The Limits to Gossip: Second-order Shared Knowledge of all Secrets is Unsatisfiable

Hans van Ditmarsch[1][0000−0003−4526−8687] and Malvin Gattinger[2][0000−0002−2498−5073]

[1] Open University of the Netherlands
hans.vanditmarsch@ou.nl
[2] University of Amsterdam, the Netherlands
malvin@w4eg.eu

Abstract. It is known that without synchronization via a global clock one cannot obtain common knowledge by communication. Moreover, it is folklore that without exchanging higher-level information arbitrary higher-level shared knowledge cannot be achieved.
Here we make this result precise. We use epistemic logic to formally define "everyone knows that everyone knows that everyone knows all secrets" and then prove that this statement is unsatisfiable.

Keywords: Epistemic logic · Gossip · Higher-order knowledge

1 Introduction

Consider a group of agents that initially each know a unique secret. The agents then make one-to-one phone calls during which they always share all the secrets they know. So-called *gossip protocols* in this simple model provide an efficient way to spread information using peer-to-peer communication.

Even when the gossiping agents only exchange secrets and no additional information, they also obtain higher-order knowledge: they learn what other agents know, or even what other agents know about yet other agents, and so on. But it is known that without synchronization via a global clock one cannot obtain common knowledge by peer-to-peer communication [5]. It is folklore that only the first level of shared knowledge can be achieved, but not the higher levels. In this article we make this result precise using epistemic logic.

We start with two examples that illustrate interesting cases of when and how agents can obtain higher-order knowledge.

Example 1. Suppose we have a set of four agents $\{a, b, c, d\}$. Consider the sequence of calls $ab.cd.ac.bd.ad.bc.ab.cd$ and the results shown in Table 1. After the fourth call bd everyone is an expert, i.e. they know all secrets. In both of the last two calls ab and cd the two agents involved become so-called super experts, i.e. they know that everyone knows all secrets. The example shows in particular that two agents can become super experts at the same time.

	a	b	c	d	initial state
$\stackrel{ab}{\rightarrow}$	ab	ab	c	d	
$\stackrel{cd}{\rightarrow}$	ab	ab	cd	cd	
$\stackrel{ac}{\rightarrow}$	abcd A C	ab	abcd A C	cd	
$\stackrel{bd}{\rightarrow}$	abcd A C	abcd B D	abcd A C	abcd B D	everyone is an expert
$\stackrel{ad}{\rightarrow}$	abcd A CD	abcd B D	abcd A C	abcd AB D	
$\stackrel{bc}{\rightarrow}$	abcd A CD	abcd BCD	abcd ABC	abcd AB D	
$\stackrel{ab}{\rightarrow}$	abcd ABCD	abcd ABCD	abcd ABC	abcd AB D	a, b become super experts
$\stackrel{cd}{\rightarrow}$	abcd ABCD	abcd ABCD	abcd ABCD	abcd ABCD	c, d become super experts

Table 1. Results of *ab.cd.ac.bd.ad.bc.ab.cd*. A lower case y in column x means x knows the secret of y; an upper case Y means x knows that y is an expert. Therefore, "abcd" denotes an expert and "ABCD" denotes a super expert.

	a	b	c	d	initial state
$\stackrel{ac}{\rightarrow}$	a c	b	a c	d	
$\stackrel{ad}{\rightarrow}$	a cd	b	a c	a cd	
$\stackrel{ac}{\rightarrow}$	a cd	b	a cd	a cd	a learns d from c
$\stackrel{bc}{\rightarrow}$	a cd	abcd BC	abcd BC	a cd	
$\stackrel{ac}{\rightarrow}$	abcd ABC	abcd BC	abcd ABC	a cd	a is *lucky* about b

Table 2. Results of *ac.ad.ac.bc.ac* including a lucky call.

More interestingly, an agent may learn that another agent is an expert without calling that agent. We will call this a lucky call and this notion plays a key role in the proof of our main result. In a synchronous setting where agents observe that other calls happen, lucky calls happen frequently. But in the asynchronous setting we consider here agents only observe their own calls. This limits what agents learn about each other and makes lucky calls noteworthy.

Example 2. Again suppose we have a set of four agents $\{a, b, c, d\}$ and consider the call sequence *ac.ad.ac.bc.ac*, with results shown in Table 2. Here agent a learns in the final call ac that a, b, c are experts. Because b is not involved in this call we say that this a *lucky call* and say that *a is lucky about b*. We will show that this sequence is typical: in any call sequence each agent can only once be lucky and at most $n - 2$ agents out of n agents (for $n \geq 4$) can be lucky.

Our main contribution here is the proof that "everyone knows that everyone knows that everyone knows all secrets" is unsatisfiable. On the way to this result we also provide a new definition of the "causal cone" of an agent or a subset of agents in a call sequence. This article is structured as follows. We discuss related work in Section 2 and provide definitions in Section 3. In Sections 4 we define causal cones and in Section 5 we use them to characterise lucky calls. Section 6 contains our main result.

2 Related Work

The "gossip problem" as introduced above is also known as the "telephone problem" and goes back (at least) to the article [9] from 1971. The main classical result is that only a linear amount of calls ($2n-4$ if we have n agents) is needed to ensure that *everyone knows all secrets*. We refer to [6] for a survey of variants of the gossip problem, for example over different graphs or using broadcasting.

Most of the classical results assume a central scheduler, i.e. an authority that decides in which order calls should be made. More recently, decentralised gossip has been studied, where agents decide on their own whom they should call next, and multiple logics have been developed to analyse the gossip problem and different protocols [1,2,3]. Some of these logics include not only statements to say that agents know a secret, but they also provide general "an agent knows that φ" modalities common in epistemic logics, and thereby allow us to discuss the higher-order knowledge effects of gossip.

The question which higher-order knowledge can be achieved by communication between agents goes back to the classic problem of the Byzantine Generals [8]. In general it is impossible to achieve *common knowledge* in an asynchronous distributed system, as shown in [5]. However, *shared knowledge* can be achieved with messages of the form "I know that this other agents knows that ...". This has been studied in [7] where agents always tell each other *all* they know. Concretely, among n agents shared knowledge of level k can be achieved with $(k+1)(n-2)$ many calls. Here we only allow agents to exchange secrets and do not allow them to exchange any other kind of information.

Our result is most related to [1,3] who both investigate when the truth of formulas stabilises during gossip protocol execution, including the case of the most general gossip protocol where any call can be made at any time and where agents only observe their own calls (setting $\langle \bullet, \Diamond, \beta \rangle$ in [3], also called asynchronous ANY in [4]). This is the same setting as ours.

By different methods the authors of [1] and those of [3] demonstrate that making new calls no longer affects the truth of epistemic formulas at some stage. Although the objectives of these publications were different, namely decidability of logics for gossip or correctness of gossip protocols, there is some overlap in methods. They show that in any (fairly scheduled) call sequence, with the standard call semantics that only secrets are exchanged in a call, at some stage further calls have no informational effect — such calls are redundant. This was relevant to observe for gossip protocols where the goal was that all agents became experts, because it showed that, in principle, even if one were to consider epistemic goals such as knowing that others are experts, things would eventually come to a stop. But they did not consider any specific epistemic goals.

In [1] only epistemic formulas of depth 1 were considered (the crucial result is [1, Lemma 21]), but here we focus on depth 2 and higher which in [1] is only mentioned as a generalization for future work.

The authors of [3] considered arbitrary epistemic formulas. However, there the comparison to our result stops: [3, Prop. 5.5, Cor. 5.6] shows that formulas of any epistemic depth remain true forever or false forever after call sequences

of certain length (bounded by a polynomial in terms of the number of agents). Hence a formula like $EEExp_A$ must remain true forever or false forever after further extending call sequences. But the authors of [3] did not investigate specific formulas. Here we show that $EEExp_A$ remains false forever.

3 Syntax and Semantics

We assume a finite set of at least four agents $A = \{a, b, c, d, \dots\}$ throughout this article. This assumption is needed for our main results. Some lemmas may also hold for less than four agents, but these are boundary cases of little interest.

Each agent holds a single secret. The agents communicate with each other through telephone calls. During a call between two agents x and y, they exchange all the secrets that they knew before the call.

A *call* is a pair of agents $(x, y) \in A \times A$ for which we write xy. Agent x is the *caller* and agent y is the *callee*. Given call xy, call yx is the *dual call*. An agent x is *involved* in a call yz iff $y = x$ or $z = x$. In this contribution the direction of the call does not matter, so it only matters if an agent is involved in a call. We will therefore arbitrarily write xy or yx for the call between x and y, where we often prefer the lexicographic order of agents. A *call sequence* is defined by induction: the empty sequence ϵ is a call sequence. If σ is a call sequence and xy is a call, then $\sigma.xy$ is a call sequence. We write $|\sigma|$ to denote the length of a call sequence.

Given call sequences τ, σ, by induction on the length of σ we further define that τ is a *subsequence* of σ. This is the inductive definition: $\epsilon \subseteq \epsilon$, and if $\tau \subseteq \sigma$ then $\tau, \tau.ab \subseteq \sigma.ab$ and $\tau, ab.\tau \subseteq ab.\sigma$.

If $\sigma = \rho.\tau$, then ρ is a *prefix* of σ, denoted as $\rho \sqsubseteq \sigma$, and τ is the *complement* of ρ in σ, where τ is also denoted $\sigma \setminus \rho$.

Definition 1 (Language). *For a finite set of agents A the* language $\mathcal{L}$ *is given by* $\varphi ::= b_a \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi$ *where $a, b \in A$. Let $\rightarrow$ and $\vee$ be defined as usual.*

The φ are called *formulas*. The atomic formula b_a reads as 'agent a has the secret of b' or 'agent a knows the secret of b'. The formula $K_a\varphi$ reads 'agent a knows that φ is true'. We also define $E\varphi := \bigwedge_{a \in A} K_a\varphi$ and read it as 'everyone knows that φ' ($E\varphi$ is also known as *shared* or *mutual* knowledge of φ).

Agent a is an *expert* if she knows all the secrets, formally $\bigwedge_{b \in A} b_a$, abbreviated as Exp_a. *Everyone is an expert* is represented by the formula $Exp_A := \bigwedge_{a \in A} \bigwedge_{b \in A} b_a$. Agent a is a *super expert* if she knows that everyone is an expert, formally $K_a Exp_A$. Therefore, $EExp_A$ means that everyone is a super expert and $EEExp_A$ means that everyone knows that everyone is a super expert. The main result we prove in this article is that the latter cannot be achieved, i.e. that $EEExp_A$ is not satisfiable.

Let $B \subseteq A$. If $a \in B$ knows the secrets of all the agents in B, a is a B *expert*, and if a knows that all the agents in B know the secrets of all the agents in B, a is a B *super expert*. We find it convenient to have some additional notation for such matters: $\bigwedge_{b \in B} K_b\varphi$ is denoted $E_B\varphi$, and '$a \in B$ knows the secrets of all the

agents in B' is denoted $Exp_a(B)$ with similar notational variations as before, so that $E_B Exp_B(B)$ denotes that all in B are B super experts (and for $B = A$ we omit the parameter).

The epistemic relation defined below models that agents only observe the calls they are involved in. In particular there is no global clock and the conditions are asynchronous, meaning agents do not know how many calls have taken place.

Definition 2 (Epistemic relation). *Let $a \in A$. The* epistemic relation $\sim_a$ *is the smallest equivalence relation between call sequences such that:*

- $\epsilon \sim_a \epsilon$
- *if $\sigma \sim_a \tau$ and $a \notin \{b, c\}$, then $\sigma.bc \sim_a \tau$*
- *if $\sigma \sim_a \tau$, and for all c, $\sigma \models c_b$ iff $\tau \models c_b$, then $\sigma.ab \sim_a \tau.ab$*

Definition 3 (Semantics). *Let call sequence σ and formula $\varphi \in \mathcal{L}$ be given. We define $\sigma \models \varphi$ by induction on the structure of φ. Moreover, we define the valuation of atoms by induction on σ, for any $a, b \in A$ with $a \neq b$.*

$$
\begin{array}{lll}
\epsilon \models a_b & \text{iff } a = b & \\
\sigma.ab \models c_a & \text{iff } \sigma \models c_a \text{ or } \sigma \models c_b & \text{for all } c \in A \\
\sigma.ab \models c_b & \text{iff } \sigma \models c_a \text{ or } \sigma \models c_b & \text{for all } c \in A \\
\sigma.ab \models c_d & \text{iff } \sigma \models c_d & \text{for all } c, d \in A \text{ with } d \notin \{a, b\} \\
\sigma \models \neg\varphi & \text{iff } \sigma \not\models \varphi & \\
\sigma \models \varphi \wedge \psi & \text{iff } \sigma \models \varphi \text{ and } \sigma \models \psi & \\
\sigma \models K_a\varphi & \text{iff } \tau \models \varphi \text{ for all } \tau \text{ such that } \sigma \sim_a \tau &
\end{array}
$$

A formula φ is valid*, notation $\models \varphi$, iff for all call sequences σ we have $\sigma \models \varphi$. We abbreviate the set of secrets know by a after σ with $a(\sigma) := \{c \in A \mid \sigma \models c_a\}$.*

If in call ab agent a or b becomes an expert, then the other agent must also be an expert after this call. In contrast, if in a call ab agent a or b becomes a super expert, then the other agent does not have to be a super expert after this call.

4 Causal Relation and Causal Cone

We now introduce additional notation for specific subsequences of calls. The goal is to make it easy to select and reason about those calls that are relevant for a specific subset of agents.

Definition 4 (Causal relation). *For any sequence σ and calls $ab, cd \in \sigma$, we write $ab < cd$ iff σ has shape $\sigma_1.ab.\sigma_2.cd.\sigma_3$ (i.e. this occurrence of the call ab in σ is* before *this occurrence of the call cd in σ). We define the relation $\ll_0$ over call occurrences in σ by $ab \ll_0 cd \;:\Leftrightarrow\; ab < cd$ and $\{a, b\} \cap \{c, d\} \neq \varnothing$. Let $\ll$ be the reflexive transitive closure of $\ll_0$. Calls ab and cd are* causally related *iff $ab \ll cd$.*

Definition 5 (Causal cone). *Given a call sequence σ and a set of agents $B \subseteq A$, the* causal cone $\sigma^B_{\ll}$ *is defined by induction:*

$$\epsilon^B_{\ll} := \epsilon \qquad (\sigma.ab)^B_{\ll} := \begin{cases} \sigma^{B\cup\{a,b\}}_{\ll}.ab & \text{if } a \in B \text{ or } b \in B \\ \sigma^B_{\ll} & \text{otherwise} \end{cases}$$

For $\sigma^{\{a,b\}}_{\ll}$ we write $\sigma^{ab}_{\ll}$. Furthermore, let $\sigma^B_{\not\ll}$ be the complement of $\sigma^B_{\ll}$ in σ.

Intuitively, given a call sequence $\sigma.ab$, the sequence $\sigma^{ab}_{\ll}$ is the subsequence of σ of which all calls are *causally related* to the final call ab (it is the causal cone of ab), in other words $\sigma^{ab}_{\ll}$ is the subsequence consisting of all $cd \in \sigma$ such that $cd \ll ab$ in $\sigma.ab$. Also note that we can identify $\sigma^a_{\ll}$ with the causal cone of the last call in σ involving a. In other words, $\sigma^a_{\ll}$ determines what a knows after σ.

The complement $\sigma^B_{\not\ll}$ are the calls that do *not* determine what a knows after σ. For example, $\sigma^{ab}_{\not\ll}$ consists of all $cd \in \sigma$ that do not determine what a or b know after σ.

The causal relations between calls are only interesting when there are at least four agents. If there are two or three agents, all calls are causally related.

Lemma 1. *Let call sequence σ, group $B \subseteq A$ of agents, and agent $a \in B$ be given. Then (i) $\sigma \sim_a \sigma^B_{\ll}$ and also (ii) $\sigma \sim_a \sigma^B_{\ll}.\sigma^B_{\not\ll}$.*

Proof. We prove (i) by induction on the length of σ. Note that we declared B and $a \in B$ after σ, hence these may occur differently in our inductive assumption. For $B \cup \{a\}$ we write $B + a$ and similarly for $B \cup \{b, c\}$ we write $B + bc$.

Case $\sigma = \epsilon$. This is by definition as $\epsilon^B_{\ll} = \epsilon$.

Case $\sigma = \tau.ab$. By the definition of $\ll$ we have $(\tau.ab)^B_{\ll} = \tau^{B+b}_{\ll}.ab$ (we recall that $a \in B$). By inductive assumption we have $\tau \sim_a \tau^{B+b}_{\ll}$ and also $\tau \sim_b \tau^{B+b}_{\ll}$, so that agent b holds the same secrets after both. Therefore, by definition of $\sim_a$, we have $\tau.ab \sim_a \tau^{B+b}_{\ll}.ab$. Combining this we obtain $\tau.ab \sim_a (\tau.ab)^B_{\ll}$.

Case $\sigma = \tau.bc$ with $b, c \neq a$ and $b \in B$ or $c \in B$. By the definition of $\ll$ and because b or c is in B, $(\tau.bc)^B_{\ll} = \tau^{B+bc}_{\ll}.bc$. By induction, $\tau \sim_a \tau^{B+bc}_{\ll}$, and from that, the fact that $b, c \neq a$, and the definition of $\sim_a$ we also obtain $\tau.bc \sim_a \tau^{B+bc}_{\ll}.bc$, and therefore $\tau.bc \sim_a (\tau.bc)^B_{\ll}$.

Case $\sigma = \tau.bc$ with $b, c \notin B$. We have that $\tau.bc \sim_a \tau$ by the definition of $\sim_a$, because $b, c \neq a$. By inductive assumption, $\tau \sim_a \tau^B_{\ll}$. By the definition of $\ll$ and because $b, c \neq a$, $\tau^B_{\ll} = (\tau.bc)^B_{\ll}$. Combining all this we obtain $\tau.bc \sim_a (\tau.bc)^B_{\ll}$.

We now prove (ii). From $\sigma \sim_a \sigma^B_{\ll}$ it follows that $\sigma \sim_a \sigma^B_{\ll}.\sigma^B_{\not\ll}$ by the definition of $\sim_a$ and the observation that a does not occur in any call in $\sigma^B_{\not\ll}$. □

An instantiation of Lemma 1 is that $\sigma \sim_a \sigma^a_{\ll}.\sigma^a_{\not\ll}$ and $\sigma \sim_a \sigma^a_{\ll}$: a considers it possible that all not causally related calls, if any, take place after her last call.

Moreover, note that we have $(\sigma.ab)^{ab}_{\ll} = \sigma^{ab}_{\ll}.ab$ and $(\sigma.ab)^{ab}_{\not\ll} = \sigma^{ab}_{\not\ll}$ by Definition 5. Hence Lemma 1 also implies the following corollary.

Corollary 1. *Let call sequence σ and call ab be given. Then $\sigma.ab \sim_a \sigma^{ab}_{\ll}.ab.\sigma^{ab}_{\not\ll}$ and also $\sigma.ab \sim_a \sigma^{ab}_{\ll}.ab$.*

5 Lucky Calls

In the introduction we informally introduced the notion of a lucky call. The following definition makes this notion precise and the goal of this section is to characterise when lucky calls can happen.

Definition 6. *An agent a is lucky in a call ab if she learns in that call that another agent c is an expert. Formally, given a call sequence $\sigma.ab$ and $c \notin \{a, b\}$, agent a is* lucky *about c in ab iff we have $\sigma \not\models K_a Exp_c$ and $\sigma.ab \models K_a Exp_c$.*

In subsequent proofs we show and use that only $A - b$ super experts, who know that all but one agent b know the secrets of all but one agent b, can be lucky. We stress that our results only concern asychrony. In a synchronous setting, agents are lucky all the time and this is nothing special.

Lemma 2. *When two agents become experts in a call, they cannot be lucky.*

Proof. Let σ and a, b be given such that $\sigma \not\models Exp_a$, $\sigma \not\models Exp_b$, and $\sigma.ab \models Exp_a$ as well as $\sigma.ab \models Exp_b$. We show that for all $c \neq a, b$, $\sigma.ab \not\models K_a Exp_c$.

From Corollary 1 it follows that $\sigma.ab \sim_a \sigma^{ab}_{\ll}.ab$. No call in $\sigma^{ab}_{\ll}$ may contain an agent who is an expert, as the causal relation would then have made a or b an expert before call ab. Therefore $\sigma^{ab}_{\ll} \not\models Exp_c$ for all $c \in A$ other than a or b, and because of that and the semantics of calls also $\sigma^{ab}_{\ll}.ab \not\models Exp_c$. From that and $\sigma.ab \sim_a \sigma^{ab}_{\ll}.ab$ it follows that $\sigma.ab \models \neg K_a Exp_c$.

Similarly we show (replace $\sim_a$ by $\sim_b$) that $\sigma.ab \not\models K_b Exp_c$. □

Lemma 3. *When two agents become experts, neither becomes a super expert.*

Proof. We recall that there at least four agents. By Lemma 2, after the call wherein two agents become experts, they both remain uncertain whether the two or more agents not involved in the call are expert. □

We already showed (Lemma 2) that when two agents become experts in a call, they cannot be lucky. Now we consider the case where one of the agents was already an expert before the call. We want to characterise when the other agent who becomes expert can be lucky.

As an example, recall the call sequence $ac.ad.ac.bc.ac$ from Example 2 wherein a learns that b and c are expert in the final call ac. In other words, a's final call ac is lucky. After the prefix $ac.ad.ac$, a is a super expert for all agents but one (for $\{a, c, d\}$): a knows that a, c, d know all the secrets of a, c, d. This allows a to learn in the final call ac that *someone* from those all but one agents a, c, d must have called the one agent b and clearly it was not herself. Agent a learns in the final call that c or d called b, not that c called b.

We will now show that this typical case is also the only case.

Although we assume that there are at least four agents, the next result also holds for three agents. In call sequence $ac.bc.ac$, the third call is lucky and a then learns that c and b are experts. In the first call, a becomes a $\{a, c\}$ super expert. In the second call ac, agent a that c must have called b, and thus learns that b and c are experts, and thus becomes a $\{a, b, c\}$ super expert.

Lemma 4. *An agent becoming an expert can only be lucky if she is a super expert for all agents but one. Formally: Let $a, b, c \in A$ and σ be given and suppose $\sigma \not\models Exp_a$ and $\sigma.ac \models Exp_a$. Then $\sigma.ac \models K_a Exp_b$ iff $\sigma \models K_a Exp_{A-b}(A-b)$.*

Proof. Suppose we have $\sigma \not\models Exp_a$ and $\sigma.ac \models Exp_a$. We show the two directions. ($\Rightarrow$): We show the requested by contraposition: if a is not a $A-b$ super expert ($\sigma \not\models K_a Exp_{A-b}(A-b)$), then a cannot be lucky about b (i.e. $\sigma.ac \not\models K_a Exp_b$).

Intuitively, there are two ways in which non-expert a can be not a $A-b$ super expert: when she does not know enough or when she knows too much. In the first case *a does not know the secret of b* but she is not a super expert. In the second case *a knows the secret of b*, as this is, in a way, 'more' than being an $A-b$ super expert who is not an expert, which implies ignorance of b.

Case a does not know b. If $\sigma \models \neg b_a$ then a considers it possible that b has not yet made a call and thus only knows its own secret. In that case, as a is not a $A-b$ super expert, then a considers it possible that there is an agent $d \in A-b$ such that d does not know all of $A-b$'s secrets, that is, then there is a $e \in A-b$ (where e may be a or c) such that d does not know the secret of e. Agent a thus considers it possible that the next two calls are $db.dc$ and that b is not involved in further calls. After db, agent b is not an expert because neither b nor d know the secret of e. In dc agent d informs c of the secret of b. This can still be followed by any call sequence τ of calls between the agents of $A-ba$ making c expert before call ac. Altogether we get $\sigma.ac \sim_a \sigma.db.dc.\tau.ac$ and $\sigma.db.dc.\tau.ac \not\models Exp_b$. Therefore $\sigma.ac \not\models K_a Exp_b$.

Case a knows b. If $\sigma \models b_a$, then a cannot also be a $A-b$ super expert as this implies that a also knows all other secrets and therefore is an expert, which contradicts our assumption. It remains to show that $\sigma.ac \not\models K_a Exp_b$.

As non-expert a became expert in call ac, a learns a secret of some agent d in that call. As a already knew the secret of b, we must have $d \neq b$.

First assume there is a last call in σ between a and b. In that call ab, a therefore did not learn the secret of d. So after this call a still considers it possible that d only knows its own secret.

- If after call ab agent a also knows the secret of c, then a considers it possible that c does not know d. If then the subsequent calls are $bc.cd$ and b was not involved in further calls, then after call bc agent b still does not know d so b is not an expert. The part $bc.cd$ can still be followed by any sequence τ of calls between the agents of $A-ba$ making c expert before call ac. Altogether we get $\sigma.ac \sim_a \sigma.bc.cd.\tau.ac$ and $\sigma.bc.cd.\tau.ac \not\models Exp_b$. Therefore $\sigma.ac \not\models K_a Exp_b$.
- If after call ab agent a does *not* know the secret of c, then b also does not know c, and a considers it possible that subsequently $bd.cd$ took place. After bd agent b is not an expert (because b still does not know c). Call cd informs c of the secret of b. This can still be followed by any number of calls between the agents of $A-ba$ making c expert before call ac. Therefore $\sigma.ac \not\models K_a Exp_b$.

Second, assume there was no call in σ between a and b. Then, given that a knows b and thus knows that a call took place between b and some agent e (where e may be c) in $A-ba$, we again conclude that b did not know d after

that call nor after any call before the last call by a in σ (which happens to be the final call in $\sigma^a_{\ll}$). Agent a considers a call sequence possible (with prefix $\sigma^a_{\ll}$) wherein after her last call b did not make further calls and that b's secret was instead initially spread by agent e among the $A-ba$, and so on until c became expert. Therefore also in this case, $\sigma.ac \not\models K_a Exp_b$.

($\Leftarrow$): Suppose $\sigma \models K_a Exp_{A-b}(A-b)$. Note that we not only have $\sigma \not\models Exp_a$, but also $\sigma \models K_a \neg Exp_a$. Moreover, we claim that $\sigma \models Exp_c$. To see this, note that after σ agent a is an $A-b$ super expert, so a knows that c knows all secrets except b. Therefore, if a becomes expert in the last a call with c, then c must have learnt another secret. This can only be the secret of b. Therefore a learnt that c was an expert after any $\tau \sim_a \sigma$.

In order to show $\sigma.ac \models K_a Exp_b$, let τ be arbitrary such that $\tau.ac \sim_a \sigma.ac$. By definition of $\sim_a$ we have $\tau \sim_a \sigma$. From $\sigma \models Exp_c$ we get $c(\tau) = c(\sigma) = A$. It remains to show that $\tau \models Exp_b$. Also note that $c(\tau) = A$ implies $\tau \models b_c$.

From $\tau \sim_a \sigma$ as well as $\sigma \models K_a Exp_{A-b}(A-b)$ and $\sigma \models K_a \neg Exp_a$ we obtain $\tau \models Exp_{A-b}(A-b)$ and $\tau \models \neg Exp_a$.

These two imply that $\tau^a_{\ll} \models \neg b_d$ for any $d \neq b$ (b does not occur in $\tau^a_{\ll}$) and therefore in particular that $\tau^a_{\ll} \models \neg b_c$.

From $\tau^a_{\ll} \models \neg b_c$ whereas $\tau \models b_c$ it follows that $\tau^a_{\ll}$ must contain a call bd involving b and some agent d where either $d = c$ or there is a subsequent call ce involving c (where $d, e \neq a$). Because $d(\tau^a_{\ll}) = A-b$ for any such d, in the call bd agent b becomes an expert. (Also, ce explains how agent c became an expert.)

Therefore $\tau \models Exp_b$, and as τ was arbitrary such that $\tau.ac \sim_a \sigma.ac$ this shows that $\sigma.ac \models K_a Exp_b$. □

We can conclude from Lemma 4 that if $\sigma \models E_{A-b} Exp_{A-b}(A-b)$, so when all agents in $A-b$ are $A-b$ super experts, and if also no agent in $A-b$ is expert, all but one of those can be lucky in the same call sequence. To see this, after σ let some agent c call b. Let now all agents in $A-bc$ call c. Then they all also learn in that call that b is expert. Given $|A| = n$ agents, we therefore get $n-2$ many lucky calls.

Example 3. With 4 agents we can have 2 lucky calls. Recall $ac.ad.ac.bc.ac$ from Example 2 wherein a learns that b and c are expert in the final call ac. Now consider the expanded sequence $ac.ad.ac.cd.bc.ac.cd$: after $ac.ad.ac.cd$, all of a, c, d know that all of a, c, d know all secrets of a, c, d. In penultimate call ac, a learns that b and c are expert, and in final call cd, d learns that b and c are expert.

6 Main Result

Lemma 5. *An agent cannot become an expert and a super expert in the same call.*

Proof. Suppose a becomes expert in ac, i.e. $\sigma \not\models Exp_a$ and $\sigma.ac \models Exp_a$. We need to show that $\sigma.ac \not\models K_a Exp_A$. From $\sigma \not\models Exp_a$ we know there must be a $d \in A$ such that $\sigma \models \neg d_a$. As there are at least four agents, there must be a $b \notin \{a, c, d\}$.

From $\sigma \models \neg d_a$ we get $\sigma \not\models Exp_a(A-b)$ and thus $\sigma \not\models K_a Exp_{(A-b)}(A-b)$. Now by Lemma 4 we have $\sigma.ac \not\models K_a Exp_b$. This implies $\sigma.ac \not\models K_a Exp_A$. □

Lemma 6. *An agent becoming a super expert considers it possible that the other agent involved in that call did not become a super expert.*

Proof. Let $\sigma.ab$ be a call sequence wherein agent a becomes super expert in final call ab. From Lemma 5 we conclude that a was already expert after σ.

Suppose that b became expert in the call ab. Then also from Lemma 5 we conclude that b did not become a super expert in the call ab. As a considers the actual call sequence possible, we infer that agent a then considers it possible that agent b is not a super expert after $\sigma.ab$.

Thus we can assume that b already was an expert after σ. As a was expert before call ab, a became super expert by learning in call ab that b is expert.

First suppose a considers possible that b made a lucky call in σ. The lucky call was not with a, as a would then already have known that b is an expert. So let that call be bc with some agent $c \neq a$. If in that call b also learns that d is expert, where $d \neq a, c$, then a considers possible that the call instead of bc was bd. In call bd agent b only learns that d is an expert (we recall that by Lemma 4 agent b must have been an $A-d$ superexpert before that call), and does not learn that c is expert. So a considers a call sequence possible wherein b was not lucky when he became expert, and therefore when becoming expert remained uncertain whether some other agent c is expert.

We continue the argument by reasoning about this agent c.

Suppose a learnt that c is expert in lucky call ad. Then a considers possible that all further calls involving c were instead involving d except for further calls ac. After this call sequence b does not know that c is expert, so b is not a super expert, so a considers possible after $\sigma.ab$ that b is not a super expert.

Now suppose that ac is the first call in σ after which a knows that c is expert. We now distinguish four cases by whether c and b are experts before the call ac.
1 Suppose c was already expert before that call ac.
1.1 If b became expert before call ac in call bd with $d \neq c$, b was ignorant whether c is expert after that call (we assumed bd was not lucky), replace bd by $bd.cd$ in the call sequence between bd and ac, and replace all further occurrences of c except in further calls ac by d. This call sequence is indistinguishable for a and preserves that b does not know that c is expert (after final call ab). Therefore a considers possible after $\sigma.ab$ that b is not a super expert.
1.2 If b became expert after call ac in call bd with $d \neq c$, replace all subsequent occurrences of c except in calls ac by d. (We need not change any calls between ac and bd, as we assumed that b does not learn whether any other agent than d is expert in call bd, which includes agent c.) Then b does not know after final call ab that c is expert. Therefore a considers possible after $\sigma.ab$ that b is not a super expert.
2 Suppose c became expert in that call ac.
2.1 If b became expert before ac in bd with $d \neq c$, agent a considers possible that all further calls after ac involving c were instead by d, except for further

occurrences of ac. We need not change any calls involving c between bd and ac: note that this cannot have been bc as that would have made c expert before ac contrary to our assumption. This call sequence preserves that b does not know whether c is expert after final call ab. Therefore a considers possible after $\sigma.ab$ that b is not a super expert.
2.2 If b became expert after ac in bd with $d \neq c$ agent a considers possible all further calls after bd involving c were instead by d except for future occurrences of ac. Now consider the calls involving c between ac and bd: we need not change any of those as we assumed that b only learnt that d is expert in call bd. This call sequence preserves that b does not know whether c is expert after final call ab. Therefore a considers possible after $\sigma.ab$ that b is not a super expert.

We have exhaustively investigated all cases and this ends the proof. □

In Lemma 6 it is important to observe that the agent c that agent a remains uncertain about is different from the agent b involved in the call wherein she became super expert. This will be used in the final theorem.

Example 4. As an example of Lemma 6, recall the sequence $ab.cd.ac.bd.ad.bc.ab$ from Example 1 above where a and b become super experts in the last call. It is indistinguishable for a from call sequence $ab.cd.ac.bd.ad.ab$ after which b is not a super expert. Alternatively we could replace bc by bd in the actual sequence.

With Lemma 6 we can now prove our main result.

Theorem 1. *$EEExp_A$ is unsatisfiable.*

Proof. Let ρ be an arbitrary call sequence. We show that $\rho \not\models EEExp_A$. If $\rho \not\models EExp_A$, then clearly $\rho \not\models EEExp_A$. Hence we assume that $\rho \models EExp_A$.

Consider any agent a becoming a super expert in ρ, in other words choose σ and τ such that $\rho = \sigma.ab.\tau$ where in call ab agent a becomes a super expert. We will show that $\rho \not\models K_a K_b Exp_c$ for some $c \in A$.

From Lemma 6 it follows that after $\sigma.ab$ agent a considers possible a call sequence $\sigma'.ab$ after which b is not a super expert. Therefore, b considers possible a call sequence $\sigma''.ab$ that does not satisfy Exp_A, that is, $\sigma''.ab$ does not satisfy Exp_c for some $c \in A$. Clearly we must have that $c \neq a$ (because b learnt that a is an expert in the call ab). Formally, we have: $\sigma.ab \sim_a \sigma'.ab$ and $\sigma'.ab \sim_b \sigma''.ab$ and $\sigma''.ab \models \neg Exp_c$. The last implies $\sigma'' \models \neg Exp_c$.

The sequence $\sigma.ab.\tau$ is indistinguishable for a from $\sigma.ab.\tau'$ where τ' is τ without all calls involving b but not a. This is because no secrets are exchanged in any call in τ, because all agents are already experts (because a is a super expert after $\sigma.ab$).

For the same reason, $\sigma.ab.\tau'$ is indistinguishable for b from $\sigma.ab.\tau''$ where τ'' is τ' restricted to calls involving b. Call sequence τ'' is a finite and possibly empty sequence consisting only of calls ab. We can write ab^n for that, where $n \in \mathbb{N}$ is the number of occurrences of ab in τ (and τ').

First, from $\sigma.ab \sim_a \sigma.ab$ (reflexivity of $\sim_a$), $\tau \sim_a \tau'$, and the fact that all are experts before τ and τ' (so no new secrets are learnt in calls) we get

$\sigma.ab.\tau \sim_a \sigma.ab.\tau'$. Then, from the assumption $\sigma.ab \sim_a \sigma'.ab$, $\tau' \sim_a \tau'$ (reflexivity of $\sim_a$ again), and that all are experts before τ', we get $\sigma.ab.\tau' \sim_a \sigma'.ab.\tau'$. Finally, from $\sigma.ab.\tau \sim_a \sigma.ab.\tau'$ and $\sigma.ab.\tau' \sim_a \sigma'.ab.\tau'$ and transitivity of $\sim_a$ we then get $\sigma.ab.\tau \sim_a \sigma'.ab.\tau'$.

Similarly, from $\sigma'.ab \sim_b \sigma'.ab$ and $\tau' \sim_b \tau''$ we get $\sigma'.ab.\tau' \sim_b \sigma'.ab.\tau''$, from $\sigma'.ab \sim_b \sigma''.ab$ and $\tau'' \sim_b \tau''$ we get $\sigma'.ab.\tau'' \sim_b \sigma''.ab.\tau''$ (although c is not an expert in σ'', a and b are, so as above no new secrets are learnt in calls in τ''), and therefore from both we obtain $\sigma'.ab.\tau' \sim_b \sigma''.ab.\tau''$.

From $\sigma'' \not\models Exp_c$ and $\sigma''.ab.\tau'' = \sigma''.ab^{n+1}$ we obtain $\sigma''.ab.\tau'' \not\models Exp_c$. Finally, in view of the above $\sigma'.ab.\tau' \not\models K_b Exp_c$, and also $\sigma.ab.\tau \not\models K_a K_b Exp_c$.

This implies $\rho \not\models EEExp_A$ and because ρ was an arbitrary call sequence, we have shown $EEExp_A$ must be unsatisfiable. □

7 Conclusion and Open Questions

We have shown that "everyone knows that everyone knows all secrets" is the maximum level of shared knowledge that can be reached in asynchronous gossip. Formally, we can reach $EExp_A$ but the next level $EEExp_A$ is unsatisfiable.

Our results suggest at least two immediate related questions. First, it remains open whether $K_i EExp_A$ is satisfiable. That is, can at least one agent learn that everyone is a super-expert? Second, a footnote to [3]: maybe one does not need to consider arbitrary epistemic formulas, but could show that beyond a certain modal depth *any formula* of a certain kind (with positively stacked modalities, i.e. excluding negations before epistemic modalities) is unsatisfiable. That would simplify model checking gossip.

Finally, in ongoing work we study *minimal* call sequences reaching super success. We conjecture that at least $2n-3$ calls are needed to reach $K_i Exp_A$ for some agent i, and that at least $n-2+\binom{n}{2}$ calls are needed to reach $EExp_A$.

References

1. Apt, K., Wojtczak, D.: Verification of distributed epistemic gossip protocols. J. Artif. Intell. Res. **62**, 101–132 (2018). https://doi.org/10.1613/jair.1.11204
2. van Ditmarsch, H., Gattinger, M., Kuijer, L., Pardo, P.: Strengthening gossip protocols using protocol-dependent knowledge. Journal of Applied Logics - IfCoLog Journal of Logics and their Applications **6**(1), 157–203 (2019), `https://arxiv.org/abs/1907.12321`
3. van Ditmarsch, H., van der Hoek, W., Kuijer, L.: The logic of gossiping. Artificial Intelligence **286**, 103306 (2020). https://doi.org/10.1016/j.artint.2020.103306
4. van Ditmarsch, H., Gattinger, M., Ramezanian, R.: Everyone knows that everyone knows: Gossip protocols for super experts. CoRR **abs/2011.13203** (2020), `https://arxiv.org/abs/2011.13203`
5. 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

6. Hedetniemi, S.M., Hedetniemi, S.T., Liestman, A.L.: A survey of gossiping and broadcasting in communication networks. Networks **18**(4), 319—-349 (1988). https://doi.org/10.1002/net.3230180406
7. Herzig, A., Maffre, F.: How to share knowledge by gossiping. AI Communications **30**(1), 1—-17 (2017). https://doi.org/10.3233/aic-170723
8. Lamport, L., Shostak, R., Pease, M.: The byzantine generals problem. ACM Transactions on Programming Languages and Systems **4**(3), 382–401 (1982). https://doi.org/10.1145/357172.357176
9. Tijdeman, R.: On a telephone problem. Nieuw Archief voor Wiskunde **3**(19), 188–192 (1971)