Varieties of Distributed Knowledge

Rustam Galimullin

University of Bergen
Bergen, Norway

Louwe B. Kuijer

University of Liverpool
Liverpool, U.K.

Abstract

Distributed knowledge is one of the better known group knowledge modalities. While its intuitive idea is relatively clear, there is ample room for interpretation of details. We investigate 12 definitions of distributed knowledge that differ from each other in the kinds of information sharing the agents can perform in order to achieve mutual knowledge of a proposition. We then show which kinds of distributed knowledge are equivalent, and which kinds imply each other, i.e., for any two variants τ_1 and τ_2 of distributed knowledge we show whether a proposition φ being distributed knowledge under definition τ_1 implies that φ is distributed knowledge under definition τ_2.

Keywords: Epistemic Logic, Distributed Knowledge

1 Introduction

Epistemic logic (see, e.g., [7,13]) can be used to describe the knowledge of one or more agents. If multiple agents are involved, one can then study various kinds of group knowledge. On the one hand, we may consider types of group knowledge that are stronger than individual knowledge; for example, one may wonder whether a particular proposition φ is known by all members of the group, or even whether φ is so obvious (to the group members) as to be common knowledge. On the other hand, we can also consider a type of group knowledge that is weaker than individual knowledge; even if φ is not currently known by any individual group member, the group might be able to learn φ if they combine their information. For example, perhaps agent a knows that $\varphi \to \psi$ and b knows φ. Neither of them knows ψ, yet if they pool their knowledge they would be able to get to know it.

This latter kind of group knowledge is typically known as *distributed knowledge* (see, e.g., [11,13,7,16,14,4]). Distributed knowledge is a kind of hypothetical knowledge: φ is distributed knowledge among a group G if the members

E-mails: `Rustam.Galimullin@uib.no` and `Louwe.Kuijer@liverpool.ac.uk`

of G could, if they combined their knowledge, learn that φ is true. Note that we do not require the agents to actually combine their knowledge in this way, it suffices that they *could* do so and learn φ.

While the general idea of distributed knowledge is reasonably clear, there is no consensus about how to formally define it. Broadly speaking, there have historically been two main approaches.We will refer to these approaches as the *intersection* approach and the *full communication* approach, with the latter term being derived from [16]. We should note, however, that these are not standardised terms. In fact, even "distributed knowledge" is not fully standard, with other terms such as "group knowledge", "collective knowledge" and "implicit knowledge" also being used.

In both approaches, the distributed knowledge of a group G of agents depends on what information the group members possess. In epistemic logic, the information state of an agent a is generally represented by an accessibility relation $\sim_a$, and a knows a formula φ, denoted $\Box_a\varphi$, in world s if and only if φ is true in every world s' such that $s \sim_a s'$.

The intersection approach is the most common one, and is used in [11,13,7,16,10,14,4], among many others. Here, φ is distributed knowledge in world s if φ is true in every world s' such that $s \sim_G s'$, where $\sim_G = \bigcap_{a\in G} \sim_a$. The intuition behind this approach is that the group G is collectively capable of distinguishing between s and s' if any of its members can.

The full communication approach is less common, but still used in many places, including [12,16,10,14]. In this approach, φ is distributed knowledge among G if and only if the set of formulas known by any of the agents entails φ, i.e., if $\{\psi \in \mathcal{L}_0 \mid \exists a \in G : s \models \Box_a\psi\} \models \varphi$. In order to avoid circularity we do have to be careful to specify that the known formulas ψ must not reference distributed knowledge, i.e. they must be from the basic epistemic logic $\mathcal{L}_0$.

Observant readers may notice that there is significant overlap between the list of papers using the intersection approach and those using the full communication approach. This is because one of the topics studied has been the relation between the intersection and full communication versions of distributed knowledge. The outcomes of this comparison are that (i) if φ is full communication distributed knowledge then it is also intersection distributed knowledge, (ii) φ can be intersection intersection distributed knowledge without being full communication distributed knowledge, and (iii) on certain types of models, the two kinds of distributed knowledge are equivalent.

The reason the two approaches to distributed knowledge persist side by side, albeit with the intersection approach being more popular, is that they appeal to different intuitions. Specifically, the issue is whether agents share information in a way that can be expressed in epistemic logic. Suppose that a considers a world s_1 possible but s_2 impossible, while b considers s_1 to be impossible and s_2 to be possible, but that s_1 and s_2 are not distinguishable by any formula of epistemic logic. Can a and b, when working together, discover that neither s_1 nor s_2 is possible?

The full communication approach says "no, they cannot exclude s_1 and

s_2". After all, while a does not consider s_2 to be possible, there is nothing they can say to b that would communicate this impossibility, and b is likewise incapable of communicating the impossibility of s_1. The intersection approach, on the other hand, says "yes, they can exclude s_1 and s_2". Even if neither of them can express the difference (in epistemic logic, at least), a knows that s_2 is impossible and b knows that s_1 is impossible, so together they know both are impossible. Perhaps they communicate this impossibility to each other in a language other than epistemic logic, such as first-order logic. Perhaps they simply point at the worlds they consider impossible. Perhaps they perform a Vulcan mind-meld, or somehow merge their databases or neural networks. What matters, to the intersection approach, is not how the agents share their information, but only that the agents possess the required information.

In this paper, we will not try to settle the debate in favour of one variant. On the contrary, we will introduce several further variants of distributed knowledge. This is because, in addition to the form of information sharing (formulas or mind-meld) which makes the difference between the intersection and full communication variants, there are several more questions one can ask about how distributed knowledge is established.

How much information do the agents share? Do they share information simultaneously, or is there an order? Are all agents required to know that φ is true after the knowledge sharing, or does it suffice if one agent knows φ?

Different answers to these questions may lead to different notions of distributed knowledge. That is not to say that all possible combinations lead to different kinds of distributed knowledge. For example, suppose that each agent shares a single proposition known to that agent. Then it does not matter whether the agents share simultaneously or in order (Proposition 4.4), and if one of the agents can learn φ then all of them can learn it (Proposition 4.3). But in other cases, the difference does matter.

We will consider 12 possible definitions of distributed knowledge. One of these, which we label $(\cap, \epsilon, \epsilon, \forall)$, is the intersection definition of distributed knowledge. None of our definitions is exactly the same as the full communication definition, but the variant that we label $(\mathcal{L}_0, \odot, \Uparrow, \forall)$ is equivalent to full communication (Proposition 3.2). [1] To the best of our knowledge, other definitions of distributed knowledge in our taxonomy have not been considered in the literature before.

After introducing the basic technical definitions in Section 2, we will define all variants of distributed knowledge in Section 3 and compare them to the existing approaches. Then we compare the introduced definitions of distributed knowledge to each other in Section 4. Finally, we conclude and outline the directions for further research in Section 5.

[1] Furthermore, several other variants are equivalent to $(\mathcal{L}_0, \odot, \Uparrow, \forall)$, and therefore, by transitivity, also equivalent to full communication.

2 Basic Definitions

Each of the variants of distributed knowledge that we consider will use the same language, which is basic epistemic logic with an additional operator D_G, indicating distributed knowledge among group G.

Definition 2.1 Let $\mathcal{P}$ be a countable set of propositional atoms and $\mathcal{A}$ a finite set of agents. The language $\mathcal{L}$ is given by the following normal form

$$\mathcal{L} \ni \varphi ::= p \mid \neg\varphi \mid (\varphi \vee \varphi) \mid \Box_a\varphi \mid D_G\varphi$$

where $p \in \mathcal{P}$, $a \in \mathcal{A}$ and $\emptyset \neq G \subseteq \mathcal{A}$. We denote the fragment of $\mathcal{L}$ that does not contain D_G by $\mathcal{L}_0$. We omit parentheses where this should not cause confusion.

We use $\wedge, \rightarrow, \leftrightarrow$ and $\Diamond_a$ in the usual way as abbreviations. Similarly, we use $\bigwedge$ and $\bigvee$ for n-ary conjunction and disjunction, respectively.

Because we are describing (distributed) *knowledge*, we will use S5 models. We should stress, however, that our results also hold for K models.

Definition 2.2 A *model* $\mathcal{M}$ is a tuple $(S, \sim, V)$, where S is a non-empty set of *worlds*, $\sim: \mathcal{A} \rightarrow 2^{S\times S}$ assigns to each agent $a \in \mathcal{A}$ an equivalence relation $\sim_a \subseteq S \times S$ and $V : \mathcal{P} \rightarrow 2^S$ is the valuation function. If necessary, we will refer to the elements of the tuple as $S^{\mathcal{M}}$, $\sim^{\mathcal{M}}$, and $V^{\mathcal{M}}$. A *pointed model* is a pair $\mathcal{M}, s$ where s is a world of $\mathcal{M}$.

All operators other than D_G are given their normal semantics. The semantics of $\mathcal{L}_0$ is therefore as follows.

Definition 2.3 Let $\mathcal{M} = (S, \sim, V)$ be a model and $s \in S$. Then

$\mathcal{M}, s \models p$	iff	$s \in V(p)$
$\mathcal{M}, s \models \neg\varphi$	iff	$\mathcal{M}, s \not\models \varphi$
$\mathcal{M}, s \models \varphi \vee \psi$	iff	$\mathcal{M}, s \models \varphi$ or $\mathcal{M}, s \models \psi$
$\mathcal{M}, s \models \Box_a\varphi$	iff	$\mathcal{M}, t \models \varphi$ for all t such that $s \sim_a t$.

For D_G, the semantics will depend on the type of distributed knowledge under consideration, which we discuss in the next section.

In several of the proofs throughout this paper, we will make use of the concept of Q-bisimilarity, where $Q \subseteq \mathcal{P}$.

Definition 2.4 Let $Q \subseteq \mathcal{P}$, and $\mathcal{M} = (S^{\mathcal{M}}, \sim^{\mathcal{M}}, V^{\mathcal{M}})$ and $\mathcal{N} = (S^{\mathcal{N}}, \sim^{\mathcal{N}}, V^{\mathcal{N}})$ be models. We say that $\mathcal{M}$ and $\mathcal{N}$ are *Q-bisimilar* (denoted $\mathcal{M} \approx_Q \mathcal{N}$) if there is a non-empty relation $B \subseteq S^{\mathcal{M}} \times S^{\mathcal{N}}$, called *$Q$-bisimulation*, such that for all $B(s,t)$, the following conditions are satisfied:

Atoms for all $p \in Q$: $s \in V^{\mathcal{M}}(p)$ if and only if $t \in V^{\mathcal{N}}(p)$,

Forth for all $a \in \mathcal{A}$ and $u \in S^{\mathcal{M}}$ such that $s \sim_a^{\mathcal{M}} u$, there is a $v \in S^{\mathcal{N}}$ such that $t \sim_a^{\mathcal{N}} v$ and $B(u,v)$,

Back for all $a \in \mathcal{A}$ and $v \in S^{\mathcal{N}}$ such that $t \sim_a^{\mathcal{N}} v$, there is a $u \in S^{\mathcal{M}}$ such that $s \sim_a^{\mathcal{M}} u$ and $B(u,v)$.

We say that $\mathcal{M}, s$ and $\mathcal{N}, t$ are Q-bisimilar and denote this by $\mathcal{M}, s \approx_Q \mathcal{N}, t$ if there is a Q-bisimulation linking worlds s and t.

In the paper, we will make use of the classic result that bisimilar models satisfy the same formulas of epistemic logic.

Theorem 2.5 *Given $\mathcal{M}, s$ and $\mathcal{N}, t$, if $\mathcal{M}, s \approx_Q \mathcal{N}, t$, then for all $\varphi \in \mathcal{L}_0$ that include atoms only from Q, we have that $\mathcal{M}, s \models \varphi$ if and only if $\mathcal{N}, t \models \varphi$.*

3 Varieties of distributed knowledge

In Section 1 we mentioned a number of questions regarding the exact workings of distributed knowledge. Here we discuss these questions in more detail, and use the potential answers to define types of distributed knowledge. Before we get into these details, however, we should remark on one aspect of distributed knowledge that will hold for every variant, namely that distributed knowledge is backward-looking.

In both the intersection and full communication definitions of distributed knowledge, a proposition φ is distributed knowledge among G if, by combining their knowledge, G can discover that φ *was* true *before* they combined their knowledge.[2] In this paper we also follow this tradition. The past tense is important because φ may contain claims that certain group members are ignorant of some fact, and this ignorance can be broken when agents in G share their knowledge.

For example, let φ be the famous Moore-sentence $p \wedge \neg\Box_a p$, i.e., "p is true but a does not know that p is true." This sentence cannot be known by agent a, yet it can be distributed knowledge between a and b. Perhaps a knows that $\neg\Box_a p$ while b knows that p. When a and b combine their knowledge, they will learn that $p \wedge \neg\Box_a p$ used to be true, but that very same communication will render the formula false, since a will learn that p is true.

Because of this backward looking nature, distributed knowledge is a *static* operator, as opposed to the *dynamic* operators from dynamic epistemic logic (DEL) [17]. A dynamic take on distributed knowledge is also possible, and would likely correspond to what agents may learn through communication with each other[3]. Some of the known approaches include a single agent sharing all her information with everyone [5], a group of agents sharing everything they know among themselves [4,6], topic-based communication within a group of agents [9], and various forms of public communication by agents and their effects [2,1,3]. While such dynamic treatment of distributed knowledge is interesting, it is outside the scope of this contribution.

We now continue with a detailed discussion of each of our questions regarding the meaning of distributed knowledge.

---

[2] See [4] for a more thorough discussion of this aspect of distributed knowledge.

[3] As with static distributed knowledge, we need to account for many small but important implementation decisions, for example the extent to which agents that are not part of the group will be aware of the discussion among the group members.

Forms of information The first important question is the form of information shared by the group members in their attempt to establish knowledge of a proposition. We consider two answers to this question. Firstly, agents may be able to share *formulas of $\mathcal{L}_0$ that they know.* So if $\Box_a\psi$ holds, for some $\psi \in \mathcal{L}_0$ and $a \in G$, then in a group discussion among G, agent a can contribute ψ. The restriction to formulas that agents actually know arises from the intuition that distributed knowledge is about combining individual *knowledge* rather than arbitrary formulas.

The alternative is that the agents may be sharing information in a way that is either entirely non-lingual, or at least phrased in a language stronger than $\mathcal{L}_0$. Importantly, such information sharing is not bound to respect bisimilarity in the sense of Theorem 2.5 (see, e.g., [14]).

Note that we restrict ψ to $\mathcal{L}_0$, so the formulas that the agents can share in their deliberation cannot include the D_G operator. This is required in order to avoid vicious circularity; if we allow the D_G operator to be used during the deliberations, there are situations where agents can learn φ if they combine their knowledge, but only if that knowledge includes the fact that φ is distributed knowledge. Hence φ would be distributed knowledge if and only if... φ is distributed knowledge. As a result of this circularity, the semantics would become underdetermined, i.e., there would be pointed models where both $D_G\varphi$ and $\neg D_G\varphi$ are consistent with the semantics (see Appendix A).

Recall from the introduction that the intersection definition of distributed knowledge assumes the non-lingual answer to this question, whereas the full communication definition assumes the information being communicated is in the form of $\mathcal{L}_0$ formulas.

Amount of information The next question is how much information the agents share. In principle, any measure could be used here. For example, one could imagine a situation where each agent has, say, 5 seconds to contribute their share. Or perhaps agents are limited to statements of a given maximum complexity.

Here, however, we will restrict ourselves to a coarser distinction: agents will be able to share either a single formula, or an infinite set of formulas. Note that since we are limiting only the amount of formulas, not their complexity, it would not make sense to restrict to a given finite number of formulas, since any finite number of formulas can be combined into one using conjunctions.

Note that this distinction only makes sense if information is shared as $\mathcal{L}_0$ formulas; if information is shared non-linguistically we do not have a sensible measure of the amount of information shared.

Order and turn-taking Another consideration is whether the agents share all information simultaneously, or in some order (with agents taking a single turn if they share one formula, or multiple turns if they share a set of formulas). This distinction is relevant because agents can only share formulas that they know; if the agents share their information in some order, then the later agents may be able to contribute some formulas that they did not know initially but that they have come to know based on the information provided by the agents before them. Again, this distinction only works if agents share their information in the form of formulas.

Collective or individual success Finally, we can distinguish between a type of distributed knowledge where all group members need to learn the truth of a formula and a type where only one group member needs to learn it. In other words, if there is a possible communication between a and b that would result in a knowing φ while b remains ignorant of it, would φ be distributed knowledge?

Equivalence among variants Not all of the combinations of answers to questions from the previous section make sense. Still, the answers allow us to define 12 variations of distributed knowledge.

We should stress, however, that not all these variations are truly different. For example, suppose that, in their communication, every agent shares a single formula, and they do so simultaneously. Then it is possible for the agents to communicate in such a way that a single agent learns φ if and only if it is possible for them to communicate in a way where all agents learn φ (see Proposition 4.2).

In fact, our main contributions in this paper are (1) formal definitions of the various kinds of distributed knowledge, and (2) the results on which of the variants are equivalent to each other.

3.1 Semantics for distributed knowledge

In our taxonomy, a type of distributed knowledge can be identified by the form of information shared, the amount of information, whether there is an order, and how many agents need to learn the target formula. A type τ is therefore a tuple $\tau = (f, a, o, q)$, where $f \in \{\cap, \mathcal{L}_0\}$ indicates whether the information is presented in the form of formulas ($\mathcal{L}_0$) or not ($\cap$), $a \in \{\odot, \bullet, \epsilon\}$ indicates whether a single formula is shared ($\odot$) or a set of formulas ($\bullet$), while $a = \epsilon$ is used for the case where $f = \cap$ and therefore no formulas are shared at all. The parameter $o \in \{\Uparrow, \omega, \Omega, \epsilon\}$ indicates whether the agents share their knowledge simultaneously ($\Uparrow$), in a sequence with a length α bounded by the first infinite ordinal (ω), in a sequence that can have any ordinal α as its length (Ω), or whether $f = \cap$ and therefore the question of an order doesn't make sense (ϵ). Finally, $q \in \{\exists, \forall\}$ indicates whether at least one agent must learn φ ($\exists$) or all of them must learn it ($\forall$).

We use $f(\tau)$, $a(\tau)$, $o(\tau)$ and $q(\tau)$ to denote the values of f, a, o and q, respectively, in τ. Each type τ of distributed knowledge induces semantics for the language $\mathcal{L}$, which we denote by $\models_\tau$.

The main idea of each of the semantics is that φ is distributed knowledge among G if there is some way for G to share information among themselves that would result in them learning the truth of φ. Communication by G will change the current information state, which is encoded by the set $\sim$ of relations, into a new information state $\sim'$. Often there are different things that G could communicate, and each such possible communication will lead to a new information state. Hence we will, in general, need to consider not one new information state $\sim'$ but a set of such information states.

We will denote the set of information states that can be reached by group G, when discussing in world s, using the communication type τ, as $\mathcal{R}_{G,s,\tau}$.

The semantics for the distributed knowledge operator are then given by $\mathcal{M}, s \models_\tau D_G\varphi$ iff

$$\exists \sim' \in \mathcal{R}_{G,s,\tau} \exists a \in G : \mathcal{M}, t \models \varphi \text{ for all } t \text{ such that } s \sim'_a t$$

if $q(\tau) = \exists$, or $\mathcal{M}, s \models_\tau D_G\varphi$ iff

$$\exists \sim' \in \mathcal{R}_{G,s,\tau} \forall a \in G : \mathcal{M}, t \models \varphi \text{ for all } t \text{ such that } s \sim'_a t$$

if $q(\tau) = \forall$.

We should note that, while the agents are generally not *required* to share as much information as they can, distributed knowledge is about whether the agents are able to achieve knowledge of φ. The more information is shared, the more likely it is that the agents will learn φ.[4] In particular, if there is a unique "maximal communication", it suffices to consider only that communication.

For example, in the intersection definition of distributed knowledge, agents are capable of explaining, in a non-linguistic way, exactly which worlds they consider possible. Conceptually, it seems reasonable that agents could, instead of communicating their exact set of possible worlds, communicate a superset of it. We need not consider this possibility, however, since sharing the exact set of worlds they consider possible is the optimal strategy. In this case, it therefore suffices to consider the singleton set $\mathcal{R}_{G,s,\tau} = \{\sim'\}$ where, for every $a \in G$, $\sim'_a = \bigcap_{b\in G} \sim_b$.

In other cases there may be no single most informative communication, so we cannot restrict ourselves to a single information state in this way. For example, if every agent can communicate a single formula that is known to them, there is not, in general, a single most informative formula for them to state. For every formula $\psi_G = \bigwedge_{a\in G} \psi_a$ with the property that $\mathcal{M}, s \models \Box_a\psi_a$ for every a, we therefore need to consider the information state $\sim^{\psi_G}$.

Based on the considerations from the previous section, we can define the following 12 variants of distributed knowledge (see Figure 1 for the full list of variants and their relative strength). Recall, however, that some of these variants are equivalent to one another.

Non-linguistic sharing Suppose information is shared non-linguistically. Then our method of restricting the amount of information shared is inapplicable. Furthermore, because we do not know how information is shared we also cannot speak of an ordering in which information is presented.

The only further distinction that is available is whether one agent needs to learn the formula or all of them do. We therefore need to consider the variants $\tau = (\cap, \epsilon, \epsilon, \exists)$ and $\tau = (\cap, \epsilon, \epsilon, \forall)$, respectively.

---

[4] This does rely on the fact that we are looking at *static* communication, i.e. φ is distributed knowledge if the agents can learn that φ *used to be* true. In *dynamic* communication, where agents are trying to learn that φ *is* true, sharing as much information as possible may not be an optimal strategy since φ can contain ignorance conditions that become false when more information is shared (e.g. the Moore formula).

As stated above, for either of these cases we have $\mathcal{R}_{G,s,\alpha} = \{\sim'\}$ where $\sim'_a = \bigcap_{b\in G} \sim_b$. Note, also, that at this point it is already easy to see that $(\cap, \epsilon, \epsilon, \exists)$ and $(\cap, \epsilon, \epsilon, \forall)$ are equivalent. This is because all agents end up with the same accessibility relation, so if one of them learns φ then they all do.

Proposition 3.1 *$\mathcal{M}, s \models_{(\cap,\epsilon,\epsilon,\exists)} D_G\varphi$ if and only if $\mathcal{M}, s \models_{(\cap,\epsilon,\epsilon,\forall)} D_G\varphi$.*

Simultaneous sharing of formulas Suppose the information sharing happens by the agents stating one or more formulas each, and that this happens simultaneously. So we are considering the variants $\tau \in \{(\mathcal{L}_0, \odot, \Uparrow, \exists), (\mathcal{L}_0, \odot, \Uparrow, \forall), (\mathcal{L}_0, \bullet, \Uparrow, \exists), (\mathcal{L}_0, \bullet, \Uparrow, \forall)\}$, where the first two assume that each agent contributes a single formula while the last two let each agent contribute a (potentially infinite) set of formulas.

We specify which information states the agents can achieve by sharing information in two steps. First, we specify all of the ways the agents could share information. In effect, this acts as a set of indices used to identify the various outcome information states. Then, for each index we specify what that outcome information state is.

The important condition on information sharing is that each agent must contribute one or more formulas that they know.[5] Hence if $\tau = (\mathcal{L}_0, \odot, \Uparrow, \exists)$ or $\tau = (\mathcal{L}_0, \odot, \Uparrow, \forall)$ then

$$\mathcal{R}_{G,s,\tau} = \{\sim^{\{\psi_a | a\in G\}} \mid \forall a \in G : \psi_a \in \mathcal{L}_0 \text{ and } \mathcal{M}, s \models \Box_a \psi_a\}$$

and

$$\sim_b^{\{\psi_a|a\in G\}} = \{(x, y) \in S \times S \mid (x, y) \in \sim_b \text{ and } \mathcal{M}, y \models \bigwedge_{a\in G} \psi_a\}.$$

If each agent shares a set of formulas, i.e. if $\tau = (\mathcal{L}_0, \bullet, \Uparrow, \exists)$ or $\tau = (\mathcal{L}_0, \bullet, \Uparrow, \forall)$, then they need to know each of the formulas they provide, so

$$\mathcal{R}_{G,s,\tau} = \{\sim^{\{\Psi_a|a\in G\}} \mid \forall a \in G : \Psi_a \subseteq \mathcal{L}_0 \text{ and } \forall \psi \in \Psi_a : \mathcal{M}, s \models \Box_a \psi\}$$

and

$$\sim_b^{\{\Psi_a|a\in G\}} = \{(x, y) \in S \times S \mid (x, y) \in \sim_b \text{ and } \forall a \forall \psi \in \Psi_a : \mathcal{M}, y \models \psi\}.$$

Taking turns Suppose that, as in the previous case, agents share one or more formulas, but now they do so sequentially. The crucial difference with simultaneous communication is that in the sequential case agents can state a formula that they only know because of the information provided to them by the previous speakers.

Assume, for example, that a knows p and b knows $p \to q$. If a speaks first and tells b that p is true, b can then, when it is their turn to speak, say that q is true, which they only know because a told them p.

[5] Note that we can assume without loss of generality that every agent shares at least one formula because agents can always use the uninformative formula $\top$ that is known by everyone.

If every agent shares exactly one formula, then the sequential sharing of information means every agent takes a single turn, where the later agents can use the information provided by the earlier ones.

If each agent provides a set of formulas, we still need to specify the order among the agents, but this will generally have to be an infinite order. Note that we do not have to assume that this order is "fair", since agents can skip their turn by providing the trivial formula $\top$. What is potentially important, however, is whether the turn-taking is limited to ω rounds (where ω is the first infinite ordinal), or whether any ordinal can be used. This gives us six variants: $\tau \in \{(\mathcal{L}_0, \odot, \omega, \exists), (\mathcal{L}_0, \odot, \omega, \forall), (\mathcal{L}_0, \bullet, \omega, \exists), (\mathcal{L}_0, \bullet, \omega, \forall), (\mathcal{L}_0, \bullet, \Omega, \exists), (\mathcal{L}_0, \bullet, \Omega, \forall)\}$. Observe that we do not consider types $(\mathcal{L}_0, \odot, \Omega, \forall)$ and $(\mathcal{L}_0, \odot, \Omega, \exists)$ since a finite number of agents communicating one formula each will never step on the trans-finite territory.

Before we can define the possible effects of sequential communication, we first need a little bit more notation. We want to consider finite sequences, infinite sequences, and even trans-finite sequences of statements. Therefore, let α be any ordinal. At each ordinal $\delta < \alpha$, one of the agents will state the truth of one formula; let us write $f(\delta)$ for the agent and $g(\delta)$ for the formula. As α can be considered to be identical to the set of all ordinals less than it, this means f and g are functions of type $f : \alpha \to G$ and $g : \alpha \to \mathcal{L}_0$.

Agent $f(\delta)$ needs to know formula $g(\delta)$ at time δ, since otherwise they would not be able to state the truth of the formula. As such, we need to keep track of the information state at each point in the process. Formally, this means that we are interested in the final information state $\sim^{\alpha,f,g}$, but we also need to define $\sim^{\alpha,f,g,\delta}$ for $\delta < \alpha$, which represents the information state immediately after the announcement that takes place at time δ. We do this by defining

$$\sim_a^{\alpha,f,g,0} = \{(x, y) \in S \times S \mid (x, y) \in \sim_a \text{ and } \mathcal{M}, y \models g(0)\}$$

and

$$\sim_a^{\alpha,f,g,\delta} = \{(x, y) \in S \times S \mid (x, y) \in \bigcap_{\epsilon<\delta} \sim_a^{\alpha,f,g,\epsilon} \text{ and } \mathcal{M}, y \models g(\delta)\}$$

for $0 < \delta < \alpha$. Finally, we define $\sim_a^{\alpha,f,g} = \bigcap_{\delta<\alpha} \sim_a^{\alpha,f,g,\delta}$.

Note that $\sim^{\alpha,f,g,\delta}$ is the information state *after* the communication at time δ, and that there is such a communication at every time $\delta < \alpha$. Hence $\sim_a^{\alpha,f,g,0}$ is generally not identical to $\sim_a$, since the latter represents the information state before any communication takes place.

Furthermore, communication only happens at $\delta < \alpha$. Hence, in particular, if $\alpha = \omega$ then communication takes place at every finite time step, but there is no "infinity-th" communication at ω.

Importantly, the definition of $\sim^{\alpha,f,g}$ does not check whether agent $f(\delta)$ actually knows $g(\delta)$ at time δ. So while the definition determines the effect that a given communication sequence would have, it does not determine whether the agents are actually capable of saying the formulas included in the sequence. For this, we need to look at another property that we will refer to as *correctness*. We say that α, f and g are *correct* for group G and world s if

$$\forall s' : \text{ if } (s, s') \in \sim_{f(\delta)} \cap \bigcap_{\epsilon<\delta} \sim^{\alpha,f,g,\epsilon}_{f(\delta)}, \text{ then } \mathcal{M}, s' \models g(\delta).$$

In other words, if s' was accessible originally (for agent $f(\delta)$) and has not been excluded by any of the preceding statements at times $\epsilon < \delta$, then $g(\delta)$ must be true in s'.

Now, we can formally define $\mathcal{R}_{G,s,\tau}$ for the sequential types of communication. If $\tau = (\mathcal{L}_0, \odot, \omega, \exists)$ or $\tau = (\mathcal{L}_0, \odot, \omega, \forall)$, then we require every agent to have exactly one turn, and hence $\alpha = |G|$. Moreover, f is a bijection, so

$$\mathcal{R}_{G,s,\tau} = \{\sim^{|G|,f,g} \mid f \text{ is a bijection and } |G|, f \text{ and } g \text{ are correct for } G, s\}.$$

If $\tau = (\mathcal{L}_0, \bullet, \omega, \exists)$ or $\tau = (\mathcal{L}_0, \bullet, \omega, \forall)$, then we take $\alpha = \omega$, i.e., we allow infinite statements, but there is no "infinity-th statement". We could demand that the turn-taking by the agents is "fair" in some way, but that is pointless; if there is an unfair turn-taking the agents could use for their communication, this can be transformed into a fair one where some agents give the trivial statement $\top$. As such, we get

$$\mathcal{R}_{G,s,\tau} = \{\sim^{\omega,f,g} \mid \omega, f \text{ and } g \text{ are correct for } G, s\}.$$

Finally, we can allow any ordinal number α of statements. If $\tau = (\mathcal{L}_0, \bullet, \Omega, \exists)$ or $\tau = (\mathcal{L}_0, \bullet, \Omega, \forall)$, then

$$\mathcal{R}_{G,s,\tau} = \{\sim^{\alpha,f,g} \mid \alpha, f \text{ and } g \text{ are correct for } G, s\}.$$

Note that all $\mathcal{R}_{G,s,\tau}$'s are sets, and we essentially quantify over all possible sequences of announcements, and, according to the definition of distributed knowledge, it is enough that at least one knowledge state induced by any sequence satisfies φ.

3.2 Connections to the traditional definitions

Our variant $(\cap, \epsilon, \epsilon, \forall)$ of distributed knowledge is, modulo some notation, identical to the traditional definition of distributed knowledge based on intersection. Our variant that most closely matches the full communication definition of distributed knowledge is $(\mathcal{L}_0, \odot, \Uparrow, \forall)$.

In fact, as mentioned in the introduction, $(\mathcal{L}_0, \odot, \Uparrow, \forall)$ is equivalent to the full communication definition, but this is not entirely obvious and therefore requires a short proof.

Proposition 3.2 *We have*

$$\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\forall)} D_G\varphi$$

if and only if

$$\{\psi \in \mathcal{L}_0 \mid \exists a \in G : \mathcal{M}, s \models \Box_a\psi\} \models \varphi.$$

Proof. Suppose $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\forall)} D_G\varphi$. Then there are $\{\psi_b \mid b \in G\}$ such that (1) for all $b \in G$, $\mathcal{M}, s \models \Box_b\psi_b$ and (2) for all $a \in G$ and every s', if $s \sim_a^{\{\psi_b \mid b \in G\}} s'$ then $\mathcal{M}, s' \models \varphi$. Furthermore, $s \sim_a^{\{\psi_b \mid b \in G\}} s'$ holds if and only if $s \sim_a s'$ and $\mathcal{M}, s' \models \bigwedge\{\psi_b \mid b \in G\}$.

This implies that for every s', if $s \sim_a s'$, then $\mathcal{M}, s' \models \bigwedge\{\psi_b \mid b \in G\} \to \varphi$. Furthermore, since $\mathcal{M}, s \models \Box_a\psi_a$, we also have $\mathcal{M}, s' \models \psi_a$. It follows that $\mathcal{M}, s \models \Box_a(\psi_a \wedge (\bigwedge\{\psi_b \mid b \in G\} \to \varphi))$.

Now, note that $\{\psi_a \wedge (\bigwedge\{\psi_b \mid b \in G\} \to \varphi) \mid a \in G\} \models \varphi$. As such, $\{\psi \mid \exists a \in G : \mathcal{M}, s \models \Box_a\psi\} \models \varphi$.

For the other direction, suppose that $\{\psi \mid \exists a \in G : \mathcal{M}, s \models \Box_a\psi\} \models \varphi$. Since epistemic logic is compact, there is a finite $\Psi \subseteq \{\psi \mid \exists a \in G : \mathcal{M}, s \models \Box_a\psi\}$ such that $\Psi \models \varphi$. For every $a \in G$, let $\psi_a = \bigwedge\{\psi \in \Psi \mid \mathcal{M}, s \models \Box_a\psi\}$.

We now have $\mathcal{M}, s \models \Box_a\psi_a$ for every $a \in G$. Furthermore, since $\Psi \models \varphi$, we also have $\mathcal{M}, s' \models \bigwedge\{\psi_a \mid a \in G\} \to \varphi$ for every $s' \in S$ such that $s \sim_a s'$. It follows that $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\forall)} D_G\varphi$. □

Note that the proof critically depends on the compactness of epistemic logic. If we used a non-compact base logic, such as epistemic logic with common knowledge, the equivalence would not hold.

4 Relative Strength

Now that we have formally defined the various types of distributed knowledge that we are interested in, we can investigate their properties. In particular, we are interested in which variants imply each other.

In most cases, it is clear that one variant τ_1 is at least as strong as another variant τ_2, in the sense that $\mathcal{M}, s \models_{\tau_1} D_G\varphi$ implies $\mathcal{M}, s \models_{\tau_2} D_G\varphi$. In particular, it is easy to see that the following hold.

Proposition 4.1

- *For every τ, $\mathcal{M}, s \models_\tau D_G\varphi$ implies $\mathcal{M}, s \models_{(\cap,\epsilon,\epsilon,\forall)} D_G\varphi$.*
- *For every f, a and o, $\mathcal{M}, s \models_{(f,a,o,\forall)} D_G\varphi$ implies $\mathcal{M}, s \models_{(f,a,o,\exists)} D_G\varphi$.*
- *For all o and q, $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,o,q)} D_G\varphi$ implies $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,o,q)} D_G\varphi$.*
- *For all a and q, $\mathcal{M}, s \models_{(\mathcal{L}_0,a,\Uparrow,q)} D_G\varphi$ implies $\mathcal{M}, s \models_{(\mathcal{L}_0,a,\omega,q)} D_G\varphi$. Furthermore, $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\omega,q)} D_G\varphi$ implies $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Omega,q)} D_G\varphi$.*

Furthermore, it follows from [16] that $(\cap, \epsilon, \epsilon, \forall)$ and $(\cap, \epsilon, \epsilon, \exists)$ do not imply $(\mathcal{L}_0, \odot, \Uparrow, \forall)$.This, however, leaves many comparisons open, which we solve here.

4.1 Single formula

First, let us compare variants that use a single formula. In particular, we will show that all the four variants, i.e. $(\mathcal{L}_0, \odot, \Uparrow, \exists)$, $(\mathcal{L}_0, \odot, \Uparrow, \forall)$, $(\mathcal{L}_0, \odot, \omega, \exists)$, and $(\mathcal{L}_0, \odot, \omega, \forall)$, are equivalent

Proposition 4.2 *$\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\exists)} D_G\varphi$ if and only if $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\forall)} D_G\varphi$.*

Proof. We know from Proposition 4.1 that "all" implies "single", so it suffices to show that the reverse also holds. Suppose, therefore, that $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\exists)}$

$D_G\varphi$. Let $\{\psi_a \mid a \in G\}$ be the witnessing formulas that, if communicated among the group, would make one agent, let's call them x, learn that φ is true.

This means that for every s', if $(s, s') \in \sim_x$ and $\mathcal{M}, s' \models \bigwedge_{a\in G} \psi_a$, then $\mathcal{M}, s' \models \varphi$. This implies that for every $(s, s') \in \sim_x$, we have $\mathcal{M}, s' \models \bigwedge_{a\in G} \psi_a \to \varphi$. Furthermore, since x was able to provide the formula ψ_x, we also have $\mathcal{M}, s' \models \psi_x$ for each such s'. Hence $\mathcal{M}, s \models \Box_x(\psi_x \wedge (\bigwedge_{a\in G} \psi_a \to \varphi))$.

Consider then the alternative communication $\{\psi'_a \mid a \in G\}$ where $\psi'_a = \psi_a$ for $a \neq x$ and $\psi'_x = \psi_x \wedge (\bigwedge_{a\in G} \psi_a \to \varphi)$. For every agent $a \in G$ and every s', if $(s, s') \in \sim_a$ and $\mathcal{M}, s' \models \bigwedge_{a\in G} \psi'_a$, we then have, in particular, $\mathcal{M}, s' \models \bigwedge_{a\in G} \psi_a$ and $\mathcal{M}, s' \models \bigwedge_{a\in G} \psi_a \to \varphi$, and hence $\mathcal{M}, s' \models \varphi$. This implies that all agents learn φ, and therefore $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\forall)} D_G\varphi$. □

In effect, the single agent x that learns φ can include hypothetical reasoning in the formula that they provide to the group. Instead of saying "ψ_x is true", they can say "ψ_x is true, and if you were to tell me $\{\psi_a \mid a \in G\}$, then I would learn that φ is true". This suffices for all the other agents to learn the truth of φ, if all $a \neq x$ do indeed provide formulas ψ_a. The same trick can be used in the sequential version.

Proposition 4.3 *$\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\omega,\exists)} D_G\varphi$ if and only if $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\omega,\forall)} D_G\varphi$.*

Proof. Let $(a_1, \cdots, a_n)$ be an ordering of G such that a_i takes their turn before a_j iff $i < j$, and let $(\psi_1, \cdots, \psi_n)$ be the corresponding formulas that witness $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\omega,\exists)} D_G\varphi$, where a_x is the agent that learns φ.

Then a_x already knows, before the communication starts, that $\bigwedge\{\psi_i \mid 1 \leq i \leq n\} \to \varphi$. Furthermore, once it is their turn, they have also learned that ψ_{a_x}. Hence they could instead say $\psi'_{a_x} = \psi_{a_x} \wedge (\bigwedge\{\psi_i \mid 1 \leq i \leq n\} \to \varphi)$, which would result in all agents learning φ. Note that since we work with the static notion of distributed knowledge, rather than a dynamic one, knowledge of agents is monotonic under announcements and they can always announce their respective formulas. □

Additionally, a similar kind of hypothetical reasoning can be used to remove the reliance on sequential communication.

Proposition 4.4 *$\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\omega,\forall)} D_G\varphi$ if and only if $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\forall)} D_G\varphi$.*

Proof. By Proposition 4.1, $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\forall)} D_G\varphi$ implies $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\omega,\forall)} D_G\varphi$. Left to show is the other direction.

Suppose therefore that $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\omega,\forall)} D_G\varphi$, as witnessed by order $(a_1, \cdots, a_n)$ and formulas $(\psi_1, \cdots, \psi_n)$. Then, at stage i, the preceding communications $\{\psi_j \mid j < i\}$ suffice for agent i to learn that ψ_i holds, in the sense that for all s', if $(s, s') \in \sim_{a_i}$ and $\mathcal{M}, s' \models \bigwedge_{j<i} \psi_j$ then $\mathcal{M}, s' \models \psi_i$.

It follows that, before the communication started, a_i already knew $\bigwedge_{j<i} \psi_j \to \psi_i$. So, in the simultaneous version, a_i could provide that formula.

Furthermore, collectively, communicating $\{\psi_i \mid 1 \leq i \leq n\}$ has the same effect as communicating $\{\bigwedge_{j<i} \psi_j \to \psi_i \mid 1 \leq i \leq n\}$. So we also have $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\forall)} D_G\varphi$. □

The equivalence of other pairs of the single formula variants follows immediately from the transitivity of the equivalence relation.

4.2 Sets of formulas

Let us now consider the variants where each agent may provide a set of formulas. As a first step, we will show that none of them imply the single formula variants. For this, it suffices to show that any one of the single formula variants is not implied by the strongest set variant (i.e., the set variant that is the hardest to satisfy, which is $(\mathcal{L}_0, \bullet, \Uparrow, \forall)$).

Proposition 4.5 *$\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Uparrow,\forall)} D_G\varphi$ does not necessarily imply that $\mathcal{M}, s \models_{(\mathcal{L}_0,\odot,\Uparrow,\forall)} D_G\varphi$.*

Proof. (*Sketch; full proof in Appendix B*). Suppose that, for every $i \in \mathbb{N}$, agent a knows whether p_i holds while b knows whether q_i holds. Furthermore, suppose both agents know that r is true if and only if, for every i, $p_i \leftrightarrow q_i$ holds. Consider the case where all p_i and q_i are, in fact, false, so r holds.

Under the $(\mathcal{L}_0, \bullet, \Uparrow, \forall)$ definition, we then have $D_{\{a,b\}}r$. After all, a can tell b that all p_i are false, while b can tell a that all q_i are false. This suffices for both of them to discover that r is true.

Under the $(\mathcal{L}_0, \odot, \Uparrow, \forall)$ definition, however, we have $\neg D_{\{a,b\}}r$. This is because the agents can only learn that r is true if $p_i \leftrightarrow q_i$ for all i, which cannot be expressed in a finite set of formulas. □

Next, let us note that $(\mathcal{L}_0, \bullet, \Uparrow, \exists)$ does not imply $(\mathcal{L}_0, \bullet, \Uparrow, \forall)$.

Proposition 4.6 *$\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Uparrow,\exists)} D_G\varphi$ does not necessarily imply $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Uparrow,\forall)} D_G\varphi$.*

Proof. (*Sketch; full proof in Appendix B*). As in Proposition 4.5, suppose a knows whether p_i is true and b knows whether q_i is true. Now, however, suppose that b knows that whether r is true depends on the parity of the number of indices i such that p_i and q_i differ in value. Specifically, b knows that r is true if that number is even, while a is uncertain whether r holds if the number is even, or if it is odd. Both agents know r is false if p_i and q_i differ infinitely often. As before, suppose that all p_i and q_i happen to be false.

With the $(\mathcal{L}_0, \bullet, \Uparrow, \exists)$ definition of distributed knowledge, we then have $D_{\{a,b\}}r$. This is because, when a tells b that all p_i are false, agent b will learn that there are 0 indices where p_i and q_i differ, so r is true.

Yet r is not distributed knowledge under the $(\mathcal{L}_0, \bullet, \Uparrow, \forall)$ definition of distributed knowledge, since only b can learn that r is true. The reason a can't learn this is that "r is true iff there is an even number of i such that p_i and q_i disagree" cannot be expressed in epistemic logic. Furthermore, while b learns that r is true once the communication is complete, $(\mathcal{L}_0, \bullet, \Uparrow, \forall)$ requires simultaneous communication, so b cannot simply say that r is true. □

What does not make a difference, however, is simultaneous statements or ω-sequential ones.

Proposition 4.7 *$\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\omega,\exists)} D_G\varphi$ iff $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Uparrow,\exists)} D_G\varphi$, and $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\omega,\forall)} D_G\varphi$ iff $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Uparrow,\forall)} D_G\varphi$*

Proof. At every step in the ω-sequential communication, when ψ_i is stated, a finite set $\{\psi_j \mid j < i\}$ preceded it. We can replace sequential announcement of $\{\psi_i \mid i \in \mathbb{N}\}$ by simultaneous announcement of $\{\bigwedge_{j<i} \psi_j \to \psi_i \mid i \in \mathbb{N}\}$. □

The Ω-sequential variant, on the other hand, is strictly weaker than ω-sequential or simultaneous ones.

Proposition 4.8 *$\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Omega,\exists)} D_G\varphi$ does not necessarily imply $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Uparrow,\exists)} D_G\varphi$.*

Proof. (*Sketch; full proof in Appendix B*). Suppose that a knows, for all i and j, whether $p_{i,j}$ and $q_{i,j}$ hold. Furthermore, suppose that the value of x_i depends on the number of j such that $p_{i,j}$ differs from $q_{i,j}$, in a way known to b but not to a and c, and that this dependence cannot be expressed in epistemic logic. (See the full proof in the appendix for one way to create such an inexpressible dependence.) Similarly, y_i depends on the number of j such that $p_{i,j}$ and $q_{i,j}$ differ, in an inexpressible way that is known to c but not to a and b. Finally, all three agents know that z is true iff there is an even number of i such that x_i and y_i differ.

Then if z is true, that is distributed knowledge using the $(\mathcal{L}_0, \bullet, \Omega, \exists)$ definition, since a can tell b and c which $p_{i,j}$ and $q_{i,j}$ hold, at which point they can say which x_i and y_i are true, allowing all of them to determine that z holds.

With the $(\mathcal{L}_0, \bullet, \Uparrow, \exists)$ definition z is not distributed knowledge, however. This is because, in order for any of the three agents to learn that z is true, all three agents need to contribute their information. But b and c cannot initially contribute any non-trivial formulas, since their only private information is the way in which x_i or y_j depends on the values of $p_{i,j}$ and $q_{i,j}$, and this dependence is not expressible in epistemic logic.

It is only after a has informed b and c about the values of every $p_{i,j}$ and $q_{i,j}$ that b and c can apply their knowledge in order to determine the truth of x_i and y_j, respectively, which they can then communicate to the other agents. So a first needs to contribute at least ω formulas before b and c can get involved, which is not possible if $\tau = (\mathcal{L}_0, \bullet, \Uparrow, \exists)$. □

However, if we allow any ordinal number of communication steps, the difference between a single agent learning the formula or all of them doing so disappears.

Proposition 4.9 *$\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Omega,\exists)} D_G\varphi$ if and only if $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Omega,\forall)} D_G\varphi$.*

Proof. Suppose that $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Omega,\exists)} D_G\varphi$. So there is a sequence of communications of length α after which one agent a knows φ. Now, consider the sequence of length $\alpha + 1$ where, in the last step, agent a states that φ is true. This suffices for all agents to learn that φ is true. □

The above suffices to determine the comparative strength of each of the variants we discussed.

5 Discussion

We have considered 12 natural interpretations of the idea behind distributed knowledge. For these interpretations, we have analysed which ones are equivalent to each other, and which ones are different. The complete landscape of distributed knowledge is shown in Figure 1.

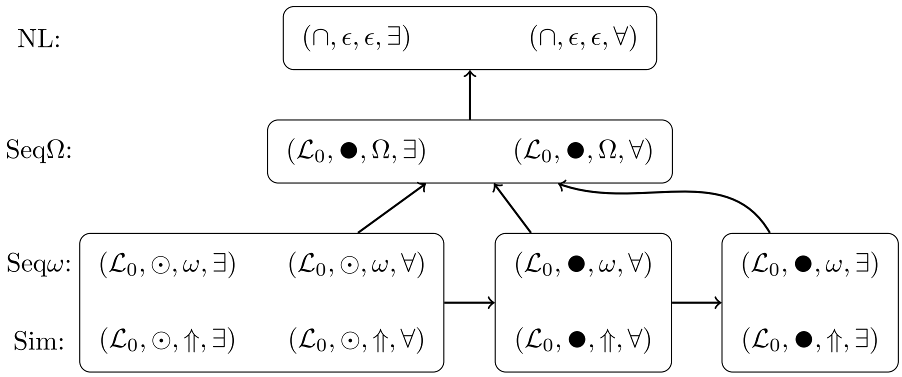

Fig. 1. The expressivity landscape of distributed knowledge, where 'NL' stands for 'non-linguistic sharing', 'SeqΩ' denotes unlimited sequential sharing, 'Seqω' stands for sequential sharing limited to ordinal ω, and 'Sim' denotes simultaneous sharing. Equivalent variations of distributed knowledge are enclosed in a box. Arrows point from stronger variants to weaker ones. Some arrows that follow from transitivity have been omitted for the sake of clarity.

Out of all variants of distributed knowledge, only the classic one (i.e. $(\cap, \epsilon, \epsilon, \forall)$ and $(\cap, \epsilon, \epsilon, \exists)$ in our taxonomy) was axiomatised [8,15]. The task of providing axiomatisations for the remaining variants seems to be both colossal and irresistibly tempting, and we thus leave it for future work.

Acknowledgements

A preliminary version of this work has benefited from discussion with Fernando R. Velázquez-Quesada. We would also like to thank the anonymous reviewers of AiML for their attention to details and constructive criticism that helped us to improve the presentation of this work.

Appendix

A Vicious Circularity

Suppose that we define the type $(\mathcal{L}, \odot, \Uparrow, \forall)$ in the same way as the type $(\mathcal{L}_0, \odot, \Uparrow, \forall)$, except that we now allow $a \in G$ to contribute any formula $\psi_a \in \mathcal{L}$ that they know, as opposed to any known formula from $\mathcal{L}_0$.

Clearly, such a definition would be circular. Here, we show that this circularity is vicious. More specifically, it is under-determined, in the sense that there are models where both $D_G\varphi$ and $\neg D_G\varphi$ are consistent with the semantics.

Proposition A.1 *There are $\mathcal{M}, s$ and φ such that both $\mathcal{M}, s \models_{(\mathcal{L}, \odot, \Uparrow, \forall)} D_G\varphi$*

and $\mathcal{M}, s \models_{(\mathcal{L},\odot,\Uparrow,\forall)} \neg D_G\varphi$ *are consistent with the* $(\mathcal{L},\odot,\Uparrow,\forall)$*-semantics.*

Proof. Consider the following model:

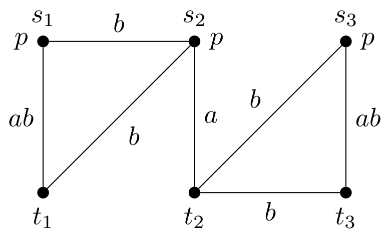

We will show that both $\mathcal{M}, s_2 \models D_{\{a,b\}}p$ and $\mathcal{M}, s_2 \not\models D_{\{a,b\}}p$ are consistent with the circular semantics. To see why this is the case, first note that the semantics are extensional, so while there are infinitely many formulas it suffices to consider only those that have different extensions. Let use denote the extension of φ by $[\![\varphi]\!]$.

Let E be the set of all extension on this model. In order for E to be consistent with the $(\mathcal{L},\odot,\Uparrow,\forall)$-semantics, it has to be "self-fulfilling", in the sense that, if we assume that E is the set of all extensions, then we should have $[\![\varphi]\!] \in E$ for all $\varphi \in \mathcal{L}$, and for every $e \in E$ there should be some $\varphi_e \in \mathcal{L}$ such that $[\![\varphi_e]\!] = e$.

We will show that there are two different sets of extensions that satisfy this criterion: we can take $E = 2^S$, in which case we have $\mathcal{M}, s_2 \models D_{\{a,b\}}p$, and we can take $E = \{\emptyset, \{s_1, s_2, s_3\}, \{t_1, t_2, t_3\}, S\}$, in which case $\mathcal{M}, s_2 \not\models D_{\{a,b\}}p$.

Suppose therefore that $E = 2^S$. It is immediate that $[\![\varphi]\!] \in E$ for all $\varphi \in \mathcal{L}$. Left to show is that every extension $e \in E$ is witnessed by some formula φ_e. To this purpose, we first note that we have $\mathcal{M}, s_2 \models D_{\{a,b\}}p$. This is because, by assumption, $E = 2^S$ is the set of extensions, so there are formulas φ_1 and φ_2 such that $[\![\varphi_1]\!] = \{s_2, t_2\}$ and $[\![\varphi_2]\!] = \{s_1, s_2, t_1\}$. Then $\mathcal{M}, s_2 \models \Box_a\varphi_1$ and $\mathcal{M}, s_2 \models \Box_b\varphi_2$, so the agents can share φ_1 and φ_2. Furthermore, by putting φ_1 and φ_2 together, the agents discover that s_2 is the only possible world. As p is true there, we have $\mathcal{M}, s_2 \models D_{\{a,b\}}p$.

In every other world, $D_{\{a,b\}}p$ is false. For t_1, t_2 and t_3 this follows from the fact that distributed knowledge is truthful and p is false in t_1, t_2 and t_3. For s_1 and s_3 it follows from the fact that there is an ab-successor (t_1 or t_3, respectively) where p is false. Since this world is an ab-successor, it can never be excluded by any formula known to a or b, so the agents cannot exclude the possibility of $\neg p$ by combining their information.

We have now shown that the formula $D_{\{a,b\}}p$ uniquely identifies the world s_2. It follows that there are also formulas uniquely identifying every other world:

s_1: $p \wedge \neg D_{\{a,b\}}p \wedge \Diamond_b D_{\{a,b\}}p$

s_3: $p \wedge \neg D_{\{a,b\}}p \wedge \neg\Diamond_b D_{\{a,b\}}p$

t_1: $\neg p \wedge \Diamond_b D_{\{a,b\}}p$

t_2**:** $\neg p \wedge \Diamond_a D_{\{a,b\}} p$

t_3**:** $\neg p \wedge \neg\Diamond_b D_{\{a,b\}} p \wedge \neg\Diamond_a D_{\{a,b\}} p$.

Let us denote the formula for any world by φ_{s_i} or φ_{t_i}. Any $e \in E$ is the extension of some disjunction of the relevant φ_{s_i} and/or φ_{t_i}.

We have now shown that $E = 2^S$ being the set of extensions is consistent with the semantics, and that we then have $\mathcal{M}, s_2 \models D_{\{a,b\}} p$. The witnessing formulas for this distributed knowledge had to have extensions $\{s_2, t_2\}$ and $\{s_1, t_1, s_2\}$, so we can take $\varphi_a = \{\varphi_{s_2} \vee \varphi_{t_2}\}$ and $\varphi_b = \{\varphi_{s_1} \vee \varphi_{s_2} \vee \varphi_{t_1}\}$.

Next, we will show that it is consistent with the semantics to have $E = \{\emptyset, \{s_1, s_2, s_3\}, \{t_1, t_2, t_3\}, S\}$, in which case $\mathcal{M}, s_2 \not\models D_{\{a,b\}} p$. In this case, it is easy to see that for every $e \in E$ there is some φ such that $e = [\![\varphi]\!]$; we have $\emptyset = [\![\bot]\!]$, $\{s_1, s_2, s_3\} = [\![p]\!]$, $\{t_1, t_2, t_3\} = [\![\neg p]\!]$ and $S = [\![\top]\!]$.

Left to show, therefore, is that for every φ, we have $[\![\varphi]\!] \in E$. Because the s_i are bisimilar to each other, as are the t_j, the bisimulation-invariance of modal logic implies that every $\psi \in \mathcal{L}_0$ will have one of the four extentions in E.

Furthermore, the D_G operator cannot break this symmetry. This is because, in every world, both a and b consider at least one t_i world and at least one s_i world possible. It follows that, for every world x, if $\mathcal{M}, x \models \Box_a\varphi$ or $\mathcal{M}, x \models \Box_b\varphi$, then $[\![\varphi]\!]$ contains at least one t_i and at least one s_i. Among the four extensions in E, the only one with this property is S.

In their communication, a nor b can therefore only contribute formulas with extension S, so neither of them provides non-trivial information. It follows that φ is distributed knowledge if and only if one of the agents already knew φ before the agents started sharing information. Hence $\mathcal{M}, x \models D_{\{a,b\}}\varphi$ if and only if $\mathcal{M}, x \models \Box_a\varphi \vee \Box_b\varphi$. It follows that, for every $\varphi \in \mathcal{L}$ there is a $\varphi_0 \in \mathcal{L}_0$ such that $[\![\varphi]\!] = [\![\varphi_0]\!]$, which implies that $[\![\varphi]\!] \in E$.

Furthermore, because $\mathcal{M}, s_2 \not\models \Box_a p \vee \Box_b p$, we have $\mathcal{M}, s_2 \not\models D_{\{a,b\}} p$. □

B Full versions of proofs

Proposition B.1 *$\mathcal{M}, s \models_{(\mathcal{L}_0, \bullet, \Uparrow, \forall)} D_G\varphi$ does not necessarily imply that $\mathcal{M}, s \models_{(\mathcal{L}_0, \odot, \Uparrow, \forall)} D_G\varphi$.*

Proof. Let $\mathcal{M} = (S, \sim, V)$, where S and $\sim$ are given as follows.

- $S = \{s_{x,y} \mid (x, y) \in [0, 1) \times [0, 1)\}$,
- $\sim_a = \{(s_{x,y}, s_{x,y'})\}$
- $\sim_b = \{(s_{x,y}, s_{x',y})\}$

Each x or y can be interpreted, written in binary notation, as a set of natural numbers. Let us write $\ulcorner x \urcorner$ for the set of natural numbers represented by x. Then we take $V(p_i) = \{s_{x,y} \mid i \in \ulcorner x \urcorner\}$, $V(q_i) = \{s_{x,y} \mid i \in \ulcorner y \urcorner\}$ and $V(r) = \{s_{x,y} \mid x = y\}$.

In every world, a knows the x-coordinate while being uncertain about the y-coordinate, while b is uncertain about the x-coordinate while knowing the y-coordinate. Since the value of p_i depends only on the x-coordinate, this means that a knows, for every i, whether p_i is true. Similarly, b knows whether q_i is true.

Furthermore, it is true throughout the model, and therefore known to both agents, that r holds if and only if $x = y$, and therefore if and only if $p_i \leftrightarrow q_i$ for all i.

We have $\mathcal{M}, s_{0,0} \models_{(\mathcal{L}_0,\bullet,\Uparrow,\forall)} D_{\{a,b\}}r$, which is witnessed by the sets $\Psi_a = \{\neg p_i \mid i \in \mathbb{N}\}$ and $\Psi_b = \{\neg q_i \mid i \in \mathbb{N}\}$. After all, the only world where all p_i and q_i are false is $s_{0,0}$, where r is true.

Suppose now, towards a contradiction, that there are formulas ψ_a and ψ_b that are known by their respective agents in $s_{0,0}$, and that would allow the agents to learn that r is true. Let Q be the set of atoms that occur in ψ_a and ψ_b. Note that this is a finite set.

Because ψ_a is known by a, it must hold in every $s_{0,y}$. Similarly, ψ_b holds on every $s_{x,0}$. Let x and y be such that $x \neq 0$, $y \neq 0$, $\ulcorner x \urcorner \cap Q = \ulcorner y \urcorner \cap Q = \emptyset$ and $x \neq y$. Now, consider the world $s_{x,y}$. We will show that it is Q-bisimilar to both $s_{x,0}$ and $s_{0,y}$.

To this end, consider the relation $\approx \subseteq S \times S$, such that $s_{x,y} \approx s_{x',y'}$ iff $s_{x,y}$ and $s_{x',y'}$ agree on $Q \cup \{r\}$. We claim that this relation is a bisimulation. Atomic agreement (when restricted to Q) between any two $\approx$-related worlds is immediate from the construction. We show forth for b, the other cases can be shown similarly.

So suppose that $s_{x,y} \approx s_{x',y'}$ and that $s_{x,y} \sim_b s_{u,v}$. Then $y = v$. Now, let $v' = y'$ and

- if $u = v$ then $u' = v'$,
- if $u \neq v$ then u' is any number such that $u' \neq v'$ and $\ulcorner u' \urcorner \cap Q = \ulcorner u \urcorner \cap Q$.

Then $s_{u,v}$ and $s_{u',v'}$ agree on all atoms in $Q \cup \{r\}$. Hence $s_{u,v} \approx s_{u',v'}$ by our assumption. Furthermore, because $y' = v'$, we also have $s_{x',y'} \sim_b s_{u',v'}$ by the construction of the model. So the forth condition is satisfied.

From this bisimilarity and the fact that $\mathcal{M}, s_{x,0} \models \psi_b$, it follows that $\mathcal{M}, s_{x,y} \models \psi_b$, and therefore also that $\mathcal{M}, s_{0,y} \models \psi_b$. For the same reason, $\mathcal{M}, s_{x,0} \models \psi_a$.

We can thus conclude that ψ_a and ψ_b cannot exclude the worlds $s_{x,0}$ and $s_{0,y}$. In both these worlds r is false, so neither a nor b learns that r is true. This contradicts our assumptions, so we have that $\mathcal{M}, s \not\models_{(\mathcal{L}_0,\odot,\Uparrow,\forall)} D_{\{a,b\}}r$. □

Proposition B.2 *$\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Uparrow,\exists)} D_G\varphi$ does not necessarily imply $\mathcal{M}, s \models_{(\mathcal{L}_0,\bullet,\Uparrow,\forall)} D_G\varphi$.*

Proof. Let

- $S = \{s_{x,y} \mid (x,y) \in [0,1) \times [0,1)\} \cup \{t_{x,y} \mid (x,y) \in [0,1) \times [0,1)\}$,
- $\sim_a = \{(u_{x,y}, v_{x,y'}) \mid u, v \in \{s,t\}, x, y, y' \in [0,1)\}$
- $\sim_b = \{(s_{x,y}, s_{x',y}) \mid x, x', y \in [0,1)\} \cup \{(t_{x,y}, t_{x',y}) \mid x, x', y \in [0,1)\}$
- $V(p_i) = \{s_{x,y} \mid i \in \ulcorner x \urcorner\} \cup \{t_{x,y} \mid i \in \ulcorner x \urcorner\}$
- $V(q_i) = \{s_{x,y} \mid i \in \ulcorner y \urcorner\} \cup \{t_{x,y} \mid i \in \ulcorner y \urcorner\}$

- $V(r) = \{s_{x,y} \mid |\ulcorner x \urcorner \triangle \ulcorner y \urcorner| \text{ is even}\} \cup \{t_{x,y} \mid |\ulcorner x \urcorner \triangle \ulcorner y \urcorner| \text{ is odd}\}$ [6]

Essentially, the model consists of two grids, one with $s_{x,y}$ worlds and another one with $t_{x,y}$ worlds. Agent a does not know which grid the current world belongs to, while agent b does. On $s_{x,y}$ worlds r is true if the number atoms on which x and y differs is even, whereas on $t_{x,y}$ worlds r is true if that number is odd. As before, a knows the value of the p_i atoms, while b knows the value of q_i. Finally, b can tell the difference between the s and t worlds, and therefore the required parity, while a does not.

We have $\mathcal{M}, s_{0,0} \models_{(\mathcal{L}_0, \bullet, \Uparrow, \exists)} D_{\{a,b\}} r$ because, as in the previous proof, a can announce which p_i hold and b can announce which q_i hold, which suffices for b to determine that r holds.

There is no way for a to learn that r is true by simultaneous statements, however. This follows from another Q-bisimilarity argument. □

Proposition B.3 *$\mathcal{M}, s \models_{(\mathcal{L}_0, \bullet, \Omega, \exists)} D_G\varphi$ does not necessarily imply $\mathcal{M}, s \models_{(\mathcal{L}_0, \bullet, \Uparrow, \exists)} D_G\varphi$.*

Proof. Let

$$
\begin{aligned}
S = \{s_{e,f,g} \mid\ & e : \mathbb{N} \times \mathbb{N} \times \{0,1\} \to \{0,1\} \\
& f : \mathbb{N} \to \mathbb{N} \\
& g : \mathbb{N} \to \mathbb{N}\}
\end{aligned}
$$

and

- $s_{e,f,g} \sim_a s_{e',f',g'}$ iff $e = e'$,
- $s_{e,f,g} \sim_b s_{e',f',g'}$ iff $f = f'$,
- $s_{e,f,g} \sim_c s_{e',f',g'}$ iff $g = g'$.

In other words, in $s_{e,f,g}$ agent a knows e, agent b knows f and agent c knows g.

Furthermore, for $i, j \in \mathbb{N}$ let

- $V(p_{i,j}) = \{s_{e,f,g} \mid e(i,j,0) = 1\}$,
- $V(q_{i,j}) = \{s_{e,f,g} \mid e(i,j,1) = 1\}$,
- $V(x_j) = \{s_{e,f,g} \mid$ the number of indices $i \in \mathbb{N}$ s.t. $e(i,j,0) \neq e(i,j,1)$ is divisible by $f(j)\}$
- $V(y_j) = \{s_{e,f,g} \mid$ the number of indices $i \in \mathbb{N}$ s.t. $e(i,j,0) \neq e(i,j,1)$ is divisible by $g(j)\}$
- $V(z) = \{s_{e,f,g} \mid$ there is an even number of indices $j \in \mathbb{N}$ s.t. exactly one of x_j, y_j holds on $s_{e,f,g}\}$

So e simply determines which $p_{i,j}$ and $q_{i,j}$ hold. The function f, meanwhile, determines for each j the number, $f(j)$, that must divide the amount of indices i on which $p_{i,j}$ and $q_{i,j}$ are different, in order for x_j to be true. Similarly,

[6] For sets X and Y, the symmetric difference $X \triangle Y$ is defined as $(X \setminus Y) \cup (Y \setminus X)$.

g determines the number of indices that must be different for y_j to be true. Finally, z holds if and only if x_j and y_j differ in value for an even number of indices j.

It is now easy to see that for any $s_{e,f,g}$ such that $\mathcal{M}, s_{e,f,g} \models z$ we have $\mathcal{M}, s_{e,f,g} \models_{(\mathcal{L}_0, \bullet, \Omega, \exists)} D_{\{a,b,c\}} z$. The communication the agents can perform is as follows:

- In the time steps before ω, agent a tells the other two exactly which $p_{i,j}$ and $q_{i,j}$ hold.
- At time ω, agents b and c know which $p_{i,j}$ and $q_{i,j}$ hold. As a consequence, b knows which x_j hold while c knows which y_j hold.
- In the time steps between ω and $2 \times \omega$, agents b and c tell the other two which x_j and y_j hold.
- At time $2 \times \omega$, all agents know exactly which x_j and y_j hold, and therefore whether z holds.
- We assumed we were in a world where z is true, so the agents learn that z is true.

Now, to show that $\mathcal{M}, s_{e,f,g} \not\models_{(\mathcal{L}_0, \bullet, \Uparrow, \exists)} D_{\{a,b,c\}} z$.

The key observation here is that b and c cannot provide any non-trivial announcements. Suppose towards a contradiction that $\mathcal{M}, s_{e,f,g} \models \Box_b \psi_b$ and $\mathcal{M}, s_{e',f',g'} \not\models \psi_b$. Let Q be the set of atoms in ψ_b. Then there is some $s_{e'',f,g''}$ that is Q-bisimilar to $s_{e',f',g'}$. But, by $\mathcal{M}, s_{e,f,g} \models \Box_b \psi_b$, we have $\mathcal{M}, s_{e'',f,g''} \models \psi_b$, which by bisimilarity implies $\mathcal{M}, s_{e',f',g'} \models \psi_b$, contradicting our assumption.

Hence the only announcements b can provide hold in every world of the model, and are therefore uninformative. That c cannot make a non-trivial announcement is shown similarly.

This means that only a can provide information, so after the communication all worlds of the form $s_{e,f',g'}$ are still accessible. Each agent then still considers both z worlds and $\neg z$ worlds to be possible, so z is not distributed knowledge. $\Box$

References

[1] Ågotnes, T., N. Alechina and R. Galimullin, *Logics with group announcements and distributed knowledge: Completeness and expressive power*, Journal of Logic, Language and Information **31** (2022), pp. 141–166.

[2] Ågotnes, T., P. Balbiani, H. van Ditmarsch and P. Seban, *Group announcement logic*, Journal of Applied Logic **8** (2010), pp. 62–81.

[3] Ågotnes, T. and R. Galimullin, *Quantifying over information change with common knowledge*, Autonomous Agents Multi-Agent Systems **37** (2023), p. 40.

[4] Ågotnes, T. and Y. N. Wáng, *Resolving distributed knowledge*, Artificial Intelligence **252** (2017), pp. 1–21.

[5] Baltag, A., *What is DEL good for?* (2010), Workshop on Logic, Rationality and Intelligent Interaction.
URL `http://ai.stanford.edu/~epacuit/lograt/esslli2010-slides/copenhagenesslli.pdf`

[6] Baltag, A. and S. Smets, *Learning what others know*, in: E. Albert and L. Kovács, editors, *Proceedings of the 23rd LPAR*, EPiC Series in Computing **73** (2020), pp. 90–119.

[7] Fagin, R., J. Y. Halpern, Y. Moses and M. Vardi, "Reasoning About Knowledge," The MIT Press, 1995.

[8] Fagin, R., J. Y. Halpern and M. Y. Vardi, *What can machines know? On the properties of knowledge in distributed systems*, Journal of the ACM **39** (1992), pp. 328–376.

[9] Galimullin, R. and F. R. Velázquez-Quesada, *(Arbitrary) partial communication*, in: N. Agmon, B. An, A. Ricci and W. Yeoh, editors, *Proceedings of the 22nd AAMAS* (2023), pp. 400–408.

[10] Gerbrandy, J., "Bisimulations on Planet Kripke," Ph.D. thesis, University of Amsterdam (1999).

[11] Halpern, J. Y. and Y. Moses, *Knolwedge and common knolwedge in a distributed environment*, Journal of the Association for Computing Machinery **37** (1990), pp. 549–587.

[12] Humberstone, I. L., *The formalities of collective omniscience*, Philosophical Studies **48** (1985), pp. 401–423.

[13] Meyer, J. C. and W. van der Hoek, "Epistemic logic for AI and computer science," Cambridge tracts in theoretical computer science **41**, CUP, 1995.

[14] Roelofsen, F., *Distributed knowledge*, Journal of Applied Non-Classical Logics **17** (2007), pp. 255–273.

[15] van der Hoek, W. and J. C. Meyer, *Making some issues of implicit knowledge explicit*, International Journal of Foundations of Computer Science **3** (1992), pp. 193–223.

[16] van der Hoek, W., B. van Linder and J.-J. Meyer, *Group knowledge is not always distributed (neither is it always implicit)*, Mathematical Social Sciences **38** (1999), pp. 215–240.

[17] van Ditmarsch, H., W. van der Hoek and B. Kooi, "Dynamic Epistemic Logic," Synthese Library **337**, Springer, 2008.