_a$	$\bar{a}h\bar{e}_a$
π_1	.2	.2	.3	.3
π_2	.3	.3	.2	.2

Consider what would happen if you observed e_a. Credence function π_1 will take this as providing some evidence that $\neg h$ is the case, since π_1 is more than 50 percent confident in $\neg a$: the claim that e_a and $\neg h$ are equivalent. Similarly, π_2 will take this as providing some evidence that h is the case. Indeed, after conditioning on e_a, your credal set would look like this:

	ahe_a	$a\bar{h}\bar{e}_a$	$\bar{a}\bar{h}e_a$	$\bar{a}h\bar{e}_a$
π_1	.4	0	.6	0
π_2	.6	0	.4	0

But what would happen if you observed $\neg e_a$ is exactly symmetrical. Credence function π_1 will take it as evidence for h, and credence function π_2 will take it as evidence against h. So you would end up with the following credal set:

	ahe_a	$a\bar{h}\bar{e}_a$	$\bar{a}\bar{h}e_a$	$\bar{a}h\bar{e}_a$
π_1	0	.4	0	.6
π_2	0	.6	0	.4

In either case, you end up with $\Pi(h|e_a) = \Pi(h|\neg e_a) = \{.4, .6\}$. And yet by assumption, $\Pi(h) = \{.5\}$. So if we let Q be the your credal set after observing the coin toss, and let Q be the credal set obtained by conditionalising on the true element of the partition $\{e_a, \neg e_a\}$, this example shows a violation of Pointwise Reflection under Certain Conditionalisation, because:

$$\Pi(h|[Q(h) = \{.4, .6\}]) = \Pi(h) = \{.5\} \not\subseteq \{.4, .6\} \tag{29}$$

Note that in this example, Pointwise Reflection fails because your updated credal set assigns two "disconnected" probability values to h. And in turn, this is because your starting credal set assigns two "disconnected" probability values to h. If we were to take as your starting credal set Π the convex hull of $\{\pi_1, \pi_2\}$, there would be no problem with this example, since then your updated credal set will have $\Pi(h|e_a) = \Pi(h|\neg e_a) = (.4, .6)$, and:

$$\Pi(h|[Q(h) = (.4, .6)]) = \Pi(h) = \{.5\} \in (.4, .6) \tag{30}$$

This suggests two ways of saving the principle. First, we could *weaken* Pointwise Reflection, and only require that your current probability value set for a, conditional on your future one being S, is a subset *of the convex hull of* S. Alternatively, we could *restrict* the intended domain of Pointwise Reflection, so that it only applies when the credal set Q is *convex*. Since Q is obtained by conditionalisation from Π, and conditionalisation preserves convexity, a natural way to ensure this is to assume that Π is convex.

I will take the second route here, both because the first one loses some of the intuitive motivation Moss gives for Pointwise Reflection, and because a number of Imprecise Bayesians already restrict their attention to convex credal sets for a variety of independent reasons (Levi, 1980). In fact, a number of popular frameworks for representing imprecise credences, such as lower probabilities and lower previsions, turn out to be equivalent to (closed and) convex credal sets.[10] But whichever route we take, Pointwise Reflection will follow from Certain Conditionalisation. The proof is the one given in (Moss, 2021)[pp.637-638] so I will not repeat it here.[11] Instead, I want to end this subsection by clarifying the relationship between Pointwise Reflection and Identity Reflection. First I will show that when we fix Pointwise Reflection either by weakening it, or by restricting it to convex credal sets, Subset Reflection will entail it. This is the content of the following proposition:

Proposition 3 (Sub-Reflection (Convex) $\Rightarrow$ Pointwise Reflection)**.** *Let* Π *be a regular credal set and* Q *the definite description of a convex and Clear credal set. If* Π *Sub-Reflects* Q*, then* Π *Pointwise-Reflects* Q*.*

Proof. See Appendix. □

Proposition 3 shows that if we fix Pointwise Reflection in such a way that it is entailed by Certain Conditionalisation, then Pointwise Reflection is no stronger than Subset Reflection, in the sense of imposing no more constraints than it. In fact we can show that under Clarity, and assuming Q is convex, Pointwise Reflection is also no weaker than Subset Reflection.

Proposition 4 (Pointwise Reflection (Convex) $\Rightarrow$ Sub-Reflection)**.** *Let* Π *be a regular credal set and* Q *the definite description of a convex and Clear credal set. If* Π *Pointwise-Reflects* Q*, then* Π *Sub-Reflects* Q*.*

Proof. See Appendix □

Together, Propositions 3 and 4 show that Subset and Pointwise Reflection impose the exact same constraints on convex credal sets, despite one being formulated in global terms and the other being formulated in local terms.

[10] See Wheeler (2022) for an overview of imprecise frameworks that are/are not equivalent to closed and convex credal sets.

[11] This also follows from Proposition 3 together with the fact that Certain Conditionalisation entails Identity Reflection (Proposition 2)

Having shown Pointwise Reflection and Subset Reflection are equivalent on convex credal sets, it is natural to wonder whether they are strictly weaker than Identity Reflection on convex credal set, i.e. whether they impose strictly weaker constraints. The next proposition shows this is in fact the case. I show this by giving an example of a convex credal set Π and a definite description of a convex credal set Q such that Π Pointwise Reflects Q, and yet Π does not Identity Reflect Q.[12]

Example 5.2. Let $\mathcal{W} = \{w_1, w_2\}$. Let $\Pi = \{\pi : \pi(w_1) = 1/2\}$ a singleton credal set. Let Q the (constant) definite description of a credal set such that:

$$Q_1 = Q_2 = \mathrm{ch}\{q, q'\}, \text{ where } q(\{w_1\}) = 0 \text{ and } q'(\{w_1\}) = 1 \tag{31}$$

Note that Q is convex at every possible world by construction, and is also trivially Clear. First I show that Π Pointwise Reflects Q. To see this, note that for every $X : \mathcal{W} \to \mathbb{R}$:

$$\begin{aligned} [Q(X) = S] \neq \emptyset \iff & S = Q_1(X) = Q_2(X) && (32)\\ \iff & S = (\min\{q(X), q'(X)\}, \max\{q(X), q'(X)\}) && (33) \end{aligned}$$

Now assume without loss of generality that $q(X) \leq q'(X)$, and let $S = (q(X), q'(X))$. Letting π be the only credence function in Π, we have:

$$\begin{aligned} \pi(X|[Q(X) = S]) &= \pi(X) && (34)\\ &= .5X(w_1) + .5X(w_2) && (35)\\ &= .5(X(w_1) + X(w_2)) + .5(X(w_1) + X(w_2)) && (36)\\ &= .5q(X) + .5q'(X) \in S = [0, 1] && (37) \end{aligned}$$

so $\Pi(X|[Q(X) = S]) \subseteq S$ whenever this conditional credal set is defined, proving that Π Pointwise Reflects Q.

It's now easy to show that Π does not Identity Reflect Q, since $[Q = \Gamma] \neq \emptyset$ iff $\Gamma = Q_1 = Q_2$, in which case $[Q = \Gamma] = \mathcal{W}$. So we have:

$$\Pi(\cdot|[Q = Q_1]) = \Pi \neq Q_1 \tag{38}$$

which shows Identity Reflection is violated.

Proposition 5 (Pointwise Reflection (Convex) $\not\Rightarrow$ Identity Reflection)**.** *There is a credal set Π and a definite description of a convex credal set Q such that Π Pointwise Reflects Q, and yet Π does not Identity Reflect Q.*

Proof. Example 5.2. □

The above results show that Pointwise Reflection is generally weaker than Identity Reflection. What should we make of this? I think the answer depends on what context (1-4) we are interested in. If we are considering deference to one's future self, and we are assuming Certain Conditionalisation, then we should not think of Pointwise Reflection as an *alternative* principle to Identity Reflection, but rather as a *local consequence* of it. In particular, Pointwise Reflection specifies a class of local constraints which

[12]Note that, in light of Proposition 2, Certain Conditionalisation must also fail in this example.

are easier to check, and which are exactly those imposed by Subset Reflection. Identity Reflection imposes all the same constraints, plus all the constraints imposed by Superset Reflection.

But while Certain Conditionalisation might be a reasonable assumption if we're interested in capturing the intuition that you should defer to your future self, it is clearly not always warranted. For example, if we are interested in giving some necessary conditions for an update to count as a genuine learning experience (2), then assuming that the update happens by conditionalising on a true proposition gives up the whole game. And similarly, only a very restrictive notion of epistemic expertise (3) would require that the expert credences be obtained by conditionalising your own on the elements of some partition. Finally, it's not clear why we should think that (4) there is some partition such that the rational credence function at every world is obtained by conditioning the same rigidly designated prior on whichever element of the partition is true at that world. So you might think that, in contexts where we don't want to assume Certain Conditionalisation, Pointwise Reflection rather than Identity Reflection might characterise the relevant notion of deference.

No matter which context (1-4) we are interested in, I am doubtful that Subset/Pointwise Reflection can tell the whole story about deference. Consider again Example 5.2. Recall that you Sub-Reflect Q whenever, conditional on Q being Γ, you make all the judgements that Γ makes. In this example you know that $Q = \Gamma$, and you Sub-Reflect Q because Γ makes *no judgements beyond the trivial ones.* For example, unless X is positive in every possible world, Q does not find X desirable, so Subset Reflection does not impose any non-trivial constraints on what gambles you should consider desirable. Indeed, any credal set Π will Sub-Reflect Q in this example!

Subset/Pointwise Reflection may well capture an important notion of deference: upon learning your superior's doxastic state, you would make all the judgements they make. But often the kind of deference we are interested in seems to require something more. To see this, it helps to look at a slightly more realistic case than Example 5.2.

Example 5.3. Let d be the proposition that the Democrats will win the next election, and let Q denote the credal set of your friend Lucy. Lucy is a staunch democrat, and only informs herself by reading highly biased pro-Democrat newspapers. You know that if the polls and other indicators strongly suggest the Democrats are going to win the election, Lucy will have learned about it from the newspapers, but if the polls and other indicators are undecisive, or if they suggest the Republicans will win, then the newspapers Lucy reads will not have reported about it.

We can model this scenario by looking at four possible worlds $\mathcal{W} = \{w_1, w_2, w_3, w_4\}$. Let $d = \{w_1, w_2\}$, and let $r = \{w_1, w_3\}$ be the proposition that the newspapers reported the polls. You know that, if the newspapers reported the polls, the Democrats are likely to win, and Lucy will be (rationally) confident in d. But if the newspapers did not report the polls, Lucy will have little to no evidence about whether d. We can formalise this by defining Lucy's credal set Q as follows:

$$Q_1 = Q_3 = \{\gamma \in \mathcal{P}_{\mathcal{W}} : \gamma(w_1) = .9, \gamma(w_3) = .1, \gamma(w_2) = \gamma(w_4) = 0)\} \tag{39}$$
$$Q_2 = Q_4 = \{\gamma \in \mathcal{P}_{\mathcal{W}} : \gamma(w_1) = \gamma(w_3) = 0, \gamma(w_2) \geq .01, \gamma(w_4) \geq .01\} \tag{40}$$

At the worlds where the newspapers report the polls, i.e. when $w_i \in r$, Lucy has $Q_i(d) = .9$. So she is very confident in d at those worlds. At the worlds where the newspaper does not report the polls, i.e. when $w_i \notin r$, Lucy has $Q_i(d) = [.01, .99]$. So she almost completely suspends judgement about whether d at those worlds, since she has almost no evidence one way or another. Note that Q is Clear, since $Q([Q = Q_i]) = \{1\}$ for every $w_i \in \mathcal{W}$.

A number of credal sets Subset Reflect Lucy in this example, but some do so in a way that seems importantly stronger. To see this, consider two agents, Martha and Stewart, with (regular) credal sets Π^M, Π^S defined as follows:

$$\Pi^M = \{\pi \in \mathcal{P}_{\mathcal{W}} : \pi(w_1) = 9/20, \pi(w_2) \geq 1/200, \pi(w_3) = 1/20, \pi(w_4) \geq 1/200\} \tag{41}$$
$$\Pi^S = \{\pi \in \mathcal{P}_{\mathcal{W}} : \pi(w_1) = 9/20, \pi(w_2) = \pi(w_4) = 1/4, \pi(w_3) = 1/20\} \tag{42}$$

Some simple calculation show that both Π^M and Π^S Subset Reflect Q. Indeed, conditional on $r = [Q = Q_1]$, both are identical to Q_1 (equivalently for Q_3). So upon learning that Lucy is confident in a Democrat victory, both Martha and Stewart are confident in a Democrat victory. But the two credal sets differ in their conditional credences, conditional on $\neg r$. In particular, we have that:

$$\Pi^M(d|\neg r) = \Pi^M(d|[Q = Q_2]) = [.01, .99] = Q_2(d) \tag{43}$$
$$\Pi^S(d)|\neg r) = \Pi^S(d|[Q = Q_2]) = \{1/2\} \in Q_2(d) \tag{44}$$

Again, as far as Subset/Pointwise Reflection are concerned, both agents defer to Lucy here. But Martha seems to defer to Lucy in a stronger way than Stewart, because upon learning that Lucy suspends judgement about whether d, Martha also suspends judgement about whether d. Stewart, on the other hand, makes many judgements about the comparative likelihood of d and other propositions, and about the desirability of gambles on d, even upon learning that the Lucy does not make such judgements. Intuitively, we would like to say that he defers to Lucy *in a weaker sense* than Martha does. If this is right, there must be a stronger notion of deference at play than the one captured by Subset Reflection.

I want to end this section by saying a bit more about what Pointwise Reflection is failing to capture in the climate scientist example. While I cannot give a definitive answer here, I want to at least sketch one way to read the example which I find insightful. By learning that Lucy assigns $Q(d) = [.01, .99]$ both Martha and Stewart are plausibly learning, among other things, that Lucy has no evidence for or against d. Martha's response to learning that Lucy has no evidence about whether d is to suspend judgement about whether d, just like Lucy does. This makes sense if Martha has no further evidence than Lucy does, and if she responds to this lack of evidence in the same way Lucy does. Borrowing the terminology of Konek (2019), we would say that Martha, like Lucy, is being *epistemically conservative* by suspending her judgement about d in the face of scarce evidence. Stewart's response to Lucy's lack of evidence is different: he adopts a credal set which makes lots of judgements about d. He either has more evidence than Lucy, or he is responding to the lack of evidence in an *epistemically liberal* way, by making many probabilistic judgements in the face of scarce evidence, despite the fact that Lucy is epistemically conservative. Under this reading, Martha defers more strongly to the expert because, unlike Stewart, both her evidence and her epistemic attitudes are aligned with those of the expert. Subset Reflection allows one to defer to an expert even when one's total evidence or epistemic attitudes are not aligned with hers, and thus cannot distinguish the stronger kind of deference Martha is displaying from the weaker one displayed by Stewart.

In the next section I look at two ways to motivate deference principles without assuming Certain Conditionalisation (but still assuming Clarity). Both strategies rely on the same assumption: that you regard someone as epistemically superior when you consider their opinions to be *at least as valuable* than yours. But they differ in the kind of value they consider. I will give a brief overview of how this assumption has been used to defend precise deference principles. Then I will show that the same assumption motivates Subset/Pointwise Reflection in the imprecise case, and propose a way to adapt this argument strategy so that it motivates the stronger principle of Identity Reflection.

6 Arguments from Epistemic and Pragmatic Value

A natural way to characterise a notion of deference in contexts (1-4) is to say that you regard the epistemically superior doxastic states to be in some way *better*, or *more valuable*:

- **Value Superiority**: you defer to a doxastic state iff you regard it to be at least as valuable as yours.

Value Superiority gives us a way to derive a deference principle from a characterisation of the value of a doxastic state. The kind of value involved in this characterisation may be either *epistemic* or *practical*. Within the precise Bayesian literature, the epistemic value of a credence function is usually captured by a strictly proper measure of inaccuracy, which measures how closely a given credence function approaches the truth (Joyce, 1998; Pettigrew, 2016). Whereas the practical value of a credence function is a measure of how well that credence function performs in guiding an agent's decisions (Good, 1967).[13]

[13]See Myrvold (2012) and Levinstein (2017) for a discussion of how these two kinds of value

Both kinds of value have been used to motivate various deference principles in the precise case (Skyrms, 1990; Huttegger, 2014; Dorst et al., 2021). I give a quick overview of these results in the next subsection. Then, in Subsection 6.2, I look at what deference principles can be derived from the Value Superiority characterisation in the imprecise case. Since at present there is not analogue of strictly proper measures of inaccuracy for imprecise credences, I will focus my attention on practical value.[14] I will also assume Clarity holds throughout. See Dorst et al. (2021) for a discussion of value-based characterisations of deference in the precise case when the Clarity assumption is dropped.

6.1 The Precise Case

One way that pragmatic value has been characterised in the precise Bayesian literature is in terms of how good a credence function is at guiding your decisions. This characterisation goes back to the work of I.J. Good 1967.[15] Here is the idea: suppose that you are facing a decision problem, i.e. you must choose from a set of options $\mathcal{X} = \{X_1, ..., X_n\}$. Each option X_j is a random vairiable mapping worlds to real values, representing a gamble that pays in linear utility. So if you choose option X_j and w_i is the actual world, you receive $X_j(w_i)$ units of utility (this is a loss of utility if $X_j(w_i) < 0$). Now consider the higher-order decision problem consisting of the following two options: either you choose an option from $\mathcal{X}$ yourself, or you let q choose from $\mathcal{X}$ for you. We say π Pragmatic Values q on $\mathcal{X}$ iff π expects letting q choose to be at least as good as choosing for itself. And π Pragmatic Values q iff π expects letting q choose to be at least as good as choosing for itself, on every decision problem $\mathcal{X}$.

To make this more precise it helps to introduce some more notation. Let $\mathcal{L}(\mathcal{W})$ be the set of all gambles on $\mathcal{W}$. If $\mathcal{X}, \mathcal{Y}$ are sets of gambles and e is a proposition, I write $-\mathcal{X}$ for the set $\{-X : X \in \mathcal{X}\}$, $\mathcal{X} + \mathcal{Y}$ for the set $\{X + Y : X \in \mathcal{X}, Y \in \mathcal{Y}\}$ (and $\mathcal{X} - \mathcal{Y}$ for $\mathcal{X} + (-\mathcal{Y})$), and $e\mathcal{X}$ for the set $\{eX : X \in \mathcal{X}\}$. To each credence function π we can associate a *choice function* $C_\pi : 2^{\mathcal{L}(\mathcal{W})} \to 2^{\mathcal{L}(\mathcal{W})}$ defined by:

$$C_\pi(\mathcal{X}) = \{X \in \mathcal{X} : X \text{ maximises } \pi(X) \text{ in } \mathcal{X}\} \tag{45}$$

Intuitively, if $C_\pi(\mathcal{X})$ is the set of options which π considers choiceworthy among those in $\mathcal{X}$.[16]

are related.

[14]Seidenfeld et al. (2012) show that there are no continuous, real-valued, strictly proper scoring rules, of the kind used to measure accuracy in the precise case, for imprecise doxastic states. See also Schoenfield (2017); Mayo-Wilson and Wheeler (2016) for similar impossibility theorems, and Konek (2023) for a recent approach to measuring the accuracy of imprecise credences in light of these impossibility results.

[15]Although similar ideas had already been discussed by Savage (1954)

[16]Choice functions are commonly used to model the preferences of agents with imprecise credences, particularly those represented by sets of probability and utility functions. See Seidenfeld et al. (2010) for a treatment of choice functions associated to sets of probability-utility pairs, and Van Camp (2018) for choice functions not necessarily associated to such sets. Note that in this essay I am implicitly assuming that all gamble payoffs are expressed in the same linear utility scale.

Let π be a credence function and q be the definite description of a credence function which can take values $\{\gamma_1, ..., \gamma_m\}$. Let $\mathcal{X}$ be a set of gambles. What is the expected value of choosing for yourself? If $w \in \mathcal{W}$ is the case, choosing for yourself will pay off $X(w)$, where X is a choiceworthy gamble for you (i.e. $X \in C_\pi(\mathcal{X})$). Note that the set $C_\pi(\mathcal{X})$ may contain multiple gambles, whose values at each world can be different. So "choosing for yourself" does not uniquely define a payoff for every possible world, as the payoffs depend on which choiceworthy gamble you actually pick. However, since every gamble in $C_\pi(\mathcal{X})$ must have the same expectation (as it must maximise π's expectation on $\mathcal{X}$), we may take this to be the expectation of choosing for yourself.

What is the expected value letting q choose for you? If $w \in [q = \gamma_j]$, you will receive utility $X(w)$ for some $X \in C_{\gamma_j}(\mathcal{X})$. But note that, just like π, γ_j can also have multiple maximal options in $\mathcal{X}$. So the payoffs resulting from letting Q choose for you can vary depending on which of its maximal options γ_j will pick. The payoff of letting Q choose for you might correspond to the payoff of any gamble in the following set:

$$\sum_{j=1}^{m} [q = \gamma_j] C_{\gamma_j}(\mathcal{X}) \tag{46}$$

This is the set of gambles which, if $w \in [q = \gamma_h]$, pay off $X(w)$ for some $X \in C_{\gamma_j}(\mathcal{X})$.

The requirement that you expect letting q choose from $\mathcal{X}$ to be at least as good choosing yourself is therefore ambiguous between at least two possible readings. It could either mean that (i) for *every way* of adjudicating between maximal options for q, you expect q's choice to be at least as good as yours, or (ii) for *some way* of adjudicating between maximal options for q, you expect q's choice to be at least as good as yours. This ambiguity will be of central importance once we move to imprecise credences. But luckily, in the precise case, the two readings turn out to be equivalent (Dorst et al., 2021)[Lemma B.13]. So the requirement is not ambiguous after all, and we can formally define it as follows (using reading (ii)):

- **Pragmatic Value**:
  For every $\mathcal{X} = \{X_1, ..., X_k\}$ and every $X_s \in \mathcal{X}$, there is some $Y \in \sum_{j=1}^{m} [q = \gamma_j] C_{\gamma_j}(\mathcal{X})$ such that:

$$\pi(X_s) \leq \pi(Y) \tag{47}$$

when this constraint holds, we say that π Pragmatic Values q.

With this definition in place, we can prove the following result.

Proposition 6. *Let π be a credence function and q the definite description of a Clear credence function. Then π Global Reflects q iff π Pragmatic Values q.*

Proof. The earliest proof I could find for this result is due to (Huttegger, 2014). Dorst et al. (2021)[Theorem 2.2] proves a more general theorem with notation similar to the one used in this essay. ☐

Proposition 6 shows that under Clarity, if we plug a pragmatic notion of value into Value Superiority, the resulting notion of deference is equivalent to the one characterised by Global Reflection.

6.2 The Imprecise Case

It's not at all straightforward to give a good analogue of the Pragmatic Value constraint for imprecise credences. The idea behind the Pragmatic Value requirement is that you find Q to be at least as valuable as your doxastic state when you expect letting Q choose for you to be at least as good as choosing for yourself on any decision problem. And as before, we can use choice functions to make this sort of requirement more precise. To any credal set Π we can associate a *choice function* $C_\Pi : 2^{\mathcal{L}(\mathcal{W})} \to 2^{\mathcal{L}(\mathcal{W})}$ defined by:

$$C_\Pi(\mathcal{X}) = \{X \in \mathcal{X} : X \text{ maximises } \pi(X) \text{ in } \mathcal{X} \text{ for some } \pi \in \Pi\} \tag{48}$$

Intuitively, if $C_\Pi(\mathcal{X})$ is the set of options such that at least one $\pi \in \Pi$ considers them choiceworthy among those in $\mathcal{X}$. This is known as the set of *E-Admissible options* for Π among $\mathcal{X}$ (Levi, 1980; Seidenfeld et al., 2010).

As in the precise case, we can start by asking: what is the value of choosing for yourself? The set of choiceworthy gambles $C_\Pi(\mathcal{X})$ will generally contain multiple gambles, each taking different values at different possible worlds. So "choosing for yourself" does not determine a unique payoff at every possible world. But unlike the precise case, the various credence functions in your credal set will generally disagree about the expected value of the various gambles $X \in C_\Pi(\mathcal{X})$, both in the sense that each credence function may assign different expected value to different gambles in this set, and in the sense that different credence functions may assign different expected value to the same gamble in this set. So when formulating our constraint, it's not clear that we can assign an expected value or range of expected values to the option "choosing for yourself".

And what is the value of letting Q choose for you? As in the precise case, this also does not uniquely define a payoff at every possible world. If $w \in [Q = \Gamma_j]$, by letting Q choose for you you will receive value $X(w)$ for some gamble $X \in C_{\Gamma_j}(\mathcal{X})$. But Γ_j can have multiple choiceworthy options in $\mathcal{X}$, and each of them could have different values at the worlds in $[Q = \Gamma_j]$. So Q's choice could have the value of any gamble in the following set, depending on how we adjudicate between choiceworthy options for the various Γ_j's:

$$\sum_{j=1}^{m}[Q = \Gamma_j]C_{\Gamma_j}(\mathcal{X}) \tag{49}$$

So both "choosing for yourself" and "letting Q choose for you" correspond to sets of gambles. Furthermore, each credence function in your credal set may assign different expected values to different gambles within each set, and different credence functions in your credal set may assign different expected values to the same gamble in either set.

In light of this, given a problem $\mathcal{X}$, there are quite a few ways to express that you expect letting Q choose for you from $\mathcal{X}$ to be at least as good as choosing from $\mathcal{X}$ yourself, which all collapse to Pragmatic Value if both Π and Q are precise. For example:

- For every $X \in C_\Pi(\mathcal{X})$ and for every $Y \in \sum_{j=1}^m [Q = \Gamma_j] C_{\Gamma_j}(\mathcal{X})$ and for every $\pi \in \Pi$, $\pi(X) \leq \pi(Y)$.
- For every $X \in C_\Pi(\mathcal{X})$ and for every $Y \in \sum_{j=1}^m [Q = \Gamma_j] C_{\Gamma_j}(\mathcal{X})$, there is some $\pi \in \Pi$ such that $\pi(X) \leq \pi(Y)$.
- For every $X \in C_\Pi(\mathcal{X})$ there is some $Y \in \sum_{j=1}^m [Q = \Gamma_j] C_{\Gamma_j}(\mathcal{X})$ such that for every $\pi \in \Pi$, $\pi(X) \leq \pi(Y)$.
- For every $X \in C_\Pi(\mathcal{X})$ there is some $Y \in \sum_{j=1}^m [Q = \Gamma_j] C_{\Gamma_j}(\mathcal{X})$ and there is some $\pi \in \Pi$ such that $\pi(X) \leq \pi(Y)$.

all collapse to Pragmatic Value when Π and Q are precise.[17] See (Seidenfeld, 2004; Kadane et al., 2008; Bradley and Steele, 2016, 2014) for a discussion of possible characterisations of Pragmatic Value for agents with imprecise credences in the context of sequential decision-making.

A number of possible characterisations of Pragmatic Value turn out to be equivalent to our imprecise deference principles when Q denotes a *closed* credal set.[18] As mentioned earlier, many popular frameworks in the Imprecise Bayesian literature end up being equivalent to closed and convex credal sets (Wheeler, 2022), but not much attention has been devoted to the closure property for non-convex credal sets. Yet this property seems essential for the formal proofs given in this essay; while it would be desirable to relax it, I will leave it for future work to establish whether this can be done.

As an illustrative example of the characterisation of deference in terms of Pragmatic Value, I will focus on the following (comparatively weak) imprecise version of Value:

- **Weak Pragmatic Value**
  For every $\mathcal{X} = \{X_1, ..., X_k\}$, every $X \in \mathcal{X}$, and every $\pi \in \Pi$, there is some $Y \in \sum_{j=1}^m [Q = \Gamma_j] C_{\Gamma_j}(\mathcal{X})$ such that:
  $$\pi(X) \leq \pi(Y) \tag{50}$$

This requires that for every $X \in \mathcal{X}$, and every credence function in Π, there is *some way* of adjudicating between the choiceworthy options of each Γ_j such that "letting Q choose for you" is no worse than choosing X according to π. In other words, no matter what option you would pick from $\mathcal{X}$, each credence function in your credal set thinks that Q could (by adjudicating between choiceworthy options in a certain way) make at least as good of a choice. I consider this to be a plausible generalisation of Pragmatic Value for agents whose doxastic state is represented by a

[17] Note that these are not all the logical possibilities, but only "extreme ones". Instead of existentially/universally quantifiying on a set, we could existentially/universally quantify on some appropriately chosen subset to derive a stronger/weaker principle, respectively.

[18] We can think of a probability function π on $\mathcal{W} = \{w_1, ..., w_n\}$ as the vector $(\pi(w_1), ...\pi(w_n))$ in $\mathbb{R}^n$, and similarly think of a credal set as a subset of $\mathbb{R}^n$. A credal set Γ is *closed* iff the corresponding subset of $\mathbb{R}^n$ is closed with regards to the Euclidean topology. In turn this means that if $\pi \notin \Gamma$, there is some $\epsilon > 0$ such that for every $\gamma \in \Gamma$, the Euclidean distance between γ and π is greater than ϵ.

credal set. When all credal sets are singletons, it clearly coincides with Pragmatic Value.

The main formal result of this section shows that if Q is Clear and closed, Weak Pragmatic Value is equivalent to Subset Reflection.

Proposition 7. *Let Π be a regular credal set, and let Q the definite description of a Clear and closed credal set. Then Π Weak Pragmatic Values Q iff Π Sub-Reflects Q.*

Proof. See Appendix. □

This gives us a way to motivate Subset Reflection in contexts beyond those where (1) Q is your future doxastic state. The assumption that you defer to someone when you consider their credences to be at least as valuable as yours is a natural one in contexts (2-4). If value is defined as in Weak Pragmatic Value, then you defer to someone in those contexts iff you Subset Reflect them.

At the end of the previous section, I pointed out that there is intuitively a stronger notion of deference than the one characterised by Subset Reflection, since the latter allows one to defer to extremely uninformative experts while having extremely informative opinions. So it's natural to wonder whether there is a version of Pragmatic Value which yields a stronger deference principle than Subset Reflection.

We can approach this question by ranking the experts you Weak Pragmatic Value in terms of their informativeness, where an expert is more informative than another when they make strictly more judgements at every possible world. You defer to all the experts you Weak Pragmatic Value, in the weak sense that upon learning their doxastic state, you would make all the judgements they make. But you defer to the most informative ones in a stronger sense: upon learning their doxastic state, you would make no more judgements than they do. In other words, if they were any more informative, you would not Weak Pragmatic Value them anymore.

Looking at comparative informativeness allows us to capture a stronger sense of deference than the one captured by Pointwise/Subset Reflection. To get an intuitive idea of how this works, consider Example 5.3 again. Your friend Lucy's credal set is denoted by Q. Now let T denote the convex credal set of Tom, whose opinions are more informative than Lucy's at every possible world. In particular:

$$T_1 = T_3 = \{\gamma \in \mathcal{P}_{\mathcal{W}} : \gamma(w_1) = .9, \gamma(w_3) = .1, \gamma(w_2) = \gamma(w_4) = 0)\} \quad (51)$$
$$T_2 = T_4 = \{\gamma \in \mathcal{P}_{\mathcal{W}} : \gamma(w_1) = \gamma(w_3) = 0, \gamma(w_2) = \gamma(w_4) = 1/2\} \quad (52)$$

At the worlds $w_i \notin r$ where the nespapers don't report the polls, Lucy has $Q_i(d) = [.01, .99]$ whereas Tom has $T_i(d) = \{1/2\}$. This can happen, for example, if Tom reads the same newspapers as Lucy, but is less epistemically conservative than her, and so is disposed to make more judgements when facing scarce evidence. Recall that Martha was an agent who, upon learning that $Q(d) = [.01, .99]$, assigned probability interval $[.01, .99]$ to d. Martha Subset Reflects Lucy, and thus Weak Pragmatic Values her.

However, Martha *does not* Weak Pragmatic Value Tom, since:

$$\Pi^M(d|[T = T_2]) = \Pi^M(d|\neg r) = \Pi^M(d|[Q = Q_2]) = [.01, .99] \nsubseteq [.4, .6] \tag{53}$$

Indeed, Martha does not Weak Pragmatic Value any definite description which is more informative than Q. Intuitively this is because her epistemic values and total evidence are aligned with those of Lucy, so she would not endorse the further judgements made by Tom or any other more informative expert in the worlds where Lucy has no evidence about whether d. Stewart on the other hand Weak Pragmatic Values Tom, as well as many other more informative credal sets. This suggests that, when you defer to someone in the stronger sense in which Martha defers to Lucy, they are maximally informative amongst the experts you Weak Pragmatic Value.

The most informative experts among those you Weak Pragmatic Value turn out to be exactly those which you Identity Reflect. To show this, we first need to make precise our intuitive notion of informativeness of a definite description.

Definition 6.1 (Informativeness of a definite description)**.** Let Q, Q' be definite descriptions of credal sets defined on the same domain $\mathcal{W}$. Q is at least as informative as Q', written $Q \subseteq Q'$, iff for every $w_i \in \mathcal{W}$, $Q_i \subseteq Q'_i$.

Now let $\mathcal{Q}$ be a set of definite descriptions. We say Q is *maximally informative in* $\mathcal{Q}$ iff there is no $Q' \in \mathcal{Q}$ such that $Q' \subseteq Q$ and $Q \nsubseteq Q'$. Now we can show that the maximally informative definite descriptions among those that Π Weak Pragmatic Values are exactly the definite descriptions that Π Identity Reflects.

Proposition 8. *Let* Π *be a regular credal set, and let* $\mathcal{Q}$ *be the set of Clear and Closed definite descriptions that* Π *Weak Pragmatic Values. Then* Q *is maximally informative in* $\mathcal{Q}$ *iff* Π *Identity Reflects* Q*.*

Proof. This follows from Proposition 7 together with Theorem 5.8 in Grunwald and Halpern (2007). A proof using the present notation is given in the Appendix. □

We can now motivate Identity Reflection from the following extension of Value Superiority, which distinguishes between a weak and a strong notion of deference:

- **Value-Information Superiority**:
    1. You *weakly* defer to a doxastic state iff you regard it to be at least as valuable as your doxastic state.
    2. You *strongly* defer to a doxastic state iff it is maximally informative among those you defer to.

By plugging a definition of value and (comparative) informativeness of a doxastic state into Value-Information Superiority, we can derive two notions of epistemic deference for imprecise credences. If value is defined by Weak Pragmatic Value and Informativeness is defined as in Definition 6.1, then you defer to someone in the weaker sense iff you Subset Reflect them (Proposition 7), and you defer to someone in the stronger sense iff you Identity Reflect them (Proposition 8). In Example 5.3, Martha defers to Lucy in the stronger sense captured by Identity Reflection, whereas Stewart only defers to Lucy in the weaker sense captured by Subset Reflection.

7 Conclusion

This essay makes some progress towards a systematic treatment of epistemic deference for agents with imprecise credences. After outlining four different contexts where deference has been relevant for precise Bayesian epistemology, I set out to defend some deference principles for imprecise probabilities using assumptions appropriate to each context.

I have shown how two imprecise deference principles, Identity and Pointwise Reflection, are related: the latter is equivalent, for convex credal sets, to Subset Reflection, which in turn is strictly weaker than Identity Reflection. Then I have shown that when you are certain that you will update by conditionalisation, you should defer to your future self according to Identity Reflection.

When the expert is someone other than your future self, we can derive a deference principle from the assumption that you defer to someone if and only if you consider their doxastic state to be at least as valuable as your own. I have shown that, if value is appropriately cashed out in pragmatic terms, this assumption yields Subset/Pointwise Reflection. I have also argued that we may be interested in a stronger notion of deference than the one characterised by Subset Reflection. This can be defended by supplementing the value assumption with an ordering of the doxastic states in terms of their informativeness. I have shown that the doxastic states you Identity Reflect are the most informative ones among the doxastic states which you consider at least as valuable as your own.

References

Bradley, S. (2019). Imprecise probabilities. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2019 edition.

Bradley, S. and Steele, K. (2014). Uncertainty, learning, and the "problem" of dilation. *Erkenntnis*, 79:1287–1303.

Bradley, S. and Steele, K. (2016). Can free evidence be bad? Value of information for the imprecise probabilist. *Philosophy of Science*, 83(1):1–28.

Briggs, R. (2009). Distorted reflection. *Philosophical Review*, 118(1):59–85.

Campbell-Moore, C. (2021). Probability filters as a model of belief; comparisons to the framework of desirable gambles. In *International Symposium on Imprecise Probability: Theories and Applications*, pages 42–50. PMLR.

Carr, J. R. (2020). Imprecise evidence without imprecise credences. *Philosophical Studies*, 177:2735–2758.

Dorst, K. (2019). Higher-order uncertainty. *Higher-order evidence: New essays*, pages 35–61.

Dorst, K. (2020). Evidence: A guide for the uncertain. *Philosophy and Phenomenological Research*, 100(3):586–632.

Dorst, K. (2023). Bayesian ambiguity: Imprecision or uncertainty? *Unpublished Manuscript*.

Dorst, K., Levinstein, B. A., Salow, B., Husic, B. E., and Fitelson, B. (2021). Deference done better. *Philosophical Perspectives*, 35(1):99–150.

Elga, A. (2013). The puzzle of the unmarked clock and the new rational reflection principle. *Philosophical Studies*, 164:127–139.

Gaifman, H. (1988). A theory of higher order probabilities. In Skyrms, B. and Harper, W. L., editors, *Causation, Chance and Credence: Proceedings of the Irvine Conference on Probability and Causation Volume 1*, pages 191–219. Springer Netherlands, Dordrecht.

Gallow, J. D. (2023). Local and global deference. *Philosophical Studies*, 180(9):2753–2770.

Good, I. J. (1967). On the principle of total evidence. *The British Journal for the Philosophy of Science*, 17(4):319–321.

Grunwald, P. D. and Halpern, J. Y. (2007). A game-theoretic analysis of updating sets of probabilities. *arXiv preprint arXiv:0711.3235*.

Hall, N. (1994). Correcting the guide to objective chance. *Mind*, 103(412):505–517.

Huttegger, S. M. (2014). Learning experiences and the value of knowledge. *Philosophical Studies*, 171:279–288.

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

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

Kadane, J. B., Schervish, M., and Seidenfeld, T. (2008). Is ignorance bliss? *The Journal of Philosophy*, 105(1):5–36.

Keynes, J. M. (1973 [1921]). *A treatise on probability*, volume 8. Macmillian, London.

Konek, J. (2019). Epistemic conservativity and imprecise credence. *Philosophy and phenomenological research*.

Konek, J. (2023). Evaluating imprecise forecasts. In *International Symposium on Imprecise Probability: Theories and Applications*, pages 270–279. PMLR.

Levi, I. (1974). On indeterminate probabilities. *The Journal of Philosophy*, 71(13):391–418.

Levi, I. (1980). *The enterprise of knowledge: An essay on knowledge, credal probability, and chance*. MIT press.

Levinstein, B. A. (2017). A pragmatist's guide to epistemic utility. *Philosophy of Science*, 84(4):613–638.

Lewis, D. (1980). A subjectivist's guide to objective chance. In Jeffrey, R. C., editor, *Studies in Inductive Logic and Probability*, volume 2, pages 267–297. University of California Press.

Mahtani, A. (2019). Imprecise probabilities. In Pettigrew, R. and Weisberg, J., editors, *The Open Handbook of Formal Epistemology*, pages 107–130. PhilPapers Foundation.

Mayo-Wilson, C. and Wheeler, G. (2016). Scoring imprecise credences. *Philosophy and Phenomenological Research*, 93(1):55–78.

Molinari, G. (2023). Trust the evidence: two deference principles for imprecise probabilities. In *International Symposium on Imprecise Probability: Theories and Applications*, pages 356–366. PMLR.

Moss, S. (2018). *Probabilistic knowledge*. Oxford University Press.

Moss, S. (2021). Global constraints on imprecise credences: Solving reflection violations, belief inertia, and other puzzles. *Philosophy and Phenomenological Research*, 103(3):620–638.

Myrvold, W. C. (2012). Epistemic values and the value of learning. *Synthese*, 187:547–568.

Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press.

Pettigrew, R. and Titelbaum, M. G. (2014). Deference done right. *Philosophers' Imprint*, 14:1–19.

Quaeghebeur, E. (2014). Desirability. In *Introduction to Imprecise Probabilities*, chapter 1, pages 1–27. John Wiley and Sons, Ltd.

Savage, L. J. (1954). *The foundations of statistics*. Wiley, New York.

Schoenfield, M. (2012). Chilling out on epistemic rationality: A defense of imprecise credences (and other imprecise doxastic attitudes). *Philosophical Studies*, 158(2):197–219.

Schoenfield, M. (2017). The accuracy and rationality of imprecise credences. *Noûs*, 51(4):667–685.

Seidenfeld, T. (2004). A contrast between two decision rules for use with (convex) sets of probabilities: γ-maximin versus e-admissibility. *Synthese*, 140(1/2):69–88.

Seidenfeld, T., Schervish, M. J., and Kadane, J. B. (2010). Coherent choice functions under uncertainty. *Synthese*, 172:157–176.

Seidenfeld, T., Schervish, M. J., and Kadane, J. B. (2012). Forecasting with imprecise probabilities. *International Journal of Approximate Reasoning*, 53(8):1248–1261.

Skyrms, B. (1990). *The dynamics of rational deliberation*. Harvard University Press.

Smith, C. A. (1961). Consistency in statistical inference and decision. *Journal of the Royal Statistical Society: Series B (Methodological)*, 23(1):1–25.

Titelbaum, M. (2015). Rationality's fixed point (or: In defense of right reason). *Oxford studies in epistemology*, 5:253–294.

Topey, B. (2012). Coin flips, credences and the reflection principle. *Analysis*, 72(3):478–488.

Van Camp, A. (2018). *Choice Functions as a Tool to Model Uncertainty*. PhD thesis, Ghent University.

Van Fraassen, B. C. (1984). Belief and the will. *The Journal of Philosophy*, 81(5):235–256.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*, volume 42. Springer.

Weisberg, J. (2007). Conditionalization, reflection, and self-knowledge. *Philosophical Studies*, 135:179–197.

Wheeler, G. (2022). A gentle approach to imprecise probability. In *Reflections on the Foundations of Probability and Statistics: Essays in Honor of Teddy Seidenfeld*, pages 37–67. Springer.

White, R. (2010). Evidential symmetry and mushy credence. *Oxford studies in epistemology*, 3:161–186.

Williams, P. M. (1976). Indeterminate probabilities. In *Formal Methods in the Methodology of Empirical Sciences: Proceedings of the Conference for Formal Methods in the Methodology of Empirical Sciences, Warsaw, June 17–21, 1974*, pages 229–246. Springer.

Appendices

A Proof of the Results

Remark 1. Let Π be a regular credal set, let Q be the definite description of a Clear credal set, and let $\mathcal{E} = \{e_1, ..., e_m\}$ be a partition. Π respects Certain Conditionalisation with regards to Q and $\mathcal{E}$ iff it obeys the following two conditions:

i $\Pi([\Pi(X|e_j) = Q(X|e_j)]|e_j) = \{1\}$.

ii $\Pi(e_j \Leftrightarrow [Q(e_j) = \{1\}]) = \{1\}$.

Proof. For the right-to-left direction, assume Π obeys (i) and (ii). Note that $\Pi([\Pi(X|e_j) = Q(X|e_j)]|e_j)$ is equal to $\Pi(\Pi(X|e_j) = Q(X))$, since $\Pi([Q(e_j) = \{1\}]|e_j) = \{1\}$ by (ii), and thus $\Pi([Q(X) = Q(X|e_j)|e_j]) = \{1\}$. So for every X, we have:

$$\Pi([\Pi(X|e_j) = Q(X)]|e_j) = \Pi([\Pi(X|e_j) = Q(X|e_j)]|e_j) = \{1\} \tag{54}$$

This can be true only if $\Pi([\Pi(\cdot|e_j) = Q]|e_j) = \{1\}$, thus proving Certain Conditionalisation.

For the left-to-right direction, assume Π obeys Certain Conditionalisation. We can show that (i) holds as follows:

$$\Pi([\Pi(X|e_j) = Q(X|e_j)]|e_j) = \{1\} \tag{55}$$

$$\iff \Pi([\Pi(X|e_j) = \Pi(X|e_j)]|e_j) = 1 \text{ by Certain Conditionalisation} \tag{56}$$

where the last equality holds by coherence. We can show that (ii) holds as follows whenever $e_j \neq \mathcal{W}$:

$$\Pi(e_j \Leftrightarrow [Q(e_j) = \{1\}]) = \{1\} \tag{57}$$

$$\iff \pi(e_j \Leftrightarrow [Q(e_j) = \{1\}]) = 1 \text{ for all } \pi \in \Pi \tag{58}$$

$$\iff \pi(e_j \wedge [Q(e_j) = \{1\}]) + \pi(\neg e_j \wedge [Q(e_j) \neq \{1\}]) = 1 \tag{59}$$

$$\iff \pi([Q(e_j) = \{1\}]|e_j)\pi(e_j) + \sum_{e_k \neq e_j} \pi([Q(e_j) \neq \{1\}]|e_k)\pi(e_k) = 1 \tag{60}$$

$$\iff \pi([\Pi(e_j|e_j) = \{1\}]|e_j)\pi(e_j) + \sum_{e_k \neq e_j} \pi([\Pi(e_j|e_k) \neq \{1\}]|e_k)\pi(e_k) = 1 \tag{61}$$

$$\iff \pi(e_j) + \sum_{e_k \neq e_j} \pi(e_k) = 1 \tag{62}$$

where the last equality holds by coherence. The proof is trivial when $e_j = \mathcal{W}$. □

To prove Lemma 1 it helps to first prove the following lemma:

Lemma 0. *Let Π be a regular credal set, Q a definite description of a credal set, and $\mathcal{E} = \{e_1, ..., e_m\}$ a finite partition. Assume Π respects Certain Conditionalisation w.r.t. Q and $\mathcal{E}$. Then for every $w_j, w_k \in \mathcal{W}$, $Q_i = Q_k$ iff w_i, w_k belong to the same $e_j \in \mathcal{E}$.*

Proof. Let $w_i, w_k \in \mathcal{W}$. For the right-to-left direction, assume by contradiction that $w_i, w_k \in e_j$ and yet $Q_i \neq Q_k$. By regularity, $\Pi(\cdot|e_j)$ is well-defined. So either $\Pi(\cdot|e_j) \neq Q_i$, or $\Pi(\cdot|e_j) \neq Q_k$ (or both). Assume without loss of generality that $\Pi(\cdot|e_j) \neq Q_i$. By Certain Conditionalisation we have:

$$\Pi([Q = \Pi(\cdot|e_j)]|e_j) = \{1\} \tag{63}$$

Now pick any $\pi \in \Pi$. From the above equality we have:

$$\pi([Q = \Pi(\cdot|e_j)]|e_j) = 1 \tag{64}$$
$$\iff \pi([Q = \Pi(\cdot|e_j)] \wedge e_j) = \pi(e_j) \tag{65}$$
$$\iff \pi(e_j) - \pi([Q \neq \Pi(\cdot|e_j)] \wedge e_j) = \pi(e_j) \tag{66}$$

but note that $w_i \in [Q \neq \Pi(\cdot|e_j)] \wedge e_j$, because $w_i \in e_j$ and $Q_i \neq \Pi(\cdot|e_j)$. Therefore $\pi([Q \neq \Pi(\cdot|e_j)] \wedge e_j) > 0$. So the last equation above is a contradiction.

The proof of the other direction is similar. Assume by contradiction that $w_i \in e_i$ and $w_k \in e_k$ with $e_i \neq e_k$, and yet $Q_i = Q_k$. Then either $Q_i \neq \Pi(\cdot|e_i)$ or $Q_k \neq \Pi(\cdot|e_k)$ (or both). Assume without loss of generality that $Q_i \neq \Pi(\cdot|e_i)$. Pick some $\pi \in \Pi$. As above, we have from Certain Conditionalisation that:

$$\pi([Q = \Pi(\cdot|e_i)]|e_i) = 1 \tag{67}$$
$$\iff \pi([Q = \Pi(\cdot|e_i)] \wedge e_i) = \pi(e_i) \tag{68}$$
$$\iff \pi(e_i) - \pi([Q \neq \Pi(\cdot|e_i)] \wedge e_i) = \pi(e_i) \tag{69}$$

But we know $w_i \in e_i$ and $w_i \in [Q \neq \Pi(\cdot|e_i)]$, so $\pi([Q \neq \Pi(\cdot|e_i)] \wedge e_i) > 0$, contradiction. □

We can now prove Lemma 1:

Lemma 1. *Let Π be a regular credal set, Q a definite description of a credal set, and $\mathcal{E} = \{e_1, ..., e_m\}$ a finite partition. Assume Π respects Certain Conditionalisation w.r.t. Q and $\mathcal{E}$. Then for every $e_j \in \mathcal{E}$:*

$$[Q = \Pi(\cdot|e_j)] = e_j \tag{70}$$

Proof. Let $w_i \in [Q = \Pi(\cdot|e_j)]$, and assume by way of contradiction that $w_i \notin e_j$. Then $w_i \in e_k \neq e_j$, since $\mathcal{E}$ is a partition. By Lemma 0, $e_k \subseteq [Q = Q_i] = [Q = \Pi(\cdot|e_j)]$. But since $\Pi(\cdot|e_j) \neq \Pi(\cdot|e_k)$, this means that $\Pi([Q = \Pi(\cdot|e_k)]|e_k) = \{0\}$, violating Certain Conditionalisation.

Now let $w_i \in e_j$ and assume by way of contradiction that $w_i \notin [Q = \Pi(\cdot|e_j)]$. Then by Lemma 0, $e_j \subseteq [Q = Q_i] \subseteq [Q \neq \Pi(\cdot|e_j)]$. But then $\Pi([Q \neq \Pi(\cdot|e_j)|e_j]) = \{1\}$, violating Certain Conditionalisation. □

Proposition 2 (Certain Conditionalisation $\Rightarrow$ Identity Reflection)**.** *Let Π be a regular credal set, Q a definite description of a credal set, and $\mathcal{E} = \{e_1, ..., e_m\}$ a finite partition. If Π respects Certain Conditionalisation w.r.t. Q and $\mathcal{E}$, then Π Identity Reflects Q.*

Proof. Assume by way of contradiction that $\Pi(\cdot|[Q = Q_i]) \neq Q_i$ for some w_i such that this conditional credal set is defined. Let $\pi \in \Pi$ and $e_j \in \mathcal{E}$. From Lemma 1 we know that $[Q = Q_i] = e_j$. By Certain Conditionalisation, we then have:

$$\pi([Q = \Pi(\cdot|e_j)]|e_j) = 1 \tag{71}$$
$$\iff \pi([Q = \Pi(\cdot|[Q = Q_i])] \wedge e_j) = \pi(e_j) \tag{72}$$
$$\iff \pi(e_j) - \pi([Q \neq \Pi(\cdot|[Q = Q_i])] \wedge e_j) = \pi(e_j) \tag{73}$$

but note that for every $w_k \in e_j$ we have $w_k \in [Q \neq \Pi(\cdot|[Q = Q_i])]$, since $Q_i = Q_k$ by Lemma 0 and $\Pi(\cdot|[Q = Q_i]) \neq Q_i$ by assumption. Therefore $e_j = [Q \neq \Pi(\cdot|[Q = Q_i])] \wedge e_j$, and because $\pi(e_j) > 0$ by regularity, the last equation above is a contradiction. $\square$

Proposition 3 (Sub-Reflection (Convex) $\Rightarrow$ Pointwise Reflection)**.** *Let Π be a regular credal set and Q the definite description of a Clear and convex credal set. If Π Sub-Reflects Q, then Π Pointwise-Reflects Q.*

Proof. Assume Π Sub-Reflects Q, where Q is the definite description of a convex credal set. Let $\{\Gamma_1, ..., \Gamma_m\}$ be the set of all possible values of Q. Let $X : \mathcal{W} \to \mathbb{R}$ and $S \subseteq \mathbb{R}$ be such that $0 \notin \Pi([Q(X) = S])$. Since every Γ_j is convex, this means that S must be convex, otherwise $[Q(X) = S]$ would be empty. Then for every $\pi \in \Pi$ we have:

$$\pi(X|[Q(X) = S]) = \frac{\pi(X[Q(X) = S])}{\pi([Q(X) = S])} \tag{74}$$

$$= \frac{\pi(X \bigcup_{j:\Gamma_j(X)=S}[Q = \Gamma_j])}{\pi(\bigcup_{j:\Gamma_j(X)=S}[Q = \Gamma_j])} \tag{75}$$

$$= \frac{\sum_{j:\Gamma_j(X)=S} \pi(X[Q = \Gamma_j])}{\sum_{j:\Gamma_j(X)=S} \pi([Q = \Gamma_j])} \tag{76}$$

$$= \sum_{j:\pi([Q=\Gamma_j])\neq 0, \Gamma_j(X)=S} \pi(X|[Q = \Gamma_j]) \frac{\pi([Q = \Gamma_j])}{\sum_{k:\Gamma_k(X)=S} \pi([Q = \Gamma_k])} \tag{77}$$

Let J be the set of indices $\{j : \Gamma_j(X) = S, \pi([Q = \Gamma_j]) \neq 0\}$. The equalities above show that, for every $\pi \in \Pi$, the value $\pi(X|[Q(X) = S])$ is a convex combination of the set of values $\{\pi(X|[Q = \Gamma_j]) : j \in J\}$. But since Π Sub-Reflects Q, we know that $\Pi(\cdot|[Q = \Gamma_j]) \subseteq \Gamma_j$, and consequently $\pi(X|[Q = \Gamma_j]) \in \Gamma_j(X) = S$ for every $\pi \in \Pi$ and $j \in J$. Therefore, because $\pi(X|[Q(X) = S])$ is a convex combination of values in the convex set S, it must also belong to the set S, completing the proof. $\square$

To prove Proposition 4 I first prove the following lemma:

Lemma 2. *Let Π be a regular credal set, and Q the definite description of a convex and Clear credal set. For any set of real values S, I write $S > 0$ as shorthand for "$s > 0$ for every $s \in S$". Then Π Pointwise-Reflects Q only if:*

$$\Pi(X[Q(X) > 0]) > 0 \tag{78}$$

whenever this credal set is defined.[19]

Proof. Assume Π Pointwise Reflects Q. Let X be a random variable such that $\Pi(\cdot|[Q(X) > 0])$ is defined. So $\pi([Q(X) > 0]) > 0$ for every $\pi \in \Pi$. Define the set of sets of real values $\mathcal{S}(X)$ as follows:

$$\mathcal{S}(X) = \{Q_i(X) : w_i \in \mathcal{W}\} \tag{79}$$

[19]This condition was proposed by Molinari (2023) as a deference principle for imprecise credences in cases where Q is not Clear. It is an imprecise analogue of the Total Trust principle introduced by Dorst et al. (2021).

Since $\mathcal{W}$ is finite we can write $\mathcal{S}$ as $\{S_1, ..., S_M\}$ for some $M \in \mathbb{N}$, and $M \geq 1$. For every $\pi \in \Pi$ we have:

$$\pi(X[Q(X) > 0]) > 0 \tag{80}$$

$$\iff \pi(X \bigcup_{k:S_k>0} [Q(X) = S_k]) > 0 \tag{81}$$

$$\iff \sum_{k:S_k>0} \pi(X[Q(X) = S_k]) > 0 \text{ because the } [Q(X) = S_k] \text{ are disjoint} \tag{82}$$

$$\iff \sum_{k:S_k>0, \pi([Q(X)=S_k])\neq 0} \pi(X[Q(X) = S_k]) > 0 \tag{83}$$

We can show this sum is strictly positive by showing that it has at least one addend, and that every addend is strictly positive. Since $\pi([Q(X) > 0]) > 0$, there is at least some k such that $\pi([Q(X) = S_k]) \neq 0$ and $S_k > 0$. So the sum has at least one addend. For every such k, by Pointwise Reflection we have $\pi(X[Q(X) = S_k]) > 0$ if and only if $\pi(X|[Q = S_k]) > 0$, which is true by Pointwise Reflection, as $\pi(X|[Q = S_k]) \in S_k$. This completes the proof. □

With this lemma in place, we can now prove Proposition 4:

Proposition 4 (Pointwise Reflection (Convex) $\Rightarrow$ Sub-Reflection)**.** *Let Π be a regular credal set and Q the definite description of a convex and Clear credal set. If Π Pointwise-Reflects Q, then Π Sub-Reflects Q.*

Proof. For any set of real values S, I write $S > 0$ as shorthand for "$s > 0$ for every $s \in S$". Assume Π Pointwise Reflects Q. Let Γ_j be a possible value of Q, and let X an arbitrary random variable such that $\Gamma_j(X) > 0$. From Clarity we also have $\Gamma_j(X[Q = \Gamma_j]) > 0$, since $\gamma(X[Q = \Gamma_j]) = \gamma(X)$ for every $\gamma \in \Gamma_j$. Then we have the following equality between random variables:

$$X[Q = \Gamma_j] = X[Q = \Gamma_j][Q(X[Q = \Gamma_j]) > 0] \tag{84}$$

But from Lemma 2, we have:

$$\Pi(X[Q = \Gamma_j][Q(X[Q = \Gamma_j]) > 0]) > 0 \tag{85}$$

and hence $\Pi(X[Q = \Gamma_j]) > 0$. For every $\pi \in \Pi$ we have $\pi([Q = \Gamma_j]) > 0$ by regularity, and so this entails that $\pi(X|[Q = \Gamma_j]) > 0$ for every $\pi \in \Pi$. This shows that $\Pi(X|[Q = \Gamma_j]) > 0$ for any random variable X such that $\Gamma_j(X) > 0$. Now assume by way of contradiction that $\Pi(\cdot|[Q = \Gamma_j]) \nsubseteq \Gamma_j$. Then we can find $\pi \in \Pi$ such that $\pi(\cdot|[Q = \Gamma_j]) \notin \Gamma_j$. Since Γ_j is convex, using the hyperplane separation theorem we can find a hyperplane (i.e. a random variable) X such that $\pi(X|[Q = \Gamma_j]) \leq 0$ and $\Gamma_j(X) > 0$, contradiction. Since Γ_j was an arbitrary possible value for Q, this proves that Π Sub-Reflects Q. □

Lemma 3. *Let π be a regular credence function and q a definite description of a credence function. Assume π respects Certain Conditionalisation with regards to q. Then π Local Reflects q iff π Global Reflects q.*

Proof. This follows from Propositions 3 and 4 by noting that, when all credal sets involved are singletons, Subset Reflection collapses to Global Reflection and Pointwise Reflection collapses to Local Reflection. □

The proof of Proposition 7 will rely on the following theorem by Seidenfeld et al. (2010).[20]

Theorem 1 (Seidenfeld et al. (2010), Theorem 2)**.** *Let π be a credence function and Γ a credal set. There is a finite set of gambles $\mathcal{X}_*$ such that:*

(a) $C_\Gamma(\mathcal{X}_*) = \mathcal{X}_*$ *iff* $\pi \in \Gamma$*, and*

(b) $\pi(X)$ *is constant for all* $X \in \mathcal{X}_*$*.*

Proof. Here is how to construct X_*. Let $k_\pi = \min\{\pi(w_i) : w_i \in \mathcal{W}, \pi(w_i) > 0\}$. Let X^* be the gamble that pays k_π at every possible world. For every w_i, define the gamble X_i as follows:

$$X_i(w_j) = \begin{cases} 1 & \text{if } i = j, \pi(w_j) = 0 \\ k_\pi & \text{if } i \neq j, \pi(w_j) = 0 \\ \frac{k_\pi}{\pi(w_j)} & \text{if } i = j, \pi(w_j) > 0 \\ 0 & \text{if } i \neq j, \pi(w_j) > 0 \end{cases} \tag{86}$$

Then let $X_* = \{X^*, X_1, ..., X_n\}$.

Prove (b) first. Clearly $\pi(X^*) = k_\pi$. And for every X_i, if $\pi(w_i) = 0$ then:

$$\pi(X_i) = \pi(w_i)1 + (1 - \pi(w_i))k_\pi = k_\pi \tag{87}$$

whereas if $\pi(w_i) \neq 0$:

$$\pi(X_i) = \pi(w_i)\frac{k_\pi}{\pi(w_i)} + (1 - \pi(w_i))0 = k_\pi \tag{88}$$

which proves (b).

To prove (a), first prove the right-to-left direction. Assume $\pi \in \Gamma$. By (b), every element of $\mathcal{X}_*$ maximises π, and so $C_\Gamma(\mathcal{X}_*) = \mathcal{X}_*$. To prove the left-to-right direction, we need to show that $C_\Gamma(\mathcal{X}_*) \neq \mathcal{X}_*$ unless $\pi \in \Gamma$. Let $\gamma \neq \pi$ be a probability function. Then we proceed by cases on whether $\pi_k = 1$.

Case 1: $k_\pi \neq 1$. There is some w_i such that $\gamma(w_i) > \pi(w_i)$, and we have:

$$\gamma(X_i) = \begin{cases} \gamma(w_i)\frac{k_\pi}{\pi(w_i)} & \text{if } \pi(w_i) > 0 \\ \gamma(w_1) + (1 - \gamma(w_i))k_\pi & \text{if } \pi(w_i) = 0 \end{cases} \tag{89}$$

and both are strictly greater than $\gamma(X^*) = k_\pi$ under the assumption that $k_\pi < 1$. This shows that $X^* \notin C_\Gamma(\mathcal{X}_*)$ unless $\pi \in \Gamma$.

Case 2: $k_\pi = 1$. Then $\pi(w_i) = 1$ for some w_i and $\mathcal{X}_* = \{X^*, X_i\}$, since X^* is the gamble that pays 1 at every possible world. So $\gamma(w_i) < 1$ and hence $\gamma(X_i) < 1 = \gamma(X^*)$. This shows that $X_i \notin C_\Gamma(\mathcal{X}_*)$ unless $\pi \in \Gamma$, completing the proof. □

[20]The original theorem is given for choice functions defined over horse lotteries. The statement and proof given here are straightforward analogues for choice functions defined over sets of gambles.

Before proving Proposition 7 it also helps to prove the following technical lemma about closed credal sets.

Lemma 4. *Let π a credence function and Γ a closed credal set defined on $\mathcal{W} = \{w_1, ..., w_n\}$. If $\pi \notin \Gamma$ then:*

$$d(\Gamma, \pi) := \inf\{\max\{\gamma(w_i) - \pi(w_i) : w_i \in \mathcal{W}\}\gamma \in \Gamma\} > 0 \tag{90}$$

Proof. Assume by way of contradiction that $d(\Gamma, \pi) \leq 0$. Note that for every $\gamma \in \Gamma$, $\max\{\gamma(w_i) - \pi(w_i) : w_i \in \mathcal{W}\}$ is non-negative, since both γ and π are probability functions. So $d(\Gamma, \pi) = 0$. This means that for every $\epsilon > 0$, there is some $\gamma \in \Gamma$ such that:

$$\max\{\gamma(w_i) - \pi(w_i) : w_i \in \mathcal{W}\} < \epsilon \tag{91}$$

Because γ and π are probability functions, this entails that:

$$\max\{|\gamma(w_i) - \pi(w_i)| : w_i \in \mathcal{W}\} < (n-1)\epsilon \tag{92}$$

This shows that:

$$\inf\{\max\{|\gamma(w_i) - \pi(w_i)| : w_i \in \mathcal{W}\}\gamma \in \Gamma\} = 0 \tag{93}$$

and thus π is a limit point of Γ. But Γ is closed, so $\pi \in \Gamma$, contradiction. $\square$

Proposition 7. *Let Π be a regular credal set, and let Q the definite description of a closed and Clear credal set. Then Π Weak Pragmatic Values Q iff Π Sub-Reflects Q.*

Proof. Let $\{\Gamma_1, ..., \Gamma_m\}$ the possible values of Q.

($\Rightarrow$): for the left-to-right direction, assume that Π does not Sub-Reflect Q. So there is some closed credal set Γ_t such that $\Pi(\cdot|[Q = \Gamma_t]) \not\subseteq \Gamma_t$. In particular, there is some $\pi^* \in \Pi$ such that $\pi^*(\cdot|[Q = \Gamma_t]) \notin \Gamma_t$. Note that $\pi^*(w_i|[Q = \Gamma_j])$ must be smaller than 1 for every w_i, because otherwise $[Q = \Gamma_j] = \{w_i\}$, in which case Γ will be equal to $\{\pi^*(\cdot|[Q = \Gamma_j])\}$ by Clarity.

Now construct the set of gambles $\mathcal{X}_* = \{X^*, X_1, ..., X_n\}$ as in the proof of Theorem 1, so that:

(a) $C_\Pi(\mathcal{X}_*) = \mathcal{X}_*$ iff $\pi^*(\cdot|[Q = \Gamma_t]) \in \Pi$, and

(b) $\pi^*(X|[Q = \Gamma_t])$ is constant for all $X \in \mathcal{X}_*$.

and X^* is the gamble paying some positive constant k_π at every possible world.

Start by showing that there is some $\epsilon > 0$ such that, for every $\gamma_t \in \Gamma_t$, there is some $X_i \in \mathcal{X}_*$ such that $\gamma_t(X_i) > \gamma_t(X^*) + \epsilon$. Lemma 4 shows that for some $\delta > 0$, every $\gamma_t \in \Gamma_t$ is such that:

$$\max\{\gamma(w_i) - \pi(w_i|[Q = \Gamma_t]) : w_i \in \mathcal{W}\} \geq \delta \tag{94}$$

Now for fixed $\gamma_t \in \Gamma_t$, let $w_i \in \mathcal{W}$ one of the worlds at which this maximum is achieved. Note that w_i must be in $[Q = \Gamma_t]$, for otherwise $\gamma_t(w_i) = 0 =$

$\pi(w_i|[Q = \Gamma_t])$. Let X_i the corresponding gamble in $\mathcal{X}_*$, defined by (86). We have:

$$\gamma_t(X_i) = \gamma_t(w_i)\frac{k_\pi}{\pi(w_i|[Q = \Gamma_t])} \tag{95}$$

$$\geq (\delta + \pi(w_i|[Q = \Gamma_t]))\frac{k_\pi}{\pi(w_i|[Q = \Gamma_t])} \tag{96}$$

$$= k_\pi + \frac{\delta k_\pi}{\pi(w_i|[Q = \Gamma_t])} \tag{97}$$

$$\geq k_\pi + \frac{\delta k_\pi}{\max\{\pi(w|[Q = \Gamma_t]) : w \in \mathcal{W}\}} \tag{98}$$

where the final expression is strictly greater than $\gamma_t(X^*) = k_\pi$, and independent of the choice of $\gamma_t \in \Gamma_t$. So there is some $\epsilon > 0$, such that for every $\gamma_t \in \Gamma_t$, there is some $X_i \in \mathcal{X}_*$ for which $\gamma_t(X_i) > \gamma_t(X^*) + \epsilon$.

Now construct the decision problem $\mathcal{X}'_*$ such that: for every $X \in \mathcal{X}_*$, if $X \neq X^*$ then $X[Q = \Gamma_t]$ is an element of $\mathcal{X}'_*$; $X^*[Q = \Gamma_t] + \epsilon$ is an element of $\mathcal{X}'_*$; and no other gamble is an element of $\mathcal{X}'_*$. We can now prove the following two claims:

- **Claim 1**: $C_{\Gamma_j}(\mathcal{X}'_*) = \{X^*[Q = \Gamma_t] + \epsilon\}$ whenever $\Gamma_j \neq \Gamma_t$.

To see that Claim 1 is true, let $\Gamma_j \neq \Gamma_t$. Then by Clarity, every $\gamma_j \in \Gamma_j$ has $\gamma_j([Q = \Gamma_t]) = 0$. So $\gamma_j(X^*[Q = \Gamma_t] + \epsilon) = \epsilon > 0$. If $X' \in \mathcal{X}'_*$ but $X' \neq X^*[Q = \Gamma_t] + \epsilon$, then $X' = X[Q = \Gamma_t]$ for some $X \in \mathcal{X}_*$, and hence $\gamma_j(X') = 0$. Therefore the only gamble in $C_{\Gamma_j}(\mathcal{X}'_*)$ is $X^*[Q = \Gamma_t] + \epsilon$.

- **Claim 2**: $X^*[Q = \Gamma_t] + \epsilon \notin C_{\Gamma_t}(\mathcal{X}'_*)$.

Claim 2 is true because we picked $\epsilon > 0$ small enough that, whenever $\gamma_t \in \Gamma_t$, there is some $X_t \in \mathcal{X}_*$ such that $\gamma_t(X_t) > \gamma_t(X^*) + \epsilon$. By Clarity, $\gamma_t(X) = \gamma_t(X[Q = \Gamma_t])$. Therefore:

$$\gamma_t(X_t[Q = \Gamma_t]) > \gamma_t(X^*[Q = \Gamma_t]) + \epsilon \tag{99}$$

$$= \gamma_t(X^*[Q = \Gamma_t] + \epsilon) \tag{100}$$

Since for every such X_t the gamble $X_t[Q = \Gamma_t]$ is an element of $\mathcal{X}'_*$, we have that $X^*[Q = \Gamma_t] + \epsilon$ is not an element of $C_{\Gamma_t}(\mathcal{X}'_*)$.

With this in place, we can show Weak Pragmatic Value is violated on $\mathcal{X}'_*$. From Claims 1 and 2 we have that "letting Q choose" on $\mathcal{X}'_*$ corresponds to the set of gambles:

$$\sum_{j=1}^{m}[Q = \Gamma_j]C_{\Gamma_j}(\mathcal{X}'_*) \tag{101}$$

$$= \{\sum_{j \neq t}[Q = \Gamma_j]([Q = \Gamma_t]X^* + \epsilon) + [Q = \Gamma_t]X_t : X_t \in C_{\Gamma_t}(\mathcal{X}'_*)\} \tag{102}$$

$$= \{\sum_{j \neq t}[Q = \Gamma_j]\epsilon + [Q = \Gamma_t]X_t : X_t \in C_{\Gamma_t}(\mathcal{X}'_*)\} \tag{103}$$

So for any $Y \in \sum_{j=1}^{m}[Q=\Gamma_j]C_{\Gamma_j}(\mathcal{X}'_*)$ we have:

$$\pi^*(X^*[Q=\Gamma_t]+\epsilon-Y) \tag{104}$$
$$=\pi^*(X^*[Q=\Gamma_t]+\epsilon-\sum_{j\neq t}[Q=\Gamma_j]\epsilon+[Q=\Gamma_t]X_t) \tag{105}$$
$$=\pi^*((X^*+\epsilon-X_t)[Q=\Gamma_t]) \tag{106}$$

for some $X_t[Q=\Gamma_t] \in C_{\Gamma_t}(\mathcal{X}'_*)$. This last expression must be strictly positive. This is because $\pi^*([Q=\Gamma_t])>0$ from regularity, and from (b) we know that $\pi(X^*-X_t|[Q=\Gamma_t])=0$ for every X_t, so $\pi((X^*-X_t)[Q=\Gamma_t])=0$. This shows that there is a member of $\mathcal{X}'_*-\delta$ which π^* strictly prefers to any member of $\sum_{j=1}^{m}[Q=\Gamma_j]C_{\Gamma_j}(\mathcal{X}'_*)$. So Weak Pragmatic Value is violated.

($\Leftarrow$): For the right-to-left direction, assume Π Sub-Reflects Q and let $\mathcal{X}$ be an arbitrary finite set of gambles. Let $\pi\in\Pi$, and let $X\in\mathcal{X}$ be a gamble that maximises π over $\mathcal{X}$. Now for each Γ_j, we can show that there is some $Y^j\in[Q=\Gamma_j]C_{\Gamma_j}(\mathcal{X})$ such that $\pi(X[Q=\Gamma_j])\leq\pi(Y^j)$. To see this, proceed by cases.
Case 1: $X\in C_{\Gamma_j}(\mathcal{X})$. Then let $Y^j=X[Q=\Gamma_j]$.
Case 2: $X\notin C_{\Gamma_j}(\mathcal{X})$. This means that for any $\gamma\in\Gamma_j$, there is some $X'\in C_{\Gamma_j}(\mathcal{X})$ such that $\gamma(X')>\gamma(X)$. By Subset Reflection, we have that $\Pi(\cdot|[Q=\Gamma_j])\subseteq\Gamma_j$, and hence $\pi(\cdot|[Q=\Gamma_j])\in\Gamma_j$. Therefore for some $X'\in C_{\Gamma_j}(\mathcal{X})$, $\pi(X'|[Q=\Gamma_j])>\pi(X|[Q=\Gamma_j])$, This in turn means $\pi(X'[Q=\Gamma_j])>\pi(X[Q=\Gamma_j])$, since from regularity $\pi([Q=\Gamma_j])>0$. So let $Y^j=X'[Q=\Gamma_j]$.

Then we can construct the following gamble:

$$\sum_{j=1}^{m}Y^j\in\sum_{j=1}^{m}[Q=\Gamma_j]C_{\Gamma_j}(\mathcal{X}) \tag{107}$$

and we have that:

$$\pi(X)=\sum_{j=1}^{m}\pi(X[Q=\Gamma_j])\leq\sum_{j=1}^{m}\pi(Y^j)=\pi(\sum_{j=1}^{m}Y^j) \tag{108}$$

So there is some gamble in $\sum_{j=1}^{m}[Q=\Gamma_j]C_{\Gamma_j}(\mathcal{X})$ which π expects to be at least as good as X. And since X was maximal for π on $\mathcal{X}$, and π was an arbitrary member of Π, this shows Π Weak Pragmatic Values Q. □

Proposition 8. *Let Π be a regular credal set, and let $\mathcal{Q}$ be the set of Clear and closed definite descriptions that Π Weak Pragmatic Values. Then Q is maximally informative in $\mathcal{Q}$ iff Π Identity Reflects Q.*

Proof. For the left to right direction, assume Q is maximally informative in $\mathcal{Q}$. Since $Q\in\mathcal{Q}$, we know Π Weak Pragmatic Values Q, and hence Π Sub-Reflects Q by Proposition 7. Let $\{\Gamma_1,...,\Gamma_m\}$ be the possible values of Q. Assume by way of contradiction that Π does not Identity Reflect Q. Then there must be some Γ_t such that:

$$\Pi(\cdot|[Q=\Gamma_t])\subsetneq\Gamma_t \tag{109}$$

Now define the definite description Q' as follows:

$$Q'_w = \pi(\cdot|[Q = \Gamma_j]) \text{ whenever } w \in [Q = \Gamma_j] \tag{110}$$

Note that Q' is well-defined because $\{[Q = \Gamma_j] : j = 1, ..., m\}$ is a partition. To see that Q' is at least as informative as Q, note that for every $w \in \mathcal{W}$, $w \in [Q = \Gamma_j]$ for some j, and:

$$Q'_w = \Pi(\cdot|[Q = \Gamma_j]) \subseteq \Gamma_j = Q_w \tag{111}$$

where the subset relation holds because Π Sub-Reflects Q. On the other hand, Q is not at least as informative as Q'. To show this, let $w \in [Q = \Gamma_t]$. Then:

$$Q'_w = \Pi(\cdot|[Q = \Gamma_t]) \subsetneq \Gamma_t = Q_w \tag{112}$$

so Q' is strictly more informative than Q. All we need to show now is that $Q' \in \mathcal{Q}$, contradicting the assumption that Q is maximally informative in $\mathcal{Q}$. To show this we must show that Π Sub-Reflects Q'. But this is true because Q' is obtained by conditioning Π on the true element of a partition, so Π Identity Reflects Q', and therefore Sub-Reflects Q' as well.

For the right-to-left direction, assume Π Identity Reflects Q. Let $\{\Gamma_1, ..., \Gamma_m\}$ be the possible values of Q. Since Identity Reflection entails Sub-Reflection, Q is an element of $\mathcal{Q}$. Let $Q' \in \mathcal{Q}$ be such that $Q' \subseteq Q$. We need to show that $Q' \subseteq Q$. Let $\{\Sigma_1, ..., \Sigma_k\}$ be the possible values of Q'. Now for every $w \in \mathcal{W}$ let Γ_j and Σ_s such that $w \in [Q = \Gamma_j]$ and $w \in [Q' = \Sigma_s]$, and let $A = [Q = \Gamma_j] \cap [Q' = \Sigma_s]$. Note that the proof is complete if we can show the following claim:

- **Claim 1**: $[Q = \Gamma_j] = [Q' = \Sigma_s]$

Using this claim, we can write:

$$Q_w = \Gamma_j = \Pi(\cdot|[Q = \Gamma_j]) = \Pi(\cdot|[Q' = \Sigma_s]) \subseteq \Sigma_s = Q'_w \tag{113}$$

and since w was an arbitrary world, this proves $Q \subseteq Q'$.

To prove Claim 1, note that if $[Q = \Gamma_j] \neq [Q' = \Sigma_s]$, then either (i) $[Q = \Gamma_j] \sim A \neq \emptyset$, or (ii) $[Q' = \Sigma_k] \sim A \neq \emptyset$, or both. If (i) is the case, we can find some $w' \in [Q' = \Sigma_k] \sim A$. From Clarity, we have $\Gamma_j(\{w'\}) = \{0\}$, and since $Q' \subseteq Q$, also $\Sigma_k(\{w'\}) = 0$. But because Π Sub-Reflects Q', this entails that $\Pi(\{w'\}|[Q' = \Sigma_k]) = \{0\}$. Since Π is regular, this can only occur if $w' \notin [Q' = \Sigma_k]$, contradiction. The proof is analogous when (ii) is the case. $\square$