

AGREEMENT THEOREMS FOR SELF-LOCATING BELIEF

MICHAEL CAIE

University of Pittsburgh

Abstract. In this paper, I first outline Aumann’s famous “no agreeing to disagree” theorem, and a second related theorem. These results show that if two or more agents, who have epistemic and credal states that are defined over algebras that do not include any self-locating propositions, have certain information about one another’s epistemic and credal states, then such agents must assign the same credence to certain propositions. I show, however, that both of these theorems fail when we consider agents who have epistemic and credal states that are defined over algebras that do include self-locating propositions. Importantly, these theorems fail for such agents even when we restrict our attention to the credences that such agents have in non-self-locating propositions. Having established this negative result, I then outline and prove three agreement theorems that hold for such agents.

Call a *centered proposition* one that may vary its truth-value not only across worlds, but also across individuals or times, and call an *uncentered proposition* one that may only vary its truth-value across worlds.¹

It is well-known that, for agents whose epistemic and credal states are defined over algebras of uncentered propositions, there are strict formal constraints on the extent to which such agents may rationally assign distinct credences to propositions, given knowledge of one another’s epistemic and credal states. Interestingly, though, if one adds to such algebras centered propositions, then such formal constraints on rational disagreement, even concerning uncentered propositions, no longer obtain.² The purpose of this paper is to explore the extent to which there are formal constraints on disagreement amongst agents whose epistemic and credal states are defined over algebras that include centered propositions, given knowledge of one another’s epistemic and credal states.

Let us say, roughly, that a set of agents are *epistemic confidants* just in case (i) it is common knowledge amongst them what each agent knows, and (ii) it is common knowledge

Received: November 10, 2015.

¹ Examples of centered propositions include the proposition that I am hungry at time t , and the proposition that it is now raining in Matera. Examples of uncentered propositions include the proposition that NN is hungry at time t , and the proposition that it is raining in Matera at time t .

One might notice that centered propositions are often expressed by sentences that contain indexicals like ‘I’ and ‘now’. Given this, one might naturally wonder whether we should also allow propositions to vary their truth-values over locations. Such propositions may be expressed by sentences involving the indexical ‘here’.

The answer is that, while we can allow for this sort of variation, we need not take it as basic. Instead, we can take such variation in truth-value over location to be *derivative* from variation over times and individuals. As we’ll see, we can represent centered propositions by sets of world, individual, time triples. Given this, we can then, for example, take the proposition that it is raining here at time t' to be the set of world, individual, time triples, $\langle w, i, t \rangle$ such that $t = t'$ and it is raining at the location of i in w at t .

² One can see this sort of phenomenon in cases discussed in Leslie (1997), Bostrom (2000), Bostrom (2002), and Lewis (2004). The formal underpinnings of these cases, though, are left unexplored by these authors.

amongst them that each agent is perfectly introspective concerning their own epistemic state.³

In §1, I'll consider two results that can be established for agents whose epistemic and credal states are defined over algebras of uncentered propositions. First, I'll outline a version of the famous agreement theorem from Aumann (1976).⁴ Second, I'll show how a similar result may be used to establish that epistemic confidants whose epistemic and credal states are defined over such algebras and who conditionalize on the same probability function must assign the same credence to every proposition.

In §2, I'll show that the analogues of both of these results fail for agents whose epistemic and credal states are defined over algebras that include centered propositions. In particular, I'll show that it is formally possible for epistemic confidants whose epistemic and credal states are defined over algebras containing centered propositions to assign different credences to the same *uncentered* proposition even if each agent sets their credences by conditionalizing the same probability function on the strongest proposition that they know.

Having shown that cases of irresolvable disagreement are at least formally possible for agents with epistemic and credal states defined over algebras that include centered propositions, I'll then outline three agreement theorems that hold for such agents. (The proofs of these theorems are provided in the Appendix.) One of these theorems shows that, despite the formal possibility of disagreement concerning uncentered propositions for epistemic confidants with epistemic states defined over such algebras, there are, nonetheless, probability functions that guarantee that *any* group of such epistemic confidants who set their credences by conditionalizing this function on their respective epistemic states will assign the same credence to any uncentered proposition. The other two theorems provide certain constraints which, together with the conditions required for a group of agents to be epistemic confidants, ensure that such agents will assign the same credence to any uncentered proposition, given that they conditionalize on the same probability function satisfying a certain natural indifference condition.

Formally, then, disagreement with respect to uncentered propositions amongst epistemic confidants who conditionalize on the same probability function is possible but not inevitable for agents whose epistemic and credal states are defined over algebras that include centered propositions.

It is worth highlighting at the outset a few features of the sorts of epistemic and credal states with which I'll be concerned in this paper. First, I'll only be concerned with agents whose epistemic and credal states are defined over *finite* algebras of propositions. Second, I'll only be concerned with agents whose epistemic and credal states are defined over algebras of intensional propositions. Such propositions may be represented by sets of

³ It is natural to think that (i) entails (ii). However, for agents with epistemic states defined over algebras that include centered propositions, this need not be so. Certainly, if it's common knowledge amongst some group what each agent in the group knows, then, for each member *i* of the group, it is common knowledge amongst the group that *i* knows exactly what *i* knows. However, assuming that *i* is able to entertain first-personal centered propositions about her own knowledge, there is a further way in which she may be introspective that need not hold even if it's common knowledge amongst the group what each agent knows. For *i* may know what *i* knows, but not know, first-personally, what *she* knows. This distinction and its importance will be fleshed out in greater detail in what follows.

⁴ For a helpful treatment of this theorem and related results see Halpern *et al.* (1995). For some generalizations of this result see Rubinstein and Wolinsky (1990) and Bacharach (1985). For a recent accessible and illuminating treatment see Lederman (2015). And for some further recent discussion Demey (2014) and Degremont and Roy (2012).

possible worlds or, if they involve self-locating information, sets of centered-worlds. While I will often leave these assumptions tacit, it should be borne in mind that the results that will be proven and the claims for which I will argue all proceed within their scope.

§1. Agreement Theorems and Non-Self-Locating Belief. We'll call an individual at a time an *agent*.⁵ Let A , then, be a finite set of agents, W a finite set of "possible worlds", and \mathcal{M}_W the set of probability measures on $\mathcal{P}(W)$.⁶ Let \star be an arbitrary object that we stipulate to be distinct from all other objects that appear in our model. And let $f^\star : \mathcal{P}(W) \rightarrow \{\star\}$.

We let $(\pi_i)_{i \in A}$ be an indexed family of functions such that $\pi_i : W \rightarrow \mathcal{M}_W \cup \{f^\star\}$. If i exists at w , then $\pi_i(w) \in \mathcal{M}_W$. In this case, $\pi_i(w)$ represents i 's "prior" credal function at w . Of course, i may not exist at a world w . In this case, then, we let $\pi_i(w) = f^\star$.

Def. We say that a probability function $Pr(\cdot)$ defined on an algebra $\mathcal{P}(W)$ is *regular* just in case for every $\phi \in \mathcal{P}(W)$ such that $\phi \neq \emptyset$, $Pr(\phi) \neq 0$.

We assume that if $\pi_i(w) \in \mathcal{M}_W$ then $\pi_i(w)$ is regular, for each $w \in W$. More generally, throughout this section, when we talk about a probability function $Pr(\cdot)$ upon which an agent may conditionalize, we will assume that $Pr(\cdot)$ is regular.

We let $(P_i)_{i \in A}$ be an indexed family of functions such that $P_i : W \rightarrow \mathcal{M}_W \cup \{f^\star\}$. If i exists at w , then $P_i(w) \in \mathcal{M}_W$. In this case, $P_i(w)$ represents i 's posterior credal function at w . Again, if i doesn't exist at w , then we let $P_i(w) = f^\star$.

Finally, we let $(K_i)_{i \in A}$ be an indexed family of functions such that $K_i : W \rightarrow \mathcal{P}(W) \cup \{\star\}$. We assume that, for each $w \in W$, $K_i(w) \neq \emptyset$. If i exists at w , then $K_i(w) \in \mathcal{P}(W)$. In this case, $K_i(w)$ represents the strongest proposition that i knows at a world w . We'll call this set i 's *epistemic state* at w . If i doesn't exist at w , then we let $K_i(w) = \star$. We'll assume, throughout, that any agent we are modeling knows a proposition ϕ at a world w just in case $K_i(w) \subseteq \phi$. Note that since, for each proposition ϕ , $\star \not\subseteq \phi$, we have, as desired, that if an agent doesn't exist at a world w , then they do not know any proposition ϕ at w .

Def. We say that $\mathcal{B} = \langle W, A, (K_i)_{i \in A} \rangle$ is a *basic agreement frame*.

Def. We say that $\mathcal{F} = \langle W, A, (K_i)_{i \in A}, (\pi_i)_{i \in A}, (P_i)_{i \in A} \rangle$ is a *full agreement frame*.

Epistemic states are, of course, subject to the following constraint:

Factivity: If an agent i knows a proposition ϕ , then i is not mistaken in virtue of believing ϕ .⁷

⁵ We need not assume that each agent in our model is located at the same time. Our models then may represent the credences of different individuals at different times and may also represent the credences of a single individual at different times.

⁶ We won't assume that W is a set of maximally specific possible worlds. Instead, we can think of each member of W as corresponding to a class of such possibilities. The members of such classes will disagree about matters about which our agents are unable to make distinctions.

⁷ Why not simply say that if an agent i knows ϕ , then ϕ is true? While this covers the case in which ϕ is an uncentered proposition, where ϕ is a centered proposition, we can't say that such a proposition is true or false simpliciter. Instead, such a proposition will be true relative to some agents and false relative to others. Our formulation, then, is meant to apply when ϕ

We need to ensure, then, that the family of functions $(K_i)_{i \in \mathbf{A}}$ in a basic or full agreement frame meet this condition. We do so as follows:

Def. We say that $K_i(\cdot)$ is *reflexive* at w just in case $w \in K_i(w)$.

Then we have that the epistemic state represented by a function $K_i(\cdot)$ at world w in a basic or full agreement frame will satisfy **Factivity** just in case $K_i(\cdot)$ is reflexive at w .⁸ We assume, then, that in any basic or full agreement frame each member of the family of functions $(K_i)_{i \in \mathbf{A}}$ is such that, for each $i \in \mathbf{A}$ and each $w \in \mathbf{W}$, if $K_i(w) \in \mathcal{P}(\mathbf{W})$, then $K_i(\cdot)$ is reflexive at w .

An important intuitive epistemic property for our purposes is that of common knowledge.⁹ We can define this inductively as follows:

Def. We say that the members of $G \subseteq \mathbf{A}$ *mutually know*₀ a proposition ϕ at w just in case, for every $i \in G$, i knows ϕ .

Def. We say that the members of $G \subseteq \mathbf{A}$ *mutually know* _{$n+1$} a proposition ϕ at w just in case, at w , the members of G mutually know₀ that the members of G mutually know _{n} ϕ .

Def. We say that the members of $G \subseteq \mathbf{A}$ *commonly know* a proposition ϕ at w just in case, for every n , the members of G *mutually know* _{n} ϕ at w .

We can represent these epistemic properties in a basic or full agreement frame as follows:

Def. Let $M_G^0(\phi) =_{df.} \{w \in \mathbf{W} : \bigcup_{i \in G} K_i(w) \subseteq \phi\}$. $M_G^0(\phi)$, then, is the proposition that is true at w just in case, for every $i \in G$, $K_i(w) \subseteq \phi$.

Def. Let $M_G^{n+1}(\phi) =_{df.} \{w \in \mathbf{W} : \bigcup_{i \in G} K_i(w) \subseteq M_G^n(\phi)\}$. $M_G^{n+1}(\phi)$, then, is the proposition that is true at w just in case $M_G^0(M^n(\phi))$ is true at w .

Def. Let $\text{Com}_G(\phi) =_{df.} \bigcap_{n \in \mathbb{N}} M_G^n(\phi)$. $\text{Com}_G(\phi)$, then, is the proposition that is true at w just in case $M_G^n(\phi)$ is true at w , for all n .

We can think of $K_i(w)$ as the set of worlds that are “accessible” from w relative to i .¹⁰ We’ll say that there is a *G-accessibility chain* of length n from w_1 to w_n just in case there is a sequence of n worlds: $\langle w_1, w_2, \dots, w_n \rangle$ such that, for each $j \in \{1, \dots, n-1\}$, there is some $i \in G$ such that $w_{j+1} \in K_i(w_j)$. $\text{Com}_G(\phi)$, then, holds at a world w just in case every w' that can be reached by a finite length G-accessibility chain from w is such that ϕ is true at w' .¹¹

is an uncentered proposition and when ϕ is a centered proposition. Where ϕ is an uncentered proposition, i will be mistaken in virtue of believing ϕ just in case ϕ is false, and where ϕ is a centered proposition, i will be mistaken in virtue of believing ϕ just in case ϕ is false relative to i .

⁸ **Justification:** An agent i represented in an basic or full agreement frame knows a proposition ϕ at w just in case $K_i(w) \subseteq \phi$. And an agent at w will be mistaken in virtue of believing ϕ just in case ϕ is false at w , i.e., $w \notin \phi$. Thus, there will be some proposition ϕ that i knows at w in virtue of which ϕ is mistaken just in case $w \notin K_i(w)$.

⁹ See, for example, Lewis (1969), Aumann (1976), and Barwise (1988) for discussion and formal characterizations of common knowledge.

¹⁰ Note that given the sorts of frames characterized so far, one could define a propositional modal language and define a logic for this language in the standard way. In what follows, however, we’ll simply work directly with these frames.

¹¹ See, for example, Halpern *et al.* (1995) for a proof of this.

Given the inductive characterizations of common knowledge and $\text{Com}_G(\phi)$, then, we have that ϕ is common knowledge amongst G at w just in case $\text{Com}_G(\phi)$ is true at w .¹²

Another important intuitive epistemic property for our purposes is the following:

Def. We say that an agent i is *third-personally introspective* at a world w just in case, for any proposition ϕ , (i) if, at w , i knows ϕ , then, at w , i knows that i knows ϕ , and (ii) if, at w , i does not know ϕ , then, at w , i knows that i does not know that ϕ .

Within a basic or full agreement frame this epistemic property can be captured by the following formal property of the function $K_i(\cdot)$:

Def. We say that $K_i(\cdot)$ is *homogeneous* at w just in case $K_i(w) \in \mathcal{P}(\mathbf{W})$ and for every $w' \in K_i(w)$, $K_i(w) = K_i(w')$. We let $\text{Hom}[K_i(\cdot)] =_{df} \{w \in \mathbf{W} : K_i(w) \in \mathcal{P}(\mathbf{W}) \text{ and } \forall w' \in K_i(w), K_i(w) = K_i(w')\}$.

If, then, i is an agent whose possible epistemic states can be represented in a basic or full agreement frame, we have that i is third-personally introspective at w just in case at $K_i(\cdot)$ is homogenous at w .¹³

Def. Let $\text{Cond}[P_i(\cdot), \pi_i(\cdot), K_i(\cdot)] =_{df} \{w \in \mathbf{W} : P_i(w), \pi_i(w) \in \mathcal{M}_{\mathbf{W}} \text{ and } K_i(w) \in \mathcal{P}(\mathbf{W}) \text{ and } P_i(w)(\cdot) = \pi_i(w)(\cdot|K_i(w))\}$

Def. Let $[\pi_i = Pr(\cdot)] =_{df} \{w \in \mathbf{W} : \pi_i(w) = Pr(\cdot)\}$

Def. Let $[P_i(\phi) = x_i] =_{df} \{w \in \mathbf{W} : P_i(w)(\phi) = x_i\}$

$\text{Cond}[P_i(\cdot), \pi_i(\cdot), K_i(\cdot)]$, then, is the proposition that is true at some world w just in case i at w sets their credences to their prior credence conditional on their epistemic state. $[\pi_i = Pr(\cdot)]$ is the proposition that is true at some world w just in case i 's prior credences at w are given by $Pr(\cdot)$. And $[P_i(\phi) = x_i]$ is the proposition that is true at some world w just in case i 's credence in ϕ at w is x_i .

Aumann (1976) shows that the following holds:¹⁴

First Agreement Theorem: Let \mathcal{F} be a full agreement frame and $Pr(\cdot)$ a probability function on $\mathcal{P}(\mathbf{W})$. If, at some $w \in \mathbf{W}$, for each

¹² **Justification:** This follows given the inductive characterizations of common knowledge and $\text{Com}_G(\phi)$ together with the fact that, for each $i \in G$, i knows ϕ at w just in case $K_i(w) \subseteq \phi$.

¹³ **Justification:** First assume that $K_i(\cdot)$ is not homogenous at w . If $K_i(w) = \star$, then i does not exist at w and so i is not third-personally introspective at w . So suppose that $K_i(w) \in \mathcal{P}(\mathbf{W})$. Then for some $w' \in K_i(w)$, $K_i(w) \neq K_i(w')$. Suppose that there is some $w^* \in K_i(w')$ such that $w^* \notin K_i(w)$. In this case, while, at w , i knows $\mathbf{W} - \{w^*\}$, i does not know, at w , that i knows $\mathbf{W} - \{w^*\}$. Next suppose, that there is some $w^* \in K_i(w)$ such that $w^* \notin K_i(w')$. In this case, while, at w , i does not know $\mathbf{W} - \{w^*\}$, i does not know, at w , that i does not know $\mathbf{W} - \{w^*\}$. Thus, if $K_i(\cdot)$ is not homogenous at w , then i is not third-personally introspective at w . Next, assume that $K_i(\cdot)$ is homogenous at w . Assume that i knows ϕ at w . Then $K_i(w) \subseteq \phi$. Given that $K_i(\cdot)$ is homogenous at w , we have that for each $w' \in K_i(w)$, $K_i(w') \subseteq \phi$. And so i knows that i knows that ϕ . Assume, next, that i does not know ϕ at w . Then $K_i(w) \not\subseteq \phi$. Given that $K_i(\cdot)$ is homogenous at w , we have that for each $w' \in K_i(w)$, $K_i(w') \not\subseteq \phi$. And so i knows that i does not know that ϕ . Thus if $K_i(\cdot)$ is homogenous at w , then i is third-personally introspective at w .

¹⁴ It's worth noting that, in the frames that Aumann considers, agents are taken to exist at each world in the frame. Thus, in such frames, $P_i(w), \pi_i(w) \in \mathcal{M}_{\mathbf{W}}$ and $K_i(w) \in \mathcal{P}(\mathbf{W})$, for each $w \in \mathbf{W}$. However, nothing essential about the proof changes when we relax this assumption. For an elegant proof of Aumann's theorem that makes explicit various assumptions that are tacit in Aumann (1976) and that notes various ways in which the result may fail if certain of these assumptions are relaxed see Lederman (2015).

$i \in G \subseteq \mathbf{A}$, the following propositions are true: (a) $\text{Com}_G(\text{Hom}[K_i(\cdot)])$, (b) $\text{Com}_G(\text{Cond}[P_i(\cdot), \pi_i(\cdot), K_i(\cdot)])$, and (c) $\text{Com}_G([\pi_i = Pr(\cdot)])$, then if, for each $i \in G$, there is some $x_i \in [0, 1]$ such that $\text{Com}_G([P_i(\phi) = x_i])$ is true at w , then for each $i, j \in G$, $x_i = x_j$.

This, of course, is a purely formal result. However, given what $P_i(\cdot)$, $K_i(\cdot)$ and $\pi_i(\cdot)$ represent in a full agreement frame, it follows from this result that:

No Agreeing to Disagree: If a group of agents are such that their possible epistemic and credal states can be represented in a full agreement frame, and it is common knowledge that each is third-personally introspective, and there is some probability function $Pr(\cdot)$ such that it is common knowledge that each agent sets their credences by conditionalizing $Pr(\cdot)$ on their epistemic state, then if it's also common knowledge what each agent's credence is in some proposition ϕ , then each agent must have the same credence in ϕ .

Next we show that, for agents representable in a basic agreement frame, given sufficient knowledge of one another's epistemic states, if such agents set their credences by conditionalizing the same probability function on their respective epistemic states, then such agents will be guaranteed to assign the same credence to every proposition.

Def. Let $[K_i = \phi] =_{df} \{w \in \mathbf{W} : K_i(w) = \phi\}$.

Note that, given this definition, if $K_i(w) \in \mathcal{P}(\mathbf{W})$, then $[K_i = K_i(w)]$ is the proposition that is true at a world w' just in case $K_i(w') = K_i(w)$.

We can show that the following result holds for basic agreement frames:

Second Agreement Theorem: Let \mathcal{B} be a basic agreement frame and $Pr(\cdot)$ a probability function on $\mathcal{P}(\mathbf{W})$. If, at some $w \in \mathbf{W}$, for each $i \in G \subseteq \mathbf{A}$, the following proposition is true (d) $\text{Com}_G([K_i = K_i(w)])$, then, for each $i, j \in G$, $Pr(\cdot|K_i(w)) = Pr(\cdot|K_j(w))$.¹⁵

The short proof of this result is given in the Appendix. The key fact is that if $\text{Com}_G([K_i = K_i(w)])$ is true at w , for each $i \in G \subseteq \mathbf{A}$, then it follows that, for each $i, j \in G$, $K_i(w) = K_j(w)$.

We said, earlier, that a group of agents are epistemic confidants just in case it is common knowledge amongst them what each agent knows and it is common knowledge amongst them that each agent is perfectly introspective concerning their own epistemic state. As we will see in the next section, there is some flexibility in how we should understand the latter condition. For agents whose epistemic states are defined over algebras of possible-worlds propositions, though, being third-personally introspective is the strongest introspective property that they could instantiate.¹⁶ This, then, motivates the following terminological stipulation:

Def. We'll say that a set of agents whose epistemic states are defined over a set of possible worlds propositions are *epistemic confidants* just in case it is common knowledge amongst them what each agent's epistemic state is, and it is common knowledge amongst them that each agent is third-personally introspective.

¹⁵ See the Appendix for the proof of this result.

¹⁶ As we'll see, there is a further way for an agent with credences defined over an algebra that includes centered propositions to be introspective.

Now, for agents representable in a basic agreement frame, (d) will hold at w , for each $i \in G$, just in case it is common knowledge amongst them, at w , what each agent's epistemic state is. Moreover, given that it's common knowledge amongst the members of G what each agent's epistemic state is, it follows that it is common knowledge amongst the members of G that each $i \in G$ is third-personally introspective.¹⁷ It follows, then, that a group of agents G whose epistemic states are defined over a set of possible worlds propositions will be epistemic confidants at a world w just in case (d) holds at w , for each $i \in G$. Thus, the Second Agreement Theorem shows:

No Disagreement Amongst Confidants: If a group of agents G , whose epistemic states are defined over a set of possible worlds propositions, are epistemic confidants and each sets their credences by conditionalizing the same probability function $Pr(\cdot)$ on their epistemic state, then each member of G will assign the same credence to all propositions.

§2. Agreement Theorems and Self-Locating Belief. Basic and full agreement frames represent various possible epistemic and credal states defined over an algebra of *possible worlds* propositions. There is, however, good reason to think that, even given certain idealizing assumptions, epistemic and credal states such as ours cannot be fully represented by functions defined on such algebras. For some things that we know and some things in which we have credence would seem to be essentially self-locating. My knowledge that I am Michael Caie, for example, wouldn't seem to be reducible to my knowledge of any possible worlds proposition, and similarly for my high credence that it is now 12:30.¹⁸ While knowledge of and credence in possible worlds propositions is in the business of locating the actual world within the space of possible worlds, knowledge of and credence in propositions such as that I am Michael Caie or that it is now 12:30 would seem to be in the business of locating an individual *within* a world.

In order to represent knowledge of and credence in such propositions, then, we need a more fine-grained space than a set of possible worlds. Let \mathbf{W} , as before, be a finite set of worlds, and let \mathbf{A} , as before, be a finite set of time-bound individuals. One can think of the members of \mathbf{A} as ordered pairs of non-time-bound individuals and times. To keep notational clutter to a minimum, however, we'll suppress this additional structure and work with a

¹⁷ **Justification:** We want to show that $\text{Com}_G([K_i = K_i(w)]) \subseteq \text{Com}_G(\text{Hom}[K_i(\cdot)])$. So suppose for reductio that $\text{Com}_G([K_i = K_i(w)])$ holds at some world w but $\text{Com}_G(\text{Hom}[K_i(\cdot)])$ fails at w . Given the latter assumption we have that there is some w' that is accessible from w by a finite length G -accessibility chain such that $\text{Hom}[K_i(\cdot)]$ fails at w' . There are two possibilities (i) $K_i(w') \in \mathcal{P}(\mathbf{W})$, (ii) $K_i(w') = \star$. If (i) then there is some $w'' \in K_i(w')$ such that $K_i(w'') \neq K_i(w')$. Note though that, given that there is a finite length G -accessibility chain from w to w' and that $w'' \in K_i(w')$, it follows that there is also a finite length G -accessibility chain from w to w'' . However, given that $\text{Com}_G([K_i = K_i(w)])$ holds at w , it follows that $[K_i = K_i(w)]$ holds at all worlds that are accessible from w via a finite length G -accessibility chain. And so $[K_i = K_i(w)]$ holds at w' and w'' . Thus, we have $K_i(w') = K_i(w'') = K_i(w)$, contradicting the claim that $K_i(w'') \neq K_i(w')$. If (ii), then, there is a finite length G -accessibility chain from w to w' such that $K_i(w') = \star$. However, given that $\text{Com}_G([K_i = K_i(w)])$ holds at w , it follows that $[K_i = K_i(w)]$ holds at all worlds that are accessible from w via a finite length G -accessibility chain. And so $[K_i = K_i(w)]$ holds at w' . Thus $K_i(w') = K_i(w)$. But note that given that $\text{Com}_G([K_i = K_i(w)])$ holds at w , we have that $K_i(w) \subseteq [K_i = K_i(w)]$. And so it follows that $K_i(w) \neq \star$, contradicting the claim that $K_i(w') = \star$. It follows, then, that if $\text{Com}_G([K_i = K_i(w)])$ holds at w , then $\text{Com}_G(\text{Hom}[K_i(\cdot)])$ holds at w .

¹⁸ See, e.g., Lewis (1979) for a defense of these claims.

single parameter. We call $\mathbf{C} \subseteq \mathbf{W} \times \mathbf{A}$ a set of *centered-worlds*. As a useful convention, for any $\vec{z} \in \mathbf{C}$, we will let w_z be the world parameter of \vec{z} and a_z the agent parameter of \vec{z} . Thus $\vec{z} = \langle w_z, a_z \rangle$.

To represent epistemic and credal states that include self-locating knowledge and credence, we can take an agent's epistemic and credal states to be defined over $\mathcal{P}(\mathbf{C})$. If $\phi \in \mathcal{P}(\mathbf{C})$, then for an agent to know or believe ϕ just is for the agent to know or believe that, for some $\vec{z} \in \phi$, the actual world is w_z , and they are agent a_z , while for an agent to have credence x in ϕ just is for the agent to have credence x that, for some $\vec{z} \in \phi$, the actual world is w_z , and they are agent a_z .

Amongst the members of $\mathcal{P}(\mathbf{C})$, we can distinguish between those that are essentially self-locating and those that are not.

Def. We say that $\phi \in \mathcal{P}(\mathbf{C})$ is an *uncentered proposition* just in case for any $\vec{z}, \vec{q} \in \mathbf{C}$, such that $w_z = w_q$, $\vec{z} \in \phi$ just in case $\vec{q} \in \phi$.

Def. We say that $\phi \in \mathcal{P}(\mathbf{C})$ is a *centered proposition* just in case it not uncentered.

We can think of uncentered propositions as those members of our more fine-grained algebra that correspond to our old possible worlds propositions. In particular, for any $\phi \in \mathcal{P}(\mathbf{W})$, we let $\{\vec{z} \in \mathbf{C} : w_z \in \phi\}$ stand proxy for ϕ . Centered propositions, on the other hand, have truth-values that may vary within a world. For example, the proposition that it is raining in Matera may be true at some times and false at others. In our present framework, we can represent this by $\{\vec{z} : \text{It is raining in Matera in } w_z \text{ at the temporal location of } a_z\}$. Similarly, the proposition that I am hungry may be true for different agents at different times. In our present framework, we can represent this by $\{\vec{z} : a_z \text{ is hungry at } w_z\}$.

We can generalize both basic and full agreement frames to accommodate agents with self-locating credences as follows. Let $\mathbf{C} \subseteq \mathbf{W} \times \mathbf{A}$ be a set of centered-worlds. We'll assume that for each $a \in \mathbf{A}$, there is some $\vec{x} \in \mathbf{C}$ such that $a_x = a$. Thus $\mathbf{A} = \{a_x : \text{for some } \vec{x} \in \mathbf{C}\}$. Similarly, we assume that, for every $w \in \mathbf{W}$, there is some $\vec{x} \in \mathbf{C}$ such that $w_x = w$. We'll assume further that, for each $w \in \mathbf{W}$ and each $a \in \mathbf{A}$, a exists in w just in case there is some $\vec{x} \in \mathbf{C}$ such that $w_x = w$ and $a_x = a$. Since agents may exist in some worlds and not others it need not follow from the fact that $\langle w_1, a_1 \rangle, \langle w_2, a_2 \rangle \in \mathbf{C}$ that $\langle w_1, a_2 \rangle \in \mathbf{C}$.¹⁹

Let $\mathcal{M}_{\mathbf{C}}$ be the set of probability measures on $\mathcal{P}(\mathbf{C})$. We, again, let \star be an arbitrary object that we stipulate to be distinct from all other objects that appear in our model. And we let $f^* : \mathcal{P}(\mathbf{C}) \rightarrow \{\star\}$.

We now let $(K_i)_{i \in \mathbf{A}}$ be an indexed family of functions such that $K_i : \mathbf{C} \rightarrow \mathcal{P}(\mathbf{C}) \cup \{\star\}$, $(\pi_i)_{i \in \mathbf{A}}$ an indexed family of functions such that $\pi_i : \mathbf{C} \rightarrow \mathcal{M}_{\mathbf{C}} \cup \{f^*\}$, and $(P_i)_{i \in \mathbf{A}}$ an indexed family of functions such that $P_i : \mathbf{C} \rightarrow \mathcal{M}_{\mathbf{C}} \cup \{f^*\}$.

¹⁹ A natural question: Since we can tell from the function $K_i(\cdot)$ whether or not i exists at \vec{z} , why not simply let $\mathbf{C} = \mathbf{W} \times \mathbf{A}$? While we could do this, it seems to me that there is no reason to. For, I take it that if i doesn't exist at w , then $\langle w, i \rangle$ should not be an epistemic possibility for any agent we are modelling. Thus, if i doesn't exist at w , then no possible epistemic state that will be represented in our frames will contain $\langle w, i \rangle$. Thus, while we could include such points in our frames, they would never be in any $K_i(\vec{z})$ and so would serve no purpose.

Now one might worry that since \mathbf{W} is a set of possible worlds it may be that, even granting the sorts of idealizations we're working with in this paper, some agent may be such that $\langle w, i \rangle$ is an epistemic possibility for them. Perhaps that's right. But if one is moved by such worries then the right response is to let \mathbf{W} be a set of epistemically possible worlds. And, given this, the preceding point holds.

If i exists at w_q , then $K_i(\vec{q}) \in \mathcal{P}(\mathbf{C})$. In this case, $K_i(\vec{q})$ represents the strongest proposition that i knows at the world w_q . As before, we assume that any agent we are modeling knows all and only those propositions that are entailed by this proposition. If i doesn't exist at w_q , then we let $K_i(\vec{q}) = \star$.

If i exists at w_q , then $P_i(\vec{q}) \in \mathcal{M}_{\mathbf{C}}$. In this case, $P_i(\vec{q})$ represents i 's posterior credal function at the world w_q . If i doesn't exist at w_q , then we let $P_i(\vec{q}) = f^*$.

If i exists at w_q , then $\pi_i(\vec{q}) \in \mathcal{M}_{\mathbf{C}}$. In this case, $\pi_i(\vec{q})$ represents i 's prior credal function at the world w_q . As before, we assume that, for each $\vec{q} \in \mathbf{C}$, if $\pi_i(\vec{q}) \in \mathcal{M}_{\mathbf{C}}$, then $\pi_i(\vec{q})$ is regular. If i doesn't exist at w_q , then we let $\pi_i(\vec{q}) = f^*$.

We assume that, for every $\vec{z}, \vec{q} \in \mathbf{C}$, if $w_z = w_q$, then $K_i(\vec{z}) = K_i(\vec{q})$, $P_i(\vec{z}) = P_i(\vec{q})$, and $\pi_i(\vec{z}) = \pi_i(\vec{q})$. That is, we assume that the values of these functions are determined solely by the world parameter of a centered-world.

Def. We say that $\mathcal{B} = \langle \mathbf{C}, \mathbf{A}, (K_i)_{i \in \mathbf{A}}, \rangle$ is a *basic centered agreement frame*.

Def. We say that $\mathcal{F} = \langle \mathbf{C}, \mathbf{A}, (K_i)_{i \in \mathbf{A}}, (\pi_i)_{i \in \mathbf{A}}, (P_i)_{i \in \mathbf{A}} \rangle$ is a *full centered agreement frame*.

The central difference between agreement frames, on the one hand, and centered agreement frames, on the other, is the formal condition on the functions $(K_i)_{i \in \mathbf{A}}$ that's required in order to ensure that the epistemic states that are represented by these functions satisfy **Factivity**. **Factivity** demands that an agent's epistemic state be such that the agent is not mistaken in virtue of believing any proposition ϕ entailed by this state. In the case of agreement frames, we noted that the epistemic states represented by the functions $(K_i)_{i \in \mathbf{A}}$ will satisfy **Factivity** just in case the functions representing those states are reflexive. When we consider agents whose epistemic states can be represented in a *centered* agreement frame, however, this formal property does not serve to adequately capture the corresponding epistemic property.

Suppose, then, that we have some function $B_i(\cdot) : \mathbf{C} \rightarrow \mathcal{P}(\mathbf{C}) \cup \{\star\}$ that we take to represent an agent's belief state relative to different centered-worlds. As with $K_i(\cdot)$, we assume that $B_i(\cdot)$ is only sensitive to the world parameter so that, for each $\vec{z}, \vec{q} \in \mathbf{C}$ such that $w_z = w_q$, $B_i(\vec{q}) = B_i(\vec{z})$. It is fairly straightforward to see that even if $B_i(\vec{q}) \in \mathcal{P}(\mathbf{C})$ but $\vec{q} \notin B_i(\vec{q})$, it need not follow that at \vec{q} the agent is mistaken about some proposition that they believe.

To see this, note that if an agent i has a belief state defined over an algebra $\mathcal{P}(\mathbf{C})$, then we can represent the agent as having a true belief about who they are by taking their belief state to be such that a centered world \vec{q} is compatible with their belief state just in case $a_q = i$. Suppose, then, that, at world w , i has a true belief about who they are. Thus, we have that for any centered-world \vec{z} such that $w_z = w$, \vec{z} is compatible with the agent's belief state at \vec{z} , $B_i(\vec{z})$, only if $a_q = i$. Now let \vec{z} be such that $w_z = w$ and $a_z \neq i$. (As long as there is more than one member of \mathbf{A} , there is guaranteed to be such a point.) Then we have that $\vec{z} \notin B_i(\vec{z})$. And so $B_i(\cdot)$ is not reflexive at \vec{z} . However, surely the fact that, at \vec{z} , i has a true belief about who they are doesn't entail that there is some proposition ϕ that i believes in virtue of which they are mistaken.

We said that for an agent representable in a centered agreement frame to believe a proposition ϕ just is for them to believe that there is some $\vec{q} \in \phi$ such that the actual world is w_q and they are a_q . For an agent representable in a centered agreement frame, then, to be mistaken in virtue of believing some proposition ϕ just is for it not to be the case that there is some $\vec{q} \in \phi$ such that the actual world is w_q while they are a_q . Given this, we can see that, if $K_i(\vec{q}) \in \mathcal{P}(\mathbf{C})$, then the following formal property of a function $K_i(\cdot)$ at

point \vec{q} is both necessary and sufficient to ensure that the epistemic state that it represents satisfies the condition imposed by **Factivity**:²⁰

Def. We say that $K_i(\cdot)$ is *center-reflexive* at \vec{q} just in case there is some $\langle w_q, i \rangle \in K_i(\vec{q})$.

We assume, then, in any basic or full centered agreement frame each member of the family of functions $(K_i)_{i \in \mathbf{A}}$ is such that, for each $i \in \mathbf{A}$ and each $\vec{q} \in \mathbf{C}$, if $K_i(\vec{q}) \in \mathcal{P}(\mathbf{C})$, then $K_i(\cdot)$ is center-reflexive at \vec{q} .

Common knowledge can be captured in a centered agreement frame given the following natural characterization of propositions of the form $\mathbf{Com}_G(\phi)$ within such frames:

Def. Let $M_G^0(\phi) =_{df.} \{\vec{q} \in \mathbf{C} : \bigcup_{i \in G} K_i(\vec{q}) \subseteq \phi\}$. $M_G^0(\phi)$, then, is the proposition that is true at a centered world \vec{q} just in case, for every $i \in G$, $K_i(\vec{q}) \subseteq \phi$.

Def. Let $M_G^{n+1}(\phi) =_{df.} \{\vec{q} \in \mathbf{C} : \bigcup_{i \in G} K_i(\vec{q}) \subseteq M_G^n(\phi)\}$. $M_G^{n+1}(\phi)$, then, is the proposition that is true at a centered world \vec{q} just in case $M_G^0(M^n(\phi))$ is true at \vec{q} .

Def. Let $\mathbf{Com}_G(\phi) =_{df.} \bigcap_{n \in \mathbb{N}} M_G^n(\phi)$. $\mathbf{Com}_G(\phi)$, then, is the proposition that is true at a centered world \vec{q} just in case $M_G^n(\phi)$ is true at \vec{q} , for all n .

Then we have that ϕ is common knowledge amongst G at w_q just in case $\mathbf{Com}_G(\phi)$ is true at \vec{q} .²¹

Similarly third-personal introspection can be captured in a centered-agreement frame given the following natural characterization of propositions of the form $\mathbf{Hom}[K_i(\cdot)]$ within such frames:

Def. We say that $K_i(\cdot)$ is *homogeneous* at \vec{q} just in case $K_i(\vec{q}) \in \mathcal{P}(\mathbf{C})$ and for every $\vec{z} \in K_i(\vec{q})$, $K_i(\vec{q}) = K_i(\vec{z})$. We let $\mathbf{Hom}[K_i(\cdot)] =_{df.} \{\vec{q} \in \mathbf{C} : K_i(\vec{q}) \in \mathcal{P}(\mathbf{C}) \text{ and } \forall \vec{z} \in K_i(\vec{q}), K_i(\vec{q}) = K_i(\vec{z})\}$.

Then we also have that i is third-personally introspective at w_q just in case at $K_i(\cdot)$ is homogeneous at \vec{q} .²²

In the previous section, we noted that, for agents whose epistemic states are defined over algebras of possible-worlds propositions, being third-personally introspective is the strongest introspective property that they could instantiate. Importantly, though, the same is not true for agents whose epistemic states are defined over algebras of centered-worlds propositions.

To see this, note that corresponding to the uncentered proposition that i knows ϕ : $\{\vec{q} \in \mathbf{C} : K_i(\vec{q}) \subseteq \phi\}$, and the uncentered proposition that i does not know ϕ : $\{\vec{q} \in \mathbf{C} : K_i(\vec{q}) \not\subseteq \phi\}$, there is also the centered-proposition that one knows ϕ : $\{\vec{q} \in \mathbf{C} : K_{a_q}(\vec{q}) \subseteq \phi\}$, and the centered proposition that one does not know ϕ : $\{\vec{q} \in \mathbf{C} : K_{a_q}(\vec{q}) \not\subseteq \phi\}$. If an agent i is third-personally introspective, then for each proposition ϕ that they know, they will

²⁰ **Justification:** An agent i represented in an basic or full centered agreement frame knows a proposition ϕ at \vec{q} just in case $K_i(\vec{q}) \subseteq \phi$. And an agent at \vec{q} will be mistaken in virtue of believing ϕ just in case $\langle w_q, i \rangle \notin K_i(\vec{q})$. Thus, if $K_i(\vec{q}) \in \mathcal{P}(\mathbf{C})$, then there will be some proposition ϕ that i knows at \vec{q} in virtue of which ϕ is mistaken just in case $\langle w_q, i \rangle \notin K_i(\vec{q})$.

²¹ **Justification:** The justification is a trivial variant on the justification given in fn. 12.

²² **Justification:** The justification for this is a trivial variant of the justification given in fn. 13.

know the uncentered proposition that i knows ϕ , while if they do not know ϕ , then they will know the uncentered proposition that i does not know ϕ . It does not, however, follow from the fact that i knows either of these uncentered propositions that i knows either of the corresponding centered propositions.

Def. We say that an agent i is *first-personally introspective*, at a world w , just in case, they exist at w and, for any proposition ϕ , (i) if i knows ϕ , at w , then, at w , i knows the centered proposition that they know ϕ and (ii) if i does not know ϕ , at w , then, at w , i knows the centered proposition that they do not know ϕ .

We can characterize this epistemic property by the following formal constraint on $K_i(\cdot)$:

Def. We say that $K_i(\cdot)$ is *center homogenous* at \vec{q} just in case $K_i(\vec{q}) \in \mathcal{P}(\mathbf{C})$ and $K_i(\vec{q}) = K_{a_z}(\vec{z})$, for all $\vec{z} \in K_i(\vec{q})$. We let $\text{CHom}[K_i(\cdot)] =_{df} \{\vec{q} \in \mathbf{C} : K_i(\vec{q}) \in \mathcal{P}(\mathbf{C}) \text{ and } \forall \vec{z} \in K_i(\vec{q}), K_{a_z}(\vec{q}) = K_i(\vec{z})\}$.

Then we have that an agent i who is representable in either a basic or full centered agreement frame is first-personally introspective at w_q just in case $K_i(\cdot)$ is center homogenous at \vec{q} .²³

Def. Let $\text{Cond}[P_i(\cdot), \pi_i(\cdot), K_i(\cdot)] =_{df} \{\vec{q} \in \mathbf{C} : P_i(\vec{q}), \pi_i(\vec{q}) \in \mathcal{M}_w$ and $K_i(\vec{q}) \in \mathcal{P}(\mathbf{C}) \text{ and } P_i(\vec{q})(\cdot) = \pi_i(\vec{q})(\cdot | K_i(\vec{q}))\}$

Def. Let $[\pi_i = Pr(\cdot)] =_{df} \{\vec{q} \in \mathbf{C} : \pi_i(\vec{q}) = Pr(\cdot)\}$

Def. Let $[P_i(\phi) = x_i] =_{df} \{\vec{q} \in \mathbf{C} : P_i(\vec{q})(\phi) = x_i\}$

As before, $\text{Cond}[P_i(\cdot), \pi_i(\cdot), K_i(\cdot)]$ is the proposition that is true at some centered-world \vec{q} just in case i at w_q sets their credences to their prior credence conditional on their epistemic state. $[\pi_i = Pr(\cdot)]$ is the proposition that is true at some centered-world \vec{q} just in case i 's prior credences at w_q are given by $Pr(\cdot)$. And $[P_i(\phi) = x_i]$ is the proposition that is true at some centered-world \vec{q} just in case i 's credence in ϕ at w_q is x_i .

With all of this in place, we can now show that no result analogous to the **First Agreement Theorem**, and so no result analogous to **No Agreeing to Disagree**, can be established for full centered agreement frames. We do so by providing a full centered agreement

²³ **Justification:** First assume that $K_i(\cdot)$ is not center homogenous at \vec{q} . We will show that i is not first-personally introspective at w_q . Given that $K_i(\cdot)$ is not center homogenous at \vec{q} , we have either (i) $K_i(\vec{q}) = \star$ or (ii) there is some $\vec{z} \in K_i(\vec{q})$ such that $K_{a_z}(\vec{z}) \neq K_i(\vec{q})$. If (i) then i does not exist at w_q and so is not first-personally introspective at w_q . If (ii), then we have either (iii) there is some $\vec{x} \in K_{a_z}(\vec{z})$ such that $\vec{x} \notin K_i(\vec{q})$, or (iv) there is some $\vec{x} \in K_i(\vec{q})$ such that $\vec{x} \notin K_{a_z}(\vec{z})$. If (iii), then we have that while, at w_q , i knows $K_i(\vec{q})$, i does not know the centered proposition that they know $K_i(\vec{q})$. For this centered proposition is $\{\vec{y} : K_{a_y}(\vec{y}) \subseteq K_i(\vec{q})\}$, and we have, given (iii), that this proposition is false at one of the centered-worlds, viz., \vec{z} , compatible with i 's epistemic state at w_q . If (iv), then we have that while, at w_q , i does not know $K_{a_z}(\vec{z})$, i does not know the centered proposition that they do not know $K_{a_z}(\vec{z})$. For this centered proposition is $\{\vec{y} : K_{a_y}(\vec{y}) \not\subseteq K_{a_z}(\vec{z})\}$, and we have, given (iv), that this proposition is false at one of the centered-worlds, viz., \vec{z} , compatible with i 's epistemic state at w_q .

Next, assume $K_i(\cdot)$ is center homogenous at \vec{q} . Given that, for every $\vec{z} \in K_i(\vec{q})$, $K_{a_z}(\vec{z}) \subseteq K_i(\vec{q})$, it follows that, for any proposition ϕ that i knows at w_q , i knows $\{\vec{x} : K_{a_x}(\vec{x}) \subseteq \phi\}$, i.e., i knows the centered proposition that one knows ϕ . And, given that, for every $\vec{z} \in K_i(\vec{q})$, $K_i(\vec{q}) \subseteq K_{a_z}(\vec{z})$, it follows that any proposition ϕ that i does not know at w_q , i knows $\{\vec{x} : K_{a_x}(\vec{x}) \not\subseteq \phi\}$, i.e., i knows the centered proposition that one does not know ϕ .

frame in which, for some $\vec{q} \in \mathbf{C}$, and some probability function $Pr(\cdot)$ defined on $\mathcal{P}(\mathbf{C})$, for each $i \in \mathbf{A}$, the following propositions are true:

- (a) $\text{Com}_A(\text{Hom}[K_i(\cdot)])$
- (b) $\text{Com}_A(\text{Cond}[P_i(\cdot), \pi_i(\cdot), K_i(\cdot)])$
- (c) $\text{Com}_A([\pi_i = Pr(\cdot)])$
- (e) $\text{Com}_A(\text{CHom}[K_i(\cdot)])$

and yet there is an *uncentered* proposition ϕ such that, for some $i, j \in \mathbf{A}$, there are some $x_i \neq x_j \in [0, 1]$ such that $\text{Com}_A([P_i(\phi) = x_i])$ and $\text{Com}_A([P_j(\phi) = x_j])$ are true at \vec{q} .²⁴


Frame 1: Let \mathcal{F}_1 be such that $\mathbf{W} = \{w_1, w_2\}$, $\mathbf{A} = \{a_1, a_2\}$, and $\mathbf{C} = \mathbf{W} \times \mathbf{A}$. We let $K_{a_1}(\cdot)$ be such that, for every $\vec{x} \in \mathbf{C}$, $K_{a_1}(\vec{x}) = \{\langle w_1, a_1 \rangle, \langle w_2, a_1 \rangle\}$. And we let $K_{a_2}(\cdot)$ be such that, for every $\vec{x} \in \mathbf{C}$, $K_{a_2}(\vec{x}) = \{\langle w_1, a_2 \rangle, \langle w_2, a_2 \rangle\}$. We assume that, for each $i \in \{1, 2\}$, and each $\vec{x} \in \mathbf{C}$, $P_{a_i}(\vec{x}) = \pi_{a_i}(\cdot | K_{a_i}(\vec{x}))$. Now let $Pr(\cdot)$ be the probability function on this algebra determined by the following assignment: $Pr(\langle w_1, a_1 \rangle) = Pr(\langle w_2, a_2 \rangle) = 1/8$, and $Pr(\langle w_2, a_1 \rangle) = Pr(\langle w_1, a_2 \rangle) = 3/8$. We assume that, for each $i \in \{1, 2\}$, and each $\vec{x} \in \mathbf{C}$, $\pi_{a_i}(\vec{x}) = Pr(\cdot)$.

²⁴ It is perhaps worth pausing to ward off one potential source of confusion. In the following full centered agreement frame, the functions $K_i(\cdot)$, $\pi_i(\cdot)$, and $P_i(\cdot)$ are such that, relative to any $\vec{q} \in \mathbf{C}$, $P_i(\vec{q})(\cdot) = \pi_i(\cdot | K_i(\vec{q}))$. In this full centered agreement frame, then, agents are represented as setting their credences by conditionalizing their prior credal functions on their epistemic states.

Now, on one natural way of talking about what it is for an individual to ‘update by conditionalization’, it is clear that individuals with credences defined over sets of centered-worlds propositions often should not update by conditionalization. In particular, if by ‘updating by conditionalization’ we mean that between times t_1 and t_2 an individual adjusts their credences by conditionalizing their t_1 credences on their epistemic state at t_2 (and this is often what people mean by this phrase), then it is true that an individual with credences defined over sets of centered-worlds propositions should not always update by conditionalization. For there may be some *centered* proposition—say, the proposition that it is now t_1 —that the individual knows at t_1 , in which the individual may rationally lower their credence between t_1 and t_2 . But such a change is impossible if the individual is simply conditionalizing on their t_1 credences. For conditionalization preserves certainties. See, e.g., Arntzenius (2003) for a development of this point.

It is important, however, to see that the objections to updating by conditionalization in this sense, do not extend to an agent’s being such that they set their credences by conditionalizing an appropriate prior credal function on their epistemic state. For, note that we’ve assumed that, for each $\vec{q} \in \mathbf{C}$, $\pi_i(\vec{q})$ only assigns the necessary truth \mathbf{C} value 1. Thus, updating by conditionalizing $\pi_i(\vec{q})$ on one’s epistemic state does not force one to be dogmatically certain of any centered propositions, as updating by conditionalization over time may. And so this obvious problem with conditionalizing over time doesn’t give us reason to suppose that rational agents with credences defined over algebras that include centered propositions cannot be represented in a full centered agreement frame in which, relative to any $\vec{q} \in \mathbf{C}$, $P_i(\vec{q})(\cdot) = \pi_i(\cdot | K_i(\vec{q}))$.

Now it is, in fact, a somewhat controversial matter whether a rational individual with credences defined over an algebra that includes centered propositions may maintain the same prior over time. For example, Meacham (2008) argues that rational updating policies for such agents may require an agent to change their priors over time. But note that there is no assumption encoded in a full centered agreement frame that an individual over time will have the same prior. For our priors are assigned to agents, who are time-indexed individuals, and two agents who are related by a shift in time may have different priors at the same world in a full centered agreement frame.


 Fig. 1. \mathcal{F}_1

Some of the key features of this frame are represented in Figure 1. Here the set of centered-worlds enclosed in the dashed line corresponds to a_1 's (invariant) epistemic state, while the set of centered-worlds enclosed in the solid line corresponds to a_2 's (invariant) epistemic state. The numbers correspond to the (invariant) prior probabilities that each agent assigns to each centered-world.

First note that, for each $i \in \mathbf{A}$, $\text{Hom}[K_i(\cdot)] = \text{CHom}[K_i(\cdot)] = \text{Cond}[K_i(\cdot), \pi_i(\cdot), K_i(\cdot)] = [\pi_i = Pr(\cdot)] = \mathbf{C}$. It follows that, for each $i \in \mathbf{A}$, $\text{Com}_A(\text{Hom}[K_i(\cdot)]) = \text{Com}_A(\text{CHom}[K_i(\cdot)]) = \text{Com}_A(\text{Cond}[P_i(\cdot), \pi_i(\cdot), K_i(\cdot)]) = \text{Com}_A([\pi_i = Pr(\cdot)]) = \mathbf{C}$. And so, (a), (b), (c) and (e) hold at each $\vec{q} \in \mathbf{C}$. Next, note that the *uncentered* proposition $\phi = \{\langle w_1, a_1 \rangle, \langle w_1, a_2 \rangle\}$ is such that $[P_{a_1}(\phi) = 1/4] = [P_{a_2}(\phi) = 3/4] = \mathbf{C}$.²⁵ Thus, $\text{Com}_A([P_{a_1}(\phi) = 1/4])$ and $\text{Com}_A([P_{a_2}(\phi) = 3/4])$ are true at each $\vec{q} \in \mathbf{C}$.

Given what $K_i(\cdot)$, $\pi_i(\cdot)$ and $P_i(\cdot)$ represent in a centered agreement frame, this frame shows that a group of agents whose epistemic and credal states are defined over sets of centered-worlds propositions may be such that each sets their credences by updating the same probability function $Pr(\cdot)$ on their epistemic state, each is third-personally and first-personally introspective and these facts are common knowledge amongst the set of agents, and yet there is some proposition ϕ such that some agent assigns credence x_i to ϕ and another assigns credence x_j , and $x_i \neq x_j$ and these facts are common knowledge amongst the set of agents. In this sense, then, agents whose possible credal states are defined over algebras that include centered propositions can agree to disagree.

Now this fact, in itself, isn't terribly surprising. For, we should expect such open disagreement at least with respect to certain *centered* propositions. For example, it shouldn't be surprising that I may assign high credence to, say, the proposition that I am Michael Caie and that you may assign low credence to this same proposition, and yet our epistemic situations may be open to view in the manner characterized by (a), (b), (c), and (e).

What is perhaps initially surprising, however, is that, while there cannot be open disagreement (so defined) with respect to possible-worlds propositions amongst agents whose credences are defined only over possible-worlds propositions, there can be open disagreement (so defined) with respect to the analogues of possible-worlds propositions amongst agents whose credences are defined over centered-worlds propositions. For uncentered propositions are just those propositions that correspond, in the new algebra, to possible-worlds propositions. The reason that such open disagreement is possible, of course, is that, given certain probability functions defined over an algebra $\mathcal{P}(\mathbf{C})$, such as the one appealed to in Frame 1, there may be probabilistic correlations between uncentered propositions characterizing possible worlds and centered propositions characterizing agents. And, given such correlations, different information about which agent one is may give rise to different probabilities for how the world is.

²⁵ To see, for example, that $[P_{a_1}(\phi) = 1/4] = \mathbf{C}$, let $\vec{q} \in \mathbf{C}$. Then we have: $P_{a_1}(\vec{q})(\phi) = \pi_{a_1}(\phi|K_{a_1}(\vec{q})) = Pr(\phi)/Pr(K_{a_1}(\vec{q})) = 1/4$.

We've seen, then, that the First Agreement Theorem, and so No Agreeing to Disagree, may fail not only when we consider the credences that agents have in centered propositions, but also when we consider the credences that agents have in uncentered propositions. We can also show that the corresponding analogues of the Second Agreement Theorem and No Disagreement Amongst Confidants fail for such agents.

To see this, consider the basic centered agreement frame upon which \mathcal{F}_1 is constructed. Let \mathcal{B}_1 , then, be the basic centered agreement frame such that $\mathbf{W} = \{w_1, w_2\}$, $\mathbf{A} = \{a_1, a_2\}$, $\mathbf{C} = \mathbf{W} \times \mathbf{A}$, and such that, for every $\vec{x} \in \mathbf{C}$, $K_{a_1}(\vec{x}) = \{\langle w_1, a_1 \rangle, \langle w_2, a_1 \rangle\}$, and $K_{a_2}(\vec{x}) = \{\langle w_1, a_2 \rangle, \langle w_2, a_2 \rangle\}$. Given this frame, we have, for any $\vec{q} \in \mathbf{C}$ and each $i \in \mathbf{A}$: (d) $\text{Com}_A([K_i = K_i(\vec{q})])$ and (e) $\text{Com}_A(\text{CHom}[K_i(\cdot)])$.²⁶ To see why the former holds, note that for each $i \in \mathbf{A}$ and each $\vec{x}, \vec{y} \in \mathbf{C}$, $K_i(\vec{x}) = K_i(\vec{y})$. We've seen, though, that there is a probability function $Pr(\cdot)$ defined over $\mathcal{P}(\mathbf{C})$ and an uncentered proposition ϕ such that, for any $\vec{q} \in \mathbf{C}$, $Pr(\phi|K_{a_1}(\vec{q})) \neq Pr(\phi|K_{a_2}(\vec{q}))$. Thus, the Second Agreement Theorem fails not only when we consider the credences that agents have in centered propositions, but also when we consider the credences that agents have in uncentered propositions.

Now our informal gloss on what it is for a group of agents to be epistemic confidants is that it be common knowledge amongst them what each agent's epistemic state is, and common knowledge amongst them that each agent is perfectly introspective concerning their own epistemic state. In the case of agents whose epistemic states are defined over a set of possible worlds propositions, we noted that being third-personally introspective is the strongest introspective property that they could instantiate. We've seen, though, that this is not true for agents whose epistemic states are defined over a set of centered-worlds propositions. For such agents may be first-personally introspective, but it doesn't follow from the fact that one is third-personally introspective that one is also first-personally introspective. This, then, motivates the following additional terminological stipulation:

Def. We'll say that a set of agents whose epistemic states are defined over a set of centered-worlds propositions are *epistemic confidants* just in case it is common knowledge amongst them what each agent's epistemic state is, it is common knowledge amongst them that each agent is third-personally introspective, and it is common knowledge amongst them that each agent is first-personally introspective.

Now, given what $K_i(\cdot)$ represents in a centered agreement frame, (d) will hold at \vec{q} , for each $i \in G$, just in case each agent's epistemic state is common knowledge amongst them, at w_q . Moreover, given that each agent's epistemic state is common knowledge amongst the members of G , it follows that it is common knowledge amongst the members of G that each $i \in G$ is third-personally introspective.²⁷ It doesn't, however, follow from the fact that each agent's epistemic state is common knowledge amongst the members of G , that it is common knowledge amongst the members of G that each $i \in G$ is first-personally introspective. However, given what $K_i(\cdot)$ represents in a centered agreement frame, (e) will hold at \vec{q} , for each $i \in G$, just in case it is common knowledge amongst them, at w_q , that each agent is first-personally introspective. It follows, then, that a group of agents G

²⁶ Note that we also have for any $\vec{q} \in \mathbf{C}$ and each $i \in \mathbf{A}$: (a) $\text{Com}_A(\text{Hom}[K_i(\cdot)])$. As shown earlier, though, this follows from the fact that for any $\vec{q} \in \mathbf{C}$ and each $i \in \mathbf{A}$: (d) $\text{Com}_A([K_i = K_i(\vec{q})])$. To avoid redundancy, then, we omit this condition.

²⁷ See fn. 17 for the justification of this claim.

whose epistemic states are defined over a set of centered-worlds propositions will be epistemic confidants at a world w just in case (d) and (e) hold relative to all \vec{q} such that $w_q = w$, for each $i \in G$.²⁸

The agents represented in \mathcal{B}_1 , then, are epistemic confidants. We've seen, though, that there is a probability function $Pr(\cdot)$ defined over $\mathcal{P}(\mathbf{C})$ and an uncentered proposition ϕ such that if a_1 and a_2 set their credences by conditionalizing $Pr(\cdot)$ on their epistemic states, then they will assign different credences to ϕ . It follows, then, that a group of epistemic confidants whose epistemic and credal states are defined over a set of centered-worlds propositions may set their credences by conditionalizing the same probability function on their epistemic states and yet assign different credences to certain uncentered propositions. And so it follows that **No Disagreement Amongst Confidants** fails not only when we consider the credences that agents have in centered propositions, but also when we consider the credences that agents have in uncentered propositions.

A key feature of \mathcal{B}_1 is that a_1 and a_2 's (invariant) epistemic states are disjoint. Given this, it is clear that we can find a probability distribution such that some uncentered proposition has a different probability conditional on each of these sets. Now this feature of a_1 and a_2 's epistemic states is precluded by the assumption that $K_{a_1}(\cdot)$ and $K_{a_2}(\cdot)$ are reflexive relative to each point in our frame. However, \mathcal{B}_1 demonstrates that such disjointness of epistemic states is not ruled out by the assumption that $K_{a_1}(\cdot)$ and $K_{a_2}(\cdot)$ are *center-reflexive* relative to each point in our frame. Indeed, we noted earlier that, given that $K_{a_1}(\cdot)$ and $K_{a_2}(\cdot)$ are reflexive relative to each point in our frame, it follows from the fact that for each $i \in A = \{a_1, a_2\}$, $\text{Com}_A([K_i = K_i(w)])$ holds at w , that $K_{a_1}(w) = K_{a_2}(w)$. \mathcal{B}_1 , though, demonstrates that if $K_{a_1}(\cdot)$ and $K_{a_2}(\cdot)$ are center-reflexive relative to each point in our frame, then even if, for each $i \in A = \{a_1, a_2\}$, both $\text{Com}_A([K_i = K_i(\vec{q})])$ and $\text{Com}_A(\text{CHom}[K_i(\cdot)])$ hold at \vec{q} , nonetheless it may be the case that $K_{a_1}(\vec{q}) \cap K_{a_2}(\vec{q}) = \emptyset$. The fact, then, that **Factivity**, for agents representable in a centered agreement frame, corresponds to the functions representing their epistemic states being center-reflexive instead of reflexive, allows for the possibility that a set of epistemic confidants may set their credences by conditionalizing the same probability function on their epistemic states and yet assign distinct credences to the same uncentered proposition.

\mathcal{B}_1 , then, establishes that it is at least formally possible for a group of agents, representable in a centered agreement frame, to be epistemic confidants and to set their credences by conditionalizing the same probability function on their respective epistemic states and yet disagree about certain uncentered propositions. It's worth noting, though, that the probability functions that allow for such disagreement amongst the agents represented in \mathcal{B}_1 have a peculiar feature.

Def. We say that a probability function $Pr(\cdot)$ defined on $\mathcal{P}(\mathbf{C})$ is *center-indifferent* just in case, if $w_x = w_q$, then $Pr(\vec{x}) = Pr(\vec{q})$.

It's fairly straightforward to see that the only probability functions that allow for disagreement amongst the agents in \mathcal{B}_1 are non-center-indifferent.²⁹ Indeed, we can show something much more general:

²⁸ Note that since (d) and (e) hold relative to some \vec{q} such that $w_q = w$ just in case they hold relative to all \vec{q} such that $w_q = w$, we could equally say that a group of agents G whose epistemic states are defined over a set of centered-worlds propositions will be epistemic confidants at a world w just in case (d) and (e) hold relative to some \vec{q} such that $w_q = w$.

²⁹ To see this note that the only two contingent uncentered propositions in this algebra are $w_1 = \{\langle w_1, a_1 \rangle, \langle w_1, a_2 \rangle\}$ and $w_2 = \{\langle w_2, a_1 \rangle, \langle w_2, a_2 \rangle\}$. And, given that, $Pr(\langle w_1, a_1 \rangle)$

First Centered Agreement Theorem: Let \mathcal{B} be a basic centered agreement frame and let $Pr(\cdot)$ be a center-indifferent probability function on $\mathcal{P}(\mathbf{C})$. If, at some $\vec{q} \in \mathbf{C}$, for each $i \in \mathbf{A}$, the following propositions are true: (d) $\text{Com}_A([K_i = K_i(\vec{q})])$ and (e) $\text{Com}_A(\text{CHom}[K_i(\cdot)])$, then for every $i, j \in \mathbf{A}$ and every uncentered proposition ϕ , $Pr(\phi|K_i(\vec{q})) = Pr(\phi|K_j(\vec{q}))$.³⁰

Very roughly, the reason that this holds is the following. Say that $\langle w, i \rangle$ is a w -centered-world. And let \mathcal{K}_A be the set of epistemic states for the members of \mathbf{A} . We can show that if a set of agents \mathbf{A} have epistemic and credal states defined over an algebra $\mathcal{P}(\mathbf{C})$ such that $\mathbf{A} = \{a_x : \text{for some } \vec{x} \in \mathbf{C}\}$ and are such that (d) and (e) hold for each $i \in \mathbf{A}$, then \mathcal{K}_A will be a set of nonoverlapping subsets of \mathbf{C} , such that each member of \mathcal{K}_A has the same proportion of w -centered-worlds as any other member of \mathcal{K}_A . And, given this, it follows that, for any center-indifferent probability function $Pr(\cdot)$ and any $\xi, \gamma \in \mathcal{K}_A$, if ϕ is an uncentered proposition, then $Pr(\phi|\xi) = Pr(\phi|\gamma)$.

Since a set of agents whose epistemic and credal states are defined over a set of centered-worlds propositions will be epistemic confidants just in case their epistemic states are representable in a basic centered agreement frame in which the appropriate instances of (d) and (e) hold for each agent at each centered-world whose possible world parameter is the actual world, the First Centered Agreement Theorem shows:


Centered Agreement: If a group of agents \mathbf{A} are epistemic confidants whose epistemic states are defined over an algebra of centered-worlds propositions $\mathcal{P}(\mathbf{C})$ such that $\mathbf{A} = \{a_x : \text{for some } \vec{x} \in \mathbf{C}\}$, then if the members of \mathbf{A} set their credences by conditionalizing the same center-indifferent probability function on their epistemic states, then each member of \mathbf{A} must have the same credence in any *uncentered* proposition.

It is worth highlighting that this result depends crucially upon the assumption that it is not only common knowledge amongst \mathbf{A} what each agent's epistemic state is, but also that it is common knowledge that each agent is *first-personally* introspective. For we can provide a basic centered agreement frame in which a group of agents \mathbf{A} are such that (d) $\text{Com}_A([K_i = K_i(\vec{q})])$ holds, for each $\vec{q} \in \mathbf{C}$ and each $i \in \mathbf{A}$, and yet for some center-indifferent probability function $Pr(\cdot)$, there is some $i, j \in \mathbf{A}$, some $\vec{q} \in \mathbf{C}$ and some *uncentered* proposition ϕ such that $Pr(\phi|K_i(\vec{q})) \neq Pr(\phi|K_j(\vec{q}))$.

Frame 2: Let \mathcal{B}_2 be such that $\mathbf{W} = \{w_1, w_2\}$, $\mathbf{A} = \{a_1, a_2\}$, and $\mathbf{C} = \mathbf{W} \times \mathbf{A}$. We let $K_{a_1}(\cdot)$ be such that, for every $\vec{x} \in \mathbf{C}$, $K_{a_1}(\vec{x}) = \mathbf{C}$. And we let $K_{a_2}(\cdot)$ be such that, for every $\vec{x} \in \mathbf{C}$, $K_{a_2}(\vec{x}) = \{\langle w_1, a_1 \rangle, \langle w_1, a_2 \rangle, \langle w_2, a_2 \rangle\}$.

Now let $Pr(\cdot)$ be such that, for each $\vec{z} \in \mathbf{C}$, $Pr(\vec{z}) = 1/4$. We can represent the resulting frame and this probability distribution over it as in Figure 2. Here the set of centered-worlds enclosed in the solid line corresponds to a_2 's (invariant) epistemic state, the set of centered-worlds enclosed in the dashed line corresponds to a_1 's (invariant) epistemic state, while the numbers correspond to probability values assigned by $Pr(\cdot)$.

³⁰ See the Appendix for the proof of this result.


 Fig. 2. \mathcal{B}_2

Note that, for each $i \in \mathbf{A}$ and each $\vec{x}, \vec{y} \in \mathbf{C}$, $K_i(\vec{x}) = K_i(\vec{y}) \in \mathcal{P}(\mathbf{C})$. It follows that, for any $\vec{q} \in \mathbf{C}$ and each $i \in \mathbf{A}$, (d) $\text{Com}_A([K_i = K_i(\vec{q})])$ holds.³¹ Nonetheless, it is easy to see that there is some $\vec{q} \in \mathbf{C}$ and some *uncentered* proposition ϕ such that $Pr(\phi|K_{a_1}(\vec{q})) \neq Pr(\phi|K_{a_2}(\vec{q}))$. For example, consider $\phi = \{\langle w_2, a_1 \rangle, \langle w_2, a_2 \rangle\}$. For any $\vec{q} \in \mathbf{C}$, we have $Pr(\phi|K_{a_1}(\vec{q})) = 1/2$ and $Pr(\phi|K_{a_2}(\vec{q})) = 1/3$.

Importantly, $\text{Com}_A(\text{CHom}[K_{a_2}(\cdot)])$ is false at every $\vec{q} \in \mathbf{C}$ in \mathcal{B}_2 . Thus, while in this frame, relative to each centered-world, it is common knowledge amongst a_1 and a_2 what a_2 's epistemic state is, and so it is common knowledge amongst a_1 and a_2 that a_2 is third-personally introspective, it is not common knowledge amongst a_1 and a_2 , at any centered-world, that a_2 is first-personally introspective. For, at each centered world, a_2 's epistemic state is: $\{\langle w_1, a_1 \rangle, \langle w_1, a_2 \rangle, \langle w_2, a_2 \rangle\}$, but $K_{a_1}(\langle w_1, a_1 \rangle) = \mathbf{C}$. While, then, a_2 knows the uncentered proposition that a_2 's epistemic state is $\{\langle w_1, a_1 \rangle, \langle w_1, a_2 \rangle, \langle w_2, a_2 \rangle\}$, a_2 does not know the *centered* proposition that their epistemic state is $\{\langle w_1, a_1 \rangle, \langle w_1, a_2 \rangle, \langle w_2, a_2 \rangle\}$. This first-personal ignorance allows a_1 and a_2 's respective epistemic states to partially but not wholly overlap. And this partial overlap allows for there to be disagreement between the two with respect to certain uncentered propositions, despite the fact that each agent sets their credences by conditionalizing the same center-indifferent probability function on their respective epistemic state.

Now, it may be tempting to conclude that the First Centered Agreement Theorem shows that *any* set of epistemic confidants, whose epistemic states are defined over an algebra of centered-worlds propositions, and who set their credences by conditionalizing the same center-indifferent probability function on their respective epistemic states, must assign the same credence to every uncentered proposition over which their credences are defined. It's important to see, though, that this does not follow from the First Centered Agreement Theorem. And, indeed, we can show that this general claim is false. To see this, consider the following basic centered agreement frame.

Frame 3: Let \mathcal{B}_3 be such that $\mathbf{W} = \{w_1, w_2, w_3\}$, $\mathbf{A} = \{a_1, a_2, a_3\}$, and $\mathbf{C} = \mathbf{W} \times \mathbf{A}$. We let $K_{a_1}(\cdot)$ be such that, for every $\vec{x} \in \mathbf{C}$, $K_{a_1}(\vec{x}) = \{\langle w_1, a_1 \rangle, \langle w_2, a_1 \rangle, \langle w_3, a_1 \rangle, \langle w_3, a_2 \rangle\}$. We let $K_{a_3}(\cdot)$ be such that, for every $\vec{x} \in \mathbf{C}$, $K_{a_3}(\vec{x}) = \{\langle w_1, a_3 \rangle, \langle w_2, a_3 \rangle, \langle w_3, a_3 \rangle\}$. And we let $K_{a_2}(\cdot)$, be such that (i) $K_{a_2}(\langle w_1, a_i \rangle) = \{\langle w_1, a_2 \rangle\}$, for $i \in \{1, 2, 3\}$, (ii) $K_{a_2}(\langle w_2, a_i \rangle) = \{\langle w_2, a_2 \rangle\}$, for $i \in \{1, 2, 3\}$, and (iii) $K_{a_2}(\langle w_3, a_i \rangle) = \{\langle w_1, a_1 \rangle, \langle w_2, a_1 \rangle, \langle w_3, a_1 \rangle, \langle w_3, a_2 \rangle\}$, for $i \in \{1, 2, 3\}$.

Now let $Pr(\cdot)$ be such that, for each $\vec{z} \in \mathbf{C}$, $Pr(\vec{z}) = 1/9$. We can represent the resulting frame and this probability distribution as in Figure 3. Here the set of centered-worlds

³¹ It also follows that, for any $\vec{q} \in \mathbf{C}$ and each $i \in \mathbf{A}$, (a) $\text{Com}_A(\text{Hom}[K_i(\cdot)])$ holds. See fn. 22 for a justification of this claim.

enclosed in the solid line corresponds to a_1 's (invariant) epistemic state, as well as a_2 's epistemic state at each $\langle w, a \rangle$ such that $w = w_3$, and the set of centered-worlds enclosed in the dashed line corresponds to a_3 's (invariant) epistemic state. The numbers correspond to the probability values assigned by $Pr(\cdot)$.




Fig. 3. \mathcal{B}_3

First, note that, for $i \in \{a_1, a_3\}$, $\text{CHom}[K_i(\cdot)] = \mathbf{C}$. And, for each $i \in \{a_1, a_3\}$ and each $\vec{x}, \vec{y} \in \mathbf{C}$, $K_i(\vec{x}) = K_i(\vec{y}) \in \mathcal{P}(\mathbf{C})$. It follows that, for any $\vec{q} \in \mathbf{C}$ and each $i \in \{a_1, a_3\}$: (d) $\text{Com}_{\{a_1, a_3\}}([K_i = K_i(\vec{q})])$ and (e) $\text{Com}_{\{a_1, a_3\}}(\text{CHom}[K_i(\cdot)])$ hold.³² Nonetheless, it is easy to see that there is some $\vec{q} \in \mathbf{C}$ and some *uncentered* proposition ϕ such that $Pr(\phi|K_{a_1}(\vec{q})) \neq Pr(\phi|K_{a_3}(\vec{q}))$. For example, consider $\phi = \{\langle w_3, a_1 \rangle, \langle w_3, a_2 \rangle, \langle w_3, a_3 \rangle\}$. For any $\vec{q} \in \mathbf{C}$, we have $Pr(\phi|K_{a_1}(\vec{q})) = 1/2$ and $Pr(\phi|K_{a_3}(\vec{q})) = 1/3$.

Now a set of agents whose epistemic states are defined over a set of centered-worlds propositions will be epistemic confidants just in case their epistemic states are representable in a basic agreement frame in which the appropriate instances of (d) and (e) hold for each agent, at each centered-world whose possible world parameter is the actual world. It follows, then, that epistemic confidants whose epistemic states are defined over an algebra of centered-worlds propositions may assign different credences to certain uncentered propositions, despite the fact that each agent sets their credences by conditionalizing the same center indifferent probability function on their respective epistemic state.

To see why this result is compatible with the **First Centered Agreement Theorem**, note that, while, at each point in \mathcal{B}_3 , it is common knowledge amongst $\{a_1, a_3\}$ what each of their epistemic states is, it is not the case that, at each point in \mathcal{B}_3 , it is common knowledge amongst \mathbf{A} what each of their epistemic states is. In the epistemic situation represented by this basic agreement frame, then, only a proper subset of a group of agents \mathbf{A} whose epistemic states are defined over an algebra $\mathcal{P}(\mathbf{C})$ such that $\mathbf{A} = \{a_x : \text{for some } \vec{x} \in \mathbf{C}\}$ are epistemic confidants. But the **First Centered Agreement Theorem** only tells us that a group of agents \mathbf{A} whose epistemic states are defined over an algebra $\mathcal{P}(\mathbf{C})$ such that $\mathbf{A} = \{a_x : \text{for some } \vec{x} \in \mathbf{C}\}$ and who are epistemic confidants cannot assign different credences to any uncentered propositions given that each sets their credences by conditionalizing the same center-indifferent probability function on their epistemic state.

If an agent i has an epistemic state defined over an algebra $\mathcal{P}(\mathbf{C})$ such that $\mathbf{A} = \{a_x : \text{for some } \vec{x} \in \mathbf{C}\}$, we can think of \mathbf{A} as the set of agents that i thinks they might be

³² We also have that for any $\vec{q} \in \mathbf{C}$ and each $i \in \{a_1, a_3\}$, (a) $\text{Com}_{\{a_1, a_3\}}(\text{Hom}[K_i(\cdot)])$. But, again, this follows from the fact that for any $\vec{q} \in \mathbf{C}$ and each $i \in \{a_1, a_3\}$, (d) $\text{Com}_{\{a_1, a_3\}}([K_i = K_i(\vec{q})])$.

identical to, prior to receiving any information. Call these agents i 's *epistemic candidates*. What the **First Centered Agreement Theorem** shows is that when a set of agents G are epistemic confidants, and each $i \in G$ is such that their set of epistemic candidates is G itself, then each $i \in G$ must assign the same credence to every uncentered proposition for which their credences are defined, given that each member of G sets their credences by conditionalizing the same center-indifferent probability function on their epistemic state. \mathcal{B}_3 shows, however, that when the set of epistemic candidates for a set of epistemic confidants G is a strict superset of G , such agreement is not required.

The reason for this is as follows. When the set of epistemic candidates for a set of epistemic confidants G is just G itself, it follows that the proportion of w -centered-worlds amongst a_i 's epistemic state will be the same as the proportion of w -centered worlds amongst a_j 's epistemic state, for all epistemic confidants a_i, a_j and all $w \in \mathbf{W}$. However, when there are additional epistemic candidates who are not members of G , there may be a greater proportion of w -centered-worlds in some agents' epistemic states than in others. So, in \mathcal{B}_3 we see that there is a greater proportion of w_3 -centered-worlds in a_1 's epistemic state than in a_3 's epistemic state. And the reason for this is that while both a_1 and a_3 are uncertain about whether w_1, w_2 , or w_3 is actual, for a_1 there are two epistemic w_3 possibilities—by a_1 's lights the world may be w_3 and they may be a_1 , or the world may be w_3 and they may be a_2 —while for a_3 there is only one w_3 possibility. In order, however, for there to be this additional w_3 possibility for a_1 , without there being corresponding w_1 and w_2 possibilities, it's necessary that a_2 not be an epistemic confidant of a_1 and a_3 .

We've seen, then, that a set of epistemic confidants who set their credences by conditionalizing the same center-indifferent probability function on their respective epistemic states need not assign the same credence to each uncentered proposition. A natural question that arises, then, is whether, given an algebra $\mathcal{P}(\mathbf{C})$, there is any probability function $Pr(\cdot)$, defined over this algebra, that ensures that any group of epistemic confidants whose credal states are defined over $\mathcal{P}(\mathbf{C})$ and who set their credences by conditionalizing $Pr(\cdot)$ on their respective epistemic states will assign the same credence to any uncentered proposition. The answer to this question, it turns out, is 'yes'. For we can show:

Second Centered Agreement Theorem: Let \mathcal{B} be a basic centered agreement frame. Then there exists some probability function $Pr(\cdot) \in \mathcal{M}_{\mathbf{C}}$, such that for any $\vec{q} \in \mathbf{C}$, if every $i \in G \subseteq \mathbf{A}$ is such that the following propositions are true at \vec{q} : (d) $\text{Com}_G([K_i = K_i(\vec{q})])$ and (e) $\text{Com}_G(\text{CHom}[K_i(\cdot)])$, then for each uncentered proposition ψ and each $i, j \in G$, $Pr(\psi|K_i(\vec{q})) = Pr(\psi|K_j(\vec{q}))$.³³

Very roughly, the proof for this proceeds as follows. First, we show that there is a particular subset $\mathbf{M} \subseteq \mathcal{P}(\mathbf{C})$ and a particular partition of \mathbf{M} such that, if a set of agents $G \subseteq \mathbf{A}$, whose epistemic states are defined over the algebra $\mathcal{P}(\mathbf{C})$, are such that (d) and (e) hold at some \vec{q} , for each $i \in G$, then there is some cell Δ in this partition such that $K_i(\vec{q}) \in \Delta$, for each $i \in G$. We then show that there is a probability function $Pr(\cdot)$ such that for every ξ, γ in the same cell of this partition and any uncentered proposition ψ , $Pr(\psi|\xi) = Pr(\psi|\gamma)$.

Since a set of agents whose epistemic states are defined over a set of centered-worlds propositions will be epistemic confidants just in case their epistemic states are representable in a basic agreement frame in which the appropriate instances of (d) and (e) hold

³³ See the Appendix for the proof of this result.

for each agent, at each centered-world whose possible world parameter is the actual world, the **Second Centered Agreement Theorem** shows:

Possible Agreement: There is a probability function $Pr(\cdot)$ such that any set of epistemic confidants who set their credences by conditionalizing $Pr(\cdot)$ on their respective epistemic states will assign the same credence to any uncentered proposition.

Now, in light of the fact that epistemic confidants may assign different credences to some uncentered proposition even if they set their credences by conditionalizing the same center-indifferent probability function on their respective epistemic states, another question that arises is whether there are natural constraints that we might impose on a group of agents, in addition to those required for them to be epistemic confidants, that would ensure agreement given that all such agents set their credences by conditionalizing the same center-indifferent probability function on their respective epistemic states. The answer to this question, it turns out, is also ‘yes’.

Def. We’ll say that an agent i is *first-personally G -introspective* just in case, if i ’s epistemic state is ϕ , then i knows the centered proposition that it is common knowledge amongst G that their epistemic state is characterized by ϕ .

Def. We’ll say that a set of agents G whose epistemic states are defined over a set of centered-worlds propositions are *epistemic peers* just in case the members of G are epistemic confidants and each member of G is first-personally G -introspective.

A first point to note is that in \mathcal{B}_3 , although a_1 and a_3 are epistemic confidants, a_1 is not $\{a_1, a_3\}$ -introspective. To see this, note first that the centered proposition that it is common knowledge amongst G that one’s epistemic state is characterized by ϕ may be represented in a centered agreement frame by $\{\vec{z} : \text{Com}_G([K_{a_z} = \phi])\}$, and second, that $K_{a_1}(\{\vec{z} : \text{Com}_{a_1, a_3}([K_{a_z} = \phi])\})$ fails to hold at each point in \mathcal{B}_3 , where ϕ is a_1 ’s invariant epistemic state in this frame. The reason for this is that, given a_1 ’s (invariant) epistemic state as represented in this frame, a_1 is uncertain about whether they are a_1 or a_2 . But while a_3 knows what a_1 ’s epistemic state is, a_3 does not know what a_2 ’s epistemic state is. Thus since a_1 is uncertain whether they are a_1 or a_2 , and since a_3 knows what a_1 ’s epistemic state is but not what a_2 ’s epistemic state is, it follows that a_1 does not know the centered proposition that a_3 knows what *their* epistemic state is.

A natural thought, then, is that, if we consider groups of epistemic peers, cases such as \mathcal{B}_3 , in which conditionalizing on a center-indifferent probability function can lead to different probabilities for certain uncentered propositions, will not arise. This, it turns out, is true. For we can show:

Third Centered Agreement Theorem: Let \mathcal{B} be a basic centered agreement frame, and let $Pr(\cdot)$ be a center-indifferent probability function defined on $\mathcal{P}(\mathbf{C})$. If, at some $\vec{q} \in \mathbf{C}$, every $i \in G \subseteq \mathbf{A}$ is such that the following propositions are true at \vec{q} : (d) $\text{Com}_G([K_i = K_i(\vec{q})])$, (e) $\text{Com}_G(\text{CHom}[K_i(\cdot)])$ and (f) $K_i(\{\vec{z} : \text{Com}_G([K_{a_z} = K_i(\vec{q})])\})$, then for every uncentered proposition ψ , and each $i, j \in G$, $Pr(\psi|K_i(\vec{q})) = Pr(\psi|K_j(\vec{q}))$.³⁴

³⁴ See the Appendix for the proof of this result.

The proof of this result proceeds roughly as follows. Again, say that $\langle w, i \rangle$ is a *w-centered-world*. And let \mathcal{K}_G be the set of epistemic states for the members of $G \subseteq \mathbf{A}$. We can show that if a set of agents $G \subseteq \mathbf{A}$ have epistemic and credal states defined over an algebra $\mathcal{P}(\mathbf{C})$ and are such that (d), (e) and (f) hold for each $i \in G$, then \mathcal{K}_G will be a set of subsets of \mathbf{C} , such that each member of \mathcal{K}_G has the same proportion of *w-centered-worlds* as any other member of \mathcal{K}_G . And, given this, it follows that, for any center-indifferent probability function $Pr(\cdot)$ and any $\xi, \gamma \in \mathcal{K}_A$, if ϕ is an uncentered proposition, then $Pr(\phi|\xi) = Pr(\phi|\gamma)$.

Now it's easy to see that a set of agents whose epistemic and credal states are defined over a set of centered-worlds propositions will be epistemic peers just in case their epistemic states are representable in a basic centered agreement frame in which the appropriate instances of (d), (e) and (f) hold for each agent at each centered-world whose possible world parameter is the actual world. Given this, the **Third Centered Agreement Theorem** shows:

Peer Agreement: If a group of epistemic peers G set their credences by conditionalizing the same center-indifferent probability function on their respective epistemic states, then each member of G must have the same credence in any *uncentered* proposition.

§3. Conclusion. We've seen that certain well-known theorems, which ensure that agreement will be achieved for agents whose epistemic and credal states are defined over algebras of possible worlds propositions, given that such agents have sufficient information about one another's epistemic states and that such agents update on the same probability function, fail when we consider agents whose epistemic and credal states are defined over algebras of centered-worlds propositions. We've seen, though, that while agreement may not be guaranteed given that such agents update on the same probability function, there are some probability functions that will ensure agreement at least with respect to uncentered propositions given that such agents have sufficient information about one another's epistemic states.

In closing, let me very briefly note a way in which these results may bear on a topic of philosophical controversy. It seems to me to be quite plausible that it is at least rationally permissible for epistemic confidants to assign the same credence to any uncentered proposition. And the **Second Centered Agreement Theorem** assures us that this claim is compatible with the natural Bayesian assumption that it is rationally permissible for agents to set their credences by conditionalizing the same probability function on their respective epistemic states. It turns out, though, that this plausible principle is incompatible with certain *prima facie* plausible principles about how agents should update their credences in light of self-locating information. In particular, it can be shown that this principle is incompatible with both the standard Thirder and Halfer solutions to the famous Sleeping Beauty problem.³⁵ This, I think, provides us with some reason to reject both of these treatments of this puzzle.

One natural response to this argument, though, is to reject the preceding principle in favor of a weaker principle. Instead of saying that it is always rationally permissible for epistemic confidants to assign the same credence to any uncentered proposition, one may maintain that it is always rationally permissible for epistemic peers to assign the same

³⁵ See Caie (2016) for development and discussion of this point. See Elga (2000) and Lewis (2001) for the Thirder and Halfer responses to the Sleeping Beauty problem.

credence to any uncentered proposition, though in certain cases it may be rationally impermissible for epistemic confidants to do so. It turns out that this weaker constraint is compatible with the Thirder solution to the famous Sleeping Beauty problem, though not with the Halfer solution. Moreover, the **Third Centered Agreement Theorem** assures us that this weaker principle is compatible with the prima facie plausible claim that rational agents should always set their credences by conditionalizing on a center-indifferent probability function.

These results and their bearing on the question of how rational agents should update their credences in light of self-locating information are discussed in detail in Caie (2016).

§4. Acknowledgments. Thanks to two anonymous referees at *The Review of Symbolic Logic* for very helpful comments. Special thanks to Harvey Lederman for extremely helpful comments on an ancestor to this paper.

§A. Proofs.

A.1. Second Agreement Theorem.

CLAIM: Let \mathcal{B} be a basic agreement frame and $Pr(\cdot)$ a probability function on $\mathcal{P}(\mathbf{W})$. If, at some $w \in \mathbf{W}$, for each $i \in G \subseteq \mathbf{A}$, the following proposition is true (d) $\text{Com}_G([K_i = K_i(w)])$, then, for each $i, j \in G$, $Pr(\cdot|K_i(w)) = Pr(\cdot|K_j(w))$.

Proof. Let i and j be arbitrary members of G . We'll show that, given (d), it follows that $K_i(w), K_j(w) \in \mathcal{P}(\mathbf{W})$, and $K_i(w) = K_j(w)$, and so $Pr(\cdot|K_i(w)) = Pr(\cdot|K_j(w))$. Given that $\text{Com}_G([K_i = K_i(w)])$ holds at w , it follows that $K_i(w) \subseteq [K_i = K_i(w)]$ and $K_j(w) \subseteq [K_i = K_i(w)]$ and so $K_i(w) \neq \star$ and $K_j(w) \neq \star$. Moreover, given that $\text{Com}_G([K_i = K_i(w)])$ holds at w , we have that (i) for every $w' \in K_j(w)$, $K_i(w') = K_i(w)$. Moreover, since $K_i(w) \neq \star$, and since, for every $w' \in \mathbf{W}$, if $K_i(w') \in \mathcal{P}(\mathbf{W})$ then $K_i(\cdot)$ is reflexive at w' , it follows that (ii) for every $w' \in K_j(w)$, $w' \in K_i(w')$. Together (i) and (ii) entail (iii) $K_j(w) \subseteq K_i(w)$. Furthermore, given that $\text{Com}_G([K_j = K_j(w)])$ holds at w , we have that (iv) for every $w' \in K_i(w)$, $K_j(w') = K_j(w)$. Moreover, since $K_j(w) \neq \star$, and since, for every $w' \in \mathbf{W}$, if $K_j(w') \in \mathcal{P}(\mathbf{W})$ then $K_j(\cdot)$ is reflexive at w' , it follows that (v) for every $w' \in K_i(w)$, $w' \in K_j(w')$. Together (iv) and (v) entail (vi) $K_i(w) \subseteq K_j(w)$. And so, given (iii) and (vi), we have $K_i(w) = K_j(w)$. \square

A.2. First Centered Agreement Theorem.

CLAIM: Let \mathcal{B} be a basic centered agreement frame and let $Pr(\cdot)$ be a center-indifferent probability function on $\mathcal{P}(\mathbf{C})$. If, at some $\vec{q} \in \mathbf{C}$, for each $i \in \mathbf{A}$, the following propositions are true: (d) $\text{Com}_A([K_i = K_i(\vec{q})])$ and (e) $\text{Com}_A(\text{CHom}[K_i(\cdot)])$, then for every $i, j \in \mathbf{A}$ and every uncentered proposition ϕ , $Pr(\phi|K_i(\vec{q})) = Pr(\phi|K_j(\vec{q}))$.

Proof. Let \mathcal{B} be a basic centered agreement frame, let $Pr(\cdot)$ be a regular center-indifferent probability function on $\mathcal{P}(\mathbf{C})$, and let $\vec{q} \in \mathbf{C}$ be a point such that, for each $i \in \mathbf{A}$, (d) and (e) hold at \vec{q} . We will show that, for every $i, j \in \mathbf{A}$ and every uncentered proposition ϕ , $Pr(\phi|K_i(\vec{q})) = Pr(\phi|K_j(\vec{q}))$.

We say that there is a *finite length A-accessibility chain* from \vec{z}_1 to \vec{z}_n just in case there is a sequence of n centered-worlds: $\langle \vec{z}_1, \vec{z}_2, \dots, \vec{z}_n \rangle$ such that, for each $\vec{z}_j, \vec{z}_{j+1} \in \langle \vec{z}_1, \vec{z}_2, \dots, \vec{z}_n \rangle$, there is some $i \in \mathbf{A}$ such that $\vec{z}_{j+1} \in K_i(\vec{z}_j)$.

Let γ , then, be the set of points \vec{z} such that there is a finite length \mathbf{A} -accessibility chain from \vec{q} to \vec{z} . It follows that $\vec{q} \in \gamma$. Furthermore, given that \mathcal{B} is a basic centered agreement frame, and that, at \vec{q} , (d) and (e) hold, it follows that, for each $i \in \mathbf{A}$ and each $\vec{x} \in \gamma$, (i) $K_i(\vec{q}) = K_i(\vec{x}) \in \mathcal{P}(\mathbf{C})$, (ii) K_i is center-reflexive at \vec{x} , (iii) K_i is center homogeneous at \vec{x} . There are now two cases to consider.

(Case 1): For each $i, j \in \mathbf{A}$, $K_i(\vec{q}) = K_j(\vec{q})$. In this case, it is trivial that for each uncentered proposition ϕ , $Pr(\phi|K_i(\vec{q})) = Pr(\phi|K_j(\vec{q}))$.

(Case 2): There are some $i, j \in \mathbf{A}$, such that $K_i(\vec{q}) \neq K_j(\vec{q})$. There will, then, be a finite number $|\mathbf{A}| \geq n > 1$ of distinct propositions ψ such that $K_i(\vec{q}) = \psi$ for some $i \in \mathbf{A}$. Given some arbitrary ordering of this set, we denote each such proposition ψ^m for $m \leq n$. We let $f(\psi^m) = \{i \in \mathbf{A} : K_i(\vec{q}) = \psi^m\}$. We first show:

(A) For every $\vec{z} \in \psi^m$, $a_z = i$ for some $i \in f(\psi^m)$.

Proof of (A). Assume, for reductio, that we have some $\vec{z} \in \psi^m$, such that $a_z \neq i$ for every $i \in f(\psi^m)$. Now, we know that $\psi^m \subseteq \gamma$. But, then, since $a_z \neq i$ for every $i \in f(\psi^m)$, and since, given (i), $K_{a_z}(\cdot)$ is constant on γ , it follows that $K_{a_z}(\vec{z}) \neq \psi^m$. We have, then, $K_i(\vec{q}) = \psi^m$, for any $i \in f(\psi^m)$, $\vec{z} \in \psi^m$ and $K_{a_z}(\vec{z}) \neq \psi^m$. This entails, though, that $K_i(\cdot)$ is not center-homogeneous at \vec{q} . But, given (iii), it follows that $K_i(\cdot)$ must be center-homogeneous at \vec{q} . Hence, for every $\vec{z} \in \psi^m$, $a_z = i$ for some $i \in f(\psi^m)$. \square

Next we show:

(B) For each ψ^m, ψ^k and each $i \in f(\psi^m), j \in f(\psi^k)$ and each $w \in \mathbf{W}$, $\langle w, i \rangle \in \psi^m$ iff $\langle w, j \rangle \in \psi^k$.

Proof of (B). First, assume $\langle w, i \rangle \in \psi^m$. We have $\psi^m \subseteq \gamma$. So, since, given (i), K_j is constant over γ , we have $K_j(\langle w, i \rangle) = \psi^k$. Moreover, given (ii), we have that, for each point in $\vec{x} \in \gamma$, and each $i \in \mathbf{A}$, $\langle w_x, i \rangle \in K_i(\vec{x})$. Thus, we have, $\langle w, j \rangle \in \psi^k$. And so we have that if $\langle w, i \rangle \in \psi^m$, then, $\langle w, j \rangle \in \psi^k$. Next assume that $\langle w, j \rangle \in \psi^k$. We have $\psi^k \subseteq \gamma$. So, since, given (i), K_i is constant over γ , we have $K_i(\langle w, j \rangle) = \psi^m$. And so, given (ii), we have $\langle w, i \rangle \in \psi^m$. Thus, we have that if $\langle w, j \rangle \in \psi^k$, then, $\langle w, i \rangle \in \psi^m$, and so, $\langle w, j \rangle \in \psi^k$ iff $\langle w, i \rangle \in \psi^m$. \square

Given some $w' \in \mathbf{W}$, we can think of the uncentered proposition $\{\langle w, i \rangle : w = w'\}$ as characterizing w' . Call such a proposition a *world-proposition*. We will now show that, given that $Pr(\cdot)$ is center-indifferent, $Pr(\xi|\psi^m) = Pr(\xi|\psi^k)$, for each world-proposition ξ .

For each $w \in \mathbf{W}$, we let $w(\psi^m) = \{\langle w', i \rangle \in \psi^m : w' = w\}$. Given (A) and (B), we have, then, that, for every ψ^m, ψ^k , and every $w \in \mathbf{W}$, $w(\psi^m) \neq \emptyset$ iff $w(\psi^k) \neq \emptyset$. For arbitrary ψ^k , then, we let $\mathcal{W} = \{w \in \mathbf{W} : w(\psi^k) \neq \emptyset\}$. Note that whichever ψ^k we choose for this, the resulting set will be the same.

For our proof, there are now two cases to consider.

- (I) $w' \notin \mathcal{W}$.
- (II) $w' \in \mathcal{W}$.

For (I), we have that $Pr(\{\langle w, i \rangle : w = w'\}|\psi^k) = Pr(\{\langle w, i \rangle : w = w'\}|\psi^m) = 0$. This follows from the fact that, given (A) and (B), if $w' \notin \mathcal{W}$, then, for each ψ^k, ψ^m , $\psi^k \cap \{\langle w, i \rangle : w = w'\} = \psi^m \cap \{\langle w, i \rangle : w = w'\} = \emptyset$.

For (II), let $X_w = Pr(\langle w, i \rangle)$, for arbitrary i . Note that this is only well-defined given that $Pr(\cdot)$ is center-indifferent. Next note that, given (A) and (B), we have that, for every ψ^k ,

$w(\psi^k) = \{\langle w, i \rangle : i \in f(\psi^k)\}$. Given this, we have:

$$Pr(\{\langle w, i \rangle : w = w'\} | \psi^k) = \frac{|f(\psi^k)| \times X_{w'}}{\sum_{w \in \mathcal{W}} |f(\psi^k)| \times X_w} = \frac{X_{w'}}{\sum_{w \in \mathcal{W}} X_w}$$

And we have:

$$Pr(\{\langle w, i \rangle : w = w'\} | \psi^m) = \frac{|f(\psi^m)| \times X_{w'}}{\sum_{w \in \mathcal{W}} |f(\psi^m)| \times X_w} = \frac{X_{w'}}{\sum_{w \in \mathcal{W}} X_w}$$

Finally, note that each uncentered proposition is a union of world-propositions. Thus given that $Pr(\xi | \psi^m) = Pr(\xi | \psi^k)$ for each world-proposition ξ , it follows that $Pr(\phi | \psi^m) = Pr(\phi | \psi^k)$, for each uncentered proposition ϕ . \square

A.3. Second Centered Agreement Theorem.

CLAIM: Let \mathcal{B} be a basic centered agreement frame. Then there exists some probability function $Pr(\cdot) \in \mathcal{M}_{\mathcal{C}}$, such that for any $\vec{q} \in \mathbf{C}$, if every $i \in G \subseteq \mathbf{A}$ is such that the following propositions are true at \vec{q} : (d) $\text{Com}_G([K_i = K_i(\vec{q})])$ and (e) $\text{Com}_G(\text{CHom}[K_i(\cdot)])$, then for each uncentered proposition ψ and each $i, j \in G$, $Pr(\psi | K_i(\vec{q})) = Pr(\psi | K_j(\vec{q}))$.

Proof. Let \mathbf{M} be the set of $\phi \subseteq \mathbf{C}$ such that (i) there is some $i \in \mathbf{A}$ such that, for every $\vec{z} \in \phi$, $K_i(\vec{z}) = \phi$, and (ii) for every $\vec{z} \in \phi$, $K_{a_z}(\vec{z}) = \phi$. For any $\phi \subseteq \mathbf{C}$, let $\mathbf{w}(\phi) = \{w \in \mathbf{W} : \langle w, i \rangle \in \phi, \text{ for some } i \in \mathbf{A}\}$, and let $\delta(\phi) = \{i \in \mathbf{A} : K_i(\vec{q}) = \phi, \text{ for every } \vec{q} \in \phi\}$. Note that for each $\phi, \psi \in \mathbf{M}$, such that $\phi \neq \psi$, $\phi \cap \psi = \emptyset$.

We first show that there is a natural way of embedding any $\phi \in \mathbf{M}$ in a maximal set of sets \mathcal{D} such that: (I) $\phi \in \mathcal{D}$, (II) for every $\phi_n \in \mathcal{D}$, if $\vec{q} \in \phi_n$, then $K_{a_q}(\vec{q}) = \phi_n$, (III) for every $\phi_m, \phi_n \in \mathcal{D}$, $\mathbf{w}(\phi_m) = \mathbf{w}(\phi_n)$, (IV) $\delta(\phi_m) \neq \emptyset$, for every $\phi_m \in \mathcal{D}$. To this end, let $I = \{i \in \mathbf{A} : K_i(\vec{x}) = K_i(\vec{y}) \text{ for every } \vec{x}, \vec{y} \in \phi\}$. We let $\Gamma_\phi^0 = \{\zeta \subseteq \mathbf{C} : \text{for some } i \in I \text{ and some } \vec{x} \in \phi, K_i(\vec{x}) = \zeta\}$. We let $\Gamma_\phi^1 = \Gamma_\phi^0 - \{\zeta \in \Gamma_\phi^0 : \mathbf{w}(\zeta) \neq \mathbf{w}(\phi)\}$. Finally, we let $\Gamma_\phi = \Gamma_\phi^1 - \{\zeta \in \Gamma_\phi^1 : \text{for some } \vec{q} \in \zeta, K_{a_q}(\vec{q}) \neq \zeta\}$. It is easy to see that Γ_ϕ satisfies conditions (I) - (IV). Moreover, given this construction, Γ_ϕ is maximal in the sense that any set satisfying conditions (I) - (IV) is a subset of Γ_ϕ . It is worth noting that each member of Γ_ϕ is itself an element of \mathbf{M} .

Let $g : \mathbf{M} \rightarrow \mathcal{P}(\mathbf{M})$ be such that, for any $\phi \in \mathbf{M}$, $g(\phi) = \Gamma_\phi$. We can show that the following hold:

- (1) For every $\zeta \in \mathbf{M}$, $\zeta \in g(\zeta)$.
- (2) For every $\zeta, \gamma \in \mathbf{M}$, if $\zeta \in g(\gamma)$, then $g(\zeta) = g(\gamma)$.

Proof of (1). This is immediate from the characterization of $g(\cdot)$. \square

Proof of (2). Let $\zeta \in g(\gamma)$. We first show that $g(\gamma) \subseteq g(\zeta)$. Let $\psi \in g(\gamma)$. So ψ is a set such that (i) for every $\vec{x} \in \psi$, $K_{a_x}(\vec{x}) = \psi$, (ii) $\mathbf{w}(\psi) = \mathbf{w}(\gamma)$, and (iii) there is some $i \in \mathbf{A}$, such that for every $\vec{x} \in \gamma$, $K_i(\vec{x}) = \psi$. To show that $\psi \in g(\zeta)$ we need to show that (i') for every $\vec{x} \in \psi$, $K_{a_x}(\vec{x}) = \psi$, (ii') $\mathbf{w}(\psi) = \mathbf{w}(\zeta)$, and (iii') there is some $i \in \mathbf{A}$, such that for every $\vec{x} \in \zeta$, $K_i(\vec{x}) = \psi$. Clearly, (i) entails (i'). To see that (ii) entails (ii'), note that, given that $\zeta \in g(\gamma)$, we have $\mathbf{w}(\zeta) = \mathbf{w}(\gamma)$. To see that (iii) entails (iii'), note that, for any $i \in \mathbf{A}$ and any $\vec{x} \in \mathbf{C}$, $K_i(\cdot)$ is only sensitive to the world parameter of \vec{x} . Thus, since $\mathbf{w}(\zeta) = \mathbf{w}(\gamma)$, (iii) entails (iii').

Next, we show, mutatis mutandis, that $g(\zeta) \subseteq g(\gamma)$. Let $\psi \in g(\zeta)$. So ψ is a set such that (i') for every $\vec{x} \in \psi$, $K_{a_x}(\vec{x}) = \psi$, (ii') $\mathbf{w}(\psi) = \mathbf{w}(\zeta)$, and (iii') there is some $i \in \mathbf{A}$, such that for every $\vec{x} \in \zeta$, $K_i(\vec{x}) = \psi$. To show that $\psi \in g(\gamma)$ we need to show that (i) for every

$\vec{x} \in \psi$, $K_{a_x}(\vec{x}) = \psi$, (ii) $\mathbf{w}(\psi) = \mathbf{w}(\gamma)$, and (iii) there is some $i \in \mathbf{A}$, such that for every $\vec{x} \in \gamma$, $K_i(\vec{x}) = \psi$. Clearly (i') entails (i). And, given that $\mathbf{w}(\xi) = \mathbf{w}(\gamma)$, we also have that (ii') entails (ii) and (iii') entails (iii). \square

Let $(\phi \equiv \psi) =_{df} \psi \in g(\phi)$. Then, given the above, we can see that \equiv is an equivalence relation on \mathbf{M} . First, note that, given (1), we have $\phi \in g(\phi)$ and so we have $\phi \equiv \phi$. Next note that, given (1) and (2), if $\phi \in g(\psi)$, then $\psi \in g(\phi)$. And so we have that if $\psi \equiv \phi$, then $\phi \equiv \psi$. Finally, note that, given (2), we have that if $\phi \in g(\psi)$ and $\xi \in g(\phi)$, then $\xi \in g(\psi)$. To see this, note that, given that $\phi \in g(\psi)$, it follows from (2) that $g(\phi) = g(\psi)$, and so, given that $\xi \in g(\phi)$, we have that $\xi \in g(\psi)$. Thus, if $\psi \equiv \phi$ and $\phi \equiv \xi$, then $\psi \equiv \xi$. We can think of $g(\cdot)$, then, as partitioning \mathbf{M} into equivalence classes under the relation \equiv . We'll call the cells of this partition *g-classes*. We now show that the following holds:

(A) If, for some $\vec{q} \in \mathbf{C}$, every $i \in G \subseteq \mathbf{A}$ is such that the following propositions are true at \vec{q} : (d) $\text{Com}_G([K_i = K_i(\vec{q})])$ and (e) $\text{Com}_G(\text{CHom}[K_i(\cdot)])$, then there is some *g-class* Δ , such that for each $i \in G$, $K_i(\vec{q}) \in \Delta$.

Proof of (A). To show that there is some *g-class* Δ , such that for each $i \in G$, $K_i(\vec{q}) \in \Delta$, it suffices to show that each $i, j \in G$ is such that (i) $K_i(\vec{q}), K_j(\vec{q}) \in \mathcal{P}(\mathbf{C})$, (ii) for every $\vec{x} \in K_j(\vec{q})$, $K_i(\vec{x}) = K_i(\vec{q})$, (iii) for every $\vec{x} \in K_i(\vec{q})$, $K_{a_x}(\vec{x}) = K_i(\vec{q})$, and (iv) $\mathbf{w}(K_i(\vec{q})) = \mathbf{w}(K_j(\vec{q}))$.

(i) follows from the assumption that, for each $i \in G$, (d) $\text{Com}_G([K_i = K_i(\vec{q})])$ is true at \vec{q} . For this entails that we have $K_i(\vec{q}) \subseteq [K_i = K_i(\vec{q})]$ and $K_j(\vec{q}) \subseteq [K_i = K_i(\vec{q})]$. And these will hold only if $K_i(\vec{q}), K_j(\vec{q}) \in \mathcal{P}(\mathbf{C})$.

(ii) also follows from the assumption that, for each $i \in G$, (d) $\text{Com}_G([K_i = K_i(\vec{q})])$ is true at \vec{q} . For, if at \vec{q} it is common knowledge amongst G that $[K_i = K_i(\vec{q})]$, then, for each $j \in G$ and each $\vec{x} \in K_j(\vec{q})$, $K_i(\vec{x}) = K_i(\vec{q})$.

(iii) follows from the assumption that, for each $i \in G$, (d) $\text{Com}_G([K_i = K_i(\vec{q})])$, and (e) $\text{Com}_G(\text{CHom}[K_i(\cdot)])$ are true at \vec{q} . For, given that it is common knowledge amongst G at \vec{q} that $[K_i = K_i(\vec{q})]$, then, for each $j \in G$ and each $\vec{x} \in K_j(\vec{q})$, $K_i(\vec{x}) = K_i(\vec{q})$. And, given that it is common knowledge amongst G at \vec{q} that i is first-personally introspective, we have that, for each $j \in G$ and each $\vec{x} \in K_j(\vec{q})$, $K_i(\vec{x})$ is such that, for each $\vec{z} \in K_i(\vec{x})$, $K_{a_z}(\vec{z}) = K_i(\vec{x})$. But it follows from these that, for each $i \in G$, every $\vec{x} \in K_i(\vec{q})$, $K_{a_x}(\vec{x}) = K_i(\vec{q})$.

Finally, (iv) follows from the assumption that, for each $i \in G$, (d) $\text{Com}_G([K_i = K_i(\vec{q})])$ is true at \vec{q} . For given that it's common knowledge amongst G that $[K_i = K_i(\vec{q})]$, then for every $\vec{x} \in K_j(\vec{q})$, $K_i(\vec{x}) = K_i(\vec{q}) \in \mathcal{P}(\mathbf{C})$. But since, at each $\vec{x} \in K_j(\vec{q})$, $K_i(\vec{x}) \in \mathcal{P}(\mathbf{C})$, we have that $K_i(\cdot)$ is center-reflexive at \vec{x} . And so we have that for any $\vec{x} \in K_j(\vec{q})$, $\langle w_x, a_i \rangle \in K_i(\vec{x})$, and so, since $K_i(\vec{x}) = K_i(\vec{q})$, $\langle w_x, a_i \rangle \in K_i(\vec{q})$. Thus, we have that, for any $i, j \in G$, if $\vec{x} \in K_j(\vec{q})$, then $\langle w_x, a_i \rangle \in K_i(\vec{q})$, and thus, for each $i, j \in G$, $\mathbf{w}(K_i(\vec{q})) = \mathbf{w}(K_j(\vec{q}))$. \square

Having established this, we now show:

(B) There is some probability function $Pr(\cdot)$ such that for any *g-class* Δ , any $\phi, \xi \in \Delta$, and any uncentered proposition ψ , $Pr(\psi|\phi) = Pr(\psi|\xi)$.

Proof of (B). First, note that for any *g-class* Δ , there is some $Pr(\cdot)$ such that, for any $\phi, \xi \in \Delta$, and any uncentered proposition ψ , $Pr(\psi|\phi) = Pr(\psi|\xi)$. To see this, note (i)

that for any $\phi, \xi \in \Delta$, $\mathbf{w}(\phi) = \mathbf{w}(\xi)$, and (ii) for any $\phi, \xi \in \Delta$, if $\phi \neq \xi$, then $\phi \cap \xi = \emptyset$. (i) follows from the characterization of \mathbf{M} . (ii) follows from the more general fact that for every $\gamma_1, \gamma_2 \in \mathbf{M}$, if $\gamma_1 \cap \gamma_2 \neq \emptyset$, then $\gamma_1 = \gamma_2$. For suppose that $\gamma_1 \cap \gamma_2 \neq \emptyset$. We have, then, that there is some \vec{x} such that $\vec{x} \in \gamma_1$ and $\vec{x} \in \gamma_2$. But since both $\gamma_1, \gamma_2 \in \mathbf{M}$, we have that $K_{a_x}(\vec{x}) = \gamma_1$ and $K_{a_x}(\vec{x}) = \gamma_2$, and so $\gamma_1 = \gamma_2$. Given (i) and (ii), we can let $Pr(\cdot)$ be such that, for any $w \in \mathbf{W}$, $Pr(\{\langle w', a \rangle \in \phi : w = w'\}) = Pr(\{\langle w', a \rangle \in \xi : w = w'\})$. For, given (i), if $w \notin \mathbf{w}(\phi)$, then $\{\langle w', a \rangle \in \phi : w = w'\} = \{\langle w', a \rangle \in \xi : w = w'\} = \emptyset$, and so for any $Pr(\cdot)$, $Pr(\{\langle w', a \rangle \in \phi : w = w'\}) = Pr(\{\langle w', a \rangle \in \xi : w = w'\}) = 0$. And, given (ii), if $w \in \mathbf{w}(\phi)$, then $\{\langle w', a \rangle \in \phi : w = w'\} \cap \{\langle w', a \rangle \in \xi : w = w'\} = \emptyset$, and so we can set $Pr(\cdot)$ such that $Pr(\{\langle w', a \rangle \in \phi : w = w'\}) = Pr(\{\langle w', a \rangle \in \xi : w = w'\})$. But, given that we have $Pr(\{\langle w', a \rangle \in \phi : w = w'\}) = Pr(\{\langle w', a \rangle \in \xi : w = w'\})$, it follows that for any uncentered proposition ψ , $Pr(\psi|\phi) = Pr(\psi|\xi)$. For the fact that $Pr(\{\langle w', a \rangle \in \phi : w = w'\}) = Pr(\{\langle w', a \rangle \in \xi : w = w'\})$ assures us that for each world-proposition $\omega = \{\langle w', a \rangle : w = w'\}$, $Pr(\omega|\phi) = Pr(\omega|\xi)$. And, given that two probability functions agree on all world-propositions, they will agree on every uncentered proposition, for these are simply unions of world propositions.

Given the nature of g-classes, though, it follows from the fact that for any g-class Δ , there is some $Pr(\cdot)$ such that, for any $\phi, \xi \in \Delta$, and any uncentered proposition ψ , $Pr(\psi|\phi) = Pr(\psi|\xi)$, that there is some probability function $Pr(\cdot)$ such that, for any g-class Δ , any $\phi, \xi \in \Delta$, and any uncentered proposition ψ , $Pr(\psi|\phi) = Pr(\psi|\xi)$. For, given a g-class Δ , let Pr_Δ be an arbitrary function such that for any $\phi, \xi \in \Delta$, and any uncentered proposition ψ , $Pr_\Delta(\psi|\phi) = Pr_\Delta(\psi|\xi)$. Given that the members of distinct g-classes don't overlap, it follows that there is some probability function $Pr(\cdot)$ such that, for each g-class Δ , and each $\phi \in \Delta$, $Pr(\cdot|\phi) = Pr_\Delta(\cdot|\phi)$. And so, there is some probability function $Pr(\cdot)$ such that, for any g-class Δ , and $\phi, \xi \in \Delta$, and any uncentered proposition ψ , $Pr(\psi|\phi) = Pr(\psi|\xi)$. \square

(A) and (B), together, entail our desired result. For (A) tells us that if, for some \vec{q} , every $i \in G \subseteq \mathbf{A}$ is such that (d) $\text{Com}_G([K_i = K_i(\vec{q})])$ and (e) $\text{Com}_G(\text{CHom}[K_i(\cdot)])$ are true at \vec{q} , then there is some g-class Δ , such that for each $i \in G$, $K_i(\vec{q}) \in \Delta$. And (B) tells us that there is some probability function $Pr(\cdot)$ such that for any g-class Δ , any $\phi, \xi \in \Delta$, and any uncentered proposition ψ , $Pr(\psi|\phi) = Pr(\psi|\xi)$. It follows, then, that there is some probability function $Pr(\cdot) \in \mathcal{M}_C$, such that for any $\vec{q} \in \mathbf{C}$, if every $i \in G \subseteq \mathbf{A}$ is such that (d) $\text{Com}_G([K_i = K_i(\vec{q})])$ and (e) $\text{Com}_G(\text{CHom}[K_i(\cdot)])$ are true at \vec{q} , then for each uncentered proposition ψ and each $i, j \in G$, $Pr(\psi|K_i(\vec{q})) = Pr(\psi|K_j(\vec{q}))$. \square

A.4. Third Centered Agreement Theorem.

Claim: Let \mathcal{B} be a basic centered agreement frame, and let $Pr(\cdot)$ be a center-indifferent probability function defined on $\mathcal{P}(\mathbf{C})$. If, at some $\vec{q} \in \mathbf{C}$, every $i \in G \subseteq \mathbf{A}$ is such that the following propositions are true at \vec{q} : (d) $\text{Com}_G([K_i = K_i(\vec{q})])$, (e) $\text{Com}_G(\text{CHom}[K_i(\cdot)])$ and (f) $K_i(\{\vec{z} : \text{Com}_G([K_{a_z} = K_i(\vec{q})])\})$, then for every uncentered proposition ψ , and each $i, j \in G$, $Pr(\psi|K_i(\vec{q})) = Pr(\psi|K_j(\vec{q}))$.

Proof. The proof relies on the following two claims:

(A) Given that, at \vec{q} , (d)-(f) hold, for all $i \in G \subseteq \mathbf{A}$, it follows that, for each $i, j \in G$, $\{w_z : \vec{z} \in K_i(\vec{q})\} = \{w_z : \vec{z} \in K_j(\vec{q})\}$. We'll denote this set $W_{\vec{q}}$.

(B) Given that, at \vec{q} , (d)-(f) hold, for all $i \in G \subseteq \mathbf{A}$, it follows that, for each $i \in G$, $K_i(\vec{q}) = \{a_z : \vec{z} \in K_i(\vec{q})\} \times W_{\vec{q}}$.

To see that (A) and (B) entail our result, it suffices to see that, given (A) and (B), if $Pr(\cdot)$ is center-indifferent and ζ is a world proposition, i.e., a proposition of the form $\{\langle w, i \rangle : w = w'\}$, for some w' , then $Pr(\zeta|K_i(\vec{q})) = Pr(\zeta|K_j(\vec{q}))$.

There are two cases to consider:

- (I) $\zeta = \{\langle w, i \rangle : w = w'\}$, for some $w' \notin W_{\vec{q}}$.
- (II) $\zeta = \{\langle w, i \rangle : w = w'\}$, for some $w' \in W_{\vec{q}}$.

For (I), we have that $Pr(\zeta|K_i(\vec{q})) = Pr(\zeta|K_j(\vec{q})) = 0$. This follows from the fact that, given (A) and (B), if $\zeta = \{\langle w, i \rangle : w = w'\}$, for some $w' \notin W_{\vec{q}}$, then we have $\zeta \cap K_i(\vec{q}) = \zeta \cap K_j(\vec{q}) = \emptyset$.

For (II), let $X_w = Pr(\langle w, i \rangle)$, for arbitrary i . Note that this is only well-defined given that $Pr(\cdot)$ is center-indifferent. Given (A) and (B), we have that:

$$Pr(\{\langle w, i \rangle : w = w'\}|K_i(\vec{q})) = \frac{|\{a_z : \vec{z} \in K_i(\vec{q})\}| \times X_{w'}}{\sum_{w \in W_{\vec{q}}} |\{a_z : \vec{z} \in K_i(\vec{q})\}| \times X_w} = \frac{X_{w'}}{\sum_{w \in W_{\vec{q}}} X_w}$$

And we have:

$$Pr(\{\langle w, i \rangle : w = w'\}|K_j(\vec{q})) = \frac{|\{a_z : \vec{z} \in K_j(\vec{q})\}| \times X_{w'}}{\sum_{w \in W_{\vec{q}}} |\{a_z : \vec{z} \in K_i(\vec{q})\}| \times X_w} = \frac{X_{w'}}{\sum_{w \in W_{\vec{q}}} X_w}$$

Proof of (A). This follows from the result, proved in §A.3, that if (d) and (e) hold at \vec{q} for each $i \in G \subseteq \mathbf{A}$, then there is some g-class Δ , such that, for each $i \in G$, $K_i(\vec{q}) \in \Delta$. For, as we showed, for each g-class Δ , and each $\phi, \psi \in \Delta$, $\{w : \langle w, a \rangle \in \phi \text{ for some } a \in \mathbf{A}\} = \{w : \langle w, a \rangle \in \psi \text{ for some } a \in \mathbf{A}\}$. \square

Proof of (B). Let j be an arbitrary member of G . And let $\vec{r} \in K_j(\vec{q})$. We'll show that, for each $w \in W_{\vec{q}}$, $\langle w, a_r \rangle \in K_j(\vec{q})$. Given that, at \vec{q} , (f) holds for each $i \in G$, we have $\vec{r} \in \{\vec{z} : \text{Com}_G([K_{a_z} = K_j(\vec{q})])\}$. And, given that, at \vec{q} , (d) holds for each $i \in G$, it follows from the fact that $\vec{r} \in \{\vec{z} : \text{Com}_G([K_{a_z} = K_j(\vec{q})])\}$, together with the fact that $K_{a_r}(\cdot)$ is only sensitive to the world parameter, that, for every \vec{x} such that $w_x \in W_{\vec{q}}$, $K_{a_r}(\vec{x}) = K_j(\vec{q})$. But we have that, for each \vec{x} such that $w_x \in W_{\vec{q}}$, $\langle w_x, a_r \rangle \in K_{a_r}(\vec{x})$, given that $K_{a_r}(\cdot)$ is center-reflexive at each $\vec{v} \in \mathbf{C}$. Thus, for each \vec{x} such that $w_x \in W_{\vec{q}}$, we have $\langle w_x, a_r \rangle \in K_j(\vec{q})$. And so, for each $w \in W_{\vec{q}}$, $\langle w, a_r \rangle \in K_j(\vec{q})$. \square

BIBLIOGRAPHY

- Frank Arntzenius. (2003). Some problems for conditionalization and reflection. *Journal of Philosophy*, **100**(7), 356–370.
- Robert Aumann. (1976). Agreeing to disagree. *The Annals of Statistics*, **4**(6), 1236–1239.
- Michael Bacharach. (1985). Some extensions of a claim of aumann in an axiomatic model of knowledge. *Journal of Economic Theory*, **37**(1), 167–190.
- Jon Barwise. (1988). Three views of common knowledge. In *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge*. Morgan Kaufmann San Francisco, California, pp. 365–379.
- Nick Bostrom. (2000) Observer-relative chances in anthropic reasoning? *Erkenntnis*, **52**, 93–108.
- Nick Bostrom. (2002). *Anthropic Bias*. Routledge, New York.
- Michael Caie. (2016). Agreement and updating for self-locating belief. (ms).

- Cédric Degrémont, & Oliver Roy. (2012). Agreement theorems in dynamic-epistemic logic. *Journal of Philosophical Logic*, **41**(4), 735–764.
- Lorenz Demey. (2014). Agreeing to disagree in probabilistic dynamic epistemic logic. *Synthese*, **191**(3), 409–438.
- Adam Elga. (2000). Self-locating belief and the sleeping beauty problem. *Analysis*, **60**(266), 143–147.
- Joseph Halpern, Yoram Moses, & Moshe Vardi. (1995). *Reasoning About Knowledge*. MIT Press Cambridge, Massachusetts.
- Harvey Lederman. (2015). People with common priors can agree to disagree, *this Review*, **8**(1), 11–45.
- John Leslie. (1997). Observer-relative chances and the doomsday argument. *Inquiry*, **40**, 427–436.
- David Lewis. (1969). *Convention*. Blackwell Publishing, Oxford.
- David Lewis. (1979). Attitudes de dicto and de se. *The Philosophical Review*, **88**(4), 513–543.
- David Lewis. (2001). Sleeping beauty: Reply to Elga. *Analysis*, **61**(271), 171–176.
- David Lewis. (2004). How many lives has Schrödinger’s cat? *Australasian Journal of Philosophy*, **82**(1), 3–22.
- Christopher Meacham. (2008). Sleeping beauty and the dynamics of de se beliefs. *Philosophical Studies*, **138**(2), 245–269.
- Ariel Rubinstein, & Asher Wolinsky. (1990). On the logic of ‘agreeing to disagree’ type results. *Journal of Economic Theory*, **1**, 184–193.

UNIVERSITY OF PITTSBURGH
DEPARTMENT OF PHILOSOPHY
1001 CATHEDRAL OF LEARNING
PITTSBURGH, PA, 15260
E-mail: caiemike@gmail.com