Cautious Distributed Belief

John Lindqvist, Fernando R. Velázquez-Quesada and Thomas Ågotnes

July 16, 2022

Cautious distributed belief

John Lindqvist, Fernando R. Velázquez-Quesada, and Thomas Ågotnes

Department of Information Science and Media Studies, Universitetet i Bergen.
{John.Lindqvist, Fernando.VelazquezQuesada, Thomas.Agotnes}@uib.no

Abstract This manuscript introduces and studies a notion of *cautious* distributed belief. Different from the standard distributed belief, the cautious distributed belief of a group is inconsistent only when *all* group members are individually inconsistent. The paper presents basic results about cautious distributed belief, investigates whether it inherits properties from individual belief, and compares it with standard distributed belief. Although both notions are equivalent in the class of reflexive models, this is not the case in general. The propositional language extended only with cautious distributed belief is strictly less expressive than the propositional language extended only with standard distributed belief. We, finally, identify a minimal extension of the language making the former as expressive as the latter.

Keywords: cautious distributed belief · distributed belief · epistemic logic · expressivity · bisimilarity

1 Introduction

Epistemic logic (*EL*; [12]) is a simple and yet powerful framework for representing the knowledge of a set of agents. Semantically, it typically relies on relational 'Kripke' models, assigning to each agent a binary *indistinguishability* relation over possible worlds (i.e., possible states of affairs). Syntactically, it uses the agent's indistinguishability to define her *knowledge*: at a world w an agent i knows that φ is the case if and only if φ holds in all the situations that are, for her, indistinguishable from w. Despite its simplicity, *EL* has become a widespread tool, contributing to the formal study of complex multi-agent epistemic phenomena in philosophy [9], computer science [6,14] and economics [4,15].

One of the most attractive features of *EL* is that one can reason not only about individual knowledge, but also about different forms of knowledge for groups. A historically important example is the notion of *common knowledge* [13], which is known to be crucial in social interactions.[1] Another important epistemic notion for groups, key in distributed systems, is that of *distributed knowledge* [11,7,8]. Intuitively, a group has distributed knowledge of φ if and only if φ follows from the combination of the individual knowledge of all its members. In *EL* (which, recall, uses uncertainty to define knowledge), this intuition has a

[1] A group has common knowledge of φ if and only if everybody in the group knows φ, everybody in the group knows that everybody in the group knows φ, and so on.

natural representation: at a world w a group G has distributed knowledge of φ if and only if φ holds in all the situations that *all the members of the group* consider indistinguishable (i.e., if and only if φ holds in all the situations *no one in the group* can distinguish) from w. In other words, the indistinguishability relation for the distributed knowledge of a group G corresponds to the *intersection* of the indistinguishability relation of G's members.

Since distributed knowledge is the result of combining the individual knowledge of different agents, one can wonder whether agents might have inconsistent distributed knowledge (i.e., whether it is possible for a set of agents to know $\bot$ distributively). When one works with a *truthful* notion of knowledge (semantically, when all indistinguishability relations are required to be reflexive), distributed knowledge does not have this problem: all indistinguishability relations contain the reflexive edges, and thus their intersection will never be empty. However, when one works with weaker notions of information, counterintuitive situations might occur. For example, if one works with a notion of beliefs (typically represented by using a serial, transitive and Euclidean relation; see, e.g., [12]), it is possible for all agents to be consistent (i.e., no one of them believes contradictions), and yet their distributed beliefs might contain $\bot$.

This paper introduces and studies a notion of *cautious* distributed belief (modality: $D^{\forall}$). It has the property that it does not become inconsistent in the case of mutual inconsistency, picking instead a form of maximally consistent combined information. The intuition behind it is that, although a group G as a whole might be inconsistent at some world w (i.e., the set of worlds *everybody* in G considers possible from w is empty), there might be consistent subgroups among which the maximal ones become important. Considering notions of maximal consistency is a standard approach in non-monotonic reasoning for resolving potential conflicts.[2] As its name suggest, $D^{\forall}$ uses these maximally consistent subgroups of agents in a cautious way: at a world w a group G has cautious distributed belief that φ if and only if *every* maximally consistent subgroup of G has distributed belief that φ.[3]

The manuscript is organised as follows. Section 2 recalls the definition of a relational 'Kripke' model as well as that of the standard distributed belief operator D. Then it introduces the notion of cautious distributed belief, using a relatively simple example to compare the two notions, and presenting some basic results about it. Section 3 studies whether this notion of belief for groups inherits properties from the individual beliefs of the group's members. Section 4 compares the expressive power of both modalities, showing that a modal language with only $D^{\forall}$ is strictly less expressive than a modal language with only D; it does so

[2] Think, e.g., about the *extensions* of a theory in default logic [16], or the *maximally admissible* (i.e., preferred) sets of arguments in abstract argumentation theory [5]. The idea has been also used within epistemic logic (e.g., by [2] in the context of evidence-based beliefs) and also for distributed beliefs (by [10], in the context of *explicit* beliefs defined via belief bases).

[3] This corresponds to the *skeptical* reasoner in non-monotonic reasoning. There is also an alternative that matches the *credulous* reasoner, discussed briefly in Section 5.

by providing a notion of bisimulation for $D^{\forall}$. Yet, the paper identifies what it is that D can see but $D^{\forall}$ cannot. Finally, Section 5 summarises the results and discusses further research lines.

2 Basic definitions

Throughout this text, let A be a finite non-empty set of agents and P be a countable non-empty set of atomic propositions. The basic propositional language (using $\neg$ and $\wedge$ as primitive operators) is denoted by $\mathcal{L}$. (Its semantic interpretation is as usual.) Then, $\mathcal{L}_{X_1,\ldots,X_n}$ is the language extending $\mathcal{L}$ with the operators $X_1, \ldots, X_n$. In particular, $\mathcal{L}_D$ is $\mathcal{L}$ with the additional use of D_G for $\varnothing \neq G \subseteq A$, and $\mathcal{L}_{D^{\forall}}$ is $\mathcal{L}$ with the additional use of $D^{\forall}_G$ for $\varnothing \neq G \subseteq A$.

Definition 1 (Belief model) *A* belief model *is a tuple $\mathcal{M} = \langle W, R, v\rangle$ where W is a non-empty set of possible worlds (also denoted as $\mathrm{D}(\mathcal{M})$), $R = \{R_a \subseteq W \times W \mid a \in A\}$ assigns an arbitrary* accessibility *relation to each agent $a \in A$, and $v : P \to 2^S$ is a valuation function. A* pointed belief model *is a pair $(\mathcal{M}, s)$ with $\mathcal{M}$ a belief model and $s \in \mathrm{D}(\mathcal{M})$ a world in it. The class of all belief models is denoted as* **M**. *Given $\langle W, R, v\rangle$ in* **M**, *$a \in A$ and $s \in W$, the set $C_a(s) := \{s' \in W \mid sR_as'\}$ is called* a's conjecture set at s. *The generalisation to a set of agents $G \subseteq A$, called* G's (combined) conjecture set at s, *is defined as $C_G(s) := \bigcap_{a\in G} C_a(s)$.* ◀

Belief models are nothing but multi-agent Kripke (relational) models. Thus, they allow us to represent not only the beliefs each individual agent has, but also different belief notions for groups. As discussed in the introduction, the focus here is the novel notion of *cautious distributed beliefs* ($D^{\forall}$), together with its relationship with the well-known notion of *distributed beliefs* (D). For the semantic interpretation of the first, the following definitions will be useful.

Definition 2 (Consistency and maximal consistency) *Let $\langle W, R, v\rangle$ be in* **M**. *Take sets of agents $\varnothing \subset G' \subseteq G \subseteq A$ and a world $s \in W$. The set G' is* consistent at s *if and only if $C_{G'}(s) \neq \varnothing$. It is* maximally consistent at s w.r.t. G *(notation: $G' \subseteq^{max}_s G$) if and only if it is consistent at s and, additionally, every H satisfying $G' \subset H \subseteq G$ is inconsistent (i.e., $C_H(s) = \varnothing$). Finally, the set $C^{\forall}_G(s) := \bigcup_{G' \subseteq^{max}_s G} C_{G'}(s)$ (the* consistent (combined) conjecture set *of G at s) contains the worlds that are relevant for the maximally consistent subgroups of G at world s. The* cautious distributed belief relation *$R^{\forall}_G \subseteq \mathrm{D}(\mathcal{M}) \times \mathrm{D}(\mathcal{M})$, given by $sR^{\forall}_G t$ iff $t \in C^{\forall}_G(s)$, will simplify some later work.* ◀

Here is the semantic interpretation of the two operators, D and $D^{\forall}$, together with the standard operator for individual belief B. We also present the semantics of an additional constant $\asymp_G$, which will be useful later. Languages using these operators will be discussed in Section 4.

Definition 3 (Two types of distributed belief) *Let $(\mathcal{M}, s)$ be a pointed belief model with $\mathcal{M} = \langle W, R, v\rangle$; take $a \in A$ and a non-empty $G \subseteq A$. Then,*

$\mathcal{M}, s \models B_a\,\varphi$ *iff* $\forall s' \in C_a(s)$: $\mathcal{M}, s' \models \varphi$,

$\mathcal{M}, s \models D_G\,\varphi$ *iff* $\forall s' \in C_G(s)$: $\mathcal{M}, s' \models \varphi$,

$\mathcal{M}, s \models D^{\forall}_G\,\varphi$ *iff* $\forall G' \subseteq^{max}_s G$, $\forall s' \in C_{G'}(s)$: $\mathcal{M}, s' \models \varphi$
(equivalently, $\forall s'$ *such that* $sR^{\forall}_G s'$: $\mathcal{M}, s' \models \varphi$*)*,[4]

$\mathcal{M}, s \models \asymp_G$ *iff* $C_G(s) = \varnothing$.

A formula φ *is* valid in a class of belief models $\mathbf{C}$ *(notation:* $\mathbf{C} \models \varphi$*) when* φ *is true in every world of every model in* $\mathbf{C}$*. A formula is* valid *(notation:* $\models \varphi$*) when* $\mathbf{M} \models \varphi$. ◀

Note the difference between D_G and $D^{\forall}_G$. On the one hand, $D_G\,\varphi$ holds at s when every world in the conjecture set of G satisfies φ.[5] On the other hand, $D^{\forall}_G\,\varphi$ holds at s when every world in the conjecture set of *every maximally consistent subgroup of* G satisfies φ. In other words, $D^{\forall}_G\,\varphi$ holds at s if and only if *every* maximally consistent subgroup of G has distributed belief of φ. Note also how $\asymp_G$ simply expresses the fact that the conjecture set of G is inconsistent.

Here is an simple example showing the differences between D and $D^{\forall}$.

Example 1 *Consider the belief model* $\mathcal{M}$ *below.*[6] *Note how, at* w_1*,* a *believes* p *to be true and* q *to be false (*$\mathcal{M}, w_1, \models B_a\,p \wedge B_a\,\neg q$*). Nevertheless,* b *is uncertain about* p *but believes* q *to be true (*$\mathcal{M}, w_1 \models (\neg B_b\,p \wedge \neg B_b\,\neg p) \wedge B_b\,q$*). Finally,* c *believes* p *but is uncertain about* q *(i.e.,* $\mathcal{M}, w_1, \models B_c\,p \wedge (\neg B_c\,q \wedge \neg B_c\,\neg q)$*).*

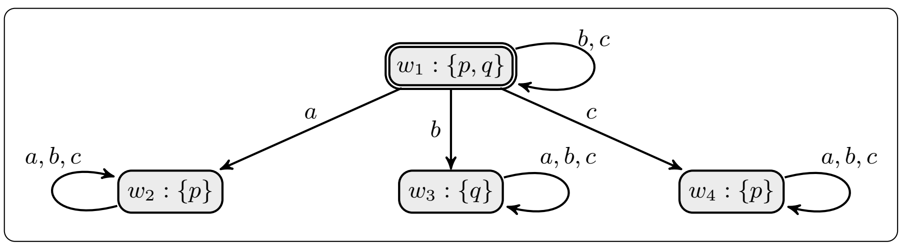

Consider first the group $G_1 = \{a, b\}$*. On the one hand, both members of* G_1 *are individually consistent at* w_1 *and yet* $C_{G_1}(w_1) = \varnothing$*; thus, at* w_1*, the maximally consistent subgroups are* $\{a\}$ *and* $\{b\}$*. Their conjecture sets are* $C_a(w_1) = \{w_2\}$ *and* $C_b(w_1) = \{w_1, w_3\}$*, and hence* G_1*'s consistent conjecture set is* $C^{\forall}_{G_1}(w_1) = \{w_1, w_2, w_3\}$*. Thus,* $\mathcal{M}, w_1 \models \neg D^{\forall}_{G_1}\,p \wedge \neg D^{\forall}_{G_1}\,q$*. On the other hand, when we consider standard distributed belief, we see that* $\mathcal{M}, w_1 \models D_{G_1}\,p \wedge D_{G_1}\,q$*. This is*

[4] The two definitions are equivalent. The first makes explicit the two quantification steps; the second, given in terms of the group's cautious distributed belief relation, reveals that $D^{\forall}_G$ is in fact a normal modality.

[5] In particular, individual belief operators B_a can be defined in terms of D, as $D_{\{a\}}\,\varphi$ (abbreviated as $D_a\,\varphi$) holds in a world s if and only if $\mathcal{M}, s' \models \varphi$ for all $s' \in C_a(s)$.

[6] Note: the individual relations are serial, transitive and Euclidean. While the paper uses the term "belief" in a rather loose way, these three properties are the ones commonly associated to a belief operator.

however due to the fact that $C_{G_1}(w_1) = \varnothing$ *and we end up quantifying over an empty set. Thus, we also get* $\mathcal{M}, w_1 \vDash D_{G_1} \bot$.

Now c *joins the group,* $G_2 = \{a, b, c\}$. *On the one hand, at* w_1 *both* b *and* c *are consistent (i.e., they can 'consistently combine information'); still,* a *and* c *are not. Thus, the maximally consistent sets are* $\{a\}$ *and* $\{b, c\}$. *The relevant conjecture sets are now* $C_a(w_1) = \{w_2\}$ *and* $C_{\{b,c\}}(w_1) = \{w_1\}$, *so* $C^{\forall}_{G_2}(w_1) = \{w_1, w_2\}$. *Then,* $\mathcal{M}, w_1 \vDash D^{\forall}_{G_2}\, p \wedge \neg D^{\forall}_{G_2}\, q$ *(the latter because, even though* b *and* c *together believe* q, *agent* a *remains 'a loner' and still believes that* q *is false). On the other hand, the situation with standard distributed belief remains as for* G_1: $\mathcal{M}, w_1 \vDash D_{G_2}\, p \wedge D_{G_2}\, q \wedge D_{G_2} \bot$. ◀

Some basic results about $D^{\forall}_G$**.** The standard notion of distributed belief, D_G, can be inconsistent even when every agent in G is consistent. The first result here shows that this is not the case for cautious distributed belief: it is inconsistent if and only if *all* agents in G are inconsistent.

Proposition 1 *For every non-empty* $G \subseteq A$ *we have* $\vDash D^{\forall}_G \bot \leftrightarrow \bigwedge_{a \in G} B_a \bot$.

Proof. Take any $\mathcal{M}$, *any* $s \in \mathrm{D}(\mathcal{M})$ *and any non-empty* $G \subseteq A$. *(*$\Rightarrow$*) If* $\mathcal{M}, s \vDash D^{\forall}_G \bot$ *then, because no world satisfies* $\bot$, *either* $C_{G'}(s) = \varnothing$ *for all* $G' \subseteq^{max}_s G$, *or there is no* G' *satisfying* $G' \subseteq^{max}_s G$. *But, by definition, no* G' *satisfying* $G' \subseteq^{max}_s G$ *is s.t.* $C_{G'}(s) = \varnothing$. *Hence, there is no* G' *satisfying* $G' \subseteq^{max}_s G$, *which means every* $G' \subseteq G$ *is s.t.* $C_{G'}(s) = \varnothing$. *In particular, all singletons* $\{a\}$ *for* $a \in G$ *are s.t.* $C_a(s) = \varnothing$, *and thus* $\mathcal{M}, s \vDash \bigwedge_{a \in G} B_a \bot$. *(*$\Leftarrow$*) If* $\mathcal{M}, s \vDash \bigwedge_{a \in G} B_a \bot$ *then* $C_a(s) = \varnothing$ *for every* $a \in G$. *Hence, every non-empty* $G' \subseteq G$ *is s.t.* $C_{G'}(s) = \varnothing$, *so there is no* G' *satisfying* $G' \subseteq^{max}_s G$. *Thus,* $\mathcal{M}, s \vDash D^{\forall}_G \bot$. ■

For another basic result, recall that individual belief operators (B_a for $a \in A$) can be expressed using the distributed belief operator for singleton groups (D_a). The same can be done with cautious distributed belief. For any world s and any agent a, there is at most one maximally consistent subgroup of $\{a\}$, namely $\{a\}$ itself. Then, $C_a(s) = C^{\forall}_a(s)$ and hence agent a's individual belief and $\{a\}$'s cautious distributed belief coincide.

Proposition 2 $\vDash B_a \varphi \leftrightarrow D^{\forall}_{\{a\}}\, \varphi$. ■

Finally, an important property of standard distributed belief is *coalition monotonicity*: if a group $H \subseteq A$ has standard distributed belief that φ, then so does any extension $G \supseteq H$ (thus, $H \subseteq G \subseteq A$ implies $\vDash D_H\, \varphi \to D_G\, \varphi$). This is not the case for cautious distributed belief. This is because the agents that join the group might not be consistent with any of the ones that were there before. In such cases, when consistent, they will be part of a different maximally consistent subgroup, which might not have the distributed belief φ. This is shown in Example 1, where $\mathcal{M}, w_1 \vDash D^{\forall}_{\{b\}}\, q$ and yet $\mathcal{M}, w_1 \nvDash D^{\forall}_{\{a,b\}}\, q$. Thus,

Fact 1 $\nvDash D^{\forall}_H\, \varphi \to D^{\forall}_G\, \varphi$ *for* $H \subseteq G \subseteq A$. ■

Frame condition	Characterising formula
seriality (l): $\forall s \in W\ \exists t \in W\,.\, sSt$	*consistency*: $\Box\varphi \to \Diamond\varphi$
reflexivity (r): $\forall s \in W\,.\, sSs$	*truthfulness of knowledge/belief*: $\Box\varphi \to \varphi$
transitivity (t): $\forall s, t, u \in W\,.\, ((sSt\ \&\ tSu) \Rightarrow sSu)$	*positive introspection*: $\Box\varphi \to \Box\Box\varphi$
symmetry (s): $\forall s, t \in W\,.\, (sSt \Rightarrow tSs)$	*truthfulness of possible knowledge/belief*: $\Diamond\Box\varphi \to \varphi$
Euclidicity (e): $\forall s, t, u \in W\,.\, ((sSt\ \&\ sSu) \Rightarrow tSu)$	*negative introspection*: $\neg\Box\varphi \to \Box\neg\Box\varphi$

Table 1: Relational properties and their well-known characterising formula.

3 Inheriting relational properties

When one studies a notion of knowledge/belief for groups, it is interesting to find out whether it inherits the properties of the knowledge/beliefs of the individuals. For example, suppose that the individual knowledge of all agents in a group is truthful and both positively and negatively introspective. Then, it is well-known that, while the group's common knowledge inherits all these properties, the group's *general* knowledge[7] inherits only truthfulness (i.e., it might not be positively or negatively introspective). Similar studies have been made for notions of belief [1].

This section studies which properties of individual belief are inherited by cautious distributed belief. The discussion is rather semantic, focussing on whether certain frame conditions on individual indistinguishability relations are inherited by the relation that defines cautious distributed belief (see Footnote 4). The connection between these conditions and the properties of knowledge/belief is made thanks to the well-known correspondence between the frame conditions and the validity of certain modal formulas [3, Chapter 3]. Using S for an arbitrary binary relation and $\Box$ ($\Diamond$) for its corresponding normal universal (existential) modality, Table 1 lists some of these frame conditions, together with the formulas that characterise them (and its intuitive epistemic/doxastic reading).[8]

Here are, then, the needed definitions.

[7] A group has general knowledge of φ if and only if everybody in the group knows φ.

[8] More precisely, a *frame* (a model without the valuation) has the given relational property if and only if the formula is valid in the frame (i.e., it is true in any world of the model under any valuation).

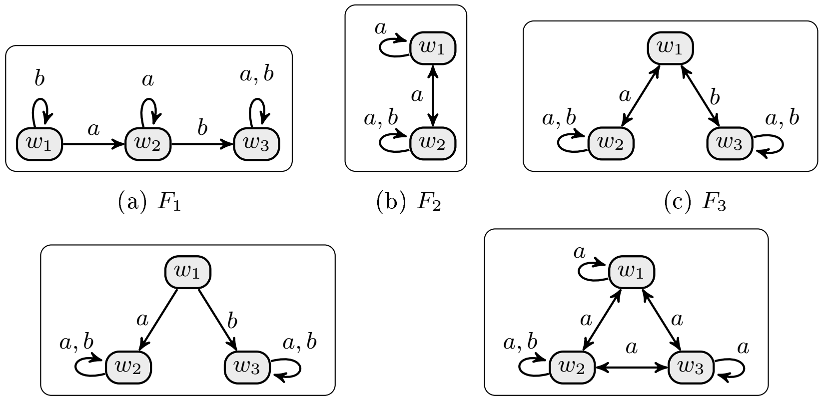

(a) F_1 (b) F_2 (c) F_3

(d) F_4 (e) F_5

Figure 1: Counterexamples for the proof of Proposition 3

Definition 4 (Inheriting properties) *Let $x \in \{l, r, t, s, e\}$ be a frame condition, and let $\mathcal{F} \subseteq \{l, r, t, s, e\}$ be a collection of them. Let $G \subseteq A$ be a non-empty set of agents, each one of them associated to a binary relation under a given domain W. A relation $S_G \subseteq W \times W$ defined in terms of the individual relations for agents in G (e.g., their union/intersection)* inherits the condition x under the additional conditions in $\mathcal{F}$ *if and only if S_G has the property x whenever* all *the relations in $\{R_i \mid i \in G\}$ have all the properties in $\mathcal{F} \cup \{x\}$.* ◀

For singleton groups, all properties are preserved. This is because if G is a singleton $\{a\}$, then the cautious distributed belief relation $R^{\forall}_{\{a\}}$ is identical to a's individual relation R_a.

Proposition 3 *Given a collection of relations $\{R_a \subseteq W \times W \mid a \in G\}$ for a group $G \subseteq A$ with at least two agents, the relation $R^{\forall}_G \subseteq W \times W$*

(1) *inherits* seriality *under $\mathcal{F} = \varnothing$;*
(2) *inherits* reflexivity *under $\mathcal{F} = \varnothing$;*
(3) (a) *does not inherit* transitivity *under any $\mathcal{F} \subseteq \{l, e\}$;*
  (b) *inherits* transitivity *under any $\mathcal{F} \supseteq \{r\}$ (also under any $\mathcal{F} \supseteq \{l, s\}$[9]);*
  (c) *inherits* transitivity *under any $\mathcal{F} \supseteq \{s\}$;*
(4) (a) *does not inherit* symmetry *under any $\mathcal{F} \subseteq \{t, e\}$;*
  (b) *does not inherit* symmetry *under any $\mathcal{F} \subseteq \{l, e\}$;*
  (c) *inherits* symmetry *under any $\mathcal{F} \supseteq \{r\}$ (also under any $\mathcal{F} \supseteq \{l, t\}$);*
(5) (a) *does not inherit* Euclidicity *under any $\mathcal{F} \subseteq \{l, s\}$;*
  (b) *does not inherit* Euclidicity *under any $\mathcal{F} \subseteq \{l, t\}$;*

[9] Inheritance under any $\mathcal{F} \subseteq \{l, s\}$ follows immediately from inheritance under any $\mathcal{F} \subseteq \{r\}$, since seriality, transitivity and symmetry together imply reflexivity. The same applies for the properties in (3)(c) and (5)(d) below.

(c) *does not inherit* Euclidicity *under any* $\mathcal{F} \subseteq \{t, s\}$;
(d) *inherits* Euclidicity *under any* $\mathcal{F} \supseteq \{r\}$ *(also under any* $\mathcal{F} \supseteq \{l, t, s\}$*)*.

Proof.

(1) *Pick any* $s \in W$. *Every relation in* $\{R_i \mid i \in G\}$ *is serial so, since* $G \neq \varnothing$, *there is* $a \in G$ *such that* R_a *is serial, and* a *is consistent at* s $(C_a(s) \neq \varnothing)$. *Thus,* G *has at least one subgroup* G' *that is maximally consistent at* s *(one containing* a*), and hence there is* $t \in C_{G'}(s) \subseteq C_G^{\forall}(s)$. *Then,* $R_G^{\forall}$ *is serial.*

(2) *Pick any* $s \in W$. *Every relation in* $\{R_i \mid i \in G\}$ *is reflexive, so* $s \in C_a(s)$ *for every* $a \in G$, *and then the only maximally consistent subgroup is* G *itself. Thus,* $C_G(s) = C_G^{\forall}(s)$ *and therefore* $s \in C_G^{\forall}(s)$. *Then,* $R_G^{\forall}$ *is reflexive.*

(3) *(a)* *In frame* F_1 *(Figure 1a), relations* R_a *and* R_b *are transitive, serial and Euclidean. Still,* $R_{\{a,b\}}^{\forall} = \{(w_1, w_1), (w_1, w_2), (w_2, w_2), (w_2, w_3), (w_3, w_3)\}$ *is not transitive.*

(b) *Pick any* $s, t, u \in W$ *such that* $sR_G^{\forall}t$ *and* $tR_G^{\forall}u$. *By reflexivity,* G *is the only maximally consistent subgroup at both* s *and* t, *so* $C_G(s) = C_G^{\forall}(s)$ *and* $C_G(t) = C_G^{\forall}(t)$. *Then,* sR_it *and* tR_iu *for every* $i \in G$, *which by transitivity implies* sR_iu *for all such* i. *Thus,* $u \in C_G^{\forall}(s)$ *and hence* $sR_G^{\forall}u$. *Then,* $R_G^{\forall}$ *is transitive.*

(c) *Pick any* $s, t, u \in W$ *such that* $sR_G^{\forall}t$ *and* $tR_G^{\forall}u$. *Then, there are* $H_1 \subseteq_s^{max} G$ *and* $H_2 \subseteq_t^{max} G$ *such that* $t \in C_{H_1}(s)$ *and* $u \in C_{H_2}(t)$. *By individual symmetry,* $s \in C_{H_1}(t)$ *and* $t \in C_{H_2}(u)$; *then, by individual transitivity,* $t \in C_{H_1}(t)$ *and* $t \in C_{H_2}(t)$. *But then,* $H_1 \cup H_2$ *is consistent at* t *and, since* H_2 *is maximally consistent at* t, *then* $(H_1 \cup H_2) \subseteq H_2$, *that is,* $H_1 \subseteq H_2$. *Hence, the previous* $u \in C_{H_2}(t)$ *implies* $u \in C_{H_1}(t)$ *which, together with* $t \in C_{H_1}(s)$ *and individual transitivity implies* $u \in C_{H_1}(s)$. *Finally, since* H_1 *is maximally consistent at* s *w.r.t.* G, $u \in C_G^{\forall}(s)$, *and hence* $sR_G^{\forall}u$.

(4) *(a)* *In frame* F_2 *(Figure 1b), relations* R_a *and* R_b *are symmetric, transitive and Euclidean. Still,* $R_{\{a,b\}}^{\forall} = \{(w_1, w_1), (w_1, w_2), (w_2, w_2)\}$ *is not symmetric.*

(b) *In frame* F_3 *(Figure 1c), relations* R_a *and* R_b *are symmetric, serial and Euclidean. Still,* $R_{\{a,b\}}^{\forall} = \{(w_1, w_2), (w_1, w_3), (w_2, w_2), (w_3, w_3)\}$ *is not symmetric.*

(c) *Pick any* $s, t \in W$ *such that* $sR_G^{\forall}t$. *By reflexivity,* G *is the only maximally consistent subgroup at both* s *and* t, *so* $C_G(s) = C_G^{\forall}(s)$ *and* $C_G(t) = C_G^{\forall}(t)$. *Then,* sR_it *for every* $i \in G$, *which by symmetry implies* tR_is *for all such* i. *Thus,* $s \in C_G^{\forall}(t)$ *and hence* $tR_G^{\forall}s$. *Then,* $R_G^{\forall}$ *is symmetric.*

(5) *(a)* *In frame* F_3 *(Figure 1c), relations* R_a *and* R_b *are Euclidean, serial and symmetric. Still,* $R_{\{a,b\}}^{\forall} = \{(w_1, w_2), (w_1, w_3), (w_2, w_2), (w_3, w_3)\}$ *is not Euclidean.*

(b) *In frame* F_4 *(Figure 1d), relations* R_a *and* R_b *are Euclidean, serial and transitive. Still,* $R_{\{a,b\}}^{\forall} = \{(w_1, w_2), (w_1, w_3), (w_2, w_2), (w_3, w_3)\}$ *is not Euclidean.*

(c) *In frame* F_5 *(Figure 1e), relations* R_a *and* R_b *are Euclidean, symmetric and transitive. Still,* $R_{\{a,b\}}^{\forall} = (W \times W) \setminus \{(w_2, w_1), (w_2, w_3)\}$ *is not Euclidean.*

(d) *Pick any $s, t, u \in W$ such that $sR^{\forall}_{G}t$ and $sR^{\forall}_{G}u$. By reflexivity, G is the only maximally consistent subgroup at both s and t, so $C_G(s) = C^{\forall}_{G}(s)$ and $C_G(t) = C^{\forall}_{G}(t)$. Then, sR_it and sR_iu for every $i \in G$, which by Euclidicity implies tR_iu for all such i. Thus, $u \in C^{\forall}_{G}(t)$ and hence $tR^{\forall}_{G}u$. Then, $R^{\forall}_{G}$ is Euclidean.* ■

Thus, seriality and reflexivity are each inherited without additional assumptions. Symmetry and Euclidicity are both inherited in the presence of reflexivity; transitivity is inherited in the presence of reflexivity, but also in the presence of symmetry. Thus, just as with *individual* belief, cautious distributed belief is factive in reflexive models, and it is consistent in serial models. However, it does not need to be introspective (neither positively nor negatively), even when the model has the frame condition (transitivity and Euclidicity, respectively).

These results are quite different from the corresponding ones for the standard notion of distributed belief. In fact, with the exception of reflexive models (in which cautious and standard distributed belief coincide; see Proposition 4 below), the behaviour of cautious distributed belief is, in this respect, the opposite of that of standard distributed belief. For the latter, transitivity, symmetry and Euclidicity are each inherited without additional assumptions, while seriality is is inherited only in the presence of reflexivity [1].

4 Relationship between D_G and $D^{\forall}_{G}$

This section discusses the relationship between standard and cautious distributed belief. The following definitions will be useful.

Definition 5 *Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two languages whose formulas can be evaluated over pointed belief models.*

- *$\mathcal{L}_2$ is* at least as expressive *as $\mathcal{L}_1$ (notation: $\mathcal{L}_1 \preccurlyeq \mathcal{L}_2$) if and only if every formula in $\mathcal{L}_1$ has a semantically equivalent formula in $\mathcal{L}_2$: for every $\alpha_1 \in \mathcal{L}_1$ there is $\alpha_2 \in \mathcal{L}_2$ s.t., for every pointed belief model $(\mathcal{M}, s)$, we have $\mathcal{M}, s \vDash \alpha_1$ if and only if $\mathcal{M}, s \vDash \alpha_2$.*[10]
- *$\mathcal{L}_1$ and $\mathcal{L}_1$ are* equally expressive *(notation: $\mathcal{L}_1 \approx \mathcal{L}_2$) if and only if $\mathcal{L}_1 \preccurlyeq \mathcal{L}_2$ and $\mathcal{L}_2 \preccurlyeq \mathcal{L}_1$.*
- *$\mathcal{L}_2$ is* strictly more expressive *than $\mathcal{L}_1$ (notation: $\mathcal{L}_1 \prec \mathcal{L}_2$) if and only if $\mathcal{L}_1 \preccurlyeq \mathcal{L}_2$ and $\mathcal{L}_2 \not\preccurlyeq \mathcal{L}_1$.*[11] ◀

The proposition below provides some connections between D_G and $D^{\forall}_{G}$. First, $D^{\forall}_{G}$ is definable in terms of D_G and Boolean operators. Second, both notions coincide when the indistinguishability relations are reflexive.

[10] A typical strategy for proving $\mathcal{L}_1 \preccurlyeq \mathcal{L}_2$ is to give a translation $tr : \mathcal{L}_1 \to \mathcal{L}_2$ such that for every $(\mathcal{M}, s)$ we have $\mathcal{M}, s \vDash \alpha_1$ iff $\mathcal{M}, s \vDash tr(\alpha_1)$. The crucial cases are those for the operators in $\mathcal{L}_1$ that do not occur in $\mathcal{L}_2$.

[11] A typical strategy for proving $\mathcal{L}_1 \not\preccurlyeq \mathcal{L}_2$ is to find two pointed models that satisfy exactly the same formulas in $\mathcal{L}_2$, and yet can be distinguished by a formula in $\mathcal{L}_1$.

Proposition 4

(1) $\models D^{\forall}_{G}\,\varphi \leftrightarrow \bigwedge_{G'\subseteq G}\Big((\neg D_{G'}\bot \wedge \bigwedge_{G'\subset H\subseteq G} D_H\bot) \to D_{G'}\varphi\Big)$.

(2) *Let* $\mathbf{T}$ *be the class of all belief models whose accessibility relations are all reflexive. Then,* $\mathbf{T} \models D^{\forall}_{G}\,\varphi \leftrightarrow D_G\varphi$.

Proof.

(1) *Suppose* $\mathcal{M}, s \models D^{\forall}_{G}\,\varphi$. *By definition, this is the case if and only if every* $G' \subseteq^{max}_{s} G$ *is such that* $\mathcal{M}, s \models D_{G'}\varphi$. *But the fact that* $G' \subseteq^{max}_{s} G$ *(i.e.,* G' *is a maximally consistent subgroup of* G *at* s*) is equivalently stated as* $\mathcal{M}, s \models \neg D_{G'}\bot \wedge \bigwedge_{G'\subset H\subseteq G} D_H\bot$.[12] *Then, the previous is the case if and only if* $\mathcal{M}, s \models \bigwedge_{G'\subseteq G}\big((\neg D_{G'}\bot \wedge \bigwedge_{G'\subset H\subseteq G} D_H\bot) \to D_{G'}\varphi\big)$.

(2) *Immediate, as* $C_G(s) = C^{\forall}_{G}(s)$ *holds for any reflexive belief model* $\mathcal{M}$*, world* $s \in \mathrm{D}(\mathcal{M})$ *and group* $\varnothing \neq G \subseteq A$ *(see the proof of Proposition 3(2)).* ∎

Using the first part of Proposition 4, one can define a translation that takes any formula in $\mathcal{L}_{D^{\forall}}$ and returns a semantically equivalent formula in $\mathcal{L}_D$. Thus, it already establishes a connection between $\mathcal{L}_D$ and $\mathcal{L}_{D^{\forall}}$.

Corollary 1 $\mathcal{L}_D$ *is at least as expressive as* $\mathcal{L}_{D^{\forall}}$ *(in symbols:* $\mathcal{L}_{D^{\forall}} \preccurlyeq \mathcal{L}_D$*).* ∎

A question remains: is $\mathcal{L}_{D^{\forall}}$ also at least as expressive as $\mathcal{L}_D$ (so the languages are equally expressive), or is $\mathcal{L}_D$ *strictly* more expressive than $\mathcal{L}_{D^{\forall}}$ (so there are situations that $\mathcal{L}_{D^{\forall}}$ cannot tell apart, and yet they can be distinguished by $\mathcal{L}_D$)?

When discussing the relative expressivity of modal languages, it is useful to have a semantic notion guaranteeing that two pointed models cannot be distinguished by a language. A multi-agent version of the standard notion of bisimulation (see, e.g., [3, Section 2.2]) plays this role for the basic multi-agent epistemic language. When the modality for standard distributed knowledge is added (i.e., for $\mathcal{L}_D$), one rather requires the notion of *collective bisimulation* [17], which asks for the conditions of the standard bisimulation to be fulfilled by the *intersection relation* of every group. Still, the results below will show that this notion is not the adequate one for our language $\mathcal{L}_{D^{\forall}}$.

The notion of $\mathcal{L}_{D^{\forall}}$-bisimulation defined below will be shown to be the adequate one for $\mathcal{L}_{D^{\forall}}$: it implies that two pointed models cannot be distinguished by $\mathcal{L}_{D^{\forall}}$ (Proposition 5), and it exists between any image-finite pointed models that cannot be distinguished by the language (Proposition 6).

Definition 6 ($\mathcal{L}_{D^{\forall}}$-Bisimulation) *Let* $\mathcal{M} = \langle W, R, v\rangle$ *and* $\mathcal{M}' = \langle W', R', v'\rangle$ *be two belief models. A non-empty relation* $Z \subseteq \mathrm{D}(\mathcal{M}) \times \mathrm{D}(\mathcal{M}')$ *is a* $\mathcal{L}_{D^{\forall}}$-bisimulation *between* $\mathcal{M}$ *and* $\mathcal{M}'$ *if and only if* Zss' *implies all of the following.*

Atom. *For all* $p \in P$*:* $s \in v(p)$ *if and only if* $s' \in v'(p)$.

Forth. *For all* $G \subseteq A$*, for all* $t \in \mathrm{D}(\mathcal{M})$*: if there is* $H \subseteq^{max}_{s} G$ *such that* $t \in C_H(s)$*, then there are* $H' \subseteq^{max}_{s'} G$ *and* $t' \in C_{H'}(s')$ *such that* Ztt'.[13]

[12] Note: this relies on the fact that G is finite (because A is finite).

[13] Equivalently: for all $G \subseteq A$, for all $t \in \mathrm{D}(\mathcal{M})$, if $sR^{\forall}_{G}t$, then $\exists t'$ such that $s'R^{\forall}_{G}t'$ and Ztt'.

Back. *For all $G \subseteq A$, for all $t' \in \mathrm{D}(\mathcal{M}')$: if there is $H' \subseteq^{max}_{s'} G$ such that $t' \in C_{H'}(s')$, then there are $H \subseteq^{max}_{s} G$ and $t \in C_H(s)$ such that Ztt'.*[14]

Write $Z : \mathcal{M}, s \leftrightarrows_{D^\forall} \mathcal{M}', s'$ when Z is a $\mathcal{L}_{D^\forall}$-bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ with Zss'. Write $\mathcal{M}, s \leftrightarrows_{D^\forall} \mathcal{M}', s$ when there is such a bisimulation Z. ◀

A $\mathcal{L}_{D^\forall}$-bisimulation follows the idea of a standard one. First, $\mathcal{L}_{D^\forall}$-bisimilar worlds should satisfy the same atoms. Then, if one of them has a 'relevant successor' t, the other should also have a 'relevant successor' t' and, moreover, these successors should be $\mathcal{L}_{D^\forall}$-bisimilar. The only difference between a $\mathcal{L}_{D^\forall}$-bisimulation and others in the literature is what 'a relevant successor' means. In a multi-agent standard bisimulation, a 'relevant successor' is any world that can be reached through the relation R_i, for some agent $i \in A$. In a collective bisimulation, a 'relevant successor' is any world that can be reached through the *intersection of the relations* of the individuals in G, for some group $G \subseteq A$. In the just defined $\mathcal{L}_{D^\forall}$ bisimulation, a 'relevant successor' is any world that belongs to the conjecture set of some maximally consistent subgroup of G, for some non-empty set of agents $G \subseteq A$.[15] As it is shown below, this definition guarantees that every world in W that is relevant for cautious distributed belief in $(\mathcal{M}, s)$ has a 'matching' world in W' that is relevant for cautious distributed belief in $(\mathcal{M}', s')$ (and vice versa). (For an example of $\mathcal{L}_{D^\forall}$-bisimilar models see the proof of fact 2 below.)

Definition 7 ($\mathcal{L}_{D^\forall}$-equivalence) *Two pointed models $\mathcal{M}, s$ and $\mathcal{M}', s'$ are $\mathcal{L}_{D^\forall}$-equivalent (notation: $\mathcal{M}, s \leftrightsquigarrow_{D^\forall} \mathcal{M}', s'$) if and only if, for every $\varphi \in \mathcal{L}_{D^\forall}$,*

$$\mathcal{M}, s \vDash \varphi \quad \textit{if and only if} \quad \mathcal{M}', s' \vDash \varphi.$$

When the models are clear from context, we will write simply $s \leftrightsquigarrow_{D^\forall} s'$. ◀

Proposition 5 ($\mathcal{L}_{D^\forall}$-Bisimilarity implies $\mathcal{L}_{D^\forall}$-equivalence) *Let $\mathcal{M}, s$ and $\mathcal{M}', s'$ be pointed belief models. Then,*

$$\mathcal{M}, s \leftrightarrows_{D^\forall} \mathcal{M}', s' \quad \textit{implies} \quad \mathcal{M}, s \leftrightsquigarrow_{D^\forall} \mathcal{M}', s'. \qquad \blacksquare$$

Proof. *First, pull out the universal quantification over formulas hidden in $\leftrightsquigarrow_{D^\forall}$, so the statement becomes* "for every formula in $\mathcal{L}_{D^\forall}$: if two pointed models are $D^\forall$-bisimilar, then they agree on the formula's truth-value"*. Now, proceed by structural induction on formulas in $\mathcal{L}_{D^\forall}$. The case for atomic propositions follows from the* ***atom*** *clause, and those for Boolean operators (in our case, $\neg$ and $\wedge$) follow from their respective inductive hypotheses.*

For formulas expressing cautious distributed belief, work by contraposition. ($\Rightarrow$) Suppose $\mathcal{M}', s' \nvDash D^\forall_G \varphi$. Then, there are $H' \subseteq^{max}_{s'} G$ and $t' \in C_{H'}(s')$

[14] Equivalently: for all $G \subseteq A$, for all $t' \in \mathrm{D}(\mathcal{M}')$, if $s' R^\forall_G t'$, then $\exists t$ such that $s R^\forall_G t$ and Ztt'.

[15] Note then that, while a collective bisimulation requires that a group is inconsistent at any world bisimilar to one at which the group is inconsistent, this not the case for a $\mathcal{L}_{D^\forall}$-bisimulation. The models in the proof of Fact 2 below show this.

such that $\mathcal{M}', t' \not\models \varphi$*. But* $\mathcal{M}, s \rightleftharpoons_{D^\forall} \mathcal{M}', s'$ *so, by the* ***back*** *clause, there are* $H \subseteq_s^{max} G$ *and* $t \in C_H(s)$ *such that* $\mathcal{M}, t \rightleftharpoons_{D^\forall} \mathcal{M}', t'$*. By inductive hypothesis, the latter implies* $\mathcal{M}, t \leftrightsquigarrow_{D^\forall} \mathcal{M}', t'$*, so from the earlier* $\mathcal{M}', t' \not\models \varphi$ *it follows that* $\mathcal{M}, t \not\models \varphi$*. Thus,* $\mathcal{M}, s \not\models D_G^\forall \varphi$*. (*$\Leftarrow$*) Similar, using the* ***forth*** *clause instead.* ∎

A weakened version of the converse holds: if two *image-finite* pointed belief models are $\mathcal{L}_{D^\forall}$-equivalent, then there is a $\mathcal{L}_{D^\forall}$-bisimulation between them.

Proposition 6 ($\mathcal{L}_{D^\forall}$-Equivalence implies $\mathcal{L}_{D^\forall}$-bisimilarity) *Let* $\mathcal{M}, s$ *and* $\mathcal{M}', s'$ *be* image-finite *pointed belief models.*[16] *Then,*

$$\mathcal{M}, s \leftrightsquigarrow_{D^\forall} \mathcal{M}', s' \quad \textit{implies} \quad \mathcal{M}, s \rightleftharpoons_{D^\forall} \mathcal{M}', s'.$$

Proof. It will be shown that $\leftrightsquigarrow_{D^\forall}$ *is in fact a* $\mathcal{L}_{D^\forall}$*-bisimulation. To do this, take any* s *and* s' *such that* $s \leftrightsquigarrow_{D^\forall} s'$*; it will be shown that the three clauses of Definition 6 are satisfied.*

Atom. *It is clear that* s *and* s' *satisfy the same atomic propositions.*

Forth. *Take any* $\varnothing \subset G \subseteq A$*; suppose there are* $H \subseteq_s^{max} G$ *and* $t \in C_H(s)$*. For the sake of a contradiction, suppose there are no* $H' \subseteq_{s'}^{max} G$ *and* $t' \in C_{H'}(s')$ *such that* $t \leftrightsquigarrow_{D^\forall} t'$*; in other words, suppose that every* $H' \subseteq_{s'}^{max} G$ *and* $t' \in C_{H'}(s')$ *are such that* $t \not\leftrightsquigarrow_{D^\forall} t'$*. This means that if* $t'_i \in C_G^\forall(s')$ *then* $t \not\leftrightsquigarrow_{D^\forall} t'_i$*: for every world* $t'_i \in C_G^\forall(s')$ *there is* $\psi_i \in \mathcal{L}_{D^\forall}$ *such that* $\mathcal{M}, t \not\models \psi_i$ *and* $\mathcal{M}', t'_i \models \psi_i$*.*
Now note that $C_G^\forall(s')$ *is non-empty and finite.*[17] *Thus,* $\psi := \bigvee_{t'_i \in C_G^\forall(s')} \psi_i$ *is a non-contradictory formula (as* $C_G^\forall(s')$ *is non-empty) in* $\mathcal{L}_{D^\forall}$ *(as* $C_G^\forall(s')$ *is finite). Hence,* $\mathcal{M}, t \not\models \psi$ *and yet* $\mathcal{M}', t'_i \models \psi$ *for every* $t'_i \in C_G^\forall(s')$*. Since* $H \subseteq_s^{max} G$ *and* $t \in C_H(s)$*, the former implies* $\mathcal{M}, s \not\models D_G^\forall \psi$*; nevertheless, the latter implies* $\mathcal{M}', s' \models D_G^\forall \psi$*. This contradicts the original assumption* $s \not\leftrightsquigarrow_{D^\forall} s'$*. Therefore, there is some* $H' \subseteq_{s'}^{max} G$ *and some* $t' \in C_{H'}(s')$ *such that* $t \leftrightsquigarrow_{D^\forall} t'$*.*

Back. *Analogous to the previous clause.* ∎

We have now enough tools to answer the question above.

Fact 2 $\mathcal{L}_{D^\forall}$ *is* not *at least as expressive as* $\mathcal{L}_D$ *(in symbols:* $\mathcal{L}_D \not\preccurlyeq \mathcal{L}_{D^\forall}$*).*

Proof. Consider the belief models shown below.

[16] A belief model $\mathcal{M}$ is image-finite iff $C_a(s)$ is finite for every $s \in \mathrm{D}(\mathcal{M})$ and every $a \in A$ (equivalently, iff $C_G(s)$ is finite for every $s \in \mathrm{D}(\mathcal{M})$ and every $G \subseteq A$).

[17] It is non-empty because, from $H \subseteq_s^{max} G$ and $t \in C_H(s)$, it follows that $t \in C_a(s)$ for some $a \in H \subseteq G$, and thus $\mathcal{M}, s \models \neg D_a^\forall \bot$. But $s \leftrightsquigarrow_{D^\forall} s'$, so $\mathcal{M}', s' \models \neg D_a^\forall \bot$, so a is consistent at s' in $\mathcal{M}'$. Then, since a is in G, there should be an $H' \subseteq_{s'}^{max} G$ with $a \in H'$. But, once again, a is consistent, so $C_{H'}(s') \neq \varnothing$ and thus $C_G^\forall(s') \neq \varnothing$. It is finite because the models are image-finite.

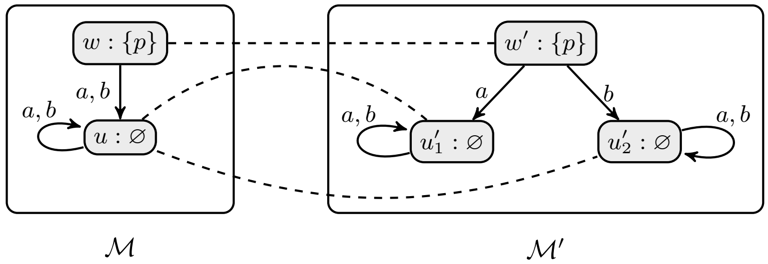

■

Use $MC_G(s)$ to denote all subgroups of G that are maximally consistent at s. The dashed edges define a bisimulation between $\mathcal{M}, w$ and $\mathcal{M}', w'$. Indeed,

- $(\boldsymbol{w, w'})$. *The **atom** clause is immediate. Now **forth**. For $G = \{a\}$, note that $MC_{\{a\}}(w) = \{\{a\}\}$ and thus $C^{\forall}_{\{a\}}(w) = \{u\}$. But then $MC_{\{a\}}(w') = \{\{a\}\}$ and thus $C^{\forall}_{\{a\}}(w') = \{u'_1\}$; moreover, Zuu'_1. The case for $G = \{b\}$ is analogous. For $G = \{a, b\}$, note that $MC_{\{a,b\}}(w) = \{\{a, b\}\}$ and thus $C^{\forall}_{\{a,b\}}(w) = \{u\}$. But then $MC_{\{a,b\}}(w') = \{\{a\}, \{b\}\}$ and thus $C^{\forall}_{\{a,b\}}(w') = \{u'_1, u'_2\}$; moreover, Zuu'_1 and Zuu'_2. The **back** clause follows a similar pattern.*
- $(\boldsymbol{u, u'_1})$. *The **atom** clause is immediate. Consider **forth**. For $G = \{a\}$, note that $MC_{\{a\}}(u) = \{\{a\}\}$ and thus $C^{\forall}_{\{a\}}(u) = \{u\}$. But then $MC_{\{a\}}(u'_1) = \{\{a\}\}$ and thus $C^{\forall}_{\{a\}}(u') = \{u'_1\}$; moreover, Zuu'_1. The case for $G = \{b\}$ is analogous. For $G = \{a, b\}$, note that $MC_{\{a,b\}}(u) = \{\{a, b\}\}$ and thus $C^{\forall}_{\{a,b\}}(u) = \{u\}$. But then $MC_{\{a,b\}}(u'_1) = \{\{a, b\}\}$ and thus $C^{\forall}_{\{a,b\}}(u'_1) = \{u'_1\}$; moreover, Zuu'_1. The **back** clause follows a similar pattern.*
- $(\boldsymbol{u, u'_2})$. *As the previous case.*

Thus, $M, w \leftrightarrows_{D^{\forall}} \mathcal{M}', w'$ and hence, by Proposition 5, $\mathcal{M}, w \leftrightsquigarrow_{D^{\forall}} \mathcal{M}', w'$. However, the pointed models can be distinguished by a formula in $\mathcal{L}_D$, as $\mathcal{M}, w \not\models D_{\{a,b\}} \bot$ and yet $\mathcal{M}', w' \models D_{\{a,b\}} \bot$. Therefore $\mathcal{L}_D \not\preccurlyeq \mathcal{L}_{D^{\forall}}$. ■

Note how the belief models used above are serial, transitive and Euclidean: the kind of models one normally uses for representing a proper notion of belief.

Corollary 2 *$\mathcal{L}_D$ is strictly more expressive than $\mathcal{L}_{D^{\forall}}$ (symbols: $\mathcal{L}_{D^{\forall}} \prec \mathcal{L}_D$).* ■

Thus, $\mathcal{L}_{D^{\forall}}$ can 'see' strictly less than what $\mathcal{L}_D$ can. The proposition below shows that the group inconsistency constant $\asymp_G$ introduced before is exactly what the former needs to 'see' exactly as much as the latter.

Proposition 7 *$\mathcal{L}_{D^{\forall},\asymp}$ and $\mathcal{L}_D$ are equally expressive (symbols: $\mathcal{L}_{D^{\forall},\asymp} \approx \mathcal{L}_D$).*

Proof. Clearly, $\models \asymp_G \leftrightarrow D_G \bot$. Thus, both $\asymp_G$ and $D^{\forall}_G$ are definable in $\mathcal{L}_D$ (for the latter, recall Proposition 4), so $\mathcal{L}_{D^{\forall},\asymp} \preccurlyeq \mathcal{L}_D$.

For proving $\mathcal{L}_D \preccurlyeq \mathcal{L}_{D^{\forall},\asymp}$, it is enough to show that D_G is definable in $\mathcal{L}_{D^{\forall},\asymp}$:

$$\models D_G \varphi \leftrightarrow (\asymp_G \lor D^{\forall}_G \varphi).$$

(⇒) Suppose $\mathcal{M}, s \models D_G\,\varphi$, so every $t \in C_G(s)$ is such that $\mathcal{M}, t \models \varphi$. Assume further that $\mathcal{M}, s \not\models \asymp_G$. Then, $MC_G(s) = \{G\}$ and thus $\mathcal{M}, s \models D^{\forall}_G\,\varphi$. (⇐) Proceed by contraposition: suppose $\mathcal{M}, s \not\models D_G\,\varphi$. Then, there is $t \in C_G(s)$ such that $\mathcal{M}, t \not\models \varphi$. Thus, $C_G(s) \neq \varnothing$, so $MC_G(s) = \{G\}$. From the former, $\mathcal{M}, s \not\models \asymp_G$; from $t \in C_G(s)$ and the latter, $\mathcal{M}, s \not\models D^{\forall}_G\,\varphi$. Thus, $\mathcal{M}, s \not\models \asymp_G \vee D^{\forall}_G\,\varphi$. ■

5 Summary and further work

This paper has introduced the notion of *cautious* distributed belief. While a set of agents G has distributed belief that φ ($D_G\,\varphi$) if and only if φ is true in every world in the conjecture set of the group, the group has cautious distributed belief that φ ($D^{\forall}_G\,\varphi$) if and only if φ is true in every world in the conjecture set of *every* maximally consistent subgroup of G.

The paper has discussed basic properties of $D^{\forall}$, showing, e.g., how it is inconsistent if and only if all agents in the group are inconsistent. Then, the paper has studied whether this group notion inherits properties from the individual notions of the group's members. It has been shown that consistency and truthfulness (technically, seriality and reflexivity) are inherited, and that so are both positive and negative introspection (technically, transitivity and Euclidicity) when the epistemic/doxastic notion is also truthful (technically, reflexive). This is the opposite of what happens with standard distributed belief, which inherits both positive and negative introspection (transitivity, symmetry and Euclidicity) without additional assumptions, and inherits consistency (seriality) only when the individual notions are truthful (reflexive). The final part of the paper has focussed on the relationship between $D^{\forall}_G$ and D_G. It has been show that, while they coincide in reflexive models (i.e., cautions distributed *knowledge* coincides with standard distributed *knowledge*), in general the latter (D) is strictly more expressive than the former ($D^{\forall}$). This difference in expressivity has been proved by providing a notion of structural equivalence that, within image-finite models, characterises modal equivalence w.r.t to $\mathcal{L}_{D^{\forall}}$ (a language extending the propositional one with $D^{\forall}$). Finally, the paper has identified the 'missing piece' that makes a language with $D^{\forall}$ as expressive as one with D.

Among the questions that still need answer, the main ones are an axiom system for the language $\mathcal{L}_{D^{\forall}}$ and a study of its complexity profile. Among the further research lines, the idea of dealing with potential group inconsistencies by looking at maximally consistent subgroups leads to another interesting alternative: a group has *bold* distributed belief that φ (say, $D^{\exists}_G\,\varphi$) if and only if φ is true in every world in the conjecture set of *some* maximally consistent subgroup of G. The quantification pattern of this alternative notion ($\exists\forall$) suggest that, different from $D^{\forall}$, the bold distributed belief operator is not a normal modal operator. Thus, further technical tools will be needed for studying its profile.

References

1. Ågotnes, T., Wáng, Y.N.: Group belief. Journal of Logic and Computation **31**(8), 1959–1978 (2021). https://doi.org/10.1093/logcom/exaa068
2. van Benthem, J., Pacuit, E.: Dynamic logics of evidence-based beliefs. Studia Logica **99**(1), 61–92 (2011). https://doi.org/10.1007/s11225-011-9347-x
3. Blackburn, P., de Rijke, M., Venema, Y.: Modal Logic. Cambridge University Press (2002)
4. de Bruin, B.: Explaining Games: The Epistemic Programme in Game Theory, Synthese Library Series, vol. 346. Springer, Dordrecht (2010)
5. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial Intelligence **77**(2), 321–358 (1995). https://doi.org/10.1016/0004-3702(94)00041-X
6. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Reasoning about knowledge. The MIT Press, Cambridge, Mass. (1995)
7. Halpern, J.Y., Moses, Y.: A guide to the modal logics of knowledge and belief: Preliminary draft. In: Joshi, A.K. (ed.) Proceedings of the 9th International Joint Conference on Artificial Intelligence. Los Angeles, CA, USA, August 1985. pp. 480–490. Morgan Kaufmann (1985), http://ijcai.org/Proceedings/85-1/Papers/094.pdf
8. Halpern, J.Y., Moses, Y.: Knowledge and common knowledge in a distributed environment. Journal of the ACM **37**(3), 549–587 (1990). https://doi.org/10.1145/79147.79161
9. Hendricks, V.F.: Introduction: 8 bridges between mainstream and formal epistemology. Philosophical Studies **128**(1), 1–5 (2006). https://doi.org/10.1007/s11098-005-4068-5
10. Herzig, A., Lorini, E., Perrotin, E., Romero, F., Schwarzentruber, F.: A logic of explicit and implicit distributed belief. In: ECAI 2020 - 24th European Conference on Artificial Intelligence. Frontiers in Artificial Intelligence and Applications, vol. 325, pp. 753–760. IOS Press (2020). https://doi.org/10.3233/FAIA200163
11. Hilpinen, R.: Remarks on personal and impersonal knowledge. Canadian Journal of Philosophy **7**(1), 1–9 (1977). https://doi.org/10.1080/00455091.1977.10716173
12. Hintikka, J.: Knowledge and Belief: An Introduction to the Logic of the Two Notions. Cornell University Press, Ithaca, N.Y. (1962)
13. Lewis, D.: Convention: A Philosophical Study. Harvard University Press, Cambridge, Massachusetts (1969)
14. Meyer, J.J.C., van der Hoek, W.: Epistemic Logic for AI and Computer Science. Cambridge University Press, New York, N.Y., U.S.A. (1995). https://doi.org/10.1017/CBO9780511569852
15. Perea, A.: Epistemic Game Theory: Reasoning and Choice. Cambridge University Press, Cambridge (2012)
16. Reiter, R.: A logic for default reasoning. Artificial Intelligence **13**(1-2), 81–132 (1980). https://doi.org/10.1016/0004-3702(80)90014-4
17. Roelofsen, F.: Bisimulation and distributed knowledge revisited. Journal of Applied Non-Classical Logics **17**, 255–273 (01 2007). https://doi.org/10.3166/jancl.17.255-273