Agreement Theorems in Dynamic-Epistemic Logic
Extended Abstract

Cédric Dégremont
Institute for Logic, Language and Computation
Universiteit van Amsterdam
The Netherlands
cedric.uva@gmail.com

Olivier Roy
Faculty of Philosophy
University of Groningen
The Netherlands
o.roy@rug.nl

Abstract

In this paper we study Aumann's Agreement Theorem in dynamic-epistemic logic. We show that common *belief* of posteriors is sufficient for agreements in "epistemic-plausibility models", under common and well-founded priors, from which the usual form of agreement results follows, using common knowledge. We do not restrict ourselves to the finite case, and show that in countable structures such results hold if *and only if* the underlying "plausibility ordering" is well-founded. We look at these results from a syntactic point of view, showing that neither well-foundedness nor common priors are expressible in a commonly used language, but that the static agreement result is finitely derivable in an extended modal logic. We finally consider "dynamic" agreement results, show they have a counterpart in epistemic-plausibility models, and provide a new form of agreements via "public announcements". Comparison of the two types of dynamic agreement reveals that they can indeed be different.

1 Introduction

In this paper we study Aumann's Agreement Theorem [1] and some of its subsequent extensions [16] and generalizations [13, 2] in dynamic-epistemic logic [5, 14]. We show that common *belief* of posteriors is sufficient for agreements in "epistemic-plausibility models", under common and well-founded priors, from which the usual form of agreement results follows, using common knowledge. We do not restrict ourselves to the finite case, which thus represents an improvement on known qualitative agreement theorems [2], and show that in countable structures such results hold if *and only if* the underlying "plausibility ordering" is well-founded. We then look at these results from a syntactic point of view, showing that neither well-foundedness nor common priors are expressible in the language proposed in [3], even if it is extended with a common belief operator, but we also show a finitary syntactic derivation of the static agreement result in an extended modal language. We finally consider "dynamic" agreement results. We show that "agreements via dialogues" [13, 2] have a counterpart in epistemic-plausibility models, and that one also gets agreements via "public announcements", a type of belief update that has so far not been considered in the agreement literature—see [10] and [19]. Comparison of the two types of dynamic agreements reveals that in some situations they are indeed different.

These technical results answer an "internal" question in dynamic-epistemic logic, namely whether agreement results hold in this framework, but they also offer new insights into the contribution of agreement theorems to interactive epistemology. That common belief of posteriors is sufficient for agreements, under common and well-founded priors, strengthens one of the key lessons of agreement theorems, viz. that first-order information is closely dependent on higher-order information in situations of interaction [10]. Our inexpressibility results, on the other hand, support a qualm already voiced in the literature concerning the difficulty for agents to reason about static agreements [21]. The two dynamic results not only make a sharp distinction between two forms of belief changes, they also allow one to capture more adequately the idea that agreements are reached via *public* dialogues. Bringing agreement theorems to dynamic-epistemic logic is thus important both technically and conceptually, and it helps to bridge the existing literature on agreements with the logical approaches to knowledge, beliefs and the dynamics of information.

In this extended abstract all the proofs are omitted, as well as some auxiliary definitions. The reader interested

TARK '09, July 6-8, 2009, California
ISBN: 978-1-60558-560-4...$10.00

in these details can communicate with the authors.

2 Definitions

In this section we introduce the models in which we study the various agreement results, and the logical language used in [3] to describe them.

2.1 Epistemic Plausibility Models

An epistemic plausibility model [5] is a qualitative representation of the agents' beliefs as well as first- and higher-order information in a given interactive situation.

Definition 2.1. [Epistemic Plausibility Model] Given a countable set of atomic propositions PROP, an *epistemic plausibility model* $\mathcal{M} = \langle W, (\leq_i)_{i\in I}, (\sim_i)_{i\in I}, V\rangle$ has $W \neq \emptyset$ and countable, $I = \{1, 2, \dots, n\}$ is a finite set of *agents*, and for each $i \in I$, $\leq_i$ is a total (plausibility) pre-order on W, $\sim_i$ is a binary equivalence relation on W, and $V : \text{PROP} \to \wp(W)$. An epistemic plausibility *frame* $\mathcal{F}$ is an epistemic plausibility model with the valuation V omitted. ⋊

The total plausibility pre-order $\leq_i$ induces i's *priors*, and can be viewed as a qualitative counterpart to a prior probability distribution on W. If $w \leq_i w'$ we say that i considers w' at least as plausible as w. Given a set $X \subseteq W$, we say that $w \in X$ is $\leq_i$-minimal in X if $w \leq_i w'$ for all $w' \in X$. The relation $\sim_i$ induces i's *information partition* W. We write $\mathcal{K}_i[w]$ to denote the cell of this partition $\{v \in W \mid w \sim_i v\}$ to which w belongs. $\mathcal{K}_i[w]$ should be regarded as i's (private) information at w. We write $|\mathcal{M}| = W$ for the domain of $\mathcal{M}$.

The next two assumptions are crucial in the following.

Definition 2.2. [(Local) well-foundedness] A plausibility pre-order satisfies:

- **Local well-foundedness**. If for all $w \in W$ and $i \in I$, for all $X \subseteq \mathcal{K}_i[w]$, X has $\leq_i$-minimal elements.
- **Well-foundedness**. If for all $X \subseteq W$ and $i \in I$, X has $\leq_i$-minimal elements.

$\mathcal{M}$ satisfies (*Local*) *Well-foundedness* if every plausibility pre-order has the corresponding property. ⋊

Definition 2.3. [(A priori/ a posteriori) Most plausible elements]

- For all $X \subseteq W$, let $\beta_i(X) = min_{\leq_i}(X) = \{w : w \text{ is } \leq_i \text{-minimal in } X\}$.
- For all $w \in W$, let $\mathcal{B}_i[w] = \beta_i(\mathcal{K}_i[w])$.

We write $w \rhd_i^{\mathcal{B}} v$ iff $v \in \mathcal{B}_i[w]$, and $w \to_i^X v$ iff $v \in \beta_i(\mathcal{K}_i[w] \cap X)$. ⋊

Intuitively $\beta_i(X)$ are the *a priori* most plausible elements of a set, ignoring the information partitions. $\mathcal{B}_i[w]$ gives the states i considers most plausible, conditional on the information he possesses at w, i.e. conditional on $\mathcal{K}_i[w]$. The relation $w \to_i^X v$ maps w to all states i considers most plausible, conditional on the information he possesses at w *and* on a given subset X. Observe that the set $\{v : w \to_i^X v\}$ might be empty for a given w and a given X, if $X \cap \mathcal{K}_i[w] = \emptyset$ or, in words, if X is already excluded by i's information at w.

Observe that β_i is well-defined if the plausibility preorder is well-founded, while local well-foundedness is sufficient for $\mathcal{B}_i$ to be well-defined. To draw an analogy with the probabilistic case, this means that local well-foundedness ensures that the conditional beliefs of an agent i are well-defined for all "events" that have a non-empty intersection with the agent's information partition. Well-foundedness, on the other hand, requires i's conditional beliefs to be well-defined for any non-empty subsets of W.

Definition 2.4. [Common Prior] There is *common prior beliefs* among group G in an epistemic plausibility model $\mathcal{M}$ when $\leq_i = \leq_j$ for all $i, j \in G$. ⋊

The reflexive-transitive closure of the union of the epistemic accessibility relations $\sim_i$ for all agents i in a group G is the model-theoretic counterpart of the notion of "common knowledge" in G [15, 14]. We define "common belief" analogously.

Definition 2.5. [Common knowledge] For each $G \subseteq I$, let $\sim_G^*$ be the reflexive-transitive closure of $\bigcup_{i\in G} \sim_i$. Let $[w]_G^* = \{w' \in W \mid w \sim_G^* w'\}$. ⋊

Definition 2.6. [Common belief] For each $G \subseteq I$, let $\rhd_G^*$ be the reflexive-transitive closure of $\bigcup_{i\in G} \rhd_i^{\mathcal{B}}$. ⋊

2.2 Doxastic-Epistemic Logic

The logical language used in [3] to describe epistemic-plausibility models is a propositional modal language with three families of modal operators, which we extend here with "common belief" operators.

Definition 2.7. [Epistemic Doxastic Language] The language $\mathcal{L}_{EDL}$ is defined as follows:

$$\phi := p \mid \neg\phi \mid \phi \wedge \phi \mid K_i\phi \mid B_i^{\phi}\phi \mid C_G\phi \mid CB_G\phi$$

where i ranges over N, p over a countable set of proposition letters PROP and $\emptyset \neq G \subseteq I$. ⋊

The propositional fragment of this language is standard, and we write $\bot$ for $p \wedge \neg p$ and $\top$ for $\neg\bot$. A formula $K_i\phi$

should be read as "i knows that ϕ", $C_G\phi$ as "it is common knowledge among group G that ϕ", $CB_G\phi$ as "it is common belief among group G that ϕ." The formula $B_i^{\phi}\psi$, should be read " conditional on ϕ, i believes that ψ." These formulas are interpreted in epistemic plausibility models as follows:

Definition 2.8. [Truth definition] We write $||\phi||^{\mathcal{M}}$ for $\{w \in |\mathcal{M}| \ : \ \mathcal{M}, w \Vdash \phi\}$. We omit $\mathcal{M}$ when it is clear from the context.

$\mathcal{M}, w \Vdash p$	iff	$w \in V(p)$
$\mathcal{M}, w \Vdash \neg\phi$	iff	$\mathcal{M}, w \nVdash \phi$
$\mathcal{M}, w \Vdash \phi \wedge \psi$	iff	$\mathcal{M}, w \Vdash \phi$ and $\mathcal{M}, w \Vdash \psi$
$\mathcal{M}, w \Vdash K_i\phi$	iff	$\forall v$ (if $w \sim_i v$ then $\mathcal{M}, v \Vdash \phi$)
$\mathcal{M}, w \Vdash B_i^{\psi}\phi$	iff	$\forall v$ (if $w \rightarrow_i^{\|\|\psi\|\|^{\mathcal{M}}} v$ then $\mathcal{M}, v \Vdash \phi$)
$\mathcal{M}, w \Vdash C_G\phi$	iff	$\forall v$ (if $w \sim_G^* v$ then $\mathcal{M}, v \Vdash \phi$)
$\mathcal{M}, w \Vdash CB_G\phi$	iff	$\forall v$ (if $w \rhd_G^* v$ then $\mathcal{M}, v \Vdash \phi$)

$\rtimes$

Simple belief conditional only on i's information at a state w can be defined using the conditional belief operator: $B_i\phi = B_i^{\top}\phi$, since: $\mathcal{M}, w \Vdash B_i^{\top}\phi$ iff $\forall v$ (if $w \rhd_i^{\mathcal{B}} v$ then $\mathcal{M}, v \Vdash \phi$).

3 Static Agreements and Well-foundedness

We first show that well-foundedness is sufficient for agreement on the posteriors under common priors and common *beliefs* of the posteriors. More precisely, we show that if an epistemic plausibility model is well-founded, then common belief that agent i believes that ϕ while j does not believe that ϕ implies that i and j have different priors, which is the contrapositive form of the agreement theorem.

Theorem 3.1 (Agreement theorem - Common Belief)**.** *If a* well-founded *epistemic plausibility model $\mathcal{M}$ satisfies $\mathcal{M}, w \Vdash CB_{\{i,j\}}(B_ip \wedge \neg B_jp)$ for some $w \in W$, then i and j have different priors in $\mathcal{M}$.*

This immediately implies the "common knowledge" agreement result below, because $C_G\phi \rightarrow CB_G\phi$ is a valid implication in epistemic plausibility models. Note, however, that this result can also have been shown independently, by application of Bacharach's [2] result on qualitative "decision functions", modulo generalization to the countable case.

Corollary 3.2 (Agreement theorem - Common Knowledge)**.** *If an epistemic plausibility model $\mathcal{M}$ satisfies* well-foundedness *and $\mathcal{M}, w \Vdash C_{\{i,j\}}(B_ip \wedge \neg B_jp)$ for one $w \in W$, then i and j have different priors in $\mathcal{M}$.*

Well-foundeness is not only sufficient for common priors to exclude the possibility of disagreements when the posterior are common beliefs, it is also necessary, as the Proposition 3.3 shows. The model behind this result is drawn in figure 1.

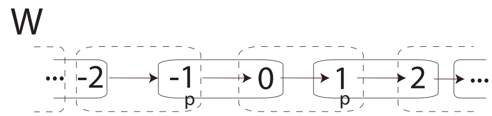

Figure 1: The epistemic plausibility model constructed in the proof of Proposition 3.3. The solid and dotted rectangles represent 1's and 2's information partitions on W, respectively. The arrows represent their common plausibility ordering.

Proposition 3.3. *There exists a pointed epistemic plausibility model $\mathcal{M}, w$ which satisfies* local well-foundedness *and* common prior *such that $\mathcal{M}, w \Vdash C_{\{1,2\}}(B_1p \wedge \neg B_2p)$.*

Well-foundedness is thus necessary for agreement results to hold, and furthermore cannot be weakened to *local well-foundedness*. This condition on the plausibility ordering is thus *the* safeguard against common knowledge of disagreement, once we drop the assumption that the state space is finite.

4 Expressive Power and Syntactic Proofs

$\mathcal{L}_{EDL}$ is a natural choice of language for talking about epistemic-plausibility models, and but we show here that it cannot express Theorem 3.1 nor Corollary 3.2, because it cannot express two of their key assumptions, common prior and well-foundedness.

Fact 4.1. *The class of epistemic plausibility frames that satisfies common prior is not definable in $\mathcal{L}_{EDL}$.*

Figure 2: The two epistemic plausibility model constructed in the proof of Fact 4.1. 1's and 2's information partitions on W are represented as in figure 1. The arrow in W represents their common plausibility ordering, while in W' the solid arrow and dotted arrows represent 1's and 2's orderings, respectively.

3

This result, which rests on the two small models drawn in figure 2, confirms the idea that to reason about (common) priors the agents must make "inter-[information]-state comparisons" [21], which they cannot do because their reasonings in $\mathcal{L}_{EDL}$ are local, i.e. they are bounded by the "hard information" [8] they have. This limitation also makes well-foundeness inexpressible, and with it the two static agreement results.

Fact 4.2. *There is no formula ϕ of $\mathcal{L}_{EDL}$ which is true in a pointed epistemic plausibility model $\mathcal{M}, w$ iff Theorem 3.1 or Corollary 3.2 holds in $\mathcal{M}, w$.*

The syntactical counterpart of the model-theoretic agreement results thus resides in more expressive languages. In the full version of the paper we present a finite syntactic derivation of Corollary 3.2 in $\mathcal{H}(@, \downarrow, C_G, \geq_j, \sim_j)$[1], which extends the hybrid language $\mathcal{H}(@, \downarrow, \geq_j, \sim_j)$ with a common knowledge modality C_G. Formally the language is the following:

$$\phi := \; p \mid i \mid x \mid \neg\phi \mid \phi \wedge \phi \mid K_j\phi \mid \langle \geq_j \rangle \phi \mid \\ @_i\phi \mid @_x\phi \mid \downarrow x.\phi \mid C_G\phi$$

Note that it allows one to scan the plausibility relation directly. The hybrid semantics draws on assignation functions that maps states variables to states and which allows the language to bind a variable to the current state and to refer to it. A detailed presentation of this language and its semantics, together with the syntactic derivation and proof of soundness of the axioms we are using is given in the full version of the paper.

On the positive side this language is able to axiomatize (converse) well-foundedness of the plausibility relation. On the negative side, the satisfiability problem for this language on the class of conversely well-founded frames is Σ^1_1-hard [12], ruling out any finite axiomatization of its validities. The derivation we show, however, is finite and uses only sound axioms. At the time of writing we still do not know whether the agreement results of Section 3 could be derived in a less complex language. The fact that the syntactic derivation reported here pertains to such an expressive language nevertheless shows that reasoning explicitly about agreement results requires onerous expressive resources.

5 Agreements via dialogues

In this section we turn to "agreements-via-dialogues" [16, 2], which analyze how agents can reach agreement in the process of exchanging information about their beliefs by updating the latter accordingly.

[1] In fact we use a language with more primitives, but, as we prove, these are entirely definable in the restricted language.

5.1 Agreements via Conditioning

We first consider agreements by repeated belief conditioning. It is known that if agents repeatedly exchange information about each others' posterior beliefs about a certain event, and update these posteriors accordingly, the posteriors will eventually converge [16, 2]. We show here that this result also holds for the "qualitative" form of beliefs conditionalization in epistemic plausibility models.

We call a *conditioning dialogue about* ϕ [16] at a state w of an epistemic plausibility model $\mathcal{M}$ a sequence of belief conditioning, for each agent, on all other agents' beliefs about ϕ. This sequence can be intuitively described as follows. It starts with the agents' simple belief about ϕ, i.e. for all i: $B_i\phi$ if $\mathcal{M}, w \Vdash B_i\phi$ and $\neg B_i\phi$ otherwise. Agent i's beliefs about ϕ at the next stage is defined by taking his beliefs about ϕ, conditional upon learning the others' belief about ϕ at that stage. Syntactically, this gives, $\mathbb{B}_{1,i} = B_i\phi$ if $\mathcal{M}, w \Vdash B_i\phi$ and $\mathbb{B}_{1,i} = \neg B_i\phi$ otherwise and, for two agents i, j, $\mathbb{B}_{n+1,i} = B_i^{\mathbb{B}_{n,j}\phi}\phi$ if $\mathcal{M}, w \Vdash B_i^{\mathbb{B}_{n,j}\phi}\phi$ and $\neg B_i^{\mathbb{B}_{n,j}\phi}\phi$ otherwise. This syntactic rendering is only intended to fix intuitions, though, since in countable models the limit of this sequence exceeds the finitary character of $\mathcal{L}_{EDL}$. We thus focus on model-theoretic conditioning.

Conditioning on a given event $A \subseteq W$ boils down to refining an agent's information partition by removing "epistemic links" connecting A and non-A states.

Definition 5.1. [Conditioning by a subset] Given an epistemic plausibility model $\mathcal{M}$, the collection of epistemic equivalence relation of the agents is an element of $\wp(W \times W)^I$. Given a group $G \subseteq I$, the function $f_G : \wp(W) \to (\wp(W \times W)^I \to \wp(W \times W)^I)$ is a *conditioning function for* G whenever:

$$(w, v) \in f_G(A)(i)(\{\sim_i\}_{i \in I}) =$$

$$\begin{cases} (w, v) \in \sim_i \text{ and } (w \in A \text{ iff } v \in A) & \text{if } i \in G \\ (w, v) \in \sim_i & \text{otherwise} \end{cases}$$

Given $\mathcal{M} = \langle W, (\leq_i)_{i \in I}, (\sim_i)_{i \in I}, V \rangle$ we write $f_G(A)(\mathcal{M})$ for the model $\langle W, (\leq_i)_{i \in I}, f_G(A)((\sim_i)_{i \in I}), V \rangle$.

⋊

It is easy to see that the relations $\sim_i$ in $f_G(A)(\mathcal{M})$ are equivalence relations. Here we are interested in cases where the agents condition their beliefs upon learning in which *belief state* the others are.

Definition 5.2. [Belief states] Let $\mathcal{M}$ an epistemic plausibility model and $A \subseteq W$, we write

$B_j^{\mathcal{M}}(A)$ for $\{w : \beta_j(\mathcal{K}_j^{\mathcal{M}}[w]) \subseteq A\}$ and

$\neg B_j^{\mathcal{M}}(A)$ for $W \setminus B_j^{\mathcal{M}}(A)$

We define $\mathbb{B}_j^{\mathcal{M},w}(A)$ as follows:

$$\mathbb{B}_j^{\mathcal{M},w}(A) = \begin{cases} B_j^{\mathcal{M}}(A) & \text{if } w \in B_j^{\mathcal{M}}(A) \\ \neg B_j^{\mathcal{M}}(A) & \text{otherwise} \end{cases}$$

$\rtimes$

Observation 5.3. *For any plausibility epistemic model $\mathcal{M}$ indexed by a finite set of agents I, $\langle \wp(W \times W)^I, \subseteq \rangle$ is a chain complete poset. Moreover for all $A \subseteq W$, $w \in W$ and $G \subseteq I$, $f_G(A)$ is deflationary.*

Taking $f_I(\bigcap_{j \in I} \mathbb{B}_j^{\mathcal{M},w}(||\phi||^{\mathcal{M}}))$ as a mapping on models, it is easy to see from the preceding observation that conditioning by agents' beliefs about some event is deflationary with respect to the relation of epistemic-submodel. It follows then by the Bourbaki-Witt fixed-point theorem [11] that conditioning by agents' beliefs has a fixed point.

Theorem 5.4 (Bourbaki-Witt [11])**.** *Let X be a chain complete poset. If $f : X \to X$ is inflationary (deflationary), then f has a fixed point.*

Given an initial pointed model $\mathcal{M}, w$ and some event $A \subseteq W$, we can construct its fixed point under conditioning by agents' beliefs as the limit of a sequence of models, which are the model-theoretic counterpart of the dialogues described above.

Definition 5.5. A *conditioning dialogue about ϕ* at the pointed plausibility epistemic model $\mathcal{M}, w$, with $\mathcal{M} = \langle W, (\leq_i)_{i \in I}, (\sim_i)_{i \in I}, V \rangle$ is the sequence of pointed epistemic plausibility models $(\mathcal{M}_n, w)$ with

$$(\mathcal{M}_0, w) = \mathcal{M}, w$$

$$(\mathcal{M}_{\beta+1}, w) = f_I(\bigcap_{j \in I} \mathbb{B}_j^{\mathcal{M}_\beta, w}(||\phi||^{\mathcal{M}}))(\mathcal{M}_\beta), w$$

$$(\mathcal{M}_\lambda, w) = \bigcap_{\beta < \lambda} (\mathcal{M}_\beta, w) \text{ for limit ordinals } \lambda$$

$\rtimes$

This extends to the countable case the standard representation of a dialogue about ϕ in the literature on dynamic agreements [16, 2]. By observation 5.3 we know that dialogues cannot last forever, i.e. that each such sequence has a limit.

Corollary 5.6. *For any pointed epistemic plausibility model $\mathcal{M}, w$ and $\phi \in \mathcal{L}_{EDL}$ there is a α^f such that, for all $i \in I$, $w \in W$ and $\alpha > \alpha^f$, $\mathcal{K}_{\alpha,i}[w] = \mathcal{K}_{\alpha^f,i}[w]$.*

Once the agents have reached this fixed-point α^f—possibly after transfinitely many steps—they have eliminated all higher-order uncertainties concerning the posteriors about ϕ of the others, viz. these posteriors are then common knowledge:

Theorem 5.7 (Common knowledge of beliefs about ϕ)**.** *At the fixed-point α^f of a conditioning dialogue about ϕ we have that for all $w \in W$ and $i \in I$, if $w \in B_i^{\mathcal{M}_{\alpha^f}, w}(||\phi||^{\mathcal{M}})$ then $w' \in B_i^{\mathcal{M}_{\alpha^f}, w}(||\phi||^{\mathcal{M}})$ for all $w' \in [w]^*_{\alpha^f, I}$, and similarly if $w \notin B_i^{\mathcal{M}_{\alpha^f}, w}(||\phi||^{\mathcal{M}})$.*

With this in hand we can directly apply the static agreement result for common knowledge (Corollary 3.2, Section 3) to find that the agents do indeed reach agreements at the fixed-point of a dialogue about ϕ.

Corollary 5.8 (Agreement via conditioning dialogue)**.** *Take any dialogue about ϕ with common and well-founded priors, and α^f as in Corollary 5.6. Then for all w in W, either $[w]^*_{\alpha^f, I} \subseteq \bigcap_{i \in I} B_i^{\mathcal{M}_{\alpha^f}, w}(||\phi||^{\mathcal{M}})$ or $[w]^*_{\alpha^f, I} \subseteq \bigcap_{i \in I} \neg B_i^{\mathcal{M}_{\alpha^f}, w}(||\phi||^{\mathcal{M}})$.*

This result brings qualitative dynamic agreement results [13, 2] to epistemic plausibility models, and show that agents can indeed reach agreement via iterated conditioning, even when the finite model assumption is dropped.

5.2 Agreements via Public Announcements

In this section we show that iterated "public announcements" lead to agreements, thus introducing a distinct form of information update to the agreement literature. Public announcements are "epistemic actions" [14] by which truthful, hard information is made public to the members of a group by a trusted source, in such a way that no member is in doubt about whether the others received the same piece of information as he did.

One extends a given logical language with public announcements by operators of the form $[\phi!]\psi$, meaning "after the announcement of ϕ, ψ holds" [20, 17]. A dialogue about ϕ via public announcements among the members of a group G thus starts, as before, with i simple beliefs about ϕ, for all $i \in I$. The agents' beliefs about ϕ at the next stage are then defined as those they would have after a public announcement of all agents' beliefs about ϕ at the first stage. Syntactically, this gives: $\mathbb{B}_{1,i}$ as in Section 5.1, and $\mathbb{B}_{n+1,i}$, as $[\bigcap_{j \in I} \mathbb{B}_{n,j}\phi!]B_i\phi$ if $\mathcal{M}, w \Vdash [\bigcap_{j \in I} \mathbb{B}_{n,j}\phi!]B_i\phi$ and as $[\bigcap_{j \in I} \mathbb{B}_{n,j}\phi!]\neg B_i\phi$ otherwise. For the same reason as in the previous section, we now move our analysis to the level of models.

The A-generated submodel of a given epistemic plausibility model $\mathcal{M}$ is the model that results after the public announcement of A in $\mathcal{M}$. We write $Sub(\mathcal{M}) = \{\mathcal{M}'$ is the A-generated submodel of $\mathcal{M} \mid A \subseteq |\mathcal{M}|\}$ and $\mathcal{M}' \sqsubseteq \mathcal{M}$ whenever $\mathcal{M}' \in Sub(\mathcal{M})$.

Definition 5.9. [Relativization by agents beliefs] Let $\mathbb{B}_i(\mathcal{M}, w, \phi)$ be defined as follows:

$$\mathbb{B}_i(\mathcal{M}, w, \phi) = \begin{cases} ||B_i\phi||^{\mathcal{M}} & \text{if } \mathcal{M}, w \Vdash B_i\phi \\ ||\neg B_i\phi||^{\mathcal{M}} & \text{otherwise} \end{cases}$$

Then given an epistemic-plausibility model $\mathcal{M} = \langle W, (\leq_i)_{i\in I}, (\sim_i)_{i\in I}, V\rangle$, the relativization $!B^{\phi}_{w}$ by agents' beliefs about ϕ at w (where $w \in |\mathcal{M}|$), takes $\mathcal{M}$ to $!B^{\phi}_{w}(\mathcal{M})$. Here $!B^{\phi}_{w}(\mathcal{M})$ is the $\bigcap_{i\in I} \mathbb{B}_i(\mathcal{M}, w, \phi)$-generated submodel $!B^{\phi}_{w}(\mathcal{M}) = \langle W^{!B^{\phi}_{w}}, \leq_i^{!B^{\phi}_{w}}, \sim_i^{!B^{\phi}_{w}}, V^{!B^{\phi}_{w}}\rangle$ of $\mathcal{M}$ such that:

- $W^{!B^{\phi}_{w}} = \bigcap_{i\in I} \mathbb{B}_i(\mathcal{M}, w, \phi)$

and for each $i \in I$

- $\leq_i^{!B^{\phi}_{w}} = \leq_i \cap (W^{!B^{\phi}_{w}} \times W^{!B^{\phi}_{w}})$
- $\sim_i^{!B^{\phi}_{w}} = \sim_i \cap (W^{!B^{\phi}_{w}} \times W^{!B^{\phi}_{w}})$
- For each $v \in W^{!B^{\phi}_{w}}$, $v \in V^{!B^{\phi}}(p)$ iff $v \in V(p)$

$\rtimes$

Note that by construction above the actual state w is never eliminated.

Observation 5.10. *For any plausibility epistemic model $\mathcal{M}$ indexed by a finite set of agents I, $\langle Sub(\mathcal{M}), \sqsubseteq\rangle$ is a chain complete poset. Moreover, for all $\phi \in \mathcal{L}_{EDL}$, $w \in W$, $!B^{\phi}$ is deflationary.*

It follows then by the Bourbaki-Witt [11] Theorem (see previous subsection) that the process of public announcement of beliefs has a fixed point. Given an initial pointed model $\mathcal{M}, w$ and some formula $\phi \in \mathcal{L}_{EDL}$, we can construct this fixed point by taking the limit of a sequence of models, which we call a public dialogue.

Definition 5.11. A *public dialogue about ϕ* starting in $\mathcal{M}, w$ is a sequence of epistemic-doxastic pointed models $\{(\mathcal{M}_n, w)\}$ such that:

- $\mathcal{M}_0 = \mathcal{M}$ is a given epistemic-plausibility model.
- $\mathcal{M}_{\beta+1} = !B^{\phi}_{w}(\mathcal{M}_{\beta})$
- $(\mathcal{M}_{\lambda})$ is the submodel of $\mathcal{M}$ generated by $\bigcap_{\beta<\lambda} |\mathcal{M}_{\beta}|$ for limit ordinals λ

$\rtimes$

It is known that such a dialogue need not stop after the first round of announcements, in e.g. the "muddy children" case [7], but by observation 5.10 we know that it will stop at some point.

Corollary 5.12 (Fixed-point)**.** *Given an epistemic-plausibility model $\mathcal{M}_0, w$ and a public dialogue about ϕ, there is a α^{ϕ} such that $(\mathcal{M}_{\alpha}, w) = (\mathcal{M}_{\alpha^{\phi}}, w)$ for all $\alpha \geq \alpha^{\phi}$.*

Moreover at $\mathcal{M}_{\alpha^{\phi}}, w$, which we call the *fixed point* of the public dialogue about ϕ, the posteriors of the agents about this formula are common knowledge, which means that they will reach an agreement on ϕ if they have common and well-founded priors.

Theorem 5.13 (Common knowledge at the fixed point)**.** *At the fixed-point of a public dialogue $\mathcal{M}_{\alpha^{\phi}}, w$ about ϕ, for all $w \in W$ and $i \in I$, if $w \in ||B_i\phi||^{\mathcal{M}_{\alpha^{\phi}}}$ then $w' \in ||B_i\phi||^{\mathcal{M}_{\alpha^{\phi}}}$ for all $w' \in [w]^{*}_{\alpha^{\phi}, I}$, and similarly if $w \notin ||B_i\phi||^{\mathcal{M}_{\alpha^{\phi}}}$.*

Corollary 5.14 (Agreements via Public Announcements)**.** *For any public dialogue about ϕ, if there is common and well-founded priors then at the fixed-point $\mathcal{M}_{\alpha^{\phi}}, w$ either all agents believe that ϕ or they all do not believe that ϕ.*

This new form of dynamic agreements result is conceptually important because it fits better than iterated conditioning the intuitive idea of a *public* dialogue, or so shall we argue in the next section, by highlighting the differences between the two processes of information exchange.

5.3 Comparing Agreements via Conditioning and Public Announcements

In this section we highlight by way of two examples that public announcements, in comparison with belief conditioning, are indeed *public*. We illustrate this first by comparing how conditioning and public announcements respectively change higher-order information, even in the case of "non-epistemic" facts. We then point out that this difference can indeed lead to different agreements, precisely in cases where the dialogues *are* about epistemic facts.

Example 5.15. *Consider the model in Figure 3. The arrows represent 1 and 2's common plausibility ordering, with $w \leq w'$ and $w' \leq w$ for all $w, w' \in W$. The solid and dotted rectangles represent 1 and 2's information partitions, respectively. Take a proposition letter p and assume that $V(p) = \{w_1, w_2\}$. Observe that the agents already agree on p at w_1, but that agent 2*

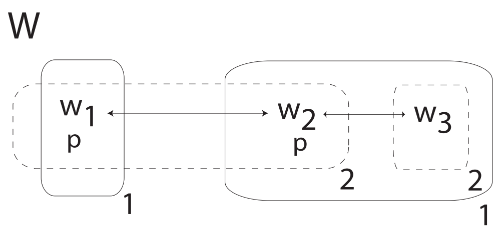

Figure 3: An epistemic plausibility model where one round of conditioning on p does not remove higher-order uncertainty about p, while a public announcement of p does.

is uncertain about 1*'s beliefs about* p*: writing* $\Diamond_2\psi$ *for* $\neg B_2\neg\psi$*, we have* $w_1 \models \Diamond_2 B_1 p \wedge \Diamond_2 \neg B_1 p$*. A single public announcement of* p *at* w_1 *suffices to remove this higher-order uncertainty:* $w_1 \models [p!]C_{1,2}p$*. Agent* 2*'s uncertainty about* 1*'s beliefs about* p*, however, remains after a single conditioning on* p*. Taking* $\Diamond_2^{\psi}\phi'$ *for* $\neg B_2^{\psi}\neg\psi'$*, we have* $w_1 \models \Diamond_2^p B_1 p \wedge \Diamond_2^p \neg B_1 p$*.*

This example illustrates the public character of announcements in comparison with the private character of conditioning. In the first case all agents know that all others have received the same piece of truthful information. This is not necessarily the case for conditioning, even if all agents condition simultaneously on the same piece of information.

Given any pointed epistemic plausibility model $\mathcal{M}, w$ and formula ϕ, the reader can check that both the dialogue about ϕ via public announcements and the dialogue about ϕ via belief conditioning at $\mathcal{M}, w$ lead to the same agreement whenever ϕ is a Boolean combination of propositional letters. This is mainly due to the fact that neither operation changes the "basic facts", i.e. the propositional valuation in a given model. They do, however, treat "informational" facts differently, as the following example shows.

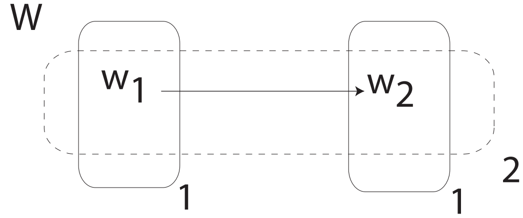

Figure 4: An epistemic plausibility model where conditioning leads to different agreements than public announcements

Example 5.16. *Consider the epistemic plausibility model in Figure 4. The arrows and rectangles are as in example 5.15. Take a proposition letter* p *and assume that* $V(p) = \{w_1\}$*. Let* $\phi := p \wedge \neg B_2 p$*, i.e. "*$p$ *but* 2 *doesn't believes that* p*". Observe that* ϕ *holds at* w_1*, that* 1 *believes it but that* 2 *does not. The conditioning dialogue and the dialogue via public announcements, both about* ϕ*, reach their fixed point* n^* *after one round in this model, where* $[w_1]_{n^*,1} = [w_1]_{n^*,2} = \{w_1\}$*. The formula* ϕ *leads to an "unsuccessful update" by public announcement [14], and at the fixed point of the dialogue neither* 1 *nor* 2 *believe that* ϕ*. In conditioning dialogue, however, both agents* do *believe that* ϕ *at the fixed point.*

This example hinges on the fact that public announcement and belief conditioning have a different influence on higher-order information. In conditioning the truth value of the formula under consideration remains fixed. If the formula contains epistemic (K_i or C_G) or doxastic (B_i, CB_G) operators, this means that the conditioning dialogue bears on the knowledge and beliefs of the agents anterior to the information exchange [3]. In dialogues via public announcements the truth value of the formula ϕ is dynamically adapted to the incoming new information, reflecting the fact that knowing that others receive the same piece of information might lead an agent to revise his higher-order information, too.

This highlights the public character of announcements in comparison with belief conditioning, and thus that the former fit well with the intuition of public dialogue that drives the dynamic agreement results.

6 Conclusion

We have studied agreement theorems from the point of view of dynamic-epistemic logic. We have shown that both static and dynamic agreement results hold in epistemic plausibility models, answering an open question in the logic literature. We pointed out the need for rather expressive logical languages to reason explicitly about static agreement results. We have furthermore improved on existing qualitative agreement results by proving that common belief in posteriors is sufficient to ensure agreement, under common and well-founded priors, and so for both finite and countable structures. Finally, we focused on the distinction between conditioning and public announcements to provide two dynamic agreement results, arguing that the later better capture the public character of dialogues. Introducing agreement theorems to dynamic-epistemic logic thus proves to be both technically and conceptually fruitful, and it bridges two important bodies of literature.

Future work should put the full generality dynamic-epistemic logic [6, 14] to use, as well as recent developments in "softer" forms of belief updates [8, 4], to analyze the possibility of agreements in a larger class of situations. It also remains open whether one can finitely axiomatize a logic which can derive the agreement results, in both their static and dynamic forms. Finally, two issues pertaining to the expressibility of the static agreement theorems should be investigated further: first, the definability of the common prior assumption via countable sets of formulas of $\mathcal{L}_{EDL}$, as shown by [18] for the probabilistic case; and second, expressibility of alternative agreement results, as e.g. the one provided in [21].

Acknowledgments

We would like to thank Johan van Benthem, Balder ten Cate, Barteld Kooi, Stig Andur Pedersen and Sonja Smets for helpful comments and discussions. We are also grateful to the referees of TARK, and to the participants at the Research Seminar in Philosophy of Linguistics and the workshop Formal Modeling in Social Epistemology, both at Tilburg University, the Dynamic Logic Working Session at the ILLC in Amsterdam, the LSE-Groningen Exchange Workshop in London and the GroLog seminar in Groningen.

References

[1] R.J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.

[2] M. Bacharach. Some extensions of a claim of Aumann in an axiomatic model of knowledge. *Journal of Economic Theory*, 37(1):167–190, October 1985.

[3] A. Baltag and S. Smets. A qualitative theory of dynamic interactive belief revision. In Giacomo Bonanno, Wiebe van der Hoek, and Michael Wooldridge, editors, *Logic and the Foundation of Game and Decision Theory (LOFT7)*, volume 3 of *Texts in Logic and Games*, pages 13–60. Amsterdam University Press, 2008.

[4] A. Baltag and S. Smets. Learning by questions and answers: From belief-revision cycles to doxastic fixed points. Under Review, 200X.

[5] A. Baltag and S. Smets. Conditional doxastic models: A qualitative approach to dynamic belief revision. In G. Mints and R. de Queiroz, editors, *Proceedings of WOLLIC 2006, Electronic Notes in Theoretical Computer Science*, volume 165, 2006.

[6] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. In *TARK 98*, 1998.

[7] J. van Benthem. One is a lonely number. In Z. Chatzidakis, P. Koepke, and W. Pohlers, editors, *Logic Colloquium '02*, Wellesley MA, 2006. ASL & A.K. Peters.

[8] J. van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-classical Logics*, 17 (2), 2007.

[9] G. Bonanno and K. Nehring. How to make sense of the common prior assumption under incomplete information. *International Journal of Game Theory*, 28(03):409–434, August 1999.

[10] G. Bonanno and K. Nehring. Agreeing to disagree: a survey. Some of the material in this paper was published in [9], 1997.

[11] N. Bourbaki. Sur le théorème de Zorn. *Archiv der Mathematik*, 2, 1949.

[12] B. ten Cate. *Model theory for extended modal languages*. PhD thesis, University of Amsterdam, 2005. ILLC Dissertation Series DS-2005-01.

[13] J. A. K. Cave. Learning to agree. *Economics Letters*, 12(2):147–152, 1983.

[14] H. van Ditmarsch, W. van de Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library Series*. Springer, 2007.

[15] R. Fagin, J.Y. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.

[16] J. Geanakoplos and H. M. Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28 (1):192–200, 1982.

[17] J. Gerbrandy. *Bisimulations on Planet Kripke*. PhD thesis, ILLC, Amsterdam, 1999.

[18] A. Heifetz. The positive foundation of the common prior assumption. *Games and Economic Behavior*, 56(1):105–120, 2006.

[19] L. Menager. *Communication, common knowledge, and consensus*. PhD thesis, Universite Paris I Pantheon-Sorbonne, 2006.

[20] J.A. Plaza. Logics of public communications. In M.L. Emrich, M.S. Pfeifer, M. Hadzikadic, and Z.W. Ras, editors, *Proceedings of the Fourth International Symposium on Methodologies for Intelligent Systems: Poster Session Program*, pages 201–216. Oak Ridge National Laboratory, 1989.

[21] D. Samet. Agreeing to disagree: The non-probabilistic case. *Games and Economic Behavior*, In Press.