Agreement Theorems and Epistemic Uniqueness

[Author Redacted]
August 2021

Abstract

Epistemic Uniqueness is the view that there is a unique rational doxastic response to any body of evidence. Epistemic Uniqueness is threatened by a well-known theorem due to Robert Aumann (1976) purporting to show that two subjects with a common prior cannot *agree to disagree*. In this paper, I will do two things. First, I will show that subjects with *imprecise credal states*—credal states that are best represented not by one probability function, but a set of probability functions—are not subject to Aumann's result. Thus, proponents of Epistemic Uniqueness can avoid the conclusion that it impossible for fully rational subjects to agree to disagree by appealing to the rationality of imprecise credences. However, in the second part of the paper, I will present a new *global agreement theorem*, showing that if two subjects know all of each other's views about a certain subject matter, and all of their evidence is about that subject matter, then they disagree about any proposition about that subject matter. The proof of this theorem, unlike the proof of Aumann's original agreement theorem, does not assume that rational doxastic states must be precise. Moreover, it relies on few assumptions about the structure of evidence. In particular, I will show that the proof of the theorem requires only the assumption that evidence is always true.

1 Introduction

Epistemic Uniqueness is the view, roughly, that there is a unique rational doxastic response to any body of evidence. *Epistemic Permissivism* is the view that sometimes, there is more than one fully rational rational doxastic response to a given body of evidence. Uniqueness is threatened by a well-known theorem due to Robert Aumann (1976) purporting to show that two subjects with a common prior cannot *agree to disagree* in the following sense. If two subjects have the same prior and commonly know their posterior credences in a given proposition H, then their posterior credences in H are equal. We can use Aumann's theorem as a premise in an argument against Epistemic Uniqueness, which runs as follows. Fully rational subjects can agree to disagree. If Epistemic Uniqueness is true, then fully rational subjects cannot agree to disagree. Therefore, Epistemic Uniqueness is false. Call this the *argument from disagreement*.[1]

[1] This way of framing the epistemological significance of Aumann's agreement theorem in terms of Epistemic Uniqueness is due to Harvey Lederman (2015). Lederman argues that the

It is fair to say that Uniqueness is not the default position among Bayesian epistemologists. Nevertheless, in recent years there have been a number of powerful arguments put forward in favor of Uniqueness. Horowitz (2015), for example, argues that certain varieties of Permissivism have a hard time accounting for the *value* of rationality within an accuracy-based framework. A second, closely related theme emerges in the work of Greco and Hedden (2016) and Levinstein (2017). These authors argue that rationality is worthy of *deference* and show that Permissivism has a hard time making sense of this norm. Finally, Schultheis (2018) argues that Permissivists have trouble vindicating plausible principles relating first-order and higher-order credences. Of course, these arguments are not conclusive; the Permissivist has a variety of responses at her disposal. Still, these arguments reveal certain attractions of Uniqueness, and as such, it is worth looking into whether the position can be defended against the argument from agreement.[2]

In this paper, I will two things. First, I will show that subjects with *imprecise credal states*—credal states that are best represented not by one probability function, but a whole set of probability functions—are not subject to Aumann's result. Specifically, I will show that if at least one of the two disagreeing subjects has an imprecise credence in the proposition that they disagree about, then it is possible for them to agree to disagree about that proposition. This provides the proponent of Epistemic Uniqueness with a way of responding to the argument from disagreement. They can reject its second premise by appealing to the rationality of imprecise credences. For many, this will seem like an especially attractive response. It is often argued that Epistemic Uniqueness is an overly demanding epistemological theory, and imprecise credences provide one way of relaxing some of those demands.

However, in the second part of the paper, I will present a new *global agreement theorem*, showing that if two subjects know all of each other's views about a certain subject matter, and all of their evidence is about that subject matter, then they cannot disagree about any proposition about that subject matter. The proof of this theorem, unlike the proof of Aumann's original agreement theorem, does not assume that rational credences are precise. Moreover, it relies on few assumptions about the structure of evidence. In particular, I will show that the proof of the theorem requires only the assume that evidence is true.

---

proponent of Epistemic Uniqueness can resist the argument by denying any of several assumptions that are not made explicit in most statements of Aumann's theorem. I will discuss these assumptions in more detail in §3 and §7.

[2]For another classic defense of Epistemic Uniqueness see White (2005). For defenses of Permissivism, see Kelly (2013), Meacham (2014), and Schoenfield (2014), among others.

Here is the plan for what follows. In §2, I will introduce the framework of *Precise Bayesianism* and state Epistemic Uniqueness in that framework. In §3, I present Aumann's agreement theorem. In §4, I introduce *Imprecise Bayesianism*. In §5, I provide a counterexample to Aumann's theorem in this framework. In §6, I defend the Imprecise Bayesian against an objection concerning the phenomenon of *dilation*. In §7, I present the global agreement theorem, and discuss where things stand for proponents of Epistemic Uniqueness in light of this result. In §8, I conclude.

2 Precise Bayesianism

According to *Precise Bayesianism*, I begin my life with an initial or *prior* credence function, P. P is a function that takes a proposition H and yields a real number in the unit interval, representing my initial degree of belief—my initial *credence*—in H. These credences are taken to be *probabilities* obeying the axioms of Kolmogorov's Probability Calculus. We define a *conditional credence* in one proposition H given another proposition E as a ratio of my unconditional credences. Specifically:

$$\text{If } P(E) > 0\text{, then } P(H|E) = \frac{P(H \wedge E}{P(E)}$$

Bayesians use these conditional credences to state a rule about how I should update my credences as I acquire evidence—the rule of *conditionalization*. Suppose that E is my total evidence at a given time, and that P is my prior credence function. Conditionalization says that my new credence function P^E should be related to my prior credence function P as follows.

> **Conditionalization**. If $P(E) > 0$, then $P^E(H) = P(H|E)$.

In the framework of Precise Bayesianism, we can state Epistemic Uniqueness as follows.

> **Precise Uniqueness**. There is some prior probability function P such that for any subject S, S is fully rational only if P is S's prior probability function.

It is easy to see that, given Precise Bayesianism, my informal statement of Epistemic Uniqueness from §1 is true just in case Precise Uniqueness is true. For suppose Precise Uniqueness were false. Then there could be two subjects S_1 and S_2 and two different prior probability functions P_1 and P_2 such that S_1 is fully rational while having P_1 as his prior and S_2 is fully rational while having P_2 as his prior. But if P_1 and P_2 are different, then there is some proposition to which they assign different probabilities, and the informal statement of Epistemic Uniqueness fails. Now suppose Precise Uniqueness is true. Then there is some prior probability function P such that, for any subject S, S is fully rational only if P is S's prior probability function. By Conditionalization, it follows that, for any body of evidence E, if E is S's total evidence, then for any proposition H, S is fully rational only if S's credence in H is equal to $P(H|E)$.

3 The Agreement Theorem

Now that we have the framework of Precise Bayesianism in place, we are ready to move on to Aumann's agreement theorem.

One often finds the theorem stated in the following way. If A and B have a common prior, and update by conditioning this prior on their (possibly different) evidence, then if A and B commonly know their posterior credences in H, then their posterior credences are equal.[3] But, as Lederman (2015) has shown, this statement of Aumann's theorem is misleading. The only assumptions that I have explicitly stated are the assumption that A and B share a common prior and that A and B update by conditionalizing on their evidence. Lederman has shown that Aumann's proof relies on several other assumptions.

First, Aumann makes assumptions about the structure of evidence. Specifically, he assumes that evidence is *partitional*. For any world w, A's evidence in w is w's cell of some partition of the sample space, and B's evidence in w is w's cell of some (possibly different) partition of the sample space.[4] Second, he assumes not just that A and B in fact have a common prior, but that they have the same common prior in all worlds in the sample space. Third, Aumann assumes not just that A and B in fact update by conditionalization, but that they update by conditionalization in all worlds in the sample space. Each of these assumptions is controversial and each goes beyond a commitment to Precise Uniqueness. However, since my main goal in this section is simply to state Aumann's theorem, I will not discuss these assumptions further here. (I will discuss the philosophical significance of these assumptions in more detail in §7.)

To make these assumptions explicit, we introduce *precise agreement frames*.

[3]See, for example, Bonnano and Nehring (1997).

[4]Note that Samet (1989) and Geanakoplos (1990) show that partitionality is not essential; they show that subjects cannot agree to disagree using weaker assumptions about the structure of information. For ease of exposition, I will stick with the assumption of partitionality in my presentation of the theorem and proof.

A precise agreement frame is a structure:

$$\langle W, E^A, E^B, P^A, P^B, p^A, p^B \rangle$$

W is a finite set of 'worlds' or 'states'. A *proposition* H is also a set of worlds, a subset of W. A and B are two subjects. $E^A : W \to \mathcal{P}(W)$ is a function that takes a world w to a set of worlds $E^A(w)$, the set of worlds consistent with A's evidence in w. A proposition H is part of A's evidence in w if and only if $E^A(w) \subseteq H$. Similarly, $E^B : W \to \mathcal{P}(W)$ is a function that takes a world w to a set of worlds $E^B(w)$, the set of worlds consistent with B's evidence in w. Aumann assumes that evidence has the following properties, where $i \in \{\text{A, B}\}$:

Truth. For any $w \in W, w \in E^i(w)$

Positive Introspection. For all $w, w' \in W$, if $w' \in E^i(w)$, then $E^i(w') \subseteq E^i(w)$

Negative Introspection. For all $w, w' \in W$: if $w' \in E^i(w)$, then $E^i(w) \subseteq E^i(w')$

Taken together, these three conditions guarantee that evidence is partitional in the sense that a subject's possible evidence propositions form a partition of the sample space. For ease of exposition, I am going to use the words 'knowledge' and 'evidence' interchangeably when presenting Aumann's theorem. So, when I say that a subject A *knows* a proposition H in a world w, I mean that H belongs to A's evidence in w. (Importantly, this is just a terminological convenience. I am introducing the term 'knowledge' so that I have an easier way to talk about the propositions that belong to a subject's evidence.)

Next, we turn to probabilities. P^A is a function that takes a world w and yields a *regular* probability function P^A_w over W. P^A_w is A's prior credence function in world w. (To say that P^A_w is *regular* is to say that for any $w \in W$, $P(\{w\}) > 0$.[5]) Likewise, P^B is a function that takes a world w and yields a regular probability function P^B_w over W—B's prior credence function in w. Aumann makes two assumptions about these prior probability functions, which, following Lederman, I will call *Constant Prior* and *Shared Prior*, respectively.

Constant Prior. For any worlds $w, w' \in W$, $P^i_w = P^i_{w'}$.

Shared Prior. For any world $w \in W : P^A_w = P^B_w$

[5]By assuming that W is finite and that probability functions are regular, I am abstracting from substantial issues about how to deal with probability zero events. See Hájek (2003) and Easwaran (2014) for discussion.

Next, we turn to the posterior probability functions. p^A is a function that takes a world w and yields A's posterior probability function in w, p^A_w. Likewise, p^B is a function that takes a world w and yields B's posterior credence function in w, p^B_w. Aumann assumes:

Global Conditionalization. For any world $w : p^i_w(\cdot) = P^i_w(\cdot|E^i(w))$.

Global Conditionalization says that A and B both update by conditionalization in every world.

Finally, to state Aumann's agreement theorem, we need to introduce the notion of common knowledge. Following Lederman (2015), say that two subjects A and B *mutually know*^{1} a proposition H in world w just in case they both know H in w. (Remember that for now we are assuming that a subject knows H in w just in case H belongs to that subject's evidence in w.) A and B mutually know^n H in w just in case they mutually know^1 that they mutually know^{n-1} that H in w. Finally, A and B *commonly know H* in w just in case, for all natural numbers n, A and B mutually know^n that H in w. In our finite setting, this definition of common knowledge is equivalent to a simpler definition in terms of *self-evident propositions*.[6] Say that a proposition F is self-evident to a subject A just in case, whenever F is true, A knows F. Say that a proposition F is *public* between A and B just in case it is self-evident to A and B. Then we can define common knowledge as follows.

> **Common Knowledge**. A and B commonly know H in a world w just in case there is some proposition F such that F is true in w, F is public between A and B, and F entails H.

This is the definition we will use in our proof of Aumann's theorem.

There is one final piece of notation that we will need. We will let $[p^A(H) = x]$ be the proposition that A's posterior credence in H is equal to x. Thus, $[p^A(H) = x] = \{w : p^A_w(H) = x\}$. Similarly, we will let $[p^B(H) = y]$ be the proposition that B's posterior credence in H is y. Thus, $[p^B(H) = y] = \{w : p^A_w(H) = y\}$.

We are now ready to state Aumann's agreement theorem.

> **Aumann's Theorem**. Let $\mathcal{F} = \langle W, E^A, E^B, P^A, P^B, p^A, p^B \rangle$ be a precise agreement frame satisfying Truth, Positive Introspection, Negative Introspection, Constant Prior, Shared Prior, and Global Conditionalization. Let w be any world in W. If there is some x and some y

[6] This definition is due to Monderer and Samet (1989). A similar idea can be found in Aumann (1976) and Lewis (1969).

such that, at w, A and B commonly know the proposition $[p^A(H) = x]$ and the proposition $[p^B(H) = y]$, then $p^A_w(H) = p^B_w(H)$.

Here is the proof. Suppose that A and B commonly know, in some world $w \in W$, that A's credence in H is x and that B's credence in H is y. Then, by our definition of common knowledge, there is a proposition F that is true in w, public between A and B, and entails both of the following.

$$[p^A(H) = x]$$
$$[p^B(H) = y]$$

To say that F entails $[p^A(H) = x]$ is to say A's credence in H is *constant at* x throughout F. That is, for any world $w \in F, p^A_w(H) = x$. Likewise, to say that F entails $[P^B(H) = y]$ is to say B's credence in H is *constant at* y throughout F. That is, for any world $w \in F, p^B_w(H) = y$.

We will first show that if A's credence in H is constant at x throughout F, then A's prior conditional probability in H given F is equal to x, and if B's credence in H is constant at y throughout F, then B's prior conditional probability in H given F is equal to y. The proof of this claim relies on the following fact.

Fact 1. If all members of a finite set of events $\mathcal{E}$ are mutually disjoint, and for every $E \in \mathcal{E} : P(H|E) = x$, then $P(H|\bigcup \mathcal{E}) = x$

Suppose that A's credence in H is constant at x throughout F. Then $p^A_w(H) = x$, for every world $w \in F$. By the definition of p^A_w, it follows that $P^A_w(H|E^A(w)) = x$ for every world $w \in F$. Let $\mathcal{A} = \{E^A(w) : w \in W, E^A(w) \subseteq F\}$. For any world w, $w \in F$ just in case $E^A(w) \in \mathcal{A}$. Thus, it follows that for all evidence propositions $E^A(w) \in \mathcal{A}$, $P^A_w(H|E^A(w)) = x$. Remember that, by Constant Prior, there is some probability function—let's call it π—such that $P^A_w = \pi$, for all $w \in W$. So, we know that for all $E^A(w) \in \mathcal{A}$, $\pi(H|E^A(w)) = x$. Since the evidence propositions in $\mathcal{A}$ are mutually disjoint, it follows from Fact 1 that $\pi(H|\bigcup \mathcal{A}) = x$. Finally, since $\mathcal{A}$ forms a partition of F, $\bigcup \mathcal{A} = F$, and hence $\pi(H|F) = x$.

We have shown that if A's credence in H is constant at x throughout F, then A's prior conditional credence in H given F is equal to x. By exactly the same reasoning we can show that if B's credence in H is constant at y throughout F, then B's prior conditional credence in H given F is equal to y. The final step of the proof appeals to Shared Prior. By Shared Prior, A's prior is identical to B's prior. So, it follows that $x = y$, that is, that A's posterior credence in H in w is equal to B's posterior credence in H in w.[7]

This completes the proof of Aumann's Theorem. Now, as Lederman emphasizes, Aumann's Theorem does not, by itself, show that if Precise Uniqueness is true, then it is impossible for rational subjects to agree to disagree. In our statement of Aumann's Theorem, we made several assumptions about the structure of evidence. It is open to the proponent of Precise Uniqueness to reject these assumptions.[8] We also assumed Constant Prior, Shared Prior, and Global Conditionalization. Again, it is open to the proponent of Precise Uniqueness to reject any one of these assumptions.

However, in this essay, I am primarily interested in exploring a different way of responding to the argument from agreement. In what follows, I will suppose that the proponent of Epistemic Uniqueness accepts (versions of) all of these assumptions, and I will ask whether she can still avoid the argument from agreement by appealing to the rationality of imprecise credences. An appeal imprecise credences will, for many, seem like an especially attractive response to the argument from agreement. It is often argued that Epistemic Uniqueness is an overly demanding epistemological theory, and imprecise credences provide one way of relaxing some of those demands.

In the next two sections, I introduce imprecise credences (§4), and show that imprecise subjects can agree to disagree (§5).

4 Imprecise Bayesianism

According to *Imprecise Bayesianism*, a subject begins her epistemic life with a *set C* of credence functions called a *prior representor*. Each credence function in the subject's representor is taken to be a probability function assigning to each proposition a real number in the unit interval. The subject's *imprecise credence* in a proposition H is then a *set* of numbers:

[7] Aumann's agreement theorem has been extended in several ways. In my informal statement of the result, I followed Aumann in assuming that the subjects' information is *partitional*. But Samet (1989) and Geanakoplos (1990) show that partitionality is not essential; they show that subjects cannot agree to disagree using weaker assumptions about the structure of information. Others have proven agreement theorems for expectations of random variables, showing that if A and B commonly know that A's expectation of the population of New York is x, and B's expectation of the population of New York is y, then $x = y$. Finally, Cave (1983) and Bacharach (1985) extend Aumann's result beyond the Bayesian setting.

[8] Note that it will not suffice to deny Negative Introspection; Samet (1989) and Geanakoplos (1990) have shown that this assumption is not essential. But denying Truth or Positive Introspection can lead to counterexamples.

$$C(H) =_{df} \{P(H) : P \in C\}$$

It is typically assumed that representors are *convex*: any probability function that can be gotten by averaging two probability functions in C is also in C.[9] If representors are convex, then for any proposition H, $C(H)$ will be an interval of real numbers, and we can characterize a subject's imprecise credence in H in terms of its upper and lower bounds. For example, we can say that a subject's imprecise credence in H is the interval $[.2, .8]$.

Proponents of Imprecise Baysianism typically assume that subjects update their credences according to the rule of *imprecise conditionalization*, which says, roughly, that each member of the representor updates by conditionalization on the new evidence. To state this more precisely, we define imprecise conditional credences as follows:

$$C(H|E) =_{df} \{P(H|E) : P \in C, P(E) > 0\}$$

Now suppose that E is my total evidence at a given time, and that C is my prior representor. Imprecise Conditionalization says that my posterior representor C^E should be related to my prior representor C as follows.

Imprecise Conditionalization. $C^E(H) = C(H|E)$

There are many reasons to accept Imprecise Bayesianism. I will not attempt to survey all of the arguments here, but I will mention a few. The reason for accepting Imprecise Bayesianism that is most commonly cited is that very often, when our evidence is sparse and unspecific, it is hard to see how we could be rationally required to assign precise credences to every proposition that we entertain. Consider a well-known case due to Elga (2010).

> A stranger approaches you on the street and starts pulling objects out of a bag. The first three objects he pulls out are a regular-sized tube of toothpaste, a live jellyfish, and a travel-sized tube of toothpaste.

Elga asks: To what degree should you believe that the next object he pulls out of his bag will be another tube of toothpaste. Now, you have no relevant statistical evidence concerning how often people carry at least three tubes of toothpaste in their bags, no relevant theories on the topic that you can appeal to. It is very

[9] Note that not all proponents of Imprecise Bayesianism accept convexity. The assumption of convexity is not essential to any of the claims that I will make in this paper. For more on convexity, see Jeffrey (1983).

natural in a situation like this to say that it is at least *permissible* for you not to have *any* precise degree of confidence in the proposition that the next object will be a tube of toothpaste. Proponents of Imprecise Bayesianism sugggest that in cases like this your credences should be imprecise. Your degree of confidence that the next object will be a tube of toothpaste should span an interval of values, such as $[.2, .8]$. A second set of motivations for imprecision concerns decision theory. For example, a number of authors suggest that imprecise credences are needed to model the phenomenon of *ambiguity aversion*.[10] As a final example, Hajek, Hawthorne, and Isaacs (forthcoming) argue that imprecise credences are the appropriate doxastic attitude to take towards non-measurable propositions (which cannot receive precise credences).

There is no settled interpretation of the formalism of Imprecise Bayesianism, no settled answer to the question of what it is to be in an imprecise credal state, and it is well beyond the scope of this essay to offer an interpretation.[11] But this fact alone is no mark against Imprecise Bayesianism. After all, there is no settled answer to the question of what precise credences are. One popular strategy is to elucidate precise credences in terms of the roles they play in other theories—our theories of rational decision, of confirmation, and so on. There is good reason to believe that we can do the same for imprecise credences. Many authors attempt to explain what imprecise credences are in terms of their place in theories of rational decision, and statistical inference, among other things.[12]

Now that we have introduced Imprecise Bayesianism, we can say what Epistemic Uniqueness comes to in this framework. In the imprecise setting, Epistemic Uniqueness is equivalent to what I will call *Imprecise Uniqueness*.

> **Imprecise Uniqueness**. Theere is some set of of probability functions C such that for any subject S, S is fully rational only if C is S's prior representor.

In the next section, I will show that imprecise subjects with a common prior representor can agree to disagree.

[10] For an overview of this literature, see Bradley (2019).

[11] One common view is the *comparativist treatment* of imprecise credences. On this picture, one takes *comparative confidence* to be fundamental, and imprecise credences are used to model subjects with *incomplete* comparative confidence orderings. See Fine (1973). For more recent defenses of this approach, see Zynda (2000), Stefansson (2017), and Konek (2019).

[12] For example, see the literature on ambiguity aversion. In a similar vein, Mahtani (2018) suggests that imprecise credences represent certain patterns of unstable betting behavior.

5 Agreement and Imprecision

When we presented Aumann's Theorem in §3, we introduced precise agreement frames. To accommodate Imprecise Bayesianism, we need to work with a more general class of frames, which I will call *imprecise agreement frames*.[13] An imprecise agreement frame is a structure:

$$\langle W, E^A, E^B, C^A, C^B, c^A, c^B \rangle$$

As before, W is a finite set of worlds. A and B are two subjects. E^A and E^B are functions from worlds to sets of worlds representing A's evidence and B's evidence, respectively. We will continue to assume that evidence satisfies Truth, Positive Introspection, and Negative Introspection. (I will also continue to use the terms 'evidence' and 'knowledge' interchangeably in this section, so that when I say that a subject knows a proposition H in w, I mean that the proposition H belongs to the subject's evidence in w.)

Now turn to probabilities. C^A is a function that takes a world w and yields A's prior representor in w, C^A_w. Likewise, C^B is a function that takes a world w and yields B's prior representor in w, C^B_w. I will assume that the prior representors contain only regular probability functions. We will assume (where $i \in \{\text{A,B}\}$):

Constant Prior Representor. For any worlds $w, w' \in W$, $C^i_w = C^i_{w'}$.

Shared Prior Representor. For any world $w \in W : C^A_w = C^B_w$

Next, c^A and c^B are functions that take a world w to A's posterior representor in w and B's posterior representor in w, respectively. We will assume:

Global Imprecise Conditionalization. For any w: $c^i_w(\cdot) = C^i_w(\cdot | E^i(w))$.

Thus, for any world w, c^A_w is the result of updating C^A_w by imprecise conditionalization on A's evidence in w, and c^B_w is the result of updating C^B_w by imprecise conditionalization on B's evidence in w.

To show that Imprecise Bayesianism is not subject to Aumann's Theorem, we provide an imprecise agreement frame $\mathcal{F} = \langle W, E^A, E^B, C^A, C^B, c^A, c^B \rangle$ satisfying Truth, Positive Introspection, Negative Introspection, Constant Prior, Shared Prior, and Global Imprecise Conditionalization, and is such that, for some $w \in W$, and some proposition H:

[13] Imprecise agreement frames are generalizations of precise agreement frames because we can define a precise agreement frame as an imprecise agreement frame in which each of C^A, C^B, c^A, and c^B is a singleton consisting of one probability function.

(1) There is some $[a, b]$ and some $[c, d]$ such that, in w, A and B commonly know $[C^A(H) = [a, b]]$ and $[C^B(H) = [c, d]]$; and

(2) $[a, b] \neq [c, d]$

We construct a frame satisfying (1) and (2) as follows. Let $W = \{w_1, w_2, w_3, w_4\}$. Let $P = \{w_1, w_2\}$. Let $E^A(w_1) = E^A(w_2) = P$. Let $E^A(w_3) = E^A(w_4) = \neg P$. Let $E^B(w_1) = E^B(w_2) = E^B(w_3) = E^B(w_4) = W$. It is easy to check that E^A and E^B both satisfy Truth, Positive Introspection, and Negative Introspection.

Let $H = \{w_1, w_3\}$. Let C be A and B's shared prior representor. (Since we are assuming Constant Prior and Shared Prior, we can leave off the subscripts and superscripts.) We will assume that C is fully precise with respect to H and with respect to P. Specifically:

- $C(H) = C(\neg H) = 1/2$
- $C(P) = C(\neg P) = 1/2$

We will assume that C is imprecise with respect to respect to *H conditional on P*. Specifically:

- $C(H|P) = [1/3, 2/3]$

It follows from these assumptions that C also has an imprecise credence in H conditional on $\neg P$. In particular, it follows $C(H|\neg P) = C(H|P)$. Thus:

- $C(H|\neg P) = [1/3, 2/3]$

A's evidence is P in w_1 and w_2 and $\neg P$ in w_3 and w_4. So, by the definition of A's posterior credence function:

- $c^A_w(H) = [1/3, 2/3]$, for all $w \in W$

B's evidence is simply W in every world. So, by the definition of B's posterior credence function:

- $c^B_w(H) = 1/2$, for all $w \in W$

A and B agree to disagree at every world in this frame. A's credence in H is constant at $[1/3, 2/3]$ throughout W. B's credence in H is constant at $1/2$ throughout W. Since W is public between A and B, it follows that, for every $w \in W$, A and B commonly know, at w, that A's imprecise credence in H is $[1/3, 2/3]$ and that B's imprecise credence in H is $1/2$.

Let's take a moment to explore why agreeing to disagree is possible for imprecise subjects. Recall that when we proved Aumann's agreement theorem for precise subjects, we relied on the following fact.

Fact 1. If all members of a finite set of events $\mathcal{E}$ are mutually disjoint, and for every $E \in \mathcal{E} : P(H|E) = x$, then $P(H|\bigcup \mathcal{E}) = x$

The important fact about imprecise subjects is that there is no analogue of Fact 1 for imprecise credences. In the frame I just presented, $C(H|P) = C(H|\neg P) = [1/3, 2/3]$ and yet $C(H|\bigcup\{P, \neg P\}) = C(H) = 1/2$. This is possible because the probability functions in C disagree about how (and whether) P is relevant to H. Some think P is *negatively relevant* to H. There's at least one probability function in C whose probability in H drops from from 1/2 to 1/3 upon learning P. Other members of C think that P is positively relevant to H. At least one increases from 1/2 to 2/3 upon learning P. Still other members of C think that H and P are independent.

Aumann's argument makes use of the following surprising consequence of Fact 1, together with our assumption that evidence is partitional. If A and B commonly each other's precise credences in H, then whatever additional evidence they have beyond what they commonly know, is irrelevant to their opinions about H. To see this, let F be the set of worlds consistent with everything that A and B commonly know in w. If A and B commonly know each other's precise credences in H, then Fact 1—together with our assumption that evidence is partitional—guarantees that A's credence in H, conditional on F *and his private evidence*, is the same as his credence in H given just F. Likewise for B. Fact 1—together with our assumption that evidence is partitional—guarantees that B's credence in H, conditional on F and *her* private evidence, is the same as her credence in H given just F. That explains why A and B must agree in the precise setting. A and B have the same prior and, in effect, the same evidence. By contrast, since there is no analogue of Fact 1 in the imprecise setting, the mere fact that A and B commonly know each other's credences does not guarantee that the private information each of A and B has is irrelevant to their opinions about H.

We have seen that it is possible for imprecise subjects to agree to disagree. This means that proponents of Epistemic Uniqueness can resist the argument from agreement by appeal to Imprecise Bayesianism, even if they accept all of the assumptions used in our proof of Aumann's Theorem mentioned in §3 (Truth, Positive Introspection, Negative Introspection, Constant Prior, Shared Prior, and Global Conditionalization).

In the next section, I will present an objection to this claim. The objection begins with the observation that agreeing to disagree in the imprecise setting is closely connected to the well-known phenomenon of *dilation* for imprecise credences. In fact, we can show that it is possible for imprecise subjects to agree to disagree only if dilation occurs. But, the objection goes, dilation is not rational; it is a pathological feature of updating with imprecise credences. But if imprecise subjects can agree to disagree only if dilation occurs, and dilation is not rational, then we have not shown that *rational* subjects with imprecise credences can agree to disagree, and so we have not, after all, found a way to respond to the argument from agreement. In the next section, I present and respond to this argument on behalf of the proponent of Imprecise Uniqueness.

6 Agreeing to Disagree and Dilation

Let's begin by saying what dilation is. Let C be a set of probability functions over W. We define a *lower probability function* $\underline{P}$, with respect to C, as follows:

$$\underline{P}(H) = \inf\{P(H) : P \in C\}$$

We define an *upper probability function*, with respect to C, in a similar way.

$$\overline{P}(H) = \sup\{P(H) : P \in C\}$$

We also have conditional lower and upper probability functions. Where $\underline{P}(E) > 0$, conditional lower and upper probabilities are defined, respectively, as follows.

$$\underline{P}(H|E) = \inf\{P(H|E) : P \in C\}$$
$$\overline{P}(H|E) = \sup\{P(H|E) : P \in C\}$$

Let $\mathcal{E}$ be a partition of the sample space W such that $\underline{P}(E) > 0$, for all $E \in \mathcal{E}$. Let H be any proposition. We will say that $\mathcal{E}$ *dilates* $C(H)$ just in case, for every $E \in \mathcal{E}$:

$$\underline{P}(H|E) < \underline{P}(H) \leq \overline{P}(H) < \overline{P}(H|E)$$

In words: $\mathcal{E}$ dilates $C(H)$ just in case C's posterior imprecise credence in H is wider than C's prior imprecise credence in H no matter which proposition in $\mathcal{E}$ is learned.

One can show that agreeing to disagree in imprecise agreement frames requires dilation. To be more precise, Let $\mathcal{F}$ be an imprecise agreement frame satisfying Truth, Positive Introspection, Negative Introspection, Shared Prior, Constant Prior, and Conditionalization. Let w be any world in W. Let A and B be two subjects, and let C be their shared prior. Suppose that A and B commonly know each other's imprecise credences in w. Then there is a proposition F that is public between A and B and that entails (1) $[c^A(H) = [a, b]]$, for some interval $[a, b]$, and (2) $[c^B(H) = [c, d]]$, for some interval $[c, d]$. Let $\mathcal{A} = \{E^A(w) : w \in W, E^A(w) \subseteq F\}$. Let $\mathcal{B} = \{E^B(w) : w \in W, E^B(w) \subseteq F\}$.

Observation 1. A and B can agree to disagree about H in w only if either $\mathcal{A}$ dilates $C(H|F)$ or $\mathcal{B}$ dilates $C(H|F)$.

I leave the proof to a footnote.[14]

I said that dilation is often thought to be a pathological feature of imprecise credences. In fact, some argue that the phenomenon provides sufficient reason to reject imprecise credences altogether.[15] To illustrate the alleged difficulties raised by dilation, let's consider an example due to Roger White.

> *Coin Game*. You haven't a clue as to whether P. But you know that I know whether P. I have a fair coin and I paint both sides so that you can't tell which side is heads and which side is tails. Then I write 'P' on one side of the coin and '$\neg P$' on the other side, with whichever is true going on the heads side and whichever is false going on the tails side. We toss the coin and see that it lands 'P'.

Before seeing how the coin lands, you have no clue whether P is true. Let's suppose that you have an imprecise credence of $[.3, .7]$ in P. You also know that the coin is fair. You have a sharp credence of .5 in the proposition that the coin landed

[14]*Proof of Observation 1.* Suppose that A and B commonly know, in w, that A's imprecise credence in H is $[a, b]$ and that B's imprecise credence in H is $[c, d]$. Then there is a public proposition F such that: A's posterior imprecise credence is constant at $[a, b]$ throughout F and B's posterior imprecise is constant at $[c, d]$. If A and B disagree, then either (1) or (2) must be true:

(1) $C(H|F) \neq [a, b]$
(2) $C(H|F) \neq [c, d]$

It's easy to show that if (1) is true, then $\mathcal{A}$ dilates $C(H|F)$, and if (2) is true, then $\mathcal{B}$ dilates $C(H|F)$. To see why, suppose, for reductio, that (1) is true and yet $\mathcal{A}$ does not dilate C(H|F). Let $C(H|F) = [x, y]$. If (1) is true, then $[a, b] \neq [x, y]$. And if $\mathcal{A}$ does not dilate $C(H|F)$, then $[a, b]$ is no wider than $[x, y]$. In that case, it follows that either $x < a$ or $b < y$. In the first case, where $x < a$, there is some probability function $P_1 \in C$ such that $P_1(H|F) = P_1(H|\bigcup\mathcal{A}) = x$ and yet for all $A \in \mathcal{A} : P_1(H|A) > x$. But that violates Fact 1, repeated below:

> **Fact 1**. If all members of a finite set of events $\mathcal{E}$ are mutually disjoint, and for every $E \in \mathcal{E} : P(H|E) = x$, then $P(H|\bigcup\mathcal{E}) = x$

(Recall that Fact 1 holds for all probability functions and P_1 is a probability function.) In the second case, where $b < y$, there is some probability function $P_2 \in C$ such that $P_2(H|E) = P_2(H|\bigcup\mathcal{A}) = y$ and yet for all $A \in \mathcal{A} : P_2(H|A) < y$. But again, this violates Fact 1.

We have now shown that if (1) is true, then $\mathcal{A}$ dilates $C(H|F)$. By parallel reasoning, we can also show that if (2) is true, then $\mathcal{B}$ dilates $C(H|F)$. But earlier we said that if A and B agree to disagree, then either (1) or (2) is true. Thus, it follows that if A and B agree to disagree, then either $\mathcal{A}$ dilates $C(H|F)$ or $\mathcal{B}$ dilates $C(H|F)$.

[15]See White (2010) for philosophical objections to dilation. The phenomenon of dilation has been thoroughly studied in the mathematics literature. See, for example, Walley (1991) and Seidenfeld and Wasserman (1993) for two of the classic papers.

heads. When you see how the coin landed, you learn that *P* is true if and only if the coin landed heads. If you update by imprecise conditionalization, your credence in the proposition that the coin lands heads must dilate to $[.3, .7]$. White, and many others, argue that this is simply absurd.

What, exactly, is so bad about dilation in this example? White offers a few different arguments. One thought he articulates is that the information you've gained seems like it just isn't relevant to your attitude towards the proposition that the coin landed heads. Your credence in heads just shouldn't change at all when you learn the biconditional. (Perhaps, as White suggests, this is because you *know* that the coin has a .5 chance of landing heads.) But the claim that, in general, one's credence in heads shouldn't change in response to learning biconditionals of this sort is not tenable. Consider a case where you are very confident in *P*—perhaps *P* is the proposition that it snowed at least once in Chicago during the month of December last year. Suppose you play the coin game with someone who knows for certain whether this is true. The coin is tossed and you see that it lands on *P*. You now know that the coin landed heads if and only if *P*. It would be eminently reasonable for you to become very confident that the coin landed heads.

A second objection to Imprecise Bayesianism based on Coin Game concerns the fact that the dilation in Coin Game is *predictable*. For if your credence in the proposition that the coin lands heads dilates from .5 to $[.3, .7]$ upon learning that the coin lands on '*P*', then it must also dilate from .5 to $[.3, .7]$ upon learning that the coin lands on '$\neg P$'. Thus, when the coin is tossed, you know that, no matter what you learn, your imprecise credence in heads will dilate from .5 to $[.3, .7]$. But, the objection continues, that can't be rational. For if you know that tomorrow you are going to adopt a certain doxastic attitude towards heads, and you know that you won't lose any information in the meantime, then you should *already* have that attitude. (This is a consequence of Bas van Fraassen's *Reflection Principle*.) But there are well-known counterexamples to principles like this—involving centered propositions, and imperfect introspection, among other things.[16] It is open to the proponent of imprecise credences to say that dilation provides another class of counterexamples.

So, I don't think that the objections to dilation based on Coin Game are compelling. Moreover, I think that there are cases in which the Imprecise Bayesian can claim that, intuitively, dilation is an appropriate response to the evidence. To see this, consider the following example.[17] One of your friends needs a kid-

[16] See Bas van Fraassen (1984) for a classic paper on the Principle of Reflection. See, among others, Arntzenius (2003) and Williamson (2000) for purported counterexamples.

[17] This case is modeled off of a case due to Robert Stalnaker (1994), who uses the case to argue

ney transplant. Adam, Ollie, and Arnie are candidate donors. Each will undergp state-of-the-art blood testing. Before testing, you ask each to tell you his blood type. Adam says that he type A. Arnie says that he is type A. Ollie says that he is type O. People are reliable about their own blood types, so you are confident in what you've been told—that Adam and Arnie are type A, and that Ollie is type O. Now suppose that they undergo testing and the test reveals that all three are the same blood type, but it does not reveal which type they are.

We will suppose that the test is always right. So, when it says that Adam, Arnie, and Ollie are all the same blood type, you trust it, and you conclude that they are all either type A or type O. Now consider your attitude towards the proposition that Adam is type A. It is very natural to think that, after learning that Adam has the same blood type as Ollie and Arnie, your attitude toward the proposition that Adam is type A should change. Before you learned this, you were confident that Adam was A; you had no reason to question his testimony. Now you aren't so sure. After all, Ollie said that he type O, and you don't know he's the one who made the mistake.

But one can show, given reasonable assumptions about the case, that this change in attitude towards the proposition that Adam is A cannot be modeled with precise credences. Specifically, suppose first that you do not trust Adam more than Ollie, that you do not trust Ollie more than Adam, and that the same is true for the other two pairs. In the precise framework, this means that you trust all of them equally, so that when each tells you his blood type, your credence that Adam is A is equal to your credence that Arnie is A, which in turn is equal to your credence that Ollie is O.[18] Second, suppose that the propositions that Adam is A, Arnie is A, and Ollie is O are mutually independent. (Whether Adam is A, for example, does not depend on what blood types Arnie and Ollie have.[19]) It follows from these two assumptions that your credence that Adam is A, conditional on the proposition that Adam is the same as Ollie and Arnie, is the same as your prior credence that Adam is A. I leave the proof to a footnote.[20]

against the principle of Rational Monotonicity in belief revision. Stalnaker's case, in turn, is an epistemic version of a case due to Matthew Ginsberg in a paper on counterfactuals. Note that Stalnaker does not discuss how this case might be modeled in a Bayesian framework. To my knowledge, the discovery that this case provides an apparent counterexample to Precise Bayesianism (and a possible argument in favor of Imprecise Bayesianism) is new.

[18] I am assuming that your initial credence—that is, your credence before Adam, Ollie, and Arnie tell you their blood types—that Adam is A is equal to your prior credence that Arnie is A, which in turn, is equal to your prior credence that Ollie is O.

[19] Importantly, we are assuming that these events are mutually independent *relative to the credence function you have before you learn that all three are the same*. The three events are not mutually independent relative to the credence function you have *after* learning this.

[20] Let *Adam A* be the proposition that Adam is type A; let *Ollie O* be the proposition that Ollie

However, if we model you with *imprecise credences*, then a change in attitude is indeed permissible upon learning that Adam is the same blood type as Ollie and Arnie. Specifically, one can show that learning that Adam is the same as Ollie and Arnie will *dilate* your credence in the proposition that Adam is A.

We will continue to assume that you do not trust Adam more than Ollie, that you do not trust Ollie more than Adam, and that the same goes for Adam and Arnie, and Ollie and Arnie. In the precise framework, this means that you trust them *equally*, and so that when each tells you his blood type, your credence that Adam is A is equal to your credence that Arnie is A, which in turn is equal to your credence that Ollie is O. Now consider the imprecise framework. Suppose that when Adam tells you that he is type A you adopt an imprecise credence—say, [.8, .9]—in the proposition that Adam is A. Likewise, when Ollie tells you that he is type O, you adopt an imprecise credence of $[.8, .9]$ in the proposition that Ollie is O, and when Arnie tells you that he is type A, you adopt an imprecise credence of $[.8, .9]$ in the proposition that Arnie is A. Moreover, suppose that you don't know exactly how to weigh Arnie's testimony against Ollie's. (This assumption seems reasonable; you don't know exactly how reliable each of Adam, Ollie, and Arnie is.) In particular, suppose that some members of your representor think that Ollie is more reliable than Arnie, while others think that Arnie is more reliable than Ollie. Given these assumptions, one can show that your credence in the proposition that Adam is A will dilate when you learn that Adam is the same

is O; let *Arnie A* be the proposition that Arnie is A; and let *All same* be the proposition that they all the same blood type. We will assume that you are sure that each of Adam, Arnie, and Ollie is either type A or type O. Let *Pr* be your credence function before learning that Adam, Ollie, and Arnie have the same blood type, but after Adam and Arnie have said that they are type A and Ollie has said that he is type O. We will make two assumptions.

(1) $Pr(Adam\ A) = Pr(Arnie\ A) = Pr(Ollie\ O) = x$

(2) $Pr(Adam\ A\ \&\ Arnie\ A\ \&\ Ollie\ O) = Pr(Adam\ A) \times Pr(Arnie\ A) \times Pr(Ollie\ O)$

With these two assumptions, we have:

$$Pr(Adam\ A|All\ same) = Pr(All\ A|All\ same) = \frac{x^2 \cdot (1-x)}{x \cdot (1-x)^2 + x^2 \cdot (1-x)} = x$$

type as Ollie and Arnie.[21]

Without getting into all of the details, let me briefly explain why. Learning that Adam is the same as Ollie and Arnie is the same as learning that Adam is the same as Ollie and that Adam is the same as Arnie. Consider those members of your representor that trust Arnie more than Ollie. Learning that Adam is the same as Ollie provides them with some evidence that Adam is O, and learning that Adam is the same as Arnie provides them with some evidence that Adam is A. But they weigh Arnie's testimony more heavily than they do Ollie's testimony. So, the evidence that they have acquired in favor of the proposition that Adam is A (that Adam is the same as Arnie) is stronger than the evidence that they have acquired in favor of the proposition that Adam is O (that Adam is the same as Ollie). Thus, the members of your representor that trust Arnie more than Ollie will become more confident that Adam is A. Now consider the members of your representor that trust Ollie more than Arnie. By similar reasoning, these members of your representor will become *less* confident that Adam is A upon learning that Adam is the same as Ollie and Arnie. In short, your credence in the proposition that Adam is A becomes more spread out—some members of your representor become less confident, while others become more confident.

I said that it is at least permissible for your attitude towards the proposition that Adam is A to change upon learning that all three have the same blood type. I have also said that, if your credences are fully precise, then such a change would not be permissible. However, if we model you with *imprecise credences*, then a change in attitude is indeed permissible. Specifically, one can show that learn-

[21]To see why this is possible, let C be your representor after Adam, Ollie, and Arnie each tells you his blood type, but before you learn that they are all the same. Suppose C contains three probability functions: Pr_1, Pr_2, Pr_3. Asumme:

- $Pr_1(Adam\ A) = Pr_1(Ollie\ O) = Pr_1(Arnie\ A) = .8$
- $Pr_2(Adam\ A) = .8$; $Pr_2(Ollie\ O) = .9$; $Pr_2(Arnie\ A) = .8$
- $Pr_3(Adam\ A) = .8$; $Pr_3(Ollie\ O) = .8$; $Pr_3(Arnie\ A) = .9$

Pr_1 trusts Adam, Arnie, and Ollie equally. Pr_2 is more confident that Ollie is O than that Arnie is A; Pr_3 is more confident that Arnie is A than that Ollie is O. Since $Pr_1(Adam\ A) = Pr_2(Adam\ A) = Pr_3(Adam\ A) = .8$, it follows that $C(Adam\ A) = .8$. We calculate conditional probabilities as follows.

- $Pr_1(Adam\ A|All\ same) = Pr_1(All\ A|All\ same) = \frac{.8\times.8\times.2}{(.8\times.8\times.2)+(.2\times.2\times.8)} = \frac{.128}{.128+.032} = .80$
- $Pr_2(Adam\ A|All\ same) = Pr_2(All\ A|All\ same) = \frac{.8\times.8\times.1}{(.8\times.8\times.1)+(.2\times.2\times.9)} = \frac{.064}{.064+.036} = .64$
- $Pr_3(Adam\ A|All\ same) = Pr_3(All\ A|All\ same) = \frac{.8\times.9\times.2}{(.8\times.9\times.2)+(.2\times.8\times.1)} = \frac{.144}{.144+.016} = .90$

So, $C(Adam\ A|All\ same) = [.64, .90]$.

ing that Adam is the same as Ollie and Arnie will *dilate* your credence in the proposition that Adam is A. Many Imprecise Bayesians assume that becoming less precise involves becoming less sure about what to think about the proposition that Adam is A.[22] If we accept this assumption, then the proponent of Imprecise Bayesianism, unlike the proponent of Precise Bayesianism, can say what seems to be the right thing about this case—that it is permissible to become less sure about Adam's blood type upon learning that Adam, Ollie, and Arnie are all the same.

Let's take stock. We started this paper with an argument against Uniqueness—the argument from agreement. This argument starts with the premise that fully rational subjects can agree to disagree. It then appeals to Aumann's agreement theorem, claiming that if Uniqueness is true, then fully rational subjects cannot agree to disagree. The conclusion of the argument is that Uniqueness must be false. I have argued that the proponent of Uniqueness can resist this argument by rejecting its second premise. Specifically, if they endorse *Imprecise Uniqueness*, rather than Precise Uniqueness, then they can say that fully rational subjects can have imprecise credences. And if fully rational subjects can have imprecise credences, then fully rational subjects can agree to disagree.

I then considered an objection to this strategy for responding to the argument from agreement by appeal to Imprecise Bayesianism. The objection says that agreeing to disagree is possible only if dilation occurs, but dilation is not rational, so agreeing to disagree can't be rational, either. I responded to this objection by defending the rationality of dilation. In the next section, I will present what I take to be a better objection. Specifically, I will present a new *global agreement theorem* that does apply to imprecise credences. Roughly, I will show that, given very minimal assumptions, if A and B know each other's views about everything, then they agree about everything.

7 Global Agreement

Let's begin by reviewing the assumptions—about the structure of evidence, and about the nature of prior and posterior representors—that we have been making throughout this paper. First, we have made two assumptions about prior representors, which are repeated below.

Constant Prior Representor. For any worlds $w, w' \in W$, $C^i_w = C^i_{w'}$.

[22] A full defense of Imprecise Bayesianism would involve saying more about how to connect the formalism of imprecise credences with more familiar attitudes like being uncertain. This is not the place to discuss those issues in detail. See Joyce (2010) and Moss (2014) for discussion.

Shared Prior Representor. For any world $w \in W : C^A_w = C^B_w$

Constant Prior and Shared Prior together guarantee that A and B commonly know, of some prior representor C, that it is their shared prior representor. Of course, that assumption is not at all innocent, and goes far beyond a commitment to Epistemic Uniqueness. Proponents of Uniqueness need not assume that A and B know what the uniquely rational prior representor is.

In that follows, I will no longer be be assuming Constant Prior Representor or Shared Prior Representor. I will assume that A and B have a common prior representor in *some* worlds, but I will not assume that they have a common prior representor in all worlds, and I will not assume that they know what their shared prior representor is. (I will continue to assume that the prior representors contain only *regular* probability functions.)

Second, we have made an assumption about how the posterior representors relate to the prior representors. Specifically, we have been assuming Global Imprecise Conditionalization, repeated below.

Global Imprecise Conditionalization. For all $w \in W$: $c^i_w(\cdot) = C^i_w(\cdot|E^i(w))$.

Global Imprecise Conditionalization guarantees that A and B commonly know that they update by conditionalization. Again, this assumption goes beyond a commitment to Epistemic Uniqueness. In what follows, I will no longer be assuming Global Conditionalization.

Finally, we have made three assumptions about the structure of evidence, namely:

Truth. For any $w \in W, w \in E^i(w)$

Positive Introspection. For all $w, w' \in W$, if $w' \in E^i(w)$, then $E^i(w') \subseteq E^i(w)$

Negative Introspection. For all $w, w' \in W$: if $w' \in E^i(w)$, then $E^i(w) \subseteq E^i(w')$

All three of these assumptions have been contested, but the second two are especially controversial. Consider, for example, the Williamsonian view evidence is knowledge. On this view, Positive Introspection is equivalent to the principle that if you know some proposition H, then you know that you know H, and Negative Introspection is equivalent to the principle that if you do not know H, then you know that you do not know H. The second principle is almost universally rejected, and the first is rejected by many epistemologists, such as those who say that knowledge requires a *margin for error* in the sense defined by Williamson (2000). There are also arguments against both Positive Introspection and Negative Introspection that do not rely on the assumption that evidence is knowledge, but I will not discuss those arguments further here.[23]

In what follows, I will make only one assumption about the structure of evidence. That assumption is Truth, which says that evidence is always true.

We are now ready to state the Global Agreement Theorem.

> **Global Agreement Theorem**. Let $\mathcal{F} = \langle \mathcal{W}, E^A, E^B, C^A, C^B, c^A, c^B \rangle$ be an imprecise agreement frame that satisfies Truth. Let w be any world in W. If (1) $C^A_w = C^B_w$, (2) $c^A_w(\cdot) = C^A_w(\cdot|E^A(w))$ and $c^B_w(\cdot) = C^B_w(\cdot|E^B(w))$, and (3) A and B both know the proposition $[c^A = c^A_w]$ and the proposition $[c^B = c^B_w]$, then $c^A_w(\cdot) = c^B_w(\cdot)$.

Suppose that, in w, A and B have a shared prior representor, and that, in w, A updates by conditioning the shared prior on his evidence in w and that B updates by conditioning the shared prior on her evidence in w. Suppose further that evidence satisfies Truth. Then it follows that if A and B know their posterior representors in w, then they agree about everything in w.

We will begin by presenting a simple lemma which we will use.

> **Lemma**. If $\mathcal{F} = \langle \mathcal{W}, E^A, E^B, C^A, C^B, c^A, c^B \rangle$ is an imprecise agreement frame satisfying Truth, then for any world $w \in W$, $[c^A = c^A_w]$ entails $E^A(w)$, and $[c^B = c^B_w]$ entails $E^B(w)$.

To streamline my presentation, I leave the proof to a footnote.[24]

We will use Lemma in a proof of our final result, the Global Agreement Theorem. Let w be any world in W. Suppose that A and B share a common prior in w, C. And suppose that A and B update according to imprecise conditionalization in w, so that $c^A_w(\cdot) = C(\cdot|E^A(w))$ and $c^B_w(\cdot) = C(\cdot|E^B(w))$. Finally, suppose that A and B both know that A's posterior representor is c^A_w and that B's posterior representor is c^B_w. This means that:

1. $c^A_w(\cdot) = c^A_w(\cdot|c^A = c^A_w \& c^B = c^B_w)$; and
2. $c^B_w(\cdot) = c^B_w(\cdot|c^A = c^A_w \& c^B = c^B_w)$

[23] See, for example, Dorst (2020).

[24] *Proof.* Suppose that $w' \in [c^A = c^A_w]$. Now suppose, for *reductio*, that $w' \notin E^A(w)$. From Truth, we know that $w' \in E^A(w')$. Hence, by regularity $c^A_{w'}(w') > 0$ and yet $c^A_w(w') = 0$, and so it follows that $c^A_{w'} \neq c^A_w$. But that is just to say that $w' \notin [c^A = c^A_w]$. We have derived a contradiction from our assumption that $w' \notin E^A(w)$, and thus it follows that $w' \in E^A(w)$.

By the definitions of $c^A_w(\cdot)$ and $c^B_w(\cdot)$, respectively, we have:

3. $c^A_w(\cdot|c^A = c^A_w \ \&\ c^B = c^B_w) = C(\cdot|E^A(w) \ \&\ c^A = c^A_w \ \&\ c^B = c^B_w)$; and
4. $c^B_w(\cdot|c^A = c^A_w \ \&\ c^B = c^B_w) = C(\cdot|E^B(w) \ \&\ c^A = c^A_w \ \&\ c^B = c^B_w)$

Lemma says that $[c^A = c^A_w]$ entails $E^A(w)$, and $[c^B = c^B_w]$ entails $E^B(w)$. Thus, we have:

5. $C(\cdot|E^A(w) \ \&\ c^A = c^A_w \ \&\ c^B = c^B_w) = C(\cdot|E^A(w) \ \&\ c^A = c^A_w \ \&\ c^B = c^B_w \ \&\ E^B(w))$
6. $C(\cdot|E^B(w) \ \&\ c^A = c^A_w \ \&\ c^B = c^B_w) = C(\cdot|E^B(w) \ \&\ c^A = c^A_w \ \&\ c^B = c^B_w \ \&\ E^A(w))$

(1) and (3) and (5) together entail (7):

7. $c^A_w(\cdot) = C(\cdot|E^A(w) \ \&\ c^A = c^A_w \ \&\ c^B = c^B_w \ \&\ E^B(w))$

And (2), (4), and (6) together entail (8):

8. $c^B_w(\cdot) = C(\cdot|E^B(w) \ \&\ c^A = c^A_w \ \&\ c^B = c^B_w \ \&\ E^A(w))$

And finally, it follows from (7) and (8) that $c^A_w(\cdot) = c^B_w(\cdot)$.

This concludes our proof of our Global Agreement Thoerem. We have shown that if A and B have a common prior in w, update their priors by conditionalization on their evidence in w, and know each other's posterior representors in w, then A and B agree about everything.

Let me take a moment to respond to an objection one might have at this stage. One might say that the Global Agreement Theorem isn't terribly worrying. Learning the identity of another person's representor (or credence function, in the case of Precise Bayesianism) involves learning their views about all propositions. But we rarely—if ever—learn another person's views about everything.

I don't think that this objection succeeds. As many authors point out, we can think of W as a set of what Savage calls *small worlds*. On this picture, W is a *subject matter*, a way of partitioning worlds according to how things stand with respect to that subject matter. The members of W—the cells of this partition—are *small* in the sense that they settle some questions—those that pertain to the subject matter—but not all. If we think of the worlds in these models as small worlds—and indeed, this seems to be how we typically understand them—then the Global Agreement Theorem can be restated as follows. Suppose that A and B share a common prior, and update by (imprecise) conditionalization. And suppose that A's evidence and B's evidence is entirely about some subject matter $\mathcal{Q}$.

Then the Global Agreement Theorem says that if A and B know all of each other's views about $\mathcal{Q}$, then A and B have all the same views about $\mathcal{Q}$.

But on the face of it, this seems wrong. If two subjects A and B know all of each other's views about $\mathcal{Q}$ yet do not have the same views about $\mathcal{Q}$, we do not conclude that at least one is irrational. To take an example from Lederman (2015), consider a group of scientists whose area of research forms a subject matter $\mathcal{Q}$. Suppose that these scientists are educated similarly, and so form the same priors after their training. In forming their views about $\mathcal{Q}$, they conditionalize their shared prior on their scientific evidence. We can imagine that these scientists exchange their views about $\mathcal{Q}$ in person (or perhaps in print). If they do not have all the same views by the end of process, then, according to the Global Agreement Theorem, they are irrational. But again, this does not seem correct.

The proof of the Global Agreement Theorem relied on very few assumptions about the structure of evidence or about the nature of rational doxastic states. First, we did not assume that credal states are fully precise; unlike Aumann's agreement theorem, imprecise subjects are subject to the Global Agreement Theorem. Second, we did not assume Constant Prior or Global Shared Prior. Thus, any two subjects A and B are subject to the Global Agreement Theorem if they in fact have a common prior, regardless of whether they know what that prior is. Third, we did not assume that evidence is positively introspective or that evidence is negatively introspective.

We did, however, assume Truth—the principle that evidence is always true. Without this assumption, we cannot derive Lemma, and without Lemma we cannot assume premises (5) and (6) in the proof above. One might take the Global Agreement Theorem as a new argument against this principle. But the decision to reject Truth should not be taken lightly. Evidence is supposed to play a certain theoretical role in our epistemological theories, and we can ask whether false evidence can play that role. Williamson (2000) forcefully argues that it cannot. For example, it is easy to explain why one would want to conform one's beliefs to the evidence if evidence is true, and harder to explain if evidence can be false. Of course, there is much more to say about this, but this is not the place to settle the issue.[25]

I conclude with this section with some thoughts about where things stand with Epistemic Uniqueness in light of the Global Agreement Theorem. All proponents of Epistemic Uniqueness—whether they accept Precise Uniqueness or Imprecise Uniqueness—are committed to the claim that if A and B have all of

[25] Fantl (2015), Drake (2018), and Comesana (2020) for defenses of the claim that evidence can be false.

the same evidence, then A and B cannot disagree about any proposition. What they might find surprising (indeed, unacceptable) about agreement theorems is the observation that, under certain conditions, A and B cannot disagree *even if they have different evidence*.

Say that A and B have a *surface disagreement* about H if they openly disagree about H but have different evidence. Say that A and B have a *deep disagreement* about H if they openly disagree about H even though they have all the same evidence. Proponents of Epistemic Uniqueness who find Aumann's result unacceptable want to drive a wedge between these two kinds of disagreement. They want to allow that are many (fully rational) surface disagreements while insisting that there are no (fully rational) deep disagreements.

I can imagine two ways to think about the philosophical significance of the Global Agreement Theorem in this context. One option is to double-down on the importance of the distinction between deep disagreements and surface disagreements, and to argue that if Truth stands in the way of making the distinction, then Truth must go. A second option—and the one that I prefer—is to use the Global Agreement Theorem to put pressure on the distinction. Given Truth, there isn't much daylight between deep disagreements and a certain class of surface disagreements—global surface disagreements. Proponents of Epistemic Uniqueness have learned to live with the impossibility of rational deep disagreements. If they accept Truth, perhaps they can learn to live with the impossibility of global surface agreements, too.

8 Conclusion

I will conclude with a brief summary of what I have done in this paper. I began with a presentation of the *argument from agreement*. The argument runs as follows. Fully rational subjects can agree to disagree. If Epistemic Uniqueness is true, then fully rational subjects cannot agree to disagree. Therefore, Epistemic Uniqueness is false. In §4-§5, I considered one way of responding to this argument on behalf of the proponent of Epistemic Uniqueness—one that appeals to the rationality of imprecise credences—and showed that imprecise subjects can agree to disagree.

I then considered an objection to this strategy for responding to the argument from agreement. The objection says that agreeing to disagree is possible only if dilation occurs, but dilation is not rational, so agreeing to disagree can't be rational, either. In §6, I responded to this objection by defending the rationality of dilation. I presented a case, and argued that, given reasonable assumptions, it is a counterexample to Precise Bayesianism. By contrast, it is not a counterexample to Imprecise Bayesianism. Imprecise Bayesians can account for the example by appeal to dilation—which, importantly, seems to be an appropriate response to your evidence in the case. I believe that this defense of dilation should be of interest to Imprecise Bayesians, regardless of what they make of the argument from agreement.

In the final section, I presented what I take to be a better objection to the strategy of defending Epistemic Uniqueness by appeal to Imprecise Bayesianism. That argument relied on a new *Global Agreement Theorem*. If A and B have a common prior, update by (imprecise) conditionalization, and A's evidence and B's evidence is entirely about some subject matter $\mathcal{Q}$, then if A and B know all of each other's views about $\mathcal{Q}$, it follows that A and B have all the same views about $\mathcal{Q}$. I showed that the Global Agreement Theorem applies to Precise and Imprecise Bayesianism, and that the proof of the theorem requires few assumptions about the structure of evidence or the nature of rational doxastic states.

References

1. Arntzenius, Frank (2003). 'Some problems for conditionalization and reflection.' *Journal of Philosophy*100 (7): pp. 356-370.
2. Aumann, Robert (1976). 'Agreeing to disagree.' *The Annals of Statistics* 4 (6): pp. 1236–1239.
3. Bonanno, Giacomo & Klaus Nehring. (1997). *Agreeing to Disagree: A Survey*.
4. Cave, Jonathan (1983). 'Learning to agree.' *Economics Letters* 12 (2): pp. 147–152.
5. Comesana, Juan (2020). 'A plea for falsehoods.' *Philosophy and Phenomenological Research* 100 (2): pp. 247-276.
6. Dorst, Kevin (2020). 'Evidence: A guide for the uncertain.' *Philosophy and Phenomenological Research* 100, pp. 586-632.
7. Drake, Jonathan (2018). 'Motivating reason to slow the factive turn in epistemology.' In Veli Mitova (ed.), *The Factive Turn in Epistemology*. Cambridge: Cambridge University Press. pp. 1-22.
8. Fantl, Jeremy (2015). 'What is it to be happy hhat P?' *Ergo: An Open Access Journal of Philosophy*.
9. van Fraassen, Bas C. (1984). 'Belief and the will.' *Journal of Philosophy* 81 (5): pp. 235-256.
10. Geanakoplos, John. (1989). 'Game theory without partitions, and applications to speculation and consensus.'
11. Ginsberg, Matthew (1986). 'Counterfactuals.' *Artificial Intelligence* 30.
12. Greco, Daniel & Brian Hedden (2016). 'Uniqueness and Metaepistemology.' *Journal of Philosophy* 113 (8): pp. 365-395.
13. Hajek, Alan (2003). 'What Conditional Probability Could Not Be.' *Synthese* 137 (3): pp. 273-323.
14. Horowitz, Sophie (2014). 'Immoderately rational.' *Philosophical Studies* 167 (1): pp. 41-56.

15. Isaacs, Yoaav, Alan Hajek, and John Hawthorne. *Non-Measurability, Imprecise Credences, and Imprecise Chances*. *Mind*, forthcoming.

16. Lederman, Harvey (2015). 'People with common priors can agree to disagree.' *Review of Symbolic Logic* 8: pp. 11-45.

17. Levinstein, Benjamin (2017). 'Permissive Rationality and Sensitivity.' *Philosophy and Phenomenological Research* 94 (2): pp. 342-370.

18. Lewis, David. (1969). *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.

19. Mahtani, Anna (2016). 'Imprecise Probabilities and Unstable Betting Behaviour.' *Nous* 52 (1): pp. 69-87.

20. Meacham, Christopher (2013). 'Impermissive Bayesianism.' *Erkenntnis* 79 (Suppl 6): pp. 1185-1217.

21. Monderer, Dov & Dov Samet (1989). 'Approximating common knowledge with common beliefs.' *Games and Economic Behavior* 1(2): pp. 170–190.

22. Schoenfield, Miriam (2014). 'Permission to believe: Why permissivism Is true and what it tells us about irrelevant influences on belief.' *Nous* 48 (2): pp. 193-218.

23. Samet, Dov (1990). 'Ignoring ignorance and agreeing to disagree.' *Journal of Economic Theory* 52(1): pp. 190–207.

24. Seidenfeld, Teddy, and Larry Wasserman (1993). 'Dilation for sets of probabilities.' *Annals of Statistics* 21: pp. 1139–1154.

25. Stalnaker, Robert (1994). 'What is a non-monotonic consequence relation?' *Fundamental Informaticae* 21: pp. 7-21.

26. Walley, Peter. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London

27. White, Roger (2005). 'Epistemic permissiveness.' *Philosophical Perspectives* 19 (1): pp. 445–459.

28. Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford University Press.