Completeness and Decidability of Protocol-Dependent Knowledge in Gossip

Hans van Ditmarsch[1], Malvin Gattinger[2], and Wouter J. Smit[2]

[1] CNRS, IRIT, University of Toulouse, France
[2] ILLC, University of Amsterdam, The Netherlands

Abstract. In the gossip problem, a group of agents aims to efficiently share information using one-to-one communication. This often occurs in decentralised systems, where agents must rely on protocols to efficiently coordinate their communication. Recent work has used epistemic logic to define gossip protocols, including protocol-dependent knowledge modalities: agent knowledge assuming common knowledge that all agents follow said protocol. While axiomatisations exist for various versions of the gossip problem, none of these include protocol-dependent knowledge. We show that protocol-dependent knowledge is strictly more expressive than standard knowledge, and we provide axiomatisations for four logics of gossip with protocol-dependent knowledge. We show that all four axiomatisations are sound and complete, as well as decidable.

1 Introduction

The gossip problem, introduced in the 1970's as the *telephone problem* [4,20] addresses how to spread information (secrets) among a group of agents by sequential pairwise communication, often called telephone calls. The goal is for all agents to know all secrets. We assume that each agent holds a single secret and that when calling each other, both agents exchange all the secrets they know, but no other information. There are many variations of the gossip problem: *one-directional exchange*, when only one agent in a call informs the other but not vice versa; multiple calls made in *parallel*; restrictions on the *network*, limiting what agents you can reach; *(a)synchronicity* determining whether agents are aware of a global clock when making calls; and more — see [16] for a survey.

Initially, research focused on finding optimal sequences of calls to achieve the goal that all agents know all secrets. Such optimal sequences require a central scheduler to plan calls. Later publications shifted to *distributed* gossip [18]. In the absence of a central scheduler, agents must rely on some protocol to coordinate their calls. Often, they rely on making calls at random. More recent developments focus on *epistemic* gossip protocols which specify pre-conditions that an agent must know are true before making a call [3,1]. There exist many such distributed epistemic protocols [1,8,9]. Other aspects of gossip protocols may also be epistemic, such as the *termination goal* (when the goal is not merely

This article is based on the Master's thesis of the third author [19].

that all secrets are known by all agents, but that they also know this) [10,11], or the *messages* being sent [7,17]. Another variation is to allow *faulty* messages [5].

To analyse how the information develops over the course of a protocol's execution we can use a dynamic epistemic logic with knowledge modalities for individual agents and a dynamic modality for calls. Various such logics have been axiomatised [14]. In these settings and even when all agents follow a given protocol, they do not assume that the other agents follow the same protocol. That is, their knowledge is not based on other agents also selecting calls on the same ground as themselves. In [12] the authors propose so-called *protocol-dependent knowledge* (see also [10,15]), which are epistemic modalities formalizing what an agent knows *given that it is common knowledge that all agents follow some protocol*. Strengthening gossip protocols with these modalities makes some protocols more successful [12]. However, the corresponding logics with protocol-dependent knowledge modalities have not yet been axiomatised.

We show that protocol-dependent knowledge is strictly more expressive than standard knowledge and provide axiomatisations for four logics of gossip with protocol-dependent knowledge modalities. To axiomatise dynamic modalities for calls, we use the technique of reduction axioms by which a formula with dynamic modalities is provably equivalent to one without. We show that all four axiomatisations are sound and complete, and that all four logics are decidable.

We proceed as follows. We first recall the definitions of gossip and protocol-dependent knowledge in Section 2 before proving our expressivity results in Section 3. We then introduce our proof systems. In Section 4 we present the logic of initial models in a call-free setting, in Section 5 we define call reductions, and in Section 6 the logic of gossip models including calls. In the latter we also show that these proof systems are decidable. Section 7 concludes the article.

2 Definitions

We assume a finite set of *agents* $a, b, \ldots \in \mathbb{A}$. Each agent knows a corresponding *secret*, also denoted $a, b, \ldots$ and $\mathbb{S}$ is the set of all secrets. We only express the secret distribution among the agents, not the content of the secrets.

A *call* is a pair of agents $a \neq b \in \mathbb{A}$. A *call sequence* $\sigma \in \Sigma$ is a list of calls. For n agents there are $n \cdot (n-1)$ different calls and all calls can always be executed. We write ab as shorthand for the call (a, b) and use a period to concatenate calls into a sequence, e.g. $ab.cd$.

Gossip has been studied under various settings [14,16]. We consider synchronous gossip with symmetric calls, i.e. agents always know how many calls took place but not necessarily which, and the caller and callee both share all secrets they know. Agents *inspect-then-merge* the secrets the other agent knows, meaning they first find out what the other agent knew before the call. they share no higher-order epistemic information or other information in a call. Agents form a total network, meaning they can always call anybody else, and calls are one-to-one.

We define the syntax with the following two mutually recursive definitions.

Definition 1 (Language). *Let $a, b \in \mathbb{A}$ be agents and $P \in \mathbb{P}$ a protocol. We define the language $\mathcal{L}^{\mathbb{P}}$ as*

$$\varphi ::= S_a b \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a^P \varphi \mid [ab]\varphi.$$

The syntax extends standard epistemic logic. The atom $S_a b$ means that *agent a knows the secret of agent b*. The protocol-dependent knowledge modality $K_a^P \varphi$ says that *assuming it is common knowledge that protocol P is followed, agent a knows that φ*. The dynamic modality $[ab]\varphi$ means that *φ holds after call ab*.

Protocols are essentially (distributed) algorithms that select calls to execute until some goal is achieved, usually that all agents know all secrets. For the purpose of this article it suffices to define a protocol by its protocol conditions, which specifies for each call when it may be executed.

Definition 2 (Protocols). *A* gossip protocol *P is a list of $n \cdot (n-1)$* protocol conditions *$P_{ab} \in \mathcal{L}^{\mathbb{P}}$ for agents $a \neq b \in \mathbb{A}$. Let $\mathbb{P}$ be the set of all protocols.*

Protocols cannot reference themselves, either directly or indirectly. This means that P_{ab} may not contain a modality K^Q such that $Q = P$ or Q references P.

The operators $\vee$, $\rightarrow$, $\top$, and $\bot$ are defined as abbreviations in the usual way. We also define the epistemic dual of K_a^P by $\widehat{K}_a^P := \neg K_a^P \neg$. While $\mathcal{L}^{\mathbb{P}}$ does not contain the standard knowledge operator K, we will use it as abbreviation for K^{ANY}, knowledge assuming the trivial protocol ANY defined in Example 3.

For any agent a and set of secrets $R \subseteq \mathbb{S}$ we furthermore define the following abbreviation to describe that *a only knows* the secrets in R.

$$O_a R := \textstyle\bigwedge_{b \in R} S_a b \wedge \bigwedge_{b \notin R} \neg S_a b \qquad \text{“}a \text{ only knows the secrets in } R\text{”}$$

We also define variations of this language. Firstly we define the basic language of gossip $\mathcal{L}$ by replacing K^P with the standard knowledge operator K.

The static fragments $\mathcal{L}_-^{\mathbb{P}}$ and $\mathcal{L}_-$ of these languages omit the dynamic modality $[ab]$. As this modality represents a call in the setting of gossip, we also use the term *call-free*.

Agents can coordinate their calls by means of protocols. A protocol stipulates whether a call is permitted at a certain point, characterised by a protocol condition. There are various well-known gossip protocols, such as *Learn New Secrets* (LNS) where agents are only allowed to make a call to an agent whose secret they do not yet know. We also use the trivial protocol *Any Call* (ANY).

Example 3 (Protocols LNS *and* ANY*).* The protocols LNS and ANY are given by the following protocol conditions $\mathsf{LNS}_{ab} := \neg S_a b$ and $\mathsf{ANY}_{ab} := \top$, respectively.

We view the protocol conditions P_{ab} as subformulas of the K^P modality, which gives rise to the following definition of modal degree.

Definition 4 (Modal Degree). *The* modal degree $d(\varphi)$ *of a formula* $\varphi \in \mathcal{L}^{\mathbb{P}}$ *is defined recursively by*

$$\begin{array}{ll} d(S_a b) & := 0, \\ d(\neg\varphi) & := d(\varphi), \\ d(\varphi \wedge \psi) & := \max\{d(\varphi), d(\psi)\}, \\ d(K_a^P \varphi) & := 1 + \max\{d(\varphi), d(P)\}, \\ d([ab]\varphi) & := d(\varphi), \end{array}$$

where $d(P) := \max\{d(P_{ab}) \mid a \neq b \in \mathbb{A}\}$ *is the* degree of protocol P.

We now define gossip models in two steps: an initial model represents the situation before any calls have happened and is then lifted to a gossip model that includes calls. Initial models with n agents are essentially $\mathbf{S5}_n$ Kripke models.

The following definition by [14] models arbitrary initial settings of gossip without protocol-dependent knowledge. These allow any secret distribution and knowledge thereof, as long as agents at least know their own secret and are aware which secrets they know. Protocols only effectively change the semantics after calls have taken place. We can therefore use the same definition for protocol-dependent knowledge too.

Definition 5 (Initial Models). *An* initial model *is a triple* $I = \langle W_0, R_0, V_0 \rangle$ *where* W_0 *is a set of worlds,* $R_0 : \mathbb{A} \to 2^{W_0 \times W_0}$ *is an equivalence relation for each agent, and* $V_0 : \mathbb{A} \times W_0 \to 2^{\mathbb{S}}$ *is a function mapping agent-world pairs to sets of secrets, such that (i)* $a \in V_0(a, w)$ *for all* $w \in W$*, and (ii) if we have* $(w_1, w_2) \in R_0(a)$ *then* $V_0(a, w_1) = V_0(a, w_2)$.

We can lift any initial model to a gossip model by inducing calls [14]. As calls are always possible, this creates for each world in the initial model an $n \cdot (n-1)$ branching tree. The resulting models are therefore forests.

Definition 6 (Gossip States). *Given an initial model* $I = \langle W_0, R_0, V_0 \rangle$*, a* gossip state *is a pair* (w, σ) *where* $w \in W_0$ *and* σ *is a call sequence. The set of gossip states induced from* I *is* $W(I) := W_0 \times \Sigma$.

Definition 7 (Valuation). *Given an initial model* $I = \langle W_0, R_0, V_0 \rangle$ *and some gossip state* $(w, \sigma) \in W(I)$*, we denote the set of secrets known by agent* a*, by* $V_a(w, \sigma)$*. We define* V *recursively as follows.*

$$\begin{array}{lll} V_a(w, \epsilon) & = V_0(a, w) & \textit{Empty sequence} \\ V_a(w, \sigma.bc) & = V_b(w, \sigma) \cup V_c(w, \sigma) & \textit{iff } a \in \{b, c\} \\ V_a(w, \sigma.bc) & = V_a(w, \sigma) & \textit{iff } a \notin \{b, c\} \end{array}$$

The semantics on gossip models are now defined as follows, with Definitions 8 to 10 being mutually recursive. The effects of protocols emerge in Definition 8 of the epistemic accessibility relation. While calls are always possible, only those calls that are P-permitted are considered by agent a under the relation $\sim_a^P$. The epistemic relation $\sim_a$ for standard gossip as in [14, Definition 3.5] instantiated for synchronous, bi-directional calls with inspect-then-merge can be retrieved by omitting the protocol conditions, marked with $(*)$.

Definition 8 (Epistemic Relation). *Given an initial model $I = \langle W_0, R_0, V_0\rangle$, we lift the initial relation R_0 for each $P \in \mathbb{P}$ and agent $a \in \mathbb{A}$ to $\sim_a^P \subseteq W(I)\times W(I)$ such that:*

$$
\begin{array}{lll}
(w_i,\epsilon) \sim_a^P (w_j,\epsilon) & \textit{iff} & (w_i,w_j) \in R_0(a); \\
(w_i,\sigma.ab) \sim_a^P (w_j,\tau.ab) & \textit{iff} & (w_i,\sigma) \sim_a^P (w_j,\tau) \\
 & & \quad \textit{and } V_b(w_i,\sigma) = V_b(w_j,\tau) \\
 & & \quad \textit{and } (w_i,\sigma) \models P_{ab} \textit{ and } (w_j,\tau) \models P_{ab}; \qquad (*) \\
(w_i,\sigma.ba) \sim_a^P (w_j,\tau.ba) & \textit{iff} & (w_i,\sigma) \sim_a^P (w_j,\tau) \\
 & & \quad \textit{and } V_b(w_i,\sigma) = V_b(w_j,\tau) \\
 & & \quad \textit{and } (w_i,\sigma) \models P_{ba} \textit{ and } (w_j,\tau) \models P_{ba}; \qquad (*) \\
(w_i,\sigma.bc) \sim_a^P (w_j,\tau.de) & \textit{iff} & (w_i,\sigma) \sim_a^P (w_j,\tau) \\
 & & \quad \textit{and } a \notin \{b,c,d,e\} \\
 & & \quad \textit{and } (w_i,\sigma) \models P_{bc} \textit{ and } (w_j,\tau) \models P_{de}. \qquad (*)
\end{array}
$$

Definition 9 (Gossip Models). *For an initial model I, we define the (induced) gossip model $M(I) := \langle W(I), \sim, V\rangle$ with $W(I)$, $\sim$, and V from Definitions 6 to 8. When the initial model is clear from context, we omit it and write M.*

Definition 10 (Semantics). *Let $\varphi, \psi \in \mathcal{L}^{\mathbb{P}}$. We define the relation $\models$ between pointed gossip models and formulas as follows.*

$$
\begin{array}{lll}
M,(w,\sigma) \models S_a b & \iff & b \in V_a(w,\sigma) \\
M,(w,\sigma) \models \neg\varphi & \iff & M,(w,\sigma) \not\models \varphi \\
M,(w,\sigma) \models \varphi \wedge \psi & \iff & M,(w,\sigma) \models \psi \textit{ and } M,(w,\sigma) \models \psi \\
M,(w,\sigma) \models K_a^P \varphi & \iff & M,(w,\sigma') \models \varphi \textit{ for all } (w,\sigma') \textit{ s.t. } (w,\sigma) \sim_a^P (w,\sigma') \\
M,(w,\sigma) \models [ab]\varphi & \iff & M,(w,\sigma.ab) \models \varphi
\end{array}
$$

When M and w are clear from context, we omit them and write $\sigma \models \varphi$.

We also define semantics for the call-free language on initial models such that $I, w \models_{\mathcal{I}} \varphi$ iff $M(I),(w,\epsilon) \models \varphi$ for all $\varphi \in \mathcal{L}_-^{\mathbb{P}}$.

Effectively, the epistemic relation for protocol P is restricted to call sequences that are P-*permitted*.

Definition 11 (P-Permitted Calls). *Given a protocol P, a call ab is* P-permitted *at state $M,(w,\sigma)$ if its protocol condition P_{ab} holds: $M,(w,\sigma) \models P_{ab}$. A call sequence σ is P-permitted if each call is P-permitted. A call or sequence is* P-illegal *if it is not P-permitted. We omit P if the protocol is clear from context.*

Protocol-dependent knowledge only differs from standard knowledge *after* calls happen. In initial models no calls have happened yet, hence knowledge does not yet depend on which protocol an agent assumes.

Lemma 12. *Let P and Q be any two protocols and let a be any agent. Then for any initial model I and world $w \in W(I)$ we have $I, w \models_{\mathcal{I}} K_a^P\varphi$ iff $I, w \models_{\mathcal{I}} K_a^Q\varphi$.*

Proof. By the semantics, in particular note that the ϵ clause of Definition 8 does not depend on the protocol but only on R_0. □

While initial models are **S5**, protocol-dependent gossip models are not: the epistemic relation for a protocol P excludes any states with P-illegal call sequences, breaking reflexivity. They are instead partial equivalence relations. As a side-effect of these semantics, the P-dependent knowledge-base in P-illegal states becomes inconsistent, meaning that for any agent a the state satisfies $K_a^P \bot$. This property is known as the *global alarm* [12].

We can describe any distribution of secrets or knowledge thereof with an initial model. However, the classic gossip problem assumes a specific initial setting: in the *root model* all agents know only their own secret and this is common knowledge among the agents.

Definition 13 (Root Model). *The* root model $I_{\mathcal{R}}$ *is the initial model with* $W_0 = \{w_{\mathcal{R}}\}$*, and* $R_0(a) = \{(w_{\mathcal{R}}, w_{\mathcal{R}})\}$ *and* $V_0(a, w_{\mathcal{R}}) = \{a\}$ *for all agents* a*. We write* $\mathcal{I}$ *for the singleton class of* $I_{\mathcal{R}}$*.*

We call the gossip model induced from $I_{\mathcal{R}}$ the *tree model*, as it contains a single tree rooted in $w_{\mathcal{R}}$. While other single-tree models exist, we use the name exclusively for this model. It is semantically equivalent to models used in literature that assume the classical setting [2,11,12,13].

Definition 14 (Tree Model). *The* tree model $M_{\mathcal{T}} := M(I_{\mathcal{R}})$ *is the gossip model induced from the root model. We write* $\mathcal{T}$ *for the singleton class of* $M_{\mathcal{T}}$*.*

We have now defined four classes of models: $\mathcal{G}$, $\mathcal{I}$, $\mathcal{R}$ and $\mathcal{T}$. We visualise the relation between them in Figure 1 below. The main class is $\mathcal{G}$. It consists of the induced gossip models of all initial models in $\mathcal{I}$, i.e. the execution trees with an initial model as their root. One particular model is the tree model $M_{\mathcal{T}} \in \mathcal{T}$ induced by the root model $I_{\mathcal{R}} \in \mathcal{R}$.

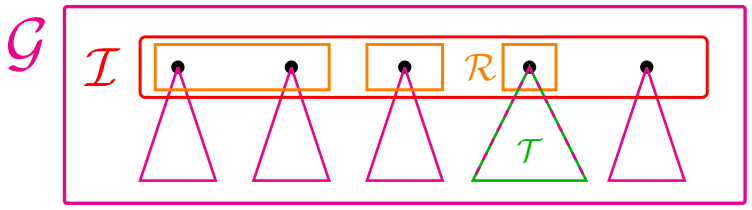

Fig. 1. The classes $\mathcal{I}$, $\mathcal{G}$, $\mathcal{R}$, and $\mathcal{T}$.

3 Expressivity of Protocol-Dependent Knowledge

We now show that $\mathcal{L}^{\mathbb{P}}$ is strictly more expressive than $\mathcal{L}$. To this end, we first show that $\mathcal{L}^{\mathbb{P}}$ can express the length of call sequences using protocol-dependent knowledge modalities. We do so by recursively defining protocols that depend on the violation of the previous.

Definition 15 (Counting Protocols). *For all natural numbers $k \geq 0$ and all agents $a \neq b$, and an arbitrary agent u we define P^k_{ab} recursively as follows.*

$$\begin{aligned} P^0_{ab} &:= \bot && \text{“Allow no calls”} \\ P^{k+1}_{ab} &:= \widehat{K}_u^{P^k} \top && \text{“The previous protocol has not been violated”} \end{aligned}$$

These protocols leverage the global alarm property. A sequence violates protocol P^n iff it exceeds length n. Recall however that a protocol can only be violated by an illegal call: the lack of a call can never cause violation. A counting protocol therefore provides an upper bound, but not a lower bound.

Lemma 16. *For all agents u, call sequences σ and all $k \geq 0$ we have*

$$M, (w, \sigma) \models \widehat{K}_u^{P^k} \top \text{ if and only if } |\sigma| \leq k.$$

Proof. Let agent u and call sequence σ be arbitrary. We use induction on k.
Base case. Suppose $k = 0$.

($\Longrightarrow$) Suppose $M, (w, \sigma) \models \widehat{K}_u^{P^0} \top$. Hence σ has not violated the protocol P^0. However, $P^0_{ab} = \bot$ for all calls ab, so no calls are P^0-permitted. Therefore we find that $\sigma = \epsilon$ and conclude that $|\sigma| \leq 0$.

($\Longleftarrow$) Suppose $|\sigma| \leq 0$. Then $\sigma = \epsilon$. The empty call sequence cannot violate any protocol, also not P^0. Hence $M, (w, \sigma) \models \neg K_u^{P^0} \bot$, i.e. $M, (w, \sigma) \models \widehat{K}_u^{P^0} \top$.
Induction Hypothesis. Let k be arbitrary and suppose we have $M, (w, \sigma) \models \widehat{K}_u^{P^k} \top$ if and only if $|\sigma| \leq k$.
Induction Step. Suppose the induction hypothesis holds for k. We show it holds for $k+1$. Let σ be an arbitrary call sequence such that $|\sigma| = k+1$ and write it as $\sigma = \tau.ab$. Using the semantics of $[ab]$ and $\widehat{K}_u$ we get the following equivalences:

$$\begin{aligned} M, (w, \tau.ab) \models \widehat{K}_u^{P^{k+1}} \top &\iff M, (w, \tau) \models [ab] \widehat{K}_u^{P^{k+1}} \top \\ &\iff M, (w, \tau) \models \neg [ab] K_u^{P^{k+1}} \bot \end{aligned}$$

We now distinguish three cases whether agent u is one of the agents a and b.

Suppose $u \notin \{a, b\}$. We get the following equivalences. The steps **Pri** and **Fnc** follow from the soundness of the axioms with the same name, that we will show below in Lemma 24. For the step at $(*)$ we use the definition of P^{k+1}_{ab} in both directions. Observe for the backwards direction that P^k is not violated at τ, so neither is P^{k+1}. This means we have $\tau \sim_u^{P^{k+1}} \tau$ which is sufficient to obtain $\tau \models \widehat{K}_u^{P^{k+1}} (P^{k+1}_{ab} \wedge [ab]\top)$. Hence τ satisfies the disjunct for $de = ab$.

$$\begin{aligned} & \tau \models \neg [ab] K_u^{P^{k+1}} \bot \\ \iff & \tau \not\models P^{k+1}_{ab} \to \textstyle\bigwedge_{d,e \neq u} K_u^{P^{k+1}} (P^{k+1}_{de} \to [de]\bot) && \textbf{(Pri)} \\ \iff & \tau \models P^{k+1}_{ab} \wedge \textstyle\bigvee_{d,e \neq u} \neg (K_u^{P^{k+1}} (P^{k+1}_{de} \to [de]\bot)) && \text{(De Morgan)} \\ \iff & \tau \models P^{k+1}_{ab} \wedge \textstyle\bigvee_{d,e \neq u} \widehat{K}_u^{P^{k+1}} (P^{k+1}_{de} \wedge \neg [de]\bot) && \text{(Sem. } \widehat{K}_u) \\ \iff & \tau \models P^{k+1}_{ab} \wedge \textstyle\bigvee_{d,e \neq u} \widehat{K}_u^{P^{k+1}} (P^{k+1}_{de} \wedge [de]\top) && \textbf{(Fnc)} \\ \iff & \tau \models \widehat{K}_u^{P^k} \top && (*) \end{aligned}$$

Suppose $u = a$. We get the following equivalences.The step $\mathbf{Obs}_1$ follows from Lemma 24 below. For the backwards direction of the step at (†), observe that there is precisely one set $Q \subseteq \mathbb{S}$ such that O_bQ. For all other sets $T \neq Q$ the conjunct is satisfied with $\neg O_bT$. The conjunct for $R = Q$ is satisfied by the knowledge operator: we use again the reflexive relation $\tau \sim^{P^{k+1}} \tau$ to obtain $\tau \models \widehat{K}_u^{P^{k+1}}(P_{ab}^{k+1} \wedge O_bQ \wedge [ab]\top)$.

$$
\begin{array}{lll}
& \tau \models \neg[ab]K_u^{P^{k+1}}\bot & \\
\iff & \tau \not\models P_{ab} \to \bigvee_{R\subseteq\mathbb{S}}(O_bR \wedge K_u^P(P_{ab} \to (O_bR \to [ab]\bot))) & (\mathbf{Obs}_1)\\
\iff & \tau \models P_{ab} \wedge \bigwedge_{R\subseteq\mathbb{S}} \neg O_bR \vee \neg K_u^P(P_{ab} \to (O_bR \to [ab]\bot)) & (\text{De Morgan})\\
\iff & \tau \models P_{ab} \wedge \bigwedge_{R\subseteq\mathbb{S}} \neg O_bR \vee \widehat{K}_u^P(P_{ab} \wedge O_bR \wedge \neg[ab]\bot) & (\text{Sem. } \widehat{K}_u)\\
\iff & \tau \models P_{ab} \wedge \bigwedge_{R\subseteq\mathbb{S}} \neg O_bR \vee \widehat{K}_u^P(P_{ab} \wedge O_bR \wedge [ab]\top) & (\mathbf{Fnc})\\
\iff & \tau \models \widehat{K}_u^{P^k}\top & (\dagger)
\end{array}
$$

Suppose $u = b$. We repeat the steps for $u = a$ and instead apply $\mathbf{Obs}_2$.

Hence $\tau.ab \models \widehat{K}_u^{P^{k+1}}\top \iff \tau \models \widehat{K}_u^{P^k}\top$. We finish by applying the induction hypothesis to find $\tau \models \widehat{K}_u^{P^k}\top \iff |\tau| \leq k \iff |\sigma| \leq k+1$. □

We negate the global alarm to invert the bound and define *counting formulas.*

Definition 17 (Counting Formulas). *We define* $\varphi_0 := \widehat{K}_u^{P^0}\top$ *and for every other natural number* $k \geq 1$*, let* $\varphi_k := \widehat{K}_u^{P^k}\top \wedge K_u^{P^{k-1}}\bot$.

Lemma 18. *For every call sequence* σ *and all* $k \in \mathbb{N}$ *we have* $\sigma \models \varphi_k$ *iff* $|\sigma| = k$.

Proof. Immediate by Lemma 16, as $K_u^{P^{k-1}}\bot$ is the negation of $\widehat{K}_u^{P^{k-1}}\top$. □

In order to show that the standard language of gossip $\mathcal{L}$ cannot express the length of call sequences, we give a standard definition of bisimulation. We can apply this definition to protocol-dependent gossip models as $\sim$ is equivalent to $\sim^{\mathsf{ANY}}$. Lemma 20 can be shown in the usual way [6, Theorem 2.20].

Definition 19 (Bisimulation). *Let* $M = \langle W, \sim, V\rangle$ *and* $M' = \langle W', \sim', V'\rangle$ *be gossip models. Two states are* bisimilar*, written* $M, s \leftrightarrows M', s'$*, if a bisimulation* Z *exists with* sZs'*. A relation* $Z \subseteq W \times W'$ *is a* bisimulation *if and only if for all gossip states* s, s' *such that* sZs' *we have:*

1. *(**Atoms**) For every agent* a *we have* $V_a(s) = V'_a(s')$*;*
2. *(**Forth**) For each agent* a*, if* $s \sim_a t$ *then there is a* $t' \in W'$ *such that* $s' \sim_a t'$ *and* tZt'*;*
3. *(**Back**) For each agent* a*, if* $s' \sim_a t'$ *then there is a* $t \in W$ *such that* $s \sim_a t$ *and* tZt'*.*

Lemma 20. *If two states are bisimilar, they satisfy the same formulas in* $\mathcal{L}$*.*

Theorem 21. $\mathcal{L}^{\mathbb{P}}$ *is strictly more expressive than* $\mathcal{L}$*.*

Proof. We show that there exist two states that satisfy the same formulas in $\mathcal{L}$ but that can be distinguished by a formula in $\mathcal{L}^{\mathbb{P}}$.

Let $I_{\mathcal{R}}$ be the initial root model for two agents a, b and $M_{\mathcal{T}}$ the gossip model induced from it. Consider call sequences $\sigma = ab$ and $\tau = ba.ab$.

We show that $M_{\mathcal{T}}, (w_{\mathcal{R}}, \sigma) \leftrightarrows M_{\mathcal{T}}, (w_{\mathcal{R}}, \tau)$. Let Z be an equivalence relation containing two equivalence classes: one containing only the empty sequence $(w_{\mathcal{R}}, \epsilon)$ and one containing the states for all other call sequences. By definition $(w_{\mathcal{R}}, \sigma) Z (w_{\mathcal{R}}, \tau)$ and we claim that Z is a bisimulation.

Observe that all atoms are satisfied after the first call, whether it is ab or ba. Next, both agents can distinguish each point in the model: they are involved in all calls. The epistemic relations of both agents contain only reflexive relations.

1. (**Atoms**) For all reflexive relations in Z, **Atoms** holds trivially. For all other $(w_{\mathcal{R}}, \sigma') Z (w_{\mathcal{R}}, \tau')$, observe that σ' and τ' must be non-empty sequences, as ϵ only occurs reflexively in Z. All secrets are already shared, so **Atoms** holds.
2. (**Forth**/**Back**) For each agent, $\sim^{\mathsf{ANY}}$ contains precisely one reflexive relation for each call sequence. Hence **Forth** and **Back** are satisfied.

Hence Z is a bisimulation. By Lemma 20 then $(w_{\mathcal{R}}, \sigma)$ and $(w_{\mathcal{R}}, \tau)$ agree on all $\varphi \in \mathcal{L}$. However, by Lemma 18 we find $\sigma \models \varphi_1$ and $\tau \not\models \varphi_1$ and $\varphi_1 \in \mathcal{L}^{\mathbb{P}}$. □

4 The Logic of Initial Models

We now provide call-free axiomatisations for the class of initial models $\mathcal{I}$ and the singleton class $\mathcal{R}$ of the root model $I_{\mathcal{R}}$. The logic for both is essentially **S5** plus axioms about agents knowing (only) their own secrets. The **PI** axiom ensures protocol invariance, which is required because protocol-dependent knowledge can only differ from standard knowledge after calls have happened.

Definition 22. *The proof system $\vdash_{\mathcal{R}}$ for the call-free language $\mathcal{L}^{\mathbb{P}}_{-}$ is defined as shown in Table 1. The system $\vdash_{\mathcal{I}}$ is obtained by omitting the* **Only** *axiom.*

Theorem 23. *The proof system $\vdash_{\mathcal{I}}$ is sound and complete for the class of all initial models $\mathcal{I}$: for all $\varphi \in \mathcal{L}^{\mathbb{P}}_{-}$ we have $\models_{\mathcal{I}} \varphi$ iff $\vdash_{\mathcal{I}} \varphi$. The system $\vdash_{\mathcal{R}}$ is sound and complete for the root model $I_{\mathcal{R}}$: for all $\varphi \in \mathcal{L}^{\mathbb{P}}_{-}$ we have $\models_{\mathcal{R}} \varphi$ iff $\vdash_{\mathcal{R}} \varphi$.*

Proof. Soundness follows directly from the semantics. Completeness can be shown using standard methods by building a canonical model [6, Section 4.2]. In particular we can define the canonical relation $R_0(a)$ over maximally consistent sets as usual using K^{ANY}. The **PI** axiom ensures that in any maximally consistent set K^P and K^Q for any two protocols P and Q agree with each other, matching Lemma 12 and therefore the semantics. □

Table 1. Rules and axioms of $\vdash_{\mathcal{R}}$. Omitting **Only** produces $\vdash_{\mathcal{I}}$.

Prop	propositional tautologies	**K**	$K_a^P(\varphi \to \psi) \to (K_a^P\varphi \to K_a^P\psi)$
MP	$\vdash \varphi, \vdash \varphi \to \psi$ imply $\vdash \psi$	**T**	$K_a^P\varphi \to \varphi$
Sub	$\vdash \varphi \leftrightarrow \psi$ implies $\vdash \chi \leftrightarrow \chi[\varphi/\psi]$	**4**	$K_a^P\varphi \to K_a^PK_a^P\varphi$
Own	S_aa	**5**	$\neg K_a^P\varphi \to K_a^P\neg K_a^P\varphi$
Only	O_aa	**Nec**	$\vdash \varphi$ implies $\vdash K_a^P\varphi$
PFi	$S_ab \to K_a^PS_ab$		
NPi	$\neg S_ab \to K_a^P\neg S_ab$	**PI**	$K^P\varphi \to K^Q\varphi$

5 Call Reductions

While in Section 3 we have seen that $\mathcal{L}^{\mathbb{P}}$ is more expressive than $\mathcal{L}$, we now show that the protocol-dependent language still shares a feature with the standard language, namely that the call operator $[ab]$ can be eliminated. We will use the established idea of *reduction axioms* to translate any formula to a call-free formula using the validities shown in Table 2. These validities originate from [14], and only the final three are adapted to protocol-dependent knowledge.

However, we will not actually stipulate these formulas as axioms for the proof systems $\vdash_{\mathcal{G}}$ and $\vdash_{\mathcal{T}}$. We only use them as rewrite rules and rely on their semantic validity to prove the soundness and completeness of $\vdash_{\mathcal{G}}$ and $\vdash_{\mathcal{T}}$.

Table 2. Call Reduction Validities on Gossip Models.

Call Basics		**Call Effects**		
Con	$[ab](\varphi \wedge \psi) \leftrightarrow ([ab]\varphi \wedge [ab]\psi)$	**Eff**	$[ab]S_cd \leftrightarrow (S_ad \vee S_bd)$	$c \in \{a, b\}$
Fnc	$[ab]\neg\varphi \leftrightarrow \neg[ab]\varphi$	**Ext**	$[ab]S_cd \leftrightarrow S_cd$	$c \notin \{a, b\}$

Calls and Protocol-Dependent Knowledge

Obs$_1$	$[ab]K_a^P\varphi \leftrightarrow (P_{ab} \to \bigvee_{R\subseteq\mathbb{S}}(O_bR \wedge K_a^P(P_{ab} \to (O_bR \to [ab]\varphi))))$	
Obs$_2$	$[ab]K_b^P\varphi \leftrightarrow (P_{ab} \to \bigvee_{R\subseteq\mathbb{S}}(O_aR \wedge K_b^P(P_{ab} \to (O_aR \to [ab]\varphi))))$	
Pri	$[ab]K_c^P\varphi \leftrightarrow (P_{ab} \to \bigwedge_{d,e\neq a} K_c^P(P_{de} \to [de]\varphi))$	$c \notin \{a, b\}$

Lemma 24. *All axioms in Table 2 are valid on the class of gossip models $\mathcal{G}$.*

Proof. The Call Basics are valid because calls are a deterministic action that can always be executed. The Call Effects are valid as they are equal to Definition 7.

To show **Obs**$_1$, let (w, σ) be an arbitrary state in some model M and let $\varphi \in \mathcal{L}^{\mathbb{P}}$ be arbitrary. We have the following chains of equivalences. Recall that $V_b(w, \sigma)$ is the set of secrets that agent b knows at state (w, σ). At step $(*)$ we use a disjunct to enumerate all possible sets of secrets $V_b(w, \sigma)$ that agent b might

know. There is precisely one set $V_b(w,\sigma) = R \subseteq \mathbb{S}$ such that $O_b R$ holds.

$$
\begin{aligned}
& (w,\sigma) \models [ab]K_a^P\varphi \\
\iff\ & (w,\sigma.ab) \models K_a^P\varphi && \text{(Sem. } [ab]) \\
\iff\ & \forall (w',\tau.de) \text{ s.t. } (w,\sigma.ab) \sim_a^P (w',\tau.de) : (w',\tau.de) \models \varphi && \text{(Sem. } K_a^P) \\
\iff\ & \forall (w',\tau) \text{ s.t. } (w,\sigma.ab) \sim_a^P (w',\tau.ab) : (w',\tau.ab) \models \varphi && \text{(Def. } \sim_a^P) \\
\iff\ & \forall (w',\tau) \text{ s.t. } (w,\sigma) \sim_a^P (w',\tau) : && \text{(Def. } \sim_a^P) \\
& \text{if } (w,\sigma) \models P_{ab} \text{ and } (w',\tau) \models P_{ab} \text{ and } V_b(w,\sigma) = V_b(w',\tau) \\
& \text{then } (w',\tau.ab) \models \varphi \\
\iff\ & \forall (w',\tau) \text{ s.t. } (w,\sigma) \sim_a^P (w',\tau) : && \text{(Semantics)} \\
& \text{If } (w,\sigma) \models P_{ab} \text{ then } (w',\tau) \models P_{ab} \to (O_b V_b(w,\sigma) \to [ab]\varphi) \\
\iff\ & (w,\sigma) \models P_{ab} \to K_a^P(P_{ab} \to (O_b V_b(w,\sigma) \to [ab]\varphi)) && \text{(Sem. } K_a^P) \\
\iff\ & (w,\sigma) \models P_{ab} \to \textstyle\bigvee_{R\subseteq\mathbb{S}}(O_b R \wedge K_a^P(P_{ab} \to (O_b R \to [ab]\varphi))) && (*)
\end{aligned}
$$

The proof for $\mathbf{Obs}_2$ is analogous to $\mathbf{Obs}_1$. For **Pri** we have the following chain of equivalences.

$$
\begin{aligned}
& (w,\sigma) \models [ab]K_c^P\varphi \\
\iff\ & (w,\sigma.ab) \models K_c^P\varphi && \text{(Sem. } [ab]) \\
\iff\ & \forall (w',\tau.de) \text{ s.t. } (w,\sigma.ab) \sim_c^P (w',\tau.de) \text{ and } c \neq d,e : && \text{(Sem. } K_c^P) \\
& (w',\tau.de) \models \varphi \\
\iff\ & \forall d,e \neq c : \forall (w',\tau) \text{ s.t. } (w,\sigma) \sim_c^P (w',\tau) : && \text{(Def. } \sim_c^P) \\
& \text{if } (w,\sigma) \models P_{ab} \text{ and } \tau \models P_{de} \text{ then } (w',\tau.de) \models \varphi \\
\iff\ & \forall d,e \neq c : \forall (w',\tau') \text{ s.t. } (w,\sigma) \sim_c^P (w',\tau) : && \text{(Semantics)} \\
& \text{if } (w,\sigma) \models P_{ab} \text{ then } (w',\tau) \models P_{de} \to [de]\varphi \\
\iff\ & \text{If } (w,\sigma) \models P_{ab} \text{ then } \forall d,e \neq c : (w,\sigma) \models K_c^P(P_{de} \to [de]\varphi) && \text{(Sem. } K_c^P) \\
\iff\ & (w,\sigma) \models P_{ab} \to \textstyle\bigwedge_{d,e\neq c} K_c^P(P_{de} \to [de]\varphi) && \square
\end{aligned}
$$

Given the validity of the call reductions, we now define and prove the following.

Definition 25 (Call Reduction). *For any formula $\varphi \in \mathcal{L}^{\mathbb{P}}$, let $\mathsf{cr}(\varphi) \in \mathcal{L}^{\mathbb{P}}_{-}$ denote the formula obtained by rewriting φ using the validities shown in Table 2 from left to right.*

Note that $\mathsf{cr}(\cdot)$ is well-defined: the rewriting terminates because after each step the subformula under the call is a strict subformula of the previous formula.

Lemma 26. *For every formula $\varphi \in \mathcal{L}^{\mathbb{P}}$ we have that $\models_{\mathcal{G}} \varphi \leftrightarrow \mathsf{cr}(\varphi)$.*

Proof. By induction on the structure of the formula, using Lemma 24. $\square$

Corollary 27. *All axioms in Table 2 are valid on the class of the tree model $\mathcal{T}$ and for every formula $\varphi \in \mathcal{L}^{\mathbb{P}}$ we have that $\models_{\mathcal{T}} \varphi \leftrightarrow \mathsf{cr}(\varphi)$.*

Proof. Immediate from $\mathcal{T} \subseteq \mathcal{G}$ and Lemmas 24 and 26. $\square$

Definition 25 also guarantees that rewriting never increases the modal degree.

Lemma 28. *Let $\varphi \in \mathcal{L}^{\mathbb{P}}$. Its call reduction $\psi \in \mathcal{L}^{\mathbb{P}}_{-}$ has degree $d(\psi) \leq d(\varphi)$.*

Proof. It suffices to show that $d(\mathsf{cr}([ab]\chi)) \leq d([ab]\chi)$. The proof is by induction on χ. We explicitly show the induction step for $\chi = K_c^P\chi'$ with $c = a$.
Induction Hypothesis. Let $\chi \in \mathcal{L}^{\mathbb{P}}$ be arbitrary. For each strict subformula χ' of χ we have $d(\mathsf{cr}(\chi')) \leq d(\chi')$.
Induction Step. Suppose $\chi = K_a^P\chi'$. Thus $\varphi = [ab]K_a^P\chi'$. Let $n = d(\varphi)$.

We have $\mathsf{cr}(\chi) = (\mathsf{cr}(P_{ab}) \to \bigvee_{R\subseteq S}(O_bR \wedge K_a^P(\mathsf{cr}(P_{ab}) \to (O_bR \to \mathsf{cr}(\chi')))))$ by Definition 25. By Definition 4, each of the strict subformulas of χ has a degree at most $n-1$. By the induction hypothesis on χ' we have thus $d(\mathsf{cr}(\chi')) \leq d(\chi') < d(\chi) = n$ and as P_{ab} is too a strict subformula of χ we have $d(\mathsf{cr}(P_{ab})) \leq d(P_{ab}) < d(\chi) = n$. Compute the degree of $\mathsf{cr}(\chi)$ to conclude that $d(\mathsf{cr}(\chi)) \leq n$. □

6 The Logic of Gossip Models

We cannot extend the axiomatisation $\vdash_{\mathcal{I}}$ and standard canonical model construction to $\mathcal{G}$ because **PI** no longer holds after calls have been made. However, we can still define provability in $\mathcal{G}$ in terms of $\mathcal{I}$, in a way that is similar to the *Tree Rule* in [14].

First, observe that it follows directly from the semantics that for any pointed gossip model and formula $\varphi \in \mathcal{L}^{\mathbb{P}}$ we have $M, (w, \sigma) \models \varphi$ iff $M, (w, \epsilon) \models [\sigma]\varphi$. We can therefore check the truth of some formula anywhere in the model already in the root state. Furthermore we can use Definition 25 to find the call reduction of $[\sigma]\varphi$. As the root states are semantically equivalent to their counterpart in the initial model, we can then use $\vdash_{\mathcal{I}}$ to determine validity of $[\sigma]\varphi$ in the roots of the gossip models, and subsequently conclude that φ always holds after sequence σ.

To find out whether φ holds after every call sequence and therefore is a validity, we only need to repeat this process for the other call sequences. However, there are infinitely many. Fortunately, we can bound the number of call sequences that we must check by showing that the number of n-bisimilarity classes is finite.

For this we now define n-bisimulation for protocol-dependent knowledge. Lemma 31 is a standard result about n-bisimulation and can be shown in the usual way, so we omit the proof [6, Proposition 2.31].

Definition 29 (n-bisimulation for Protocol-Dependent Knowledge). *Let $n \in \mathbb{N}$, and $M = \langle W, \sim, V\rangle$ and $M' = \langle W', \sim', V'\rangle$ be protocol-dependent gossip models. Two states $s \in W$ and $s' \in W'$ are* n-bisimilar*, denoted $M, s \underline{\leftrightarrow}_n M', s'$, if and only if the following conditions hold.*

1. ***(Atoms)** For every agent a we have $V_a(s) = V'_a(s')$.*

Additionally if $n > 0$ we have for each a, and for all $P \in \mathbb{P}$ with $d(P) < n$ an instance of the following two conditions.

2. ***(Forth)** For every $t \in W$ we have: if $s \sim_a^P t$ then there is a $t' \in W'$ such that $s' \sim_a'^P t'$ and $M, t \underline{\leftrightarrow}_{n-1} M', t'$.*
3. ***(Back)** For every $t' \in W$ we have: if $s' \sim_a^{P'} t'$ then there is a $t \in W$ such that $s \sim_a^P t$ and $M, t \underline{\leftrightarrow}_{n-1} M', t'$.*

Definition 30. *Two states are* n-equivalent*, written* $M,(w,\sigma) \equiv_n M',(w',\sigma')$*, iff for all* $\varphi \in \mathcal{L}^{\mathbb{P}}$ *with* $d(\varphi) \leq n$ *we have* $M,(w,\sigma) \models \varphi \iff M',(w',\sigma') \models \varphi$*.*

Lemma 31. *For any two states* $M,(w,\sigma)$ *and* $M',(w',\sigma')$ *we have*

$$M,(w,\sigma) \leftrightarrows_n M',(w',\sigma') \text{ if and only if } M,(w,\sigma) \equiv_n M',(w',\sigma').$$

Lemma 32. *The number of* $\leftrightarrows_n$*-equivalence classes* $f(n)$ *is finite.*

Proof. First we claim that there are only finitely many semantically different modal formulas of modal degree up to n. This can be shown by induction on n [6, Proposition 2.29]. For the induction step from n to $n+1$ with $K_a^P\varphi$, note that the protocol P must be defined by formulas of degree n or lower. By the induction hypothesis there are only finitely many semantically different formulas, and thus also only finitely many semantically different protocols. Hence for each degree we have only finitely many semantically different modalities. The rest is standard.

Now towards a contradiction, suppose $f(n)$ is infinite. Then there exist states $M_k,(w_k,\sigma_k)$ for $k \in \mathbb{N}$ that are all not n-bisimilar to each other. By Lemma 31 then there are formulas $\varphi_{i,j}$ that are each true at the state with index i and false at the state with index j. Moreover, infinitely many of these formulas must be semantically different, contradicting the claim. □

We now get the following result, which states that anything true after some sequence of calls is also true after some sequence of calls no longer than the number of n-bisimilarity classes.

Lemma 33. *Let* $\varphi \in \mathcal{L}^{\mathbb{P}}$ *and* $n = d(\varphi)$*. For any state* $(w,\sigma.\tau)$ *in any gossip model* M *that satisfies* φ *there is a sequence* τ' *such that* $|\tau'| \leq f(n)$ *and* $M,(w,\sigma.\tau') \models \varphi$*, where* $f(n)$ *is the number of* n*-bisimilarity classes.*

Proof. Let $M,(w,\sigma.\tau)$ and φ be arbitrary such that $M,(w,\sigma.\tau) \models \varphi$ and $|\tau| > f(n)$. Then τ must have two initial fragments $\tau_1 \neq \tau_2$ such that $(w,\tau_1) \leftrightarrows_n (w,\tau_2)$. W.l.o.g. let $|\tau_1| < |\tau_2|$ and $\tau = \tau_2.\tau_3$. Then $M,(w,\sigma.\tau_2) \models [\tau_3]\varphi$. By definition $d([\tau_3]\varphi) = n$, so by n-bisimilarity and Lemma 31 also $M,(w,\sigma.\tau_1) \models [\tau_3]\varphi$ and $M,(w,\sigma.\tau_1.\tau_3) \models \varphi$. Note that $|\tau_1.\tau_3| < |\tau|$ and we can repeat until done. □

Definition 34 (Proof System $\vdash_{\mathcal{G}}$). *Let* $\varphi \in \mathcal{L}^{\mathbb{P}}$ *and* $n = d(\varphi)$*. We define* $\vdash_{\mathcal{G}}$ *as follows, where* $f(n)$ *is the number of* n*-bisimilarity classes.*

$$\vdash_{\mathcal{G}} \varphi \iff \forall \sigma : |\sigma| \leq f(n) \text{ we have } \vdash_{\mathcal{I}} \mathsf{cr}([\sigma]\varphi)$$

Theorem 35 (Soundness and Completeness for $\mathcal{G}$). *For any* $\varphi \in \mathcal{L}^{\mathbb{P}}$*, we have* $\models_{\mathcal{G}} \varphi$ *iff* $\vdash_{\mathcal{G}} \varphi$*.*

Proof. (**Soundness**) By contraposition, suppose $\not\models_{\mathcal{G}} \varphi$. We show that $\not\vdash_{\mathcal{G}} \varphi$.

There is some model $M \in \mathcal{G}$ and some state (w,σ) such that $M,(w,\sigma) \models \neg\varphi$, i.e. we have $M,(w,\epsilon) \models \neg[\sigma]\varphi$. By Lemma 33 we can assume that $|\sigma| \leq f(d(\varphi))$ and by Lemma 26 we have the call reduction $\psi_\sigma := \mathsf{cr}([\sigma]\varphi)$.

Now suppose $\vdash_{\mathcal{G}} \varphi$, in order to reach a contradiction. Then by Definition 34 and from $|\sigma| \leq f(d(\varphi))$ we get $\vdash_{\mathcal{I}} \psi_\sigma$. By soundness of $\vdash_{\mathcal{I}}$ we then have $\models_{\mathcal{I}} \psi_\sigma$ and in particular $M, (w, \epsilon) \models \psi_\sigma$. By validity of the call reductions we get $M, (w, \epsilon) \models [\sigma]\varphi$. This contradicts $M, (w, \epsilon) \models \neg[\sigma]\varphi$, so we must have $\nvdash_{\mathcal{G}} \varphi$.

(**Completeness**) By contraposition, suppose $\nvdash_{\mathcal{G}} \varphi$. We show that $\not\models_{\mathcal{G}} \varphi$.

By Definition 34 there must be a call sequence σ such that $\nvdash_{\mathcal{I}} \mathsf{cr}([\sigma]\varphi)$. By completeness of $\vdash_{\mathcal{I}}$ we have some model and world such that $M, (w, \epsilon) \not\models \mathsf{cr}([\sigma]\varphi)$. By validity of the call reductions we get $M, (w, \epsilon) \not\models [\sigma]\varphi$. This implies $M, (w, \sigma) \not\models \varphi$ and thus $\not\models_{\mathcal{G}} \varphi$. □

We could use $\vdash_{\mathcal{I}}$ to define $\vdash_{\mathcal{G}}$, because there is a bijective relation between the two classes: each model in $\mathcal{G}$ is induced by an initial model in $\mathcal{I}$ and vice versa. Now recall that $\mathcal{T}$ is the singleton class of the tree model induced by the root model in $\mathcal{R}$. We therefore define $\vdash_{\mathcal{T}}$ analogously to $\vdash_{\mathcal{G}}$, but using $\vdash_{\mathcal{R}}$ instead of $\vdash_{\mathcal{I}}$. The proof of Theorem 37 is analogous to Theorem 35.

Definition 36 (Proof System $\vdash_{\mathcal{T}}$). *Let $\varphi \in \mathcal{L}^{\mathbb{P}}$ and $n = d(\varphi)$. We define $\vdash_{\mathcal{T}}$ as follows, where $f(n)$ is the number of n-bisimilarity classes.*

$$\vdash_{\mathcal{T}} \varphi \iff \forall \sigma : |\sigma| \leq f(n) \text{ we have } \vdash_{\mathcal{R}} \mathsf{cr}([\sigma]\varphi)$$

Theorem 37 (Soundness and Completeness for $\mathcal{T}$). *For any $\varphi \in \mathcal{L}^{\mathbb{P}}$, we have $\models_{\mathcal{T}} \varphi$ iff $\vdash_{\mathcal{T}} \varphi$.*

The definitions of both $\vdash_{\mathcal{G}}$ and $\vdash_{\mathcal{T}}$ bound the number of times they invoke $\vdash_{\mathcal{I}}$ and $\vdash_{\mathcal{R}}$ by the number of n-bisimilarity classes. As this is finite, these systems too are decidable.

Theorem 38. *Proof systems $\vdash_{\mathcal{G}}$ and $\vdash_{\mathcal{T}}$ are decidable.*

Proof. Immediate by Lemma 32. □

7 Conclusion

We have shown that the protocol-dependent knowledge modality is more expressive than the standard epistemic knowledge modality. We have defined four logics for gossip using protocol-dependent knowledge modalities and provided sound and complete proof systems that are decidable. In particular we have provided a proof system for $\mathcal{T}$, the model defined in [12]. These contributions give insight into the use of protocol-dependent knowledge modalities in epistemic logic.

An interesting avenue for future work is to generalise the use and axiomatisation of protocol-dependent knowledge modalities to domains outside of gossip.

References

1. Apt, K.R., Grossi, D., van der Hoek, W.: Epistemic protocols for distributed gossiping. In: Proceedings Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2015. EPTCS, vol. 215, pp. 51–66 (2015). https://doi.org/10.4204/EPTCS.215.5
2. Apt, K.R., Wojtczak, D.: Common knowledge in a logic of gossips. In: Proceedings Sixteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2017. EPTCS, vol. 251, pp. 10–27 (2017). https://doi.org/10.4204/EPTCS.251.2
3. Attamah, M., van Ditmarsch, H., Grossi, D., Hoek, W.: Knowledge and gossip. Frontiers in Artificial Intelligence and Applications **263**, 21–26 (2014-01). https://doi.org/10.3233/978-1-61499-419-0-21
4. Baker, B., Shostak, R.: Gossips and telephones. Discrete Mathematics **2**(3), 191–193 (1972). https://doi.org/10.1016/0012-365X(72)90001-5
5. van den Berg, L., Gattinger, M.: Dealing with unreliable agents in dynamic gossip. In: Martins, M., Sedlár, I. (eds.) Proc. of 3rd DaLí. pp. 51–67 (2020). https://doi.org/10.1007/978-3-030-65840-3_4, LNCS 12569
6. Blackburn, P., de Rijke, M., Venema, Y.: Modal Logic. No. 53 in Cambridge Tracts in Theoretical Computer Science, Cambridge University Press (2001)
7. Cooper, M., Herzig, A., Maffre, F., Maris, F., Régnier, P.: The epistemic gossip problem. Discret. Math. **342**(3), 654–663 (2019). https://doi.org/10.1016/j.disc.2018.10.041
8. van Ditmarsch, H., van Eijck, J., Pardo, P., Ramezanian, R., Schwarzentruber, F.: Epistemic protocols for dynamic gossip. Journal of Applied Logic **20**, 1–31 (2017). https://doi.org/10.1016/j.jal.2016.12.001
9. van Ditmarsch, H., van Eijck, J., Pardo, P., Ramezanian, R., Schwarzentruber, F.: Dynamic gossip. Bulletin of the Iranian Mathematical Society **45(3)**, 701–728 (2019). https://doi.org/10.1007/s41980-018-0160-4
10. van Ditmarsch, H., Gattinger, M., Ramezanian, R.: Everyone knows that everyone knows: Gossip protocols for super experts. Stud Logica **111**(3), 453–499 (2023). https://doi.org/10.1007/S11225-022-10032-3
11. van Ditmarsch, H., Gattinger, M.: You can only be lucky once: Optimal gossip for epistemic goals. Mathematical Structures in Computer Science (2024). https://doi.org/10.1017/S0960129524000082
12. van Ditmarsch, H., Gattinger, M., Kuijer, L.B., 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), http://arxiv.org/abs/1907.12321
13. van Ditmarsch, H., Gattinger, M., Ramezanian, R.: Everyone Knows that Everyone Knows: Gossip Protocols for Super Experts. Studia Logica **111**(3), 453–499 (2023). https://doi.org/10.1007/s11225-022-10032-3
14. van Ditmarsch, H., van der Hoek, W., Kuijer, L.B.: The logic of gossiping. Artificial Intelligence **286**, 103306 (2020). https://doi.org/10.1016/j.artint.2020.103306
15. Gattinger, M.: New Directions in Model Checking Dynamic Epistemic Logic. Ph.D. thesis, University of Amsterdam (2018)
16. 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

17. Herzig, A., Maffre, F.: How to share knowledge by gossiping. AI Communications **30**(1), 1–17 (2017). https://doi.org/10.3233/AIC-170723
18. Kermarrec, A.M., van Steen, M.: Gossiping in distributed systems. SIGOPS Oper. Syst. Rev. **41**(5), 2–7 (2007). https://doi.org/10.1145/1317379.1317381
19. Smit, W.J.: Axiomatising Protocol-Dependent Knowledge in Gossip. Master's thesis, University of Amsterdam (2024), https://eprints.illc.uva.nl/id/eprint/2330/
20. Tijdeman, R.: On a telephone problem. Nieuw Archief voor Wiskunde **3**(19), 188–192 (1971)