Axiomatising Protocol-Dependent Knowledge in Gossip

MSc Thesis *(Afstudeerscriptie)*

written by

Wouter Smit
(born in Amsterdam)

under the supervision of **dr. Malvin Gattinger** and **Prof. dr. Hans van Ditmarsch**, and submitted to the Examinations Board in partial fulfilment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defence:	**Members of the Thesis Committee:**
23 August 2024	dr. Benno van den Berg (chair)
	dr. Malvin Gattinger (supervisor)
	Prof. dr. Hans van Ditmarsch (co-supervisor)
	dr. Balder ten Cate
	dr. Ronald de Haan

Institute for Logic, Language and Computation

Abstract

The gossip problem is an information distribution problem where agents try to share information between each other in the most efficient way. While an optimal order of communication is easily computed, agents often find themselves in a decentralised system and therefore rely on gossip protocols for efficiently coordinating their communication. Protocol-dependent knowledge in epistemic logic expresses knowledge that an agent can deduce by assuming that a certain protocol is common knowledge. Protocol-dependent epistemic logic for the gossip problem has been studied semantically, but a sound and complete axiomatisation is yet to be found. While axiomatisations exist for various versions of the gossip problem, none of these include protocol-dependent knowledge.

In this thesis we provide a sound and complete axiomatisation for the logic of gossip with a single protocol-dependent knowledge modality and we further show that these axioms remain sound for the logic including multiple such modalities. In order to do so we analyse existing axiomatisations for non-protocol-dependent gossip and link them to existing semantic results, and we show that the language of gossip with multiple protocol modalities is considerably more expressive than the language without any such modalities.

These results aid in understanding the effects of protocols on agent knowledge in the gossip problem and provide insight in the effects of protocol-dependent knowledge modalities on expressivity of epistemic logic.

Acknowledgements

I would like to thank my supervisors Malvin Gattinger and Hans van Ditmarsch. Your enthusiasm for gossip and this project has motivated me from start to finish, and keeps me inspired even afterwards. Malvin, your unlimited positivity, energy, and patience have shown me how fulfilling research can be. Your attention to detail and style helped me to stay focused and to improve my work. You taught me many valuable lessons, both professionally and academically. Without you, neither the project nor I would have been the same. Hans, thank you for being so generous in helping out at any time throughout the project. Specifically, for dedicating so much of your scarce time in Amsterdam to working together. Those days have been pivotal to the project. More importantly, they made me feel like an equal.

I have been lucky to be surrounded by family and friends, many of whom I have met during my studies. In particular Djanira, who shared the final parts of the programme together with me and joined me in working on gossip for our side project.

Finally, I am grateful to my partner Doreen for her unconditional and unwavering support throughout. It was your care alone that helped me through the final weeks of the project. Having you by my side improved my thesis with each day, and that applies to me as a person too.

Contents

Chapter 1

Introduction

The *gossip problem*, first introduced as the *telephone problem* [Tij71], describes the distribution of information (*secrets*) between agents using peer-to-peer pairwise communication (*calls*). The goal of the problem is for all agents to learn all secrets (become *experts*), and to do so as efficiently as possible.

It is assumed that each of the finitely many agents initially holds a single secret. In each call, agents share the secrets that they know, including other secrets that they have learnt, hence the name for the problem. Agents do not share any other knowledge. In particular, they do not share higher-order epistemic information[1], such as information about who knows what.

The problem is applicable to many environments, most notably in distributed systems. The combination of information over multiple sources using limited communication resources is a common problem in this field. Examples are updating distributed databases [Kar+00] and network discovery [HLL99].

Albeit its name suggests otherwise, the gossip problem does not model socially-driven behaviour of agents, and instead only focusses on the collaborative effort of sharing information openly. However, epistemic logic can also be used to model social dynamics of gossip [Kle17].

Epistemic logic studies the knowledge of agents. The modality K encodes "knowing", such that $K_a\varphi$ means *"agent a knows that φ is true"*. This knowledge is usually based on which worlds in a Kripke frame the agent considers possible: an agent knows something if it is true in all worlds they consider possible. This relation between worlds is called the *epistemic relation* or *indistinguishability relation* of an agent.

This definition of knowledge usually makes no assumptions, but this changes with protocol-dependent knowledge [Dit+19]. The formula $K_a^P\varphi$ comes to mean that *"agent a knows that φ is true, assuming that protocol P was common knowledge"*. With the assumption that all agents behave according to a specific set of rules, the agent can often derive more knowledge and thus achieve a larger knowledge base.

While assumptions can always be encoded in the epistemic relation for a knowledge modality K_a, the novelty lies in the arbitrary combination of protocols. In this way the language can express knowledge relative to different protocols: the formula $K_a^P\varphi \wedge \neg K_a^Q\varphi$ expresses *"agent a knows that φ is true if protocol P was followed, but does not know that φ is true if protocol Q was followed"*.

Related Work

The gossip problem goes back to 1971 and research was first focused on finding the minimum number of messages required in various settings. The most well-known result is a minimum of $2n-4$ calls for $n \geq 4$ agents [Tij71; BS72].

However, this result is in many ways overly optimistic. In practice, agents are often in a distributed setting, lacking the authority of a centralised scheduler. These settings require agents to coordinate with limited information and means of communication. This has led to research into distributed algorithms, called *gossip protocols* or *epidemic*

[1] Other settings exist where agents can share higher-order knowledge [HM17].

protocols, which often employ randomized calling [Kar+00]. More recently, distributed gossip has been studied from the perspective of epistemic logic, making use of *epistemic protocols* which leverage agent knowledge to make informed decisions on what call to execute [Att+14; AGH15]. Additionally, the formalisms from epistemic logics have led to bounds on the amount of agent knowledge that can be achieved [HM17; DG22; DG24].

Much of the epistemic logic research into gossip uses a semantic approach, as research into proof systems for the gossip problem is still limited. Even within these semantics, many studies use different languages and models, often catering to specific versions of gossip or the topic of study. This has also led to different formal languages for the gossip problem that are either limited [DG24] or extended to different levels of expressivity [DGR23; Dit+19; AW17].

There have also been different approaches modelling the gossip problem in epistemic logic. Some attempts include using action models [Att+14] or using a "knowing-whether" modality [HMP21]. Recent work has provided a more efficient method using *atomic knowing* for secret representation and knowledge transformers for modelling dynamic updates [Gat18].

So far only one study has provided axiomatisations for logics describing the gossip problem [DHK20]. In doing so, the authors attempt to unify the different semantics for gossip based on three parameters. Unfortunately, these axiomatisations have not seen practical application: more recent research has not been able to apply these axiomatisations yet, instead still resorting to semantic proofs of validities [DG24]. The vastly more modular approach to defining a gossip model that the authors of [DHK20] take, can be identified as a difficulty in applying their work effectively.

Meanwhile, the notion of *protocol-dependent knowledge* has been studied in epistemic logic outside of gossip [Dit+14]. Within gossip, this notion has been extended to a protocol-specific knowledge modality to express the knowledge of agents under the assumption of a protocol [Dit+19]. This work does not provide an axiomatisation yet, leaving it an open question what the logic of protocol-dependent gossip is.

There are three central sources that we use in this thesis, denoted in Table 1.1. While [DHK20] already provides an axiomatisation for many versions of gossip, the semantics and language do not align with [Dit+19] and [DG24]. In this thesis we attempt to bridge the gap between these three works.

Table 1.1. Three recent studies using different languages and semantics.

Property	Strengthening [Dit+19]	Logic of Gossiping [DHK20]	Limits to Gossip [DG24]
Privacy	Synchronous (◓)	All (●, ◓, ○)	Asynchronous (●)
Directionality	bidirectional ($\diamond$)	All ($\triangleright, \triangleleft, \diamond$)	bidirectional ($\diamond$)
Observance	Inspect-then-merge (β)	All (α, β)	Inspect-then-merge (β)
Graph Topology	Dynamic	Total	Total
Dynamic Modality	$[\pi]$	$[ab]$	–
Knowledge Modality	K^P	K	K
Completeness	No	Yes	No

Contribution

We propose axioms for synchronous protocol-dependent gossip as defined in [Dit+19]. We show completeness for the language and models with a single protocol, and provide axioms for extending this result to languages containing multiple protocols. We do so by employing a method similar to [DHK20]. While doing so, we make minor changes to the original axiomatisation by [DHK20] and we introduce a different model definition for protocol-dependent gossip using arbitrary initial models.

We furthermore show that the language with protocol-dependent knowledge operators has higher expressivity than the language with only a general epistemic knowledge operator.

Lastly, we identify the logic used in [DG22] by relating the semantics to those of [DHK20]. We thereby show that the results in [DG22] can be derived syntactically too.

These contributions help better understand the effects of protocols on agent knowledge in the gossip problem. More generally, they provide insight in the definition of protocol-dependent knowledge by [Dit+19] and how it might be used in other epistemic settings than the gossip problem.

Outline

In chapter 2 we give the necessary background and definitions, including the main syntax and semantics of the original gossip problem without protocol-dependency. Chapter 3 continues by explaining two existing axiomatisations due to [DHK20], to which we make minor contributions before relating one of them to [DG22].

Chapter 4 introduces protocol-dependent gossip. We provide an alternative definition of protocol-dependent gossip models, which we use to show completeness for the axiomatisation for protocol-dependent gossip with a single protocol, that we provide in chapter 5. In chapter 6 we build on this axiomatisation and suggest ways to translate this to the multi-protocol setting.

The last contribution can be found in chapter 7, which provides an analysis of the difference in expressivity between the languages for basic gossip and protocol-dependent gossip. We discuss the results in chapter 8 and conclude this thesis in chapter 9.

Chapter 2

Preliminaries

We first give general definitions of the gossip problem. We then introduce the language and models for gossip based on [DHK20]. These do not include the notion of protocols and we will therefore call them *basic gossip*.

2.1 General Definitions for Gossip

Throughout this thesis we will use the following definitions and notation.

Definition 2.1 (Agents)**.** *Let Ag be the finite set of agents. We denote agents by a Roman lowercase letter.*

We assume that each agent possesses a personal secret. The set of agents and set of secrets is therefore a bijection. Moreover, it is always assumed that the agent knows their own secret. Because we view the secrets as so intrinsically connected to the agent, we re-use the agent names to denote their secrets, rather than defining a separate set of secrets.

Definition 2.2 (Secrets)**.** *Let $\mathbb{S} = Ag$ be the set of secrets, so that each secret is defined by its agent.*

It might seem that this causes confusion, but in the gossip problem we are not interested in what the secret of some agent is, or whether this secret is true. We only care about the secret distribution: who knows whose secret. We therefore never need to talk about the secret specifically: referring to "the secret a" may as well be "the secret of agent a".

Definition 2.3 (Calls)**.** *A call (a, b) is an ordered pair of agents $a \neq b \in Ag$. We call the first agent the* caller *and the second agent the* callee*. We usually omit the brackets and simply write ab.*

There are $n \cdot (n-1)$ possible calls for n agents and each call can always be executed. In some settings however, distinction between the role of caller and callee is not important. In such settings, symmetric calls ab and ba are sometimes viewed as a single call [Gat18]. The direction of calls is important for instance when the problem includes a non-total graph [Gat18], when secrets are not shared symmetrically [DHK20], or when protocols are considered [Dit+19], like in this thesis.

Definition 2.4 (Call Sequence)**.** *A* call sequence *σ is a sequential series of calls and is denoted by a Greek lowercase letter. The empty sequence is denoted by ϵ. We write concatenation of calls with a period, such that $ab.cd$ is the call sequence consisting of ab followed by cd and $\sigma.ab.\tau$ is the sequence σ followed by call ab and followed by sequence τ. We write $|\sigma|$ for the length of a call sequence σ.*

Definition 2.5 (Involved & External Agents)**.** *We call an agent c* involved *in a call ab if they are the caller or callee, that is $c \in \{a, b\}$. An agent is involved in call sequence σ if they are involved in at least one call in σ. We denote the set of involved agents in a call sequence σ by $Ag(\sigma)$. We call an agent who is not involved* external.

Agents attempt to learn all secrets. We call them *experts* when they have achieved this goal. The goal for the gossip problem is for all agents to become experts. Other goals have been studied too, including higher-order goals [DGR23].

Definition 2.6 (Expert)**.** *An agent is an* expert *if they know the secrets of all agents.*

2.2 Call Types

Call types have been introduced to better describe the various semantic properties (of calls) in different gossip settings [DHK20]. We will limit ourselves to only the most common call types, but introduce the notion fully here in order to relate to the axiomatisation of [DHK20].

The following three parameters form the basis of the call type. We only introduce them conceptually. We formalise their meaning when defining the semantics in section 2.4.

Definition 2.7 (Privacy)**.** *The level of privacy determines how much agents can observe of the calls happening around them. There are three privacy types* p*.*

$\circ =$ *transparent / observable: agents see exactly each call that takes place.*

$\mathbin{\text{◓}} =$ *synchronous: agents notice that a call takes place, but not which.*

$\bullet =$ *asynchronous: agents notice only calls they participate in.*

Privacy type ($\circ$) is also a synchronous setting in the sense that there is a global awareness of some call taking place and many results that rely on this property hold for both privacy levels. Such a notion is also called a *global clock*. However, in all relevant literature, synchronicity is used exclusively to the weaker form ($\mathbin{\text{◓}}$), because a scenario where the specifics of all calls are fully known is in many cases unrealistic. We therefore refer to ($\circ$) as *transparent*, as the actions of any agent are visible to all other agents. The term *observable* is also used in literature [DHK20].

Definition 2.8 (Directionality)**.** *A gossip call can exchange information in three ways, which is the directionality type* d*. Throughout this thesis, we limit ourselves to directionality type* $\mathsf{d} = \diamond$*.*

$\triangleright =$ *push: only the caller shares secrets*

$\triangleleft =$ *pull: only the callee shares secrets*

$\diamond =$ *bidirectional / push-pull: both share secrets*

Definition 2.9 (Observance)**.** *Agents may observe the secrets shared with them in the call in two observation types* o*. Throughout this thesis, we limit ourselves to observance type* $\mathsf{o} = \beta$*.*

$\alpha =$ *after: merge-then-inspect. Agents observe the union of their sets of secrets.*

$\beta =$ *before: inspect-then-merge. Agents observe each other's set of secrets separately.*

Together the three parameters privacy, directionality, and observance combine into 18 different call types $\mathsf{t} = (\mathsf{p}, \mathsf{d}, \mathsf{o})$ with different semantics. The logics corresponding to each of these semantics have been each axiomatised by [DHK20].

While these call parameters are useful in explicitly defining the various design options, they are not all equally commonly used. In fact, while sources exist that use each option for each parameter, the directionality type is usually $\diamond$ and the observance type is usually β. It is only the privacy type that can be viewed as an equally studied property, and even then only the $\mathbin{\text{◓}}$ and $\bullet$ cases are relatively common.

We will therefore often omit the directionality type and observance type and implicitly assume them to be $\diamond$ and β. We write $\mathsf{t} = \mathbin{\text{◓}}$ or $\mathsf{t} = \bullet$ instead of $\mathsf{t} = (\mathbin{\text{◓}}, \diamond, \beta)$ or $\mathsf{t} = (\bullet, \diamond, \beta)$.

2.3 Language and Syntax for Basic Gossip

The first notable choice in the language is how we describe secrets. We are not interested in the values of the secrets, so we do not necessarily want to encode them in propositions. Instead we are interested in the distribution of secrets. Various attempts have been made in choosing an effective strategy. An overview of the various options and their considerations can be found in [Gat18].

We call the notion that we will use to describe secrets *atomic knowing*, where atoms in the language describe what an agent knows [Gat18]. The atomic propositions in our language are defined S_ab for agents $a, b \in Ag$, to mean *"Agent a knows the secret of agent b"*. Some authors use either b_a or F_aB instead, the latter using F for *familiar with* and capitals to differentiate the secrets (upper case) from their agents (lower case) [DHK20]. We choose to use S_ab as we wish to differentiate language constructs (upper case for the S proposition) from variables (lower case for agents and their secrets a, b, c) and *secret* is the more intuitive natural language equivalent. It is furthermore useful to distinguish atoms by their type when introducing extra atoms, such as representing phone numbers, denoted by N_ab [DG22].

The basic language contains atoms for each agent pair, the usual boolean connectives for negation and conjunction, the agents' knowledge modality and a dynamic call modality. The atoms S_ab describe the distribution of secrets, and K_a is the usual epistemic knowledge operator for each agent a. The dynamic modality $[ab]$ describes a call between two agents a and b. As agents exchange secrets, a call usually incurs factual change: modification of the truth of atomic propositions.

Definition 2.10 (Basic Language of Gossip)**.** *Let $a, b \in Ag$ be agents. We define the language $\mathcal{L}^G$ as*

$$\varphi ::= S_ab \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid [ab]\varphi.$$

The operators $\vee$, $\rightarrow$, $\top$, and $\bot$ are defined as abbreviations in the usual way. We also define the epistemic dual of K by $\widehat{K}_a := \neg K_a \neg$.

For some 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.

$$O_aR := \bigwedge_{b \in R} S_ab \wedge \bigwedge_{b \notin R} \neg S_ab$$

Throughout this thesis we will often use the static fragment of $\mathcal{L}^G$, which omits the dynamic modality $[ab]$. As this modality represents a call in the setting of gossip, we also use the term *call-free*.

Definition 2.11 (Call-Free Languages)**.** *For any language of gossip $\mathcal{L}$ we define a call-free fragment $\mathcal{L}_{[-]}$ which omits the $[ab]$ modality. For the basic language of gossip we get the following definition.*

Let $a, b \in Ag$ be agents. We define the language $\mathcal{L}^G_{[-]}$ as

$$\varphi ::= S_ab \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi.$$

2.4 Models and Semantics for Basic Gossip

We now introduce gossip models as defined by [DHK20], but we restrict ourselves to only the definitions relevant for $\mathtt{t} = (\textcircled{\bullet}, \diamond, \beta)$ and $\mathtt{t} = (\bullet, \diamond, \beta)$. Contrary to other literature, the models defined by these authors are more general and are based on arbitrary *initial models* that describe an initial setting of the gossip problem. A *gossip model* then follows from lifting this definition to include calls.

This class of such gossip models is not usually of interest. Classically, the gossip problem starts in a specific state: each agent knows only their own secret and this is common knowledge. To retrieve the desired model, we start from a specific initial model that coincides with this state. We call this initial model the *initial root model* and the gossip model that follows from it the *tree model*.

Initial Models for Basic Gossip

We only require that two basic assumptions in the gossip problem are fulfilled: agents know their own secret, and they are aware of what secrets they know.

Definition 2.12 (Initial Model for Basic Gossip)**.** *An initial model is a triple $I = \langle W_0, R_0, V_0 \rangle$ where W_0 is a set of initial worlds, $R_0 : Ag \to 2^{W_0 \times W_0}$ is an equivalence relation for each agent, and $V_0 : Ag \times W_0 \to 2^{\mathbb{S}}$ is a function mapping each agent to the initial set of secrets known in some world, which satisfies (i) $a \in V_0(a, w)$ for all $w \in W$ and (ii) if $(w_1, w_2) \in R_0(a)$ then $V_0(a, w_1) = V_0(a, w_2)$.*

An initial model may therefore contain possible worlds that (some) agents cannot distinguish, as the following example shows.

Example 2.13. *See Figure 2.1. There are three agents. Agent a knows her own secret as well as those of agent b and c. Agent c knows which secrets a knows, but only knows their own secret themselves. Agent b knows that c only knows their own secret, but does not know what secrets a knows and only knows his own secret himself. In this setting, some secrets are shared and some are not. Moreover, the distribution of secrets is not common knowledge: b is not be aware what the secret distribution is, but agent c and a are.*

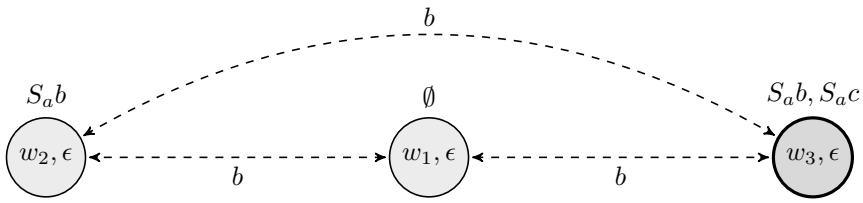

Figure 2.1. The initial model for Example 2.13. We omit $S_a a$ for all agents a which hold in all worlds, as well as reflexive relations. The actual world is w_3. Agent b does not know this, while a and c do.

Initial models are straightforward Kripke frames. We could define semantics on them in the standard way. We do not need these semantics for basic gossip, but we will define semantics on initial models for protocol-dependent gossip in chapter 4.

Lifting an Initial Model

An initial model describes the situation before any calls take place and has no notion of calls. In order to get a gossip model, we lift the definition of the initial model. This effectively induces calls into the model. Lifting the worlds and valuation is relatively straightforward, while lifting the relation is most consequential for the semantics of the resulting gossip model.

Definition 2.14 (Gossip State)**.** *Given an initial model $I = \langle W_0, R_0, V_0 \rangle$, a gossip state is a pair (w, σ) where $w \in W_0$ is an initial world and σ is a call sequence. We let W^I be the set of all gossip states relative to I and we write W when the initial model is clear from context.*

A gossip states describes exactly any situation in a gossip model: given an initial setting, the effects of calls are fully deterministic.

The secrets known by an agent, and thus the valuation of atoms for that agent, is defined as follows. We implicitly assume *bidirectional* calls ($\diamond$): both the caller and callee share all secrets that they know.

Definition 2.15. *Given an initial model $I = \langle W_0, R_0, V_0 \rangle$, we denote the set of secrets known by a, given some gossip state $(w, \sigma) \in W$, by $V(a, w, \sigma)$ or $V_a(w, \sigma)$ for short. We define V recursively as follows.*

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

We finally need to lift the relation. It is only at this point that the privacy type p affects the definition. Here we define the synchronous (◓) and asynchronous (●) version. For a complete overview, we refer to [DHK20].

Synchronicity or asynchronicity is solely determined by whether or not agents can necessarily distinguish call sequences of different length. In order to make the semantics synchronous, we will require that call sequences of related states are of the same length. We do not make this requirement in the asynchronous version. The last case of the definition reflects this difference.

As a result, agents in the asynchronous case are oblivious to any calls taking place that they are not involved in. In particular this means they consider it possible from the very start that any arbitrary number of calls has taken place without their involvement.

Definition 2.16 (Basic Epistemic Relation)**.** *Given an initial model $I = \langle W_0, R_0, V_0 \rangle$, we lift the initial relation $R_0(a) \subseteq W_0 \times W_0$ for some agent a to $\sim_a \subseteq W \times W$ such that:*

$$\begin{aligned} (w_i, \epsilon) \sim_a (w_j, \epsilon) &\textit{ iff } (w_i, w_j) \in R_0 \\ (w_i, \sigma.ab) \sim_a (w_j, \tau.ab) &\textit{ iff } (w_i, \sigma) \sim_a (w_j, \tau); \\ &\qquad \textit{and } V_b(w_i, \sigma) = V_b(w_j, \tau); \\ (w_i, \sigma.ba) \sim_a (w_j, \tau.ba) &\textit{ iff } (w_i, \sigma) \sim_a (w_j, \tau) \\ &\qquad \textit{and } V_b(w_i, \sigma) = V_b(w_j, \tau); \end{aligned}$$

and for the synchronous privacy type ◓ we let

$$\begin{aligned} (w_i, \sigma.bc) \sim_a (w_j, \tau.de) &\textit{ iff } (w_i, \sigma) \sim_a (w_j, \tau) \\ &\qquad \textit{and } a \notin \{b, c, d, e\}; \end{aligned}$$

while for the asynchronous privacy type ● we let

$$\begin{aligned} (w_i, \sigma.bc) \sim_a (w_j, \tau) &\textit{ iff } (w_i, \sigma) \sim_a (w_j, \tau) \\ &\qquad \textit{and } a \notin \{b, c\}. \end{aligned}$$

Remark 2.17. *In other literature there are no (arbitrary) initial models and therefore no definition of R_0 [DG22; DGR23; Dit+19]. In these cases, the relations are defined in the same way but the initial worlds w_i, w_j are omitted and the base case is simply $\epsilon \sim \epsilon$.*

Gossip Models

With all elements of the induced model defined, all that remains is to tie them together. A gossip model is an initial model that is lifted to include specific call semantics. Given a certain call type, the gossip model is fully determined by its initial model.

Definition 2.18 (Basic Gossip Model)**.** *Let $I := \langle W_0, R_0, V_0 \rangle$ be an initial model. Given a call type $\mathsf{t} = (\mathsf{p}, \mathsf{d}, \mathsf{o})$, we define the (induced) gossip model $M^{\mathsf{t}}(I) := \langle W, \sim, V \rangle$ with W and V as defined in definitions 2.14 and 2.15*

respectively, and the epistemic relation $\sim_a$ for each agent $a \in Ag$ as defined in definition 2.16. When the initial model or call type are clear from context, we omit them and write M or M^{t}.

Remark 2.19. *As we use a fixed $\mathsf{d} = \diamond$ and $\mathsf{o} = \beta$, we will only specify the privacy type p and write $M^{◓}$ or $M^{●}$.*

After lifting an initial model to an induced model $M^{\mathsf{t}}(I)$ to include all call sequences, we find that $M^{\mathsf{t}}(I)$ forms a forest, with each initial world as the root of a call tree, depicted in Figure 2.2.

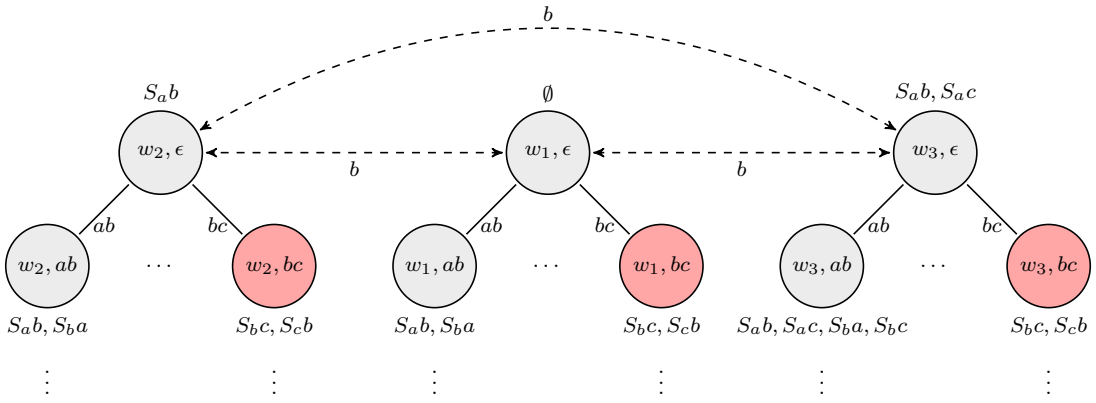

Figure 2.2. A partial gossip model for $\mathsf{p} = ◓$ induced from the initial model I from example 2.13. The three red states are b-indistinguishable from each other. Reflexive arrows are omitted.

Definition 2.20 (Semantics on Basic Gossip Models)**.** *We define $\models$ on pointed gossip models as follows. Let M be a gossip model induced from $I = \langle W_0, R_0, V_0\rangle$ and $w \in W_0$ some initial world. Finally, let σ be a call sequence.*

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

For a class of models $\mathcal{M}$, we call a formula φ valid on $\mathcal{M}$ *if for all models $M \in \mathcal{M}$ and for all states s of M we have $M, s \models \varphi$. For a single model M, we call φ* valid on M *if for all states s of M we have $M, s \models \varphi$.*

Tree Models

When the initial world is common knowledge among the agents – hence the initial distribution of secrets is common knowledge – I consists of only one world. Such an initial distribution of secrets could be arbitrary, but we often talk about the most common setting for gossip: each agent knows only their own secret and this is common knowledge [DHK20].

Definition 2.21 (Initial Root Model for Basic Gossip)**.** *Let $I_{\mathsf{root}} := \langle W_{\mathsf{root}}, R_{\mathsf{root}}, V_{\mathsf{root}}\rangle$ be an initial model, where*

- $W_{\mathsf{root}} := \{w_{\mathsf{root}}\}$*;*
- $R_{\mathsf{root}}(a) := \{(w_{\mathsf{root}}, w_{\mathsf{root}})\}$ *for all agents $a \in Ag$;*
- $V_{\mathsf{root}}(a, w_{\mathsf{root}}) := \{S_a a\}$ *for all agents $a \in Ag$.*

We call I_{root} the initial root model*. See Figure 2.3.*

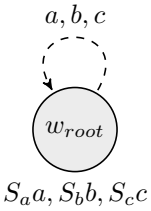

Figure 2.3. The Initial Root Model for three agents a, b, c.

We can then lift a *tree model* from I_{root}. The name refers to the structure of the model, which is now a single tree of calls. There is a tree model for each call type $\mathtt{t}$, but we only mean models lifted from I_{root} by this name.

Definition 2.22 (Tree Models for Basic Gossip)**.** *A* tree model $M^{\mathtt{t}}_{\mathsf{tree}} := M^{\mathtt{t}}(I_{\mathsf{root}})$ *is a gossip model lifted from the initial root model* I_{root} *given a call type* $\mathtt{t}$*. When the tree model that we consider is clear from context, we often omit it fully and simply refer to the call sequences, such that* $\sigma \models \varphi$ *as a shorthand for* $M^{\mathtt{t}}_{\mathsf{tree}}, (w_{\mathsf{root}}, \sigma) \models \varphi$*.*

Remark 2.23. *In literature, the term* tree model *is not used other than by [DHK20]. Other studies often consider the tree model as the only model and usually focus on one type of semantics.*

In the tree model, the initial state is easily defined by the following formula.

Definition 2.24 (Initial state – φ_ϵ)**.** *We define the following formula to describe the initial state at the empty call sequence* ϵ *in the tree model.*

$$\varphi_\epsilon := \bigwedge_{i=j} S_i j \bigwedge_{i \neq j} \neg S_i j$$

Lemma 2.25. *The following statements are equivalent.*

- $M^{\mathtt{t}}_{\mathsf{tree}}, (w_{\mathsf{root}}, \sigma) \models \varphi_\epsilon$ *for any call type* $\mathtt{t}$*.*
- $\sigma = \epsilon$
- $|\sigma| = 0$

Proof. It is clear that φ_ϵ holds in $(w_{\mathsf{root}}, \epsilon)$. After any (first) call ab, two agents necessarily exchange their secrets, so $S_a b$ and $S_b a$ hold, so φ_ϵ does not anymore. As agents do not forget secrets, this will remain true for any following calls. □

Chapter 3

Axiomatisations for Basic Gossip

In this section we discuss the axiomatisation for basic logic as proposed and shown to be sound and complete by [DHK20]. This proof system consists of three layers: a call-free proof system, call axioms, and tree rules. Each additional layer generates a larger logic.

We additionally make some contributions by proposing a few stylistic alterations to the proof systems. Specifically, we show that the call-free proof system is not minimal, and that some axioms are longer than strictly needed. Lastly, we relate these logics to the semantic definitions used in [DG24].

The axiomatisation as introduced originally is highly parameterised because it incorporates versions for all call types. We limit ourselves here to only the two most prevalent call types $(\text{◓}, \diamond, \beta)$ and $(\bullet, \diamond, \beta)$ and do not discuss any other. Throughout this section, we will therefore assume the bidirectional $(\diamond)$ and inspect-then-merge (β) types and do not specify when axioms depend on these types. See [DHK20] for a complete overview of the axioms and their dependence on the call type.

3.1 Axiomatisation for Call-Free Basic Gossip

The following axiomatisation is sound and strongly complete for call-free gossip [DHK20]. The proof system has many familiar elements: it contains the axioms of **S5** and the only additions are gossip-specific axioms **Own**, **PFi**, and **NPi**. These correspond directly to the requirements on the initial models.

Table 3.1. The rules and axioms of $\mathbb{G}$ defined by [DHK20].

Propositional		Knowledge		Secrets (static)	
Prop	propositional tautologies	**K**	$K_a(\varphi \to \psi) \to (K_a\varphi \to K_a\psi)$	**Own**	$S_a a$
MP	$\vdash \varphi, \vdash \varphi \to \psi$ imply $\vdash \psi$	**T**	$K_a\varphi \to \varphi$	**PFi**	$S_a b \to K_a S_a b$
Sub	$\vdash \varphi \leftrightarrow \psi$ implies $\vdash \chi \leftrightarrow \chi[\varphi/\psi]$	**4**	$K_a\varphi \to K_a K_a\varphi$	**NPi**	$\neg S_a b \to K_a \neg S_a b$
		5	$\neg K_a\varphi \to K_a \neg K_a\varphi$		
		Nec(K)	$\vdash \varphi$ implies $\vdash K_a\varphi$		

The system $\mathbb{G}$ is not closed under uniform substitution but instead under *substitution of equivalents*, denoted by **Sub**. In particular, the axioms for secrets cannot be uniformly substituted. We do however consider the logic a *normal modal logic* as by [HHI13] of closure under modus ponens (**MP**), the **K** axiom and necessitation (**Nec**(K)). This normality is sufficient for our goals.

PFi and **NPi** respectively provide positive and negative introspection on atoms: agents know what they do and do not know. However, while $\mathbb{G}$ includes both, only one is strictly needed. Regardless of which one we choose, the other can be derived using symmetry (**B**). We give the derivation in the following proof. While **B** is not part of the axiomatisation, it is commonly known that **S5** = **BT45** and so **B** still follows from our proof system [BRV01].

While only one of the two axioms is needed, we will still include both in the system. In the first place because both are equally intrinsic properties of the problem, so excluding one would become an arbitrary choice. Secondly, future work might consider asymmetric gossip models and accidentally lose the excluded introspection property.

Lemma 3.1. *The axiom* **NPi** *follows from the other axioms in* $\mathbb{G}$.

Proof. We assume $\neg S_a b$ to prove $K_a^P \neg S_a b$ with the following derivation. By $\to$-introduction then $\neg S_a b \to K_a^P \neg S_a b$. The indented lines contain technical steps for substitution and may be ignored for the general proof structure.

1	$\neg S_a b$	
2	$\neg S_a b \to K_a^P \widehat{K}_a^P \neg S_a b$	**B**
3	$K_a^P \widehat{K}_a^P \neg S_a b$	**MP**, 2, 1
4	$\quad \widehat{K}_a^P \neg S_a b \leftrightarrow \neg K_a^P \neg \neg S_a b$	Def. $\widehat{K}_a^P$
5	$\quad K_a^P \widehat{K}_a^P \neg S_a b \leftrightarrow K_a^P \neg K_a^P \neg \neg S_a b$	**Sub**, 4
6	$K_a^P \neg K_a^P \neg \neg S_a b$	**MP**, 3, 5
7	$\quad \neg \neg S_a b \leftrightarrow S_a b$	**Prop**
8	$\quad K_a^P \neg K_a^P \neg \neg S_a b \leftrightarrow K_a^P \neg K_a^P S_a b$	**Sub**, 7
9	$K_a^P \neg K_a^P S_a b$	**MP**, 6, 8
10	$S_a b \to K_a^P S_a b$	**PFi**
11	$(S_a b \to K_a^P S_a b) \to (\neg K_a^P S_a b \to \neg S_a b)$	**Prop** $(*)$
12	$\neg K_a^P S_a b \to \neg S_a b$	**MP**, 11, 10
13	$K_a^P (\neg K_a^P S_a b \to \neg S_a b)$	**Nec**(K), 12
14	$K_a^P (\neg K_a^P S_a b \to \neg S_a b) \to (K_a^P \neg K_a^P S_a b \to K_a^P \neg S_a b)$	**K**
15	$K_a^P \neg K_a^P S_a b \to K_a^P \neg S_a b$	**MP**, 13, 14
16	$K_a^P \neg S_a b$	**MP**, 9, 15

$(*)$ We use an instance of the tautology of contrapositives: $\vdash (p \to q) \to (\neg q \to \neg p)$. □

3.2 Adding Call Reduction Axioms

We now add axioms to $\mathbb{G}$ that describe calls. Which axioms are sound depends on the call type, but we only vary between two call types that only differ in their privacy type. Because the call axioms will contain reduction axioms, these axioms are strongly complete for the class of gossip models too.

We first introduce the axioms for *call basics* and *call effects*, which can be found in Table 3.2. Adding just these axioms to $\mathbb{G}$ does not provide any system of interest, but they do not depend on the privacy type and therefore will be part of both the synchronous and asynchronous proof system. The call basics axioms describe the interactions of calls. Specifically they ensure that the call modality is normal and functional. The call effects describe how the truth of atoms changes during calls.

There is one more set of call axioms that we must add and this is where the two proof systems start diverging. The *observance axioms* for the synchronous and asynchronous case are displayed in Tables 3.3 and 3.4 respectively. Together with the call-free proof system $\mathbb{G}$ and the previous call axioms, they form a sound and complete axiomatisation for either synchronous or asynchronous arbitrary gossip models [DHK20]. We call these systems $\mathbb{G}^{\boldsymbol{\ominus}}$ and $\mathbb{G}^{\bullet}$ respectively and obtain them by combining the axioms from Table 3.1 and Table 3.2 with those of Table 3.3 (for synchronous type $\boldsymbol{\ominus}$) or with those of Table 3.4 (for asynchronous type $\bullet$).

Table 3.2. The call axioms for gossip with directionality type $\mathsf{d} = \diamond$ and observance type $\mathsf{o} = \beta$. These axioms are used in all privacy types.

Call Basics		**Call Effects**		
K$([ab])$	$[ab](\varphi \to \psi) \to ([ab]\varphi \to [ab]\psi)$	**Eff**	$[ab]S_c d \leftrightarrow (S_a d \vee S_b d)$	$c \in \{a, b\}$
Fnc	$[ab]\neg\varphi \leftrightarrow \neg[ab]\varphi$	**Ext**	$[ab]S_c d \leftrightarrow S_c d$	$c \notin \{a, b\}$
Nec$([ab])$	$\vdash \varphi$ implies $\vdash [ab]\varphi$			

Agents are aware of the calls they are involved in regardless of the privacy type, but the privacy type does decide what they learn about other calls. The **Obs** axioms describe the knowledge of both agents in some call. They do so by considering all calls that may have happened. This is relatively straightforward in the synchronous case. For the asynchronous case, an arbitrary number of other calls may have happened immediately after the call an agent was involved in. As there are infinite possibilities, we can no longer consider all of them. Instead [DHK20] has defined a constant bound that does not depend on φ, but only on the number of agents. Specifically, $\mathcal{E}_a(n) := \{\sigma \mid |\sigma| \leq n \text{ and } \sigma \sim_a \epsilon\}$ is the set of call sequences of at most length n that are a-indistinguishable to the empty sequence[1].

While the **Pri** axioms have a different name, they are observance axioms too in the sense that they describe what external agents observe. For the synchronous privacy axiom $\mathbf{Pri}^{\boldsymbol{\ominus}}$, an agent knows something after a call (that they are not involved in) if they know it after every call they are not involved in. For the asynchronous case, we combine two privacy axioms: $\mathbf{Pri}_1^{\bullet}$ states that an agent's knowledge is not influenced by calls they do not participate in; $\mathbf{Pri}_2^{\bullet}$ describes that an agent's knowledge must persist after any call sequence that they were not involved in, because they consider it possible that the sequence has happened already.

Table 3.3. Observance axioms for $\mathsf{t} = (\boldsymbol{\ominus}, \diamond, \beta)$.

Synchronous Observance for Basic Gossip		
$\mathbf{Obs}_1^{\boldsymbol{\ominus}}$	$[ab]K_a\varphi \leftrightarrow \bigvee_{R\subseteq\mathbb{S}}(O_b R \wedge K_a(O_a R \to [ab]\varphi))$	$a \in \{a, b\}$
$\mathbf{Obs}_2^{\boldsymbol{\ominus}}$	$[ab]K_b\varphi \leftrightarrow \bigvee_{R\subseteq\mathbb{S}}(O_a R \wedge K_b(O_b R \to [ab]\varphi))$	$b \in \{a, b\}$
$\mathbf{Pri}^{\boldsymbol{\ominus}}$	$[ab]K_c\varphi \leftrightarrow \bigwedge_{d,e\neq c} K_c[de]\varphi$	$c \notin \{a, b\}$

Table 3.4. Observance axioms for $\mathsf{t} = (\bullet, \diamond, \beta)$.

Asynchronous Observance for Basic Gossip		
$\mathbf{Obs}_1^{\bullet}$	$[ab]K_a\varphi \leftrightarrow \bigvee_{R\subseteq\mathbb{S}}(O_b R \wedge K_a(O_b R \to \bigwedge_{\sigma\in\mathcal{E}_a(2\lvert Ag\rvert^3)}[ab.\sigma]\varphi))$	$a \in \{a, b\}$
$\mathbf{Obs}_2^{\bullet}$	$[ab]K_b\varphi \leftrightarrow \bigvee_{R\subseteq\mathbb{S}}(O_a R \wedge K_b(O_a R \to \bigwedge_{\sigma\in\mathcal{E}_a(2\lvert Ag\rvert^3)}[ab.\sigma]\varphi))$	$b \in \{a, b\}$
$\mathbf{Pri}_1^{\bullet}$	$[\sigma]K_c\varphi \leftrightarrow K_c\varphi$	$c \notin Ag(\sigma)$
$\mathbf{Pri}_2^{\bullet}$	$K_c\varphi \leftrightarrow K_c[\sigma]\varphi$	$c \notin Ag(\sigma)$

Remark 3.2. *The* **Obs** *axioms are completely symmetric. This should be expected, as the directionality type* $\mathsf{d} = \diamond$ *is symmetric too. For this reason, [DHK20] uses the following single but different axiom for* $\boldsymbol{\ominus}$*, which is sound but introduces redundant disjuncts. The* $\bullet$ *version exists too.*

$$\mathbf{Obs}^{\boldsymbol{\ominus}} \quad [ab]K_c^P\varphi \leftrightarrow \bigvee_{Q,R\subseteq\mathbb{S}} (O_a Q \wedge O_b R \wedge K_c^P((O_a Q \wedge O_b R) \to [ab]\varphi))$$

[1] We do not use this bound in this thesis. See corollary 5.6 in [DHK20] for more information.

These disjuncts vary over the possible sets of secrets that the agent $c \in \{a, b\}$ knows. However, as c knows what secrets she knows herself, this is not needed. In this thesis we change the axiom to avoid this redundancy. For a proof of the soundness of these versions, we refer to the proof of lemma 5.14, which shows the soundness of the protocol-dependent versions. The proof is easily adapted to the basic case.

The proof systems $\mathbb{G}^{\textcircled{\scriptsize -}}$ and $\mathbb{G}^{\bullet}$ contain *reduction axioms* for the call modality $[ab]$. These are axioms that are equivalences where the formulas within any call modality on the right hand side are strict subformulas of the formula within the call modality on the left hand side [DHK20].

These call reductions play a crucial role: because the proof system $\mathbb{G}$ is strongly complete for the call-free language on the gossip models, the call reduction axioms allow us to immediately obtain strong completeness for the full language without any additional work.

One reduction axiom is missing: the $\mathbf{K}([ab])$ axiom is an implication and not an equivalence. However, using functionality of the call modality (**Fnc**) one can show its converse to obtain the following validity [DHK20].

Fact 3.3. $\mathbb{G} \vdash [ab](\varphi \to \psi) \leftrightarrow ([ab]\varphi \to [ab]\psi)$

The following result has been proven for all call types [DHK20]. However, the authors have omitted the proofs for the two call types that we consider. We give the proof for the synchronous case here because it is similar to the protocol-dependent setting. In [DHK20], this is the call type $\mathtt{t} = (\textcircled{\scriptsize -}, \diamond, \beta)$. The asynchronous version $(\bullet, \diamond, \beta)$ is more involved and requires a different measure for the induction.

Lemma 3.4. *For every formula $\varphi \in \mathcal{L}^G$ there is a formula $\mathcal{L}^G_{[-]}$ such that $\mathbb{G}^{\textcircled{\scriptsize -}} \vdash \varphi \leftrightarrow \psi$.*

Proof. Without loss of generality we assume that φ only contains the boolean connectives $\neg$ and $\to$ as these are truth-functionally complete. Using **Sub** it furthermore suffices to consider formulas of the form $\varphi = [ab]\chi$ with $\chi \in \mathcal{L}^G_{[-]}$. We use induction on the structure of χ to show $\vdash [ab]\chi \leftrightarrow \psi$ for some $\psi \in \mathcal{L}^G_{[-]}$.

Base case. Suppose $\chi = S_c d$. If $c \notin \{a, b\}$, then **Ext** yields $\vdash [ab]S_c d \leftrightarrow S_c d$. Else we have $c \in \{a, b\}$ and **Eff** yields $\vdash [ab]S_c d \leftrightarrow (S_a d \vee S_b d)$. Both $S_c d \in \mathcal{L}^G_{[-]}$ and $(S_a d \vee S_b d) \in \mathcal{L}^G_{[-]}$, which concludes the base case.

Induction Hypothesis. Let $\chi \in \mathcal{L}^G_{[-]}$ be a formula. Suppose that for every strict subformula χ' of χ, we have some $\psi' \in \mathcal{L}^G_{[-]}$ such that $\vdash [ab]\chi' \leftrightarrow \psi'$.

Induction Step.

- Suppose $\chi = \neg\chi'$. From **Fnc** we obtain $\vdash [ab]\neg\chi' \leftrightarrow \neg[ab]\chi'$. By IH on χ' we have $\vdash [ab]\chi' \leftrightarrow \psi'$. With **Sub** we find that $\vdash [ab]\neg\chi' \leftrightarrow \neg\psi'$, the right hand of which is in $\mathcal{L}^G_{[-]}$.
- Suppose $\chi = (\chi' \to \chi'')$. By fact 3.3 we have $\vdash [ab](\chi' \to \chi'') \leftrightarrow ([ab]\chi' \to [ab]\chi'')$. By IH on χ' and χ'' we get $\vdash [ab]\chi' \leftrightarrow \psi'$ and $\vdash [ab]\chi'' \leftrightarrow \psi''$. Using **Sub** this yields $\vdash [ab](\chi' \to \chi'') \leftrightarrow (\psi' \to \psi'')$, the right hand of which is in $\mathcal{L}^G_{[-]}$.
- Suppose $\chi = K_c\chi'$. We distinguish three cases for c.

  For $c = a$ we use $\mathbf{Obs}^{\textcircled{\scriptsize -}}_1$ to get $[ab]K_c\chi' \leftrightarrow \bigvee_{R\subseteq\mathbb{S}}(O_b R \wedge K_c(O_a R \to [ab]\chi'))$. By IH on χ' we have $\vdash [ab]\chi' \leftrightarrow \psi'$. Using **Sub** we find that $[ab]K_c\chi \leftrightarrow \bigvee_{R\subseteq\mathbb{S}}(O_b R \wedge K_c(O_a R \to \psi'))$, the right hand of which is a formula in $\mathcal{L}^G_{[-]}$.

  For $c = b$ we instead use $\mathbf{Obs}^{\textcircled{\scriptsize -}}_2$ and proceed analogously.

  For $c \notin \{a, b\}$ we use $\mathbf{Pri}^{\textcircled{\scriptsize -}}$ to get $[ab]K_c\chi \leftrightarrow \bigwedge_{d,e\neq c} K_c[de]\chi'$. By IH on χ' we have $\vdash [ab]\chi' \leftrightarrow \psi'$ and $\vdash [de]\chi \leftrightarrow \psi_{de}$ for all $d, e \neq a$. Using **Sub** we find that $[ab]K_c\chi \leftrightarrow \bigwedge_{d,e\neq c} K_c\psi_{de}$, the right hand of which is a formula in $\mathcal{L}^G_{[-]}$.

This finishes the induction on χ. We conclude that for every $[ab]\chi \in \mathcal{L}^G$ we have an equivalent $\psi \in \mathcal{L}^G_{[-]}$. As every formula $\varphi \in \mathcal{L}^P$ can be written in this form, we are done. □

For each formula we can find a call-free version, which we call its *call reductions.*

Definition 3.5 (Call-Reduction)**.** *Given a formula $\varphi \in \mathcal{L}^G$, we call $\psi \in \mathcal{L}^G_{[-]}$ the* call reduction *of φ if $\vdash \varphi \leftrightarrow \psi$.*

3.3 Axiomatisations for Tree Models

At this point we have three proof systems: $\mathbb{G}$ for the call-free language, $\mathbb{G}^{\mathsf{e}}$ for synchronous arbitrary gossip models and $\mathbb{G}^{\mathsf{b}}$ for asynchronous arbitrary gossip models. However, none of these are models often used in literature, where only the specific tree models are considered. We can obtain the axiomatisation for these models by adding yet one final set of rules, which stipulates that any validities that follow from the initial root world of the tree model must be included. We call these *tree rules* and write $\mathbb{G}^{\mathsf{e}}_{\mathsf{tree}}$ and $\mathbb{G}^{\mathsf{b}}_{\mathsf{tree}}$ to denote the proof systems including them. Table 3.5 contains the synchronous tree rule and Table 3.6 lists the asynchronous tree rules.

The main challenge with the tree rules is the common knowledge in the initial world: the gossip language does not have a modality for common knowledge, but it can be approximated [DHK20]. The approach to this approximation is different for the synchronous and asynchronous case, leading to different axiomatisations.

In the synchronous case, this can be done rather straightforwardly using n-bisimulation. The asynchronous case is again more involved. Whereas with the synchronous case an agent considers only a finite number of call sequences possible, in the asynchronous case this becomes infinite: any call sequence of arbitrary length may have happened without their knowledge.

The formula root equals φ_ϵ and describes the initial distribution of secrets in the tree model. The synchronous version uses root_n to approximate the common knowledge in the initial root world up to n-bisimilarity. We will discuss this construction in more detail in chapter 5. For the rationale behind the asynchronous version we refer to [DHK20].

$$\mathsf{root}_0 := \varphi_\epsilon$$
$$\mathsf{root}_{i+1} := \mathsf{root}_i \wedge \bigwedge_{a \in Ag} K_a \mathsf{root}_i$$

Table 3.5. Tree rules for privacy type $\mathsf{p} = \mathsf{e}$, where n is the degree of φ and m is the number of n-bisimilarity classes.

Synchronous Tree Rule	
$\mathbf{Tree}^{\mathsf{e}}$	If $\vdash \mathsf{root}_n \to [\sigma]\varphi$ for all σ s.t. $\lvert\sigma\rvert \leq m$ then $\vdash \varphi$

Table 3.6. Tree rules for privacy type $\mathsf{p} = \mathsf{b}$, where n is the degree of φ and m is the number of n-bisimilarity classes.

Asynchronous Tree Rules	
$\mathbf{Tree}_1^{\mathsf{b},m}$	If $\vdash \mathsf{root} \to [\sigma]\varphi$ for all σ s.t. $\lvert\sigma\rvert \leq m \cdot (\lvert(\lvert\tau\rvert + 1)$ and $(w_{\mathsf{root}}, \sigma) \sim_a (w_{\mathsf{root}}, \tau)$ then $\vdash \mathsf{root} \to [\tau]K_a\varphi$
$\mathbf{Tree}_2^{\mathsf{b},m}$	If $\vdash \mathsf{root} \to [\sigma]\varphi$ for some σ s.t. $\lvert\sigma\rvert \leq m \cdot (\lvert(\lvert\tau\rvert + 1)$ and $(w_{\mathsf{root}}, \sigma) \sim_a (w_{\mathsf{root}}, \tau)$ then $\vdash \mathsf{root} \to [\tau]\widehat{K}_a\varphi$
$\mathbf{Tree}_3^{\mathsf{b},m}$	If $\vdash \mathsf{root} \to [\sigma]\varphi$ for all σ s.t. $\lvert\sigma\rvert \leq m$ then $\vdash \varphi$

3.4 Identifying the Logic for "Limits to Gossip"

The paper "Limits to Gossip" [DG22] and its later extension [DG24] demonstrate a limit to the level of higher order epistemic knowledge that can be achieved in the gossip problem. We show that there must exist a syntactic proof in the proof system $\mathbb{G}^{\mathsf{b}}_{\mathsf{tree}}$ because the semantics in [DG22] coincide with the tree model defined in definition 2.22.

Throughout this section, we will distinguish the operators defined in [DG22] from those defined in chapter 2 with the superscript $(.^{\mathsf{LtG}})$.

Fact 3.6. *For any agent, the epistemic relation between call sequences in [DG22] is equivalent to the asynchronous epistemic relation in definition 2.16 on the asynchronous tree model $M^{\bullet}_{\mathsf{tree}}$, that is*

$$\forall a : \sigma \sim_a^{\mathsf{LtG}} \tau \iff (w_{\mathsf{root}}, \sigma) \sim_a (w_{\mathsf{root}}, \tau).$$

Lemma 3.7. *Let $\varphi \in \mathcal{L}^G_{[-]}$ and σ be some call sequence. We have $\sigma \models^{\mathsf{LtG}} \varphi$ if and only if $M^{\bullet}_{\mathsf{tree}}, (w_{\mathsf{root}}, \sigma) \models \varphi$.*

Proof. We use induction on the structure of φ. We omit all cases except the knowledge modality, as they are not very insightful.

Induction Hypothesis. Let $\varphi \in \mathcal{L}^{P1}_{[-]}$ be arbitrary. Suppose that for all strict subformulas ψ of φ we have $\sigma \models^{\mathsf{LtG}} \psi$ if and only if $M^{\bullet}_{\mathsf{tree}}, (w_{\mathsf{root}}, \sigma) \models \varphi$.

Induction Step. Let $\varphi = K_a\psi$.

$$\begin{aligned}
& M^{\bullet}_{\mathsf{tree}}, (w_{\mathsf{root}}, \sigma) \models K_a\psi \\
\iff & \text{For all } (w', \sigma') : (w_{\mathsf{root}}, \sigma) \sim_a (w', \sigma') \text{ implies } M^{\bullet}_{\mathsf{tree}}, (w', \sigma') \models \psi && \text{(Def. semantics } K_a) \\
\iff & \text{For all } (w_{\mathsf{root}}, \sigma') : (w_{\mathsf{root}}, \sigma) \sim_a (w', \sigma') \text{ implies } M^{\bullet}_{\mathsf{tree}}, (w_{\mathsf{root}}, \sigma') \models \psi && \text{(Def. } M^{\bullet}_{\mathsf{tree}}) \\
\iff & \text{For all } (w_{\mathsf{root}}, \sigma') : \sigma \sim_a^{\mathsf{LtG}} \sigma' \text{ implies } M^{\bullet}_{\mathsf{tree}}, (w_{\mathsf{root}}, \sigma') \models \psi && \text{(By fact 3.6)} \\
\iff & \text{For all } \sigma' : \sigma \sim_a^{\mathsf{LtG}} \sigma' \text{ implies } \sigma' \models^{\mathsf{LtG}} \psi && \text{(By IH)} \\
\iff & \sigma \models^{\mathsf{LtG}} K_a\psi && \text{(Def. semantics of } K_a^{\mathsf{LtG}})
\end{aligned}$$

□

Hence, the semantics of the model in [DG22] are equivalent to those of $M^{\bullet}_{\mathsf{tree}}$, for which we have a sound and complete proof system for the language $\mathcal{L}^G$. Since $\mathcal{L}^G_{[-]}$ is a fragment of $\mathcal{L}^G$ we can use its proof system to obtain a proof in $\mathcal{L}^G$.

Theorem 3.8. *Let $\varphi \in \mathcal{L}^G_{[-]}$. If we have $\sigma \models^{\mathsf{LtG}} \varphi$ for all σ, then there exists a proof for φ in $\mathbb{G}^{\bullet}_{\mathsf{tree}}$.*

Proof. Let $\varphi \in \mathcal{L}^G_{[-]}$ be arbitrary. By lemma 3.7 the semantics on $\models^{\mathsf{LtG}}$ are equivalent to the semantics of $M^{\bullet}_{\mathsf{tree}}$. Since $\mathcal{L}^G_{[-]}$ is a fragment of $\mathcal{L}^G$, we have $\varphi \in \mathcal{L}^G$ too. Because $\mathbb{G}^{\bullet}_{\mathsf{tree}}$ is a complete proof system for $\mathcal{L}^G$ on $M^{\bullet}_{\mathsf{tree}}$, we can thus derive φ from it. □

The main result of [DG22] is a semantic proof that $\neg EEExp_{Ag}$ is a validity. It uses two abbreviations that we have not yet discussed:

- $E\varphi := \bigwedge_{a \in Ag} K_a\varphi$ for *"Everyone knows that φ is true"*;
- $Exp_A := \bigwedge_{a,b \in Ag} S_ab$ for *"Agents $A \subseteq Ag$ are experts"*.

The formula is thus read as *"Not everyone knows that everyone knows that everyone is an expert"*. Moreover, the more recent extension of this paper shows that for some agent a, the formula $\neg K_aEExp_{Ag}$ is a validity too [DG24].

Corollary 3.9. *There exists a syntactic proof for $\neg EEExp_{Ag}$ and $\neg K_aEExp_{Ag}$ in $\mathbb{G}^{\bullet}_{\mathsf{tree}}$.*

Chapter 4

Protocol-Dependent Gossip

In this section we introduce the notion of protocol-dependent knowledge in gossip models. We define gossip models with protocol-dependent knowledge in a novel way, that unites their original definition by [Dit+19] with the model construction of [DHK20].

4.1 Language and Syntax

Definition 4.1 (Protocol & Protocol Conditions)**.** *A protocol P is a set of $n \cdot (n-1)$* protocol conditions *P_{ab} for each pair of agents $a \neq b \in Ag$. We denote a protocol by a Roman capital letter and let $\mathbb{P}$ be the set of all protocols.*

Definition 4.2 (P-Permitted Calls)**.** *Given a protocol P, a call ab is* P-permitted *if the protocol condition P_{ab} is true in the state at which the call is made. A call sequence σ is P-permitted if for each call ab in σ its protocol condition P_{ab} is true when the call was made. When a call or sequence is not P-permitted, we call it* P-illegal*. In both cases, we simply say* permitted *or* illegal *if the protocol is clear from context.*

There are various well-known protocols for gossip, such as the protocol LNS *(Learn New Secrets)*. In LNS, agents are only allowed to make a call to an agents whose secret they do not yet know. We will also make use of the somewhat trivial protocol ANY *(Any Call)*, which allows any call.

Definition 4.3 (Learn New Secrets)**.** *Let LNS be the protocol defined by the following protocol conditions.*

$$\mathsf{LNS}_{ab} := \neg S_a b \qquad \text{for all } a \neq b \in Ag$$

Definition 4.4 (Any Call)**.** *Let ANY be the protocol defined by the following protocol conditions.*

$$\mathsf{ANY}_{ab} := \top \qquad \text{for all } a \neq b \in Ag$$

The language for protocol-dependent gossip is defined as follows [Dit+19]. As there are infinitely many protocols, the following language has infinitely many knowledge modalities.

Definition 4.5 (Protocol-dependent Language)**.** *Let $a, b \in Ag$ be agents and $P \in \mathbb{P}$ a protocol. We define the language $\mathcal{L}^P$ for arbitrary protocols P as*

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

We again define a static version of this language $\mathcal{L}^P_{[-]}$ by omitting the call modality.

The relation between the protocol P in definition 4.1 and the epistemic operator K_a^P in definition 4.5 may be seen as follows [Dit+19]: for n agents, let $K_a^P(P_{ab}, ..., P_{mn}, \varphi)$ be an operator with arity $n \cdot (n-1) + 1$ for the $n \cdot (n-1)$

protocol conditions and the subformula φ itself. In order to ensure that such a definition is well-founded, we do not allow self-reference of protocols: a protocol condition of P cannot contain the K^P operator (for any agent), nor can it contain any protocol that itself refers to K^P, directly or indirectly.

With this syntactic justification it becomes natural to view the protocol conditions as subformulas of K_a^P. We provide the formal definition here, which was not yet provided by [Dit+19].

Definition 4.6 (Subformula in $\mathcal{L}^P$)**.** *The set of subformulas in $\mathcal{L}^P$ is defined as follows.*

$$\begin{aligned} sub(S_a b) &:= \{S_a b\} \\ sub(\neg\varphi) &:= \{\neg\varphi\} \cup sub(\varphi) \\ sub(\varphi \wedge \psi) &:= \{\varphi \wedge \psi\} \cup sub(\varphi) \cup sub(\psi) \\ sub(K_a^P \varphi) &:= \{K_a^P \varphi\} \cup \{P_{ab} \mid a \neq b \in Ag\} \cup sub(\varphi) \\ sub([ab]\varphi) &:= \{[ab]\varphi\} \cup sub(\varphi) \end{aligned}$$

We call $sub(\varphi) \setminus \varphi$ the set of strict subformulas *of φ.*

The protocol-dependent language does not have a regular knowledge modality. However, we define $K_a := K_a^{\mathsf{ANY}}$ as an abbreviation in $\mathcal{L}^P$ because ANY is semantically equivalent to the basic epistemic modality K_a [Dit+19].

While $\mathcal{L}^P$ permits any protocol, we can also restrict this language by limiting the protocols that can be used. Various results in this thesis will be about the language restricted to only one protocol.

Definition 4.7 (Protocol-Dependent Language with a Single Protocol)**.** *Let $a, b \in Ag$ be agents and P be an arbitrary but fixed protocol. We define the language $\mathcal{L}^{P1}$ as follows.*

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

4.2 Models and Semantics

We now propose a definition for protocol-dependent gossip models. This definition differs from [Dit+19] and aims to unify the semantics of protocol-dependent knowledge with the model construction from chapter 2. In particular, we will define initial models for protocol-dependent gossip, which generalise the setting of the gossip problem.

Initial Models

We first define initial models for protocol-dependent gossip. These models generalise the problem for arbitrary distributions and agent knowledge as in section 2.4, but in the protocol-dependent setting they also generalise the concept of protocol violation, allowing that protocols may already be violated in the initial state.

Definition 4.8 (Initial Model for Protocol-Dependent Gossip)**.** *Let $\mathcal{P}$ be a set of protocols. A protocol-dependent initial model is a triple $I^{\mathcal{P}} = \langle W_0, R_0, V_0 \rangle$ where W is a set of initial worlds, $R_0 : Ag \times \mathcal{P} \to 2^{W_0 \times W_0}$ is a relation for each agent and protocol that is transitive and symmetric, and $V_0 : Ag \times W_0 \to 2^{\mathbb{S}}$ is a function mapping each agent to the initial set of secrets known in some world, which satisfies (i) $a \in V_0(a, w)$ for all $w \in W$ and (ii) if $(w_1, w_2) \in R_0(a)$ then $V_0(a, w_1) = V_0(a, w_2)$. When the set of protocols is clear from context, we omit it and write I.*

The definition above is very similar to definition 2.12. There are two differences. Firstly, R_0 is now extended to a unique relation for each $P \in \mathcal{P}$. More importantly, the relations have lost their reflexivity and therefore are no longer necessarily equivalence relations. Definition 4.8 permits reflexive relations too, making it a generalisation of definition 2.12 of initial models for basic gossip.

Remark 4.9. *These models may contain unconnected worlds. We usually relate an unconnected world to protocol violation. However, a protocol cannot be violated when no calls have happened. Since initial models describe the situation before any calls, it makes sense to require that the model is reflexive. We could model such a situation simply in the same way as definition 2.12, with one relation per agent that all are equivalence relations. Definition 4.8 is more lenient however and loses the relation between unconnectedness and protocol violation. We motivate this choice with two arguments. Most importantly, we will need this loosened definition in order to construct the canonical model and show completeness. Secondly, we are not particularly interested in gossip models starting in arbitrary states. Because we are mainly looking for an axiomatisation for the tree model, there is no problem in utilising a larger class of models to achieve this. We discuss these decisions in more detail in chapter 8.*

We can define semantics on these initial models. We will use these in the completeness proof and therefore define the semantic relation here.

Definition 4.10 (Semantics on Protocol-Dependent Initial Models)**.** *By $\models_i$ we denote the following standard semantic relation for modal logics on pointed initial models. Let I be an initial model and w a world in I.*

$$
\begin{aligned}
I, w &\models_i S_a b && \iff && b \in V_0(a, w, \sigma) \\
I, w &\models_i \neg\varphi && \iff && I, w \not\models_i \varphi \\
I, w &\models_i \varphi \wedge \psi && \iff && I, w \models_i \varphi \text{ and } I, w \models_i \psi \\
I, w &\models_i K_a^P \varphi && \iff && I, w' \models_i \varphi \text{ for all } w' \text{ s.t. } (w, w') \in R_0(a, P)
\end{aligned}
$$

Lifting the Model

The definitions for the set of states and the valuations of these states remain the same as the basic case and can be reviewed in definitions 2.14 and 2.15.

With protocol-dependent knowledge, Agents only consider states possible whose call sequences are permitted by the protocol they assume. For each protocol $P \in \mathcal{P}$ we must verify whether the call sequence σ satisfies the protocol conditions at each call.

The following definition for the synchronous case stays close to [Dit+19]. In particular, this relation is a partial equivalence and therefore is transitive and symmetric. It is an equivalence relation when restricted to P-permitted states. The asynchronous case was not defined before but is a simple modification.

Definition 4.11 (Protocol-Dependent Relation)**.** *Given an initial model $I^{\mathcal{P}} = \langle W_0, R_0, V_0 \rangle$, we lift the initial relation $R_0(a, P) \subseteq W_0 \times W_0$ for some protocol P and agent a to $\sim_a^P \subseteq W \times W$ such that:*

$$
\begin{aligned}
(w_i, \epsilon) \sim_a^P (w_j, \epsilon) \text{ iff } & (w_i, w_j) \in R_0; \\
(w_i, \sigma.ab) \sim_a^P (w_j, \tau.ab) \text{ iff } & (w_i, \sigma) \sim_a^P (w_j, \tau) \\
& \text{and } V_b(w_i, \sigma) = V_b(w_j, \tau) \\
& \text{and } (w_i, \sigma) \models P_{ab} \text{ and } (w_j, \tau) \models P_{ab}; \\
(w_i, \sigma.ba) \sim_a^P (w_j, \tau.ba) \text{ iff } & (w_i, \sigma) \sim_a^P (w_j, \tau) \\
& \text{and } V_b(w_i, \sigma) = V_b(w_j, \tau) \\
& \text{and } (w_i, \sigma) \models P_{ba} \text{ and } (w_j, \tau) \models P_{ba};
\end{aligned}
$$

and for the synchronous privacy type ◓ we let

$$
\begin{aligned}
(w_i, \sigma.bc) \sim_a^P (w_j, \tau) \text{ iff } & (w_i, \sigma) \sim_a^P (w_j, \tau) \\
& \text{and } a \notin \{b, c\} \\
& \text{and } (w_i, \sigma) \models P_{bc} \text{ and } (w_j, \tau) \models P_{de};
\end{aligned}
$$

while for the asynchronous privacy type ● *we let*

$$(w_i, \sigma.bc) \sim_a^P (w_j, \tau) \text{ iff } (w_i, \sigma) \sim_a^P (w_j, \tau) \text{ and } a \notin \{b, c\} \text{ and } (w_i, \sigma) \models P_{bc}.$$

Definition 4.12 (Protocol-Dependent Gossip Model)**.** *Let $I^{\mathcal{P}} := \langle W_0, R_0, V_0 \rangle$ be a protocol-dependent initial model for a set of protocols $\mathcal{P}$. Given a call type t, we define an (induced) gossip model $M^{\mathsf{t}}_{\mathsf{PD}}(I) := \langle W, \sim, V \rangle$ with W and V as defined in definitions 2.14 and 2.15 respectively„ and the epistemic relation $\sim_a^P$ for each agent $a \in Ag$ and protocol $P \in \mathcal{P}$ as defined in definition 4.11. When the initial model or privacy level are clear from context, we omit them and write M or $M^{\mathsf{t}}_{\mathsf{PD}}$.*

Observe that while protocols may not allow certain calls, it is always possible to execute each call. The trees are therefore identical in shape to the basic gossip models.

Remark 4.13. *We could use definition 4.8 as initial models for basic gossip models too. However, this will lead to a larger class of basic (induced) gossip models and this class of models would not satisfy* **S5***.*

We again define the semantics on these models in the standard way. The only difference with the semantics on the basic gossip is the replacement of the K_a modality for K_a^P.

Definition 4.14 (Semantics on Protocol-Dependent Gossip Models)**.** *We define $\models$ on pointed protocol-dependent gossip models as follows. Let $M = M^{\mathsf{t}}_{\mathsf{PD}}(I)$ be a gossip model induced from $I = \langle W_0, R_0, V_0 \rangle$ and $w \in W_0$ some initial world. Finally, let σ be a call sequence.*

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

For a class of models $\mathcal{M}$, we call a formula φ valid on *$\mathcal{M}$ if for all models $M \in \mathcal{M}$ and for all states s of M we have $M, s \models \varphi$. For a single model M, we call φ* valid on *M if for all states s of M we have $M, s \models \varphi$.*

Tree Models

We finally introduce the tree model of protocol-dependent gossip classes, in the same way as with the basic gossip classes. It is important to stress that the initial tree model remains identical to definition 2.21, except for the technical addition of extending the initial relation to a protocol set $\mathcal{P}$. This also means that all relations in the initial tree model are equivalence relations and satisfy **S5**. However, this property is lost for any non-trivial protocol as soon as the initial model is lifted.

Definition 4.15 (Protocol-Dependent Initial Root Model)**.** *Given a set of protocols $\mathcal{P}$, let $I^{\mathcal{P}}_{\mathsf{root}} = \langle W_{\mathsf{root}}, R_{\mathsf{root}}, V_{\mathsf{root}} \rangle$ with:*

- $W_{\mathsf{root}} := \{w_{\mathsf{root}}\}$;
- $R_{\mathsf{root}}(a, P) := \{(w_{\mathsf{root}}, w_{\mathsf{root}})\}$ *all agents $a \in Ag$ and protocols $P \in \mathcal{P}$;*
- $V_{\mathsf{root}}(a, w_{\mathsf{root}}) := \{S_a a\}$ *all agents $a \in Ag$.*

Definition 4.16 (Tree Models for Protocol-Dependent Gossip)**.** *A protocol-dependent tree model* $M^{\mathsf{t}}_{\mathsf{PDtree}} := M^{\mathsf{t}}_{\mathsf{PD}}(I^{\mathcal{P}}_{\mathsf{root}})$ *is a gossip model induced from the initial root model* $I^{\mathcal{P}}_{\mathsf{root}}$ *given a call type* t*. When the tree model is clear from context, we often omit it fully and simply refer to the call sequences, such that* $\sigma \models \varphi$ *is a shorthand for* $M^{\mathsf{t}}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \sigma) \models \varphi$*.*

4.3 Effects of Protocol-Dependent Knowledge

Protocol-dependent knowledge modalities allow us to combine protocols freely. Without it, one could still model protocol-dependent knowledge by encoding it in the regular epistemic relation. However, such a model could only ever assume one protocol. We have already given an example formula in the introduction, but let us now use concrete protocols.

$$K_a^{\mathsf{ANY}}\varphi \wedge \neg K_a^{\mathsf{LNS}}\varphi \qquad \text{“}A \text{ knows } \varphi \text{ given } \mathsf{ANY} \text{ but not given } \mathsf{LNS}\text{”}$$

This formula should never be true: ANY allows any call, so all call sequences are ANY-permitted. LNS is a more specific protocol, so the agent considers less call sequences possible. Hence the epistemic relation of $\sim_a^{\mathsf{LNS}}$ is contained in the relation $\sim_a^{\mathsf{ANY}}$. If φ holds in all ANY-related sequences, then it holds in all LNS-related sequences.

This formula is indeed unsatisfiable in the tree model for protocol-dependent gossip. We can even generalise LNS to an arbitrary protocol, because no protocol is less restrictive than ANY.

Fact 4.17. *In the tree model* $K_a^{\mathsf{ANY}}\varphi$ *implies* $K_a^P\varphi$ *for any protocol* $P \in \mathbb{P}$*.*

The P-dependent relation effectively filters out P-illegal call sequences by verifying each new call. Suppose that a protocol permits the sequence *ab.ba.cd*. If we only made the call *ab* so far, we have not yet violated any protocol condition. Hence *ab* and *ab.ba* are also still P-permitted.

Fact 4.18. *All prefixes of a* P*-permitted call sequence are* P*-permitted too.*

The initial root world is permitted by every protocol, so in the tree model we also get a stronger property.

Fact 4.19. *In the tree model, the empty sequence cannot violate any protocol.*

Recall that calls, even P-illegal ones, are still always possible. It would not make sense to forbid executing them because we never settle on one protocol, instead allowing the use of any K^P modality at any point. The modality effectively puts a constraint on the history: we limit the sequences that an agent consider possible. What happens however if we find ourselves in a P-illegal sequence?[1]

As can be seen in definition 4.11, a P-illegal sequence is excluded from the epistemic relation $\sim_a^P$. This not only means that it is unconnected from all gossip states, but also that it is not connected to itself. Since there is not a single P-dependent relation, the agent starts believing anything vacuously, even $\bot$. The agent has turned insane. We should be clear that the state itself is not inconsistent, only the knowledge base of the agent. Moreover, only their knowledge base *relative to* P. It is therefore more correct to say that they are P-insane.

In the tree model this property is stronger. Firstly, any sequence satisfying $K_a^P\bot$ must be P-illegal because no initial world can violate any protocol. The violation must come from some call in the sequence. Secondly, the initial model's relation agrees for agents and all protocols. We can therefore guarantee that the alarm is globally recognised by all agents: violation of P for a is identical to violation of P for b.

Corollary 4.20 (Global Alarm)**.** *Given the tree model. If a call sequence is* P*-illegal, then it satisfies* $K_a^P\bot$ *for any agent* a*.*

Corollary 4.21. *Given the tree model. If a call sequence is* P*-permitted, then it satisfies* $\widehat{K}_a^P\top$ *for any agent* a*.*

[1] This happens almost constantly. Consider the protocol $\mathsf{NOTANY}_{ab} := \bot$ that allows *Not Any Call*. Every gossip state with a non-empty sequence is NOTANY-illegal.

Chapter 5

Axiomatisations for Single Protocol-Dependent Gossip

In this section we propose axioms for arbitrary protocol-dependent and synchronous gossip models and protocol-dependent knowledge therein. We will do so in three layers similar to chapter 3 and show soundness and completeness results for each. Throughout this section we use P as an arbitrary fixed protocol in the language $\mathcal{L}^{P1}$ and assume all models to be synchronous.

5.1 Axiomatisation for Call-Free Protocol-Dependent Gossip

The protocol-dependent initial models defined in chapter 4 are in many ways similar to the basic initial models. The only semantic difference is the lack of reflexivity, which means that the **T** axiom is no longer sound. We do however have a transitive and symmetric relation. The axioms that correspond to these properties are **4** and **B** respectively [BRV01]. We therefore obtain proof system $\mathbb{G}$**P1** in Table 5.1. It generates a normal logic up to the exception of uniform substitution, similar to its counterpart $\mathbb{G}$ in chapter 3 for basic gossip.

Table 5.1. The rules and axioms of $\mathbb{G}$**P1** for call-free protocol-dependent gossip with a single protocol P.

Propositional		Knowledge		Secrets (static)	
Prop	propositional tautologies	**K**(P)	$K_a^P(\varphi \to \psi) \to (K_a^P\varphi \to K_a^P\psi)$	**Own**	$S_a a$
MP	$\vdash \varphi, \vdash \varphi \to \psi$ imply $\vdash \psi$	**B**(P)	$\varphi \to K_a^P \widehat{K}_a^P \varphi$	**PFi**(P)	$S_a b \to K_a^P S_a b$
Sub	$\vdash \varphi \leftrightarrow \psi$ implies $\vdash \chi \leftrightarrow \chi[\varphi/\psi]$	**4**(P)	$K_a^P\varphi \to K_a^P K_a^P \varphi$	**NPi**(P)	$\neg S_a b \to K_a^P \neg S_a b$
		Nec(P)	$\vdash \varphi$ implies $\vdash K_a^P\varphi$		

Lemma 5.1. $\mathbb{G}$**P1** *is sound for the class of protocol-dependent initial models with a single protocol.*

Proof. Let $a \in Ag$ and $I^P = \langle W_0, R_0, V_0\rangle$ be an arbitrary initial model. The soundness of **B**(P) and **4**(P) are immediate from the symmetry and transitivity of $R_0(a, P)$ in definition 4.8. **Own** is sound because we have $a \in V_0(a, w)$ for all initial worlds w by definition 4.8.

For **PFi**(P) and **NPi**(P), let $w \in W_0$ and $a \in Ag$ be arbitrary. We have $V_0(a, w) = V_0(a, w')$ for all w' satisfying $(w, w') \in R_0(a, P)$ by definition 4.8. Hence for all $b \in Ag$ we have that all $R_0(a, P)$-related worlds agree on $S_a b$. By semantics of K_a^P we therefore obtain both $w \models S_a b$ implies $w \models K_a^P S_a b$ as well as $w \models \neg S_a b$ implies $w \models K_a^P \neg S_a b$. As w and a were arbitrary, this holds globally and thus **PFi**(P) and **NPi**(P) are sound. □

Lemma 5.2. $\mathbb{G}$**P1** *is sound for the class of protocol-dependent gossip models with a single protocol.*

Proof. Let $a \in Ag$ and $M^{\mathsf{p}}_{\mathsf{PD}}(I) = \langle W, \sim, V\rangle$ be an arbitrary gossip model with $\mathsf{p} \in \{\mathsf{e}, \bullet\}$. Axioms $\mathbf{B}(P)$ and $\mathbf{4}(P)$ are sound because $\sim^P_a$ in definition 4.11 is symmetric and transitive.

Own is sound in the initial model by lemma 5.1, so for all $w \in W_0$ we have $S_a a \in V_a(w, \epsilon)$. **Own** remains sound for $M^{\mathsf{p}}_{\mathsf{PD}}(I)$ because for any sequence σ and call ab we have $V_a(w, \sigma) \subseteq V_a(w, \sigma.ab)$ by definition 2.15.

Finally, $\sim^P_a$ in definition 4.11 retains the same condition for each relation as $R_0(a, P)$ in the initial model, so $\mathbf{PFi}(P)$ and $\mathbf{NPi}(P)$ also remain sound. □

We now show that $\mathbb{G}\mathbf{P1}$ is complete for the call-free language on both the initial models and gossip models with a single protocol. The proof method is similar to that of [DHK20] and relies on defining a canonical initial model that can be lifted to a canonical gossip model.

We will use the following basic definitions.

Definition 5.3. *Given a language $\mathcal{L}$ and a logic λ, a set Δ is λ-consistent if $\lambda, \Delta \nvdash \bot$ and maximal if for all $\varphi \in \mathcal{L}$ either $\varphi \in \Delta$ or $\neg\varphi \in \Delta$. A set Δ is* maximally λ-consistent *(λ-MCS) if it is both maximal and λ-consistent. We omit the logic λ whenever it is clear from context.*

Definition 5.4. *Let $(K^P_a)^-\Delta := \{\varphi \mid K^P_a\varphi \in \Delta\}$ be the* modal projection *of Δ for protocol P and agent a.*

The canonical initial model is constructed in the standard way, with MCSs as worlds, the valuation of atoms matching with membership of the MCS, and the modal relationship based on the modal projection of the MCSs.

Definition 5.5 (Canonical Initial Model)**.** *The canonical initial model $I^c = \langle W^c_0, R^c_0, V^c_0\rangle$ is given by*

- $W^c_0 := \{\Gamma \mid \Gamma$ *is a* $\mathbb{G}\mathbf{P1}$*-MCS*$\}$*;*
- $R^c_0(a, P) := \{(\Gamma, \Delta) \in W^c \times W^c \mid (K^P_a)^-\Gamma \subseteq \Delta\}$ *for all agents $a \in Ag$;*
- $V^c_0(a, \Gamma) := \{b \mid S_a b \in \Gamma\}$ *for all agents $a \in Ag$ and worlds $\Gamma \in W^c_0$.*

While we can lift I^c to obtain a gossip model, we will first limit ourselves to the initial model I^c only. It will become apparent that this is sufficient for our completeness result.

We first show an existence lemma relating the $\mathcal{L}^{P1}_{[-]}$-MCSs to one another. The proof is standard for normal modal logics [BRV01].

Lemma 5.6 (Existence Lemma)**.** *Let Γ be a $\mathbb{G}\mathbf{P1}$-MCS. For any agent $a \in Ag$, if $\neg K^P_a\varphi \in \Gamma$ then there is an MCS Δ such that $\neg\varphi \in \Delta$ and $(K^P_a)^-\Gamma \subseteq \Delta$.*

Proof. Let a be an arbitrary agent. Let Γ be an MCS and suppose that $K^P_a\varphi \in \Gamma$. We show that there exists an MCS Δ such that $\neg\varphi \in \Delta$ and $(K^P_a)^-\Gamma \subseteq \Delta$.

Let $\Delta^- := \{\neg\varphi\} \cup (K^P_a)^-\Gamma$. We claim that Δ^- is consistent. For suppose not, then there is a finite set $\psi_1, ..., \psi_n$ such that $\vdash (\psi_1 \wedge ... \wedge \psi_n) \to \neg\neg\varphi$, because $(K^P_a)^-\Gamma$ is consistent by definition. By propositional calculus this is equivalent to $\vdash (\psi_1 \wedge ... \wedge \psi_n) \to \varphi$ and by $\mathbf{Nec}(K)$ we find that $\vdash K^P_a(\psi_1 \wedge ... \wedge \psi_n) \to K^P_a\varphi$. As $\mathbb{G}\mathbf{P1}$ is normal[1], we have $\vdash (K^P_a\psi_1 \wedge ... \wedge K^P_a\psi_n) \to K^P_a(\psi_1 \wedge ... \wedge \psi_n)$ and by **MP** we obtain $\vdash (K^P_a\psi_1 \wedge ... \wedge K^P_a\psi_n) \to K^P_a\varphi$.

We have $(K^P_a\psi_1 \wedge ... \wedge K^P_a\psi_n) \in \Gamma$ since $K^P_a\psi_1, ..., K^P_a\psi_n \in \Gamma$ and Γ is maximal, and so by **MP** we have $K^P_a\varphi \in \Gamma$. By assumption $\neg K^P_a\varphi \in \Gamma$, so we find that $\Gamma \vdash \bot$. This is a contradiction as Γ is consistent. We conclude that Δ^- must be consistent too.

Let then Δ be the MCS extending Δ^-. By construction we have $\neg\varphi \in \Delta$ and $(K^P_a)^-\Gamma \subseteq \Delta$ as required. As a was arbitrary, this holds for all agents. □

While the Existence Lemma 5.6 is phrased in syntactic terms only, we should note that the condition on Δ is equivalent to the definition of $R^c_0(a, P)$ in definition 5.5 of the canonical initial model. The lemma therefore immediately guarantees the existence of a $R^c_0(a, P)$-related world $\Delta \in W^c$ in the canonical initial model.

[1] We do not have uniform substitution, but that is not necessary for this result.

Next we show a truth lemma relating the logic $\mathbb{G}\mathbf{P1}$ to the canonical initial model I^c. Recall from definition 4.10 that $\models_i$ is the semantic relation on the initial model.

Lemma 5.7 (Truth Lemma)**.** *Let I^c be the canonical initial model. For every $\mathbb{G}\mathbf{P1}$-MCS Γ and every $\varphi \in \mathcal{L}^{P1}_{[-]}$, we have*

$$I^c, \Gamma \models_i \varphi \text{ if and only if } \varphi \in \Gamma.$$

Proof. By contraposition it suffices to show $I^c, \Gamma \not\models_i \varphi$ iff $\varphi \notin \Gamma$. We use induction on the structure of φ.

Base Case. Suppose $\varphi = S_a b$. The statement follows immediately from the definition of $V^c_0(a, \Gamma)$ and the fact that Γ is an MCS.

Induction Hypothesis. Let $\varphi \in \mathcal{L}^{P1}_{[-]}$ and suppose $I^c, \Gamma \not\models \psi$ iff $\psi \notin \Gamma$ is true for all strict subformulas ψ.

Induction Step.

- The boolean cases follow as usual.
- Suppose $\varphi = K^P_a \psi$. Forwards we use the following implications.

$$\begin{aligned}
& I^c, \Gamma \not\models_i K^P_a \psi \\
\implies & \exists \Delta \text{ s.t. } \Gamma R_a \Delta \text{ and } I^c, \Delta \models_i \neg\psi && \text{(Sem. } K^P_a\text{)} \\
\implies & \exists \Delta \text{ s.t. } (K^P_a)^- \Gamma \subseteq \Delta \text{ and } I^c, \Delta \not\models_i \psi && \text{(Def. } R_a \text{ + Sem. } \neg\text{)} \\
\implies & \exists \Delta \text{ s.t. } (K^P_a)^- \Gamma \subseteq \Delta \text{ and } \psi \notin \Delta && \text{(IH on } \psi\text{)} \\
\implies & \psi \notin (K^P_a)^- \Gamma && \text{(Def of } \subseteq\text{)} \\
\implies & K^P_a \psi \notin \Gamma && \text{(Def of } (K^P_a)^-\text{)}
\end{aligned}$$

Backwards we use the Existence Lemma 5.6 as follows.

$$\begin{aligned}
& K^P_a \psi \notin \Gamma \\
\implies & \neg K^P_a \psi \in \Gamma && \text{(Maximality of } \Gamma\text{)} \\
\implies & \exists \Delta \text{ s.t. } (K^P_a)^- \Gamma \subseteq \Delta \text{ and } \neg\psi \in \Delta && \text{(Existence Lemma 5.6)} \\
\implies & \exists \Delta \text{ s.t. } (K^P_a)^- \Gamma \subseteq \Delta \text{ and } \psi \notin \Delta && \text{(Consistency of } \Delta\text{)} \\
\implies & \exists \Delta \text{ s.t. } (K^P_a)^- \Gamma \subseteq \Delta \text{ and } I^c, \Delta \not\models_i \psi && \text{(IH on } \psi\text{)} \\
\implies & I^c, \Gamma \not\models_i K^P_a \psi && \text{(Sem. } K^P_a\text{)}
\end{aligned}$$

This completes the induction. □

Lemma 5.8. $\mathbb{G}\mathbf{P1}$ *is strongly complete for $\mathcal{L}^P_{[-]}$ on the class of protocol-dependent initial models with a single protocol.*

Proof. Let Γ be a consistent set of formulas. We show the existence of a protocol-dependent initial model I and a state w such that $I, w \models_i \Gamma$.

Let $I^c = \langle W^c_0, R^c_0, V^c_0 \rangle$ be the canonical initial model and let Γ^+ be any MCS extending Γ. By Truth lemma 5.7 we have that $I^c, \Gamma^+ \models_i \Gamma$. We verify that I^c is a protocol-dependent initial model. Recall definition 4.8.

We show that the frame of I^c is transitive. Let $u, v, w \in W^c_0$ and $a \in Ag$ be arbitrary such that $(u, v) \in R^c_0(a, P)$ and $(v, w) \in R^c_0(a, P)$. We show that $(u, w) \in R^c_0(a, P)$. Suppose some $\varphi \in w$. Then by the respective relations, $\widehat{K}^P_a \varphi \in v$ and $\widehat{K}^P_a \widehat{K}^P_a \varphi \in u$. Since u is maximal it contains $\mathbf{4}(P)$ and also $\widehat{K}^P_a \widehat{K}^P_a \varphi \to \widehat{K}^P_a \varphi$ which follows from it. By **MP** we find $\widehat{K}^P_a \in u$, thus $(u, w) \in R^c_0(a, P)$ as required.

We show that the frame is symmetric. Let $u, v \in W^c_0$ and $a \in Ag$ be arbitrary such that $(u, v) \in R^c_0(a, P)$. We show that $(v, u) \in R^c_0(a, P)$. Suppose $\varphi \in u$. Since $u \in W^c_0$, it is an MCS, so it contains the $\mathbf{B}(P)$ axiom

$\varphi \to K_a^P \widehat{K}_a^P \varphi$ and by **MP** then $K_a^P \widehat{K}_a^P \in u$. As $(u,v) \in R_0^c(a,P)$ we find $\widehat{K}_a^P \varphi \in v$ and thus we have $(v,u) \in R_0^c$ as required.

We show that agents know their own secret. Let $w \in W_0^c$ and $a \in Ag$ be arbitrary. Observe that w is an MCS. Then by **Own** we have $S_a a \in w$. By definition of V_0^c then $S_a a \in V_0^c(w)$ as required.

We finish by showing agents are aware of what secrets they know. Let $u, v \in W_0^c$ and $a \in Ag$ be arbitrary such that $(u,v) \in R_0^c(a,P)$. We show that we have $V_0^c(a,u) = V_0^c(a,v)$. Let $S_a b$ be a secret atom for arbitrary $b \in Ag$. We suppose $S_a b \in V_0^c(a,v)$ to get the following chain of implications.

$$
\begin{aligned}
& S_a b \in V_0^c(a,v) && \\
\implies & S_a b \in v && \text{(Def. } V_0^c\text{)} \\
\implies & I^c, v \models_i S_a b && \text{(Truth Lemma 5.7)} \\
\implies & I^c, u \models_i \widehat{K}_a^P S_a b && \text{(Sem. } \widehat{K}_a^P \text{ and } u R_0^c v\text{)} \\
\implies & \widehat{K}_a^P S_a b \in u && \text{(Truth Lemma 5.7)} \\
\implies & S_a b \in u && (\mathbf{NPi}(P) + \mathbf{MP}) \\
\implies & S_a b \in V_0^c(a,u) && \text{(Def. } V_0^c\text{)}
\end{aligned}
$$

In the last step we use the dual of $\mathbf{NPi}(P)$ which is $\widehat{K}_a^P S_a b \to S_a b$, resulting in $S_a b \in u$ by **MP**.

Since I^c is symmetric, we also have $(v,u) \in R_0^c(a,P)$ so we can re-apply the proof starting with $S_a b \in V_0^c(a,u)$ and using $(v,u) \in R_0^c(a,P)$ to find that $S_a b \in V_0^c(a,u) \implies S_a b \in V_0^c(a,v)$ too. Hence we conclude that $S_a b \in u \iff S_a b \in v$ and so $V_0^c(a,u) = V_0^c(a,v)$ as required. □

Lemma 5.9. $\mathbb{G}\mathbf{P1}$ *is sound and strongly complete for* $\mathcal{L}_{[-]}^{P1}$ *on the class of protocol-dependent initial models with a single protocol.*

Proof. Immediate by soundness from lemma 5.1 and completeness from lemma 5.8. □

The above completeness result is in relation to initial models. Luckily, we can lift it to a synchronous gossip model while preserving its canonicity in the root states of the induced model[2]. We can therefore use the induced gossip model $M^c := M_{\mathsf{PD}}^{\bullet}(I^c)$ as lifted canonical gossip model.

Lemma 5.10. *Let* I^c *be the canonical initial model and* $M^c = M_{\mathsf{PD}}^{\bullet}(I^c)$ *be its lifted gossip model. For any formula* $\varphi \in \mathcal{L}_{[-]}^{P1}$ *and any MCS* Γ *we have*

$$M^c, (\Gamma, \epsilon) \models \varphi \text{ if and only if } I^c, \Gamma \models_i \varphi.$$

Proof. As M^c is synchronous, by definition 4.11 a state (Γ, σ) can only be $\sim_a^P$-related to another state with a call sequence of the same length as σ. Thus root states can only be $\sim_a^P$-related to other root states. By definition this holds if and only if they were $R_0(a)$-related in I^c. Furthermore $V_a(\Gamma, \epsilon) = V_0(a, \Gamma)$ by definition. It then follows from a straightforward induction on $\varphi \in \mathcal{L}_{[-]}^{P1}$ that $M^c, (\Gamma, \epsilon) \models \varphi$ if and only if $I^c, \Gamma \models_i \varphi$. □

We can therefore define a truth lemma with respect to the root states of the canonical gossip model.

Lemma 5.11 (Truth Lemma)**.** *For every MCS* Γ *and every* $\varphi \in \mathcal{L}_{[-]}^{P1}$*, we have* $M^c, (\Gamma, \epsilon) \models \varphi$ *if and only if* $\varphi \in \Gamma$.

Proof. Immediate from Truth Lemma 5.7 and lemma 5.10. □

That M^c is a protocol-dependent gossip model is furthermore immediate from its definition: any model induced from a protocol-dependent initial model is a protocol-dependent gossip model.

[2]This only applies to synchronous models and fails with asynchronous gossip models.

Theorem 5.12. $\mathbb{G}\mathbf{P1}$ *is sound and strongly complete for* $\mathcal{L}^{P1}_{[-]}$ *on the class of protocol-dependent gossip models with a single protocol.*

Proof. Completeness is immediate by definition 4.12 of protocol-dependent gossip models and lemma 5.8. Soundness was shown in lemma 5.1. □

We have now obtained completeness for a class of gossip models, but we only used root states of the lifted canonical gossip model. Truth Lemma 5.11 guarantees that any set of formulas satisfied at any state (Δ, τ) somewhere in M^c is in fact an MCS, and therefore is already included in the canonical initial model and consequently in a root state of the lifted gossip model. We only get this property after proving the truth lemma, but using the root states is a fundamental part of the completeness proof.

5.2 Adding Synchronous Call Reduction Axioms

We now extend the proof system to include calls and call effects. We limit ourselves to only axioms that are sound for the synchronous case. Similar to chapter 3, we will ensure that these are call reduction axioms. Together with the completeness for the call-free language from Theorem 5.12, we can then immediately retrieve the completeness result for the full language.

The proof system $\mathbb{G}^{\bullet}$ for basic gossip forms a strong basis. In fact, the axioms describing the call effects in Table 3.2 remain sound. They do not use any knowledge modality, so they are expressible in the language and the call semantics have not changed.

The observance axioms do require change, as they relate to knowledge. The protocol-dependent observance axioms share many characteristics of the basic versions and can be found in Table 5.2.

Table 5.2. Observance axioms for synchronous protocol-dependent gossip.

Synchronous Observance for Protocol-Dependent Gossip		
$\mathbf{Obs}^{\bullet}_1(P)$	$[ab]K^P_a\varphi \leftrightarrow (P_{ab} \to \bigvee_{R\subseteq\mathbb{S}}(O_bR \wedge K^P_a(P_{ab} \to (O_bR \to [ab]\varphi))))$	$a \in \{a,b\}$
$\mathbf{Obs}^{\bullet}_2(P)$	$[ab]K^P_b\varphi \leftrightarrow (P_{ab} \to \bigvee_{R\subseteq\mathbb{S}}(O_aR \wedge K^P_b(P_{ab} \to (O_aR \to [ab]\varphi))))$	$b \in \{a,b\}$
$\mathbf{Pri}^{\bullet}(P)$	$[ab]K^P_c\varphi \leftrightarrow (P_{ab} \to \bigwedge_{d,e\neq a} K^P_c(P_{de} \to [de]\varphi))$	$c \notin \{a,b\}$

Remark 5.13. *Like with the non-protocol dependent observance axioms, the axioms* $\mathbf{Obs}^{\bullet}_1(P)$ *and* $\mathbf{Obs}^{\bullet}_2(P)$ *are completely symmetric. We could alternatively define both cases* $c \in \{a,b\}$ *at once with the following formula, which is sound but introduces redundant disjuncts.*

$$[ab]K^P_c\varphi \leftrightarrow (P_{ab} \to \bigvee_{Q,R\subseteq\mathbb{S}} (O_aQ \wedge O_bR \wedge K^P_c(P_{ab} \to (O_aQ \wedge O_bR \to [ab]\varphi))))$$

The protocol-dependent observance axioms are similar in structure to their original counterparts in Table 3.3. The main difference is the addition of the protocol conditions. These are introduced both for the hypothetical situations that agent c considers possible, as well as for the actual call to be permitted.

The difference between these two uses of the protocol conditions is clearest in the proof for $\mathbf{Pri}^{\bullet}(P)$, where the two usages are with respect to different variables ab and de.

Lemma 5.14. *The axioms* $\mathbf{Obs}^{\bullet}_1(P)$, $\mathbf{Obs}^{\bullet}_2(P)$, *and* $\mathbf{Pri}^{\bullet}(P)$ *are sound for the class of synchronous protocol-dependent gossip models with a single protocol.*

Proof. We omit the proof for $\mathbf{Obs}^{\bullet}_2(P)$ as it is analogous to $\mathbf{Obs}^{\bullet}_1(P)$. Let (w,σ) be an arbitrary gossip state in some gossip model M and let $\varphi \in \mathcal{L}^P$ be arbitrary.

For $\mathbf{Obs}_1^{\bullet}(P)$ we have the following chains of equivalences. Recall that $V_b(w,\sigma)$ is not a formula but the set of secrets that agent b knows at gossip 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_bR holds.

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

For $\mathbf{Pri}^{\bullet}(P)$ we have the following chain of equivalences.

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

The three axioms in Table 5.2 are call-reduction axioms for the operator K_a^P, and the other axioms in $\mathbb{G}\mathbf{P1}^{\bullet}$ still cover the other cases for atoms, implication, and negation like they did in the basic case. We can therefore again prove that each formula in $\mathcal{L}^{P1}$ is equivalent to a call reduction in $\mathcal{L}^{P1}_{[-]}$.

We should first reassure ourselves that we can again use the converse of $\mathbf{K}([ab])$.

Lemma 5.15. $\mathbb{G}\mathbf{P1}^{\bullet} \vdash [ab](\varphi \to \psi) \leftrightarrow ([ab]\varphi \to [ab]\psi)$

Proof. Immediate by fact 3.3 and the fact that $\mathbf{K}([ab])$ and **Fnc** are axioms in $\mathbb{G}\mathbf{P1}^{\bullet}$. □

The proof is then similar to lemma 3.4. The only difference is in the use of protocol-dependent axioms.

Lemma 5.16. *For every formula $\varphi \in \mathcal{L}^{P1}$ there is a formula $\psi \in \mathcal{L}^{P1}_{[-]}$ such that $\mathbb{G}\mathbf{P1}^{\bullet} \vdash \varphi \leftrightarrow \psi$.*

Proof. Without loss of generality we assume that φ only contains the boolean connectives $\neg$ and $\to$ as these are truth-functionally complete. Using **Sub** it furthermore suffices to consider formulas of the form $\varphi = [ab]\chi$ with $\chi \in \mathcal{L}^{P1}_{[-]}$. We use induction on the structure of χ to show $\vdash [ab]\chi \leftrightarrow \psi$ for some $\psi \in \mathcal{L}^{P1}_{[-]}$.

Base case. Suppose $\chi = S_cd$. If $c \notin \{a,b\}$, then **Ext** yields $\vdash [ab]S_cd \leftrightarrow S_cd$. Else we have $c \in \{a,b\}$ and **Eff** yields $\vdash [ab]S_cd \leftrightarrow (S_ad \vee S_bd)$. Both $S_cd \in \mathcal{L}^{P1}_{[-]}$ and $(S_ad \vee S_bd) \in \mathcal{L}^{P1}_{[-]}$, which concludes the base case.

Induction Hypothesis. Let $\chi \in \mathcal{L}^{P1}_{[-]}$ be a formula. Suppose that for every strict subformula χ' of χ, we have some $\psi' \in \mathcal{L}^{P1}_{[-]}$ such that $\vdash [ab]\chi' \leftrightarrow \psi'$.

Induction Step.

- Suppose $\chi = \neg\chi'$. From **Fnc** we obtain $\vdash [ab]\neg\chi' \leftrightarrow \neg[ab]\chi'$. By induction hypothesis on χ' we have $\vdash [ab]\chi' \leftrightarrow \psi'$. With **Sub** we find that $\vdash [ab]\neg\chi' \leftrightarrow \neg\psi'$, the right hand of which is in $\mathcal{L}^{P1}_{[-]}$.
- Suppose $\chi = (\chi' \to \chi'')$. By fact 3.3 we have $\vdash [ab](\chi' \to \chi'') \leftrightarrow ([ab]\chi' \to [ab]\chi'')$. By induction hypothesis on χ' and χ'' we get $\vdash [ab]\chi' \leftrightarrow \psi'$ and $\vdash [ab]\chi'' \leftrightarrow \psi''$. Using **Sub** this yields $\vdash [ab](\chi' \to \chi'') \leftrightarrow (\psi' \to \psi'')$, the right hand of which is in $\mathcal{L}^{P1}_{[-]}$.
- Suppose $\chi = K^P_c\chi'$. We consider three cases for c.

  For $c = a$ we use $\mathbf{Obs}^{\bullet}_1(P)$ to get $\vdash [ab]K^P_a\chi' \leftrightarrow (P_{ab} \to \bigvee_{R\subseteq S}(O_bR \wedge K^P_a(P_{ab} \to (O_bR \to [ab]\chi'))))$. By induction hypothesis on χ' we have $\vdash [ab]\chi' \leftrightarrow \psi$ and using **Sub** we find that $\vdash [ab]K^P_a\chi' \leftrightarrow (P_{ab} \to \bigvee_{R\subseteq S}(O_bR \wedge K^P_a(P_{ab} \to (O_bR \to \psi))))$. Recall that P_{ab} is a subformula of χ and χ is call-free. Thus the right hand of the equivalence is a formula in $\mathcal{L}^{P1}_{[-]}$.

  For $c = b$ we instead use $\mathbf{Obs}^{\bullet}_2(P)$ and proceed analogously.

  For $c \notin \{a, b\}$ we use $\mathbf{Pri}^{\bullet}(P)$ to get $\vdash [ab]K^P_c\chi' \leftrightarrow (P_{ab} \to \bigwedge_{d,e\neq a} K^P_c(P_{de} \to [de]\chi'))$. By induction hypothesis on χ' we have $\vdash [ab]\chi' \leftrightarrow \psi'$. Using **Sub** we find that $\vdash [ab]K^P_c\chi \leftrightarrow (P_{ab} \to \bigwedge_{d,e\neq a} K^P_c(P_{de} \to \psi'))$. Recall that P_{ab} and P_{de} for all $d, e \neq a$ are subformulas of χ and χ is call-free. Thus the right hand of the equivalence is a formula in $\mathcal{L}^{P1}_{[-]}$.

This finishes the induction on χ. We conclude that for every $[ab]\chi \in \mathcal{L}^{P1}$ we have an equivalent $\psi \in \mathcal{L}^{P1}_{[-]}$. As every formula $\varphi \in \mathcal{L}^{P}$ can be written in this form, we are done. □

We can now use the equivalence between formulas and their call reductions to immediately retrieve completeness for the full language $\mathcal{L}^{P1}$ on the synchronous protocol-dependent gossip models with a single protocol.

Theorem 5.17. *$\mathbb{G}\mathbf{P1}^{\bullet}$ is sound and strongly complete for $\mathcal{L}^{P1}$ on the class of synchronous protocol-dependent gossip models with a single protocol.*

Proof. Completeness follows from lemma 5.16 and Theorem 5.12. Soundness follows from lemmas 5.1 and 5.14. □

5.3 Axiomatisation for the Synchronous Tree Model

So far we have obtained soundness and completeness on the class of arbitrary models. We now extend the proof system $\mathbb{G}\mathbf{P1}^{\bullet}$ with a tree rule. We will show that the resulting proof system is sound and complete for $\mathcal{L}^{P1}$ on the synchronous tree model with a single protocol.

In order to do so, we will approximate the tree model. We define a notion of n-bisimilarity for protocol-dependent gossip models that preserves truth and additionally provides a bound on the length of call sequences that can satisfy a formula. Together with a formula to approximate the initial root world in the protocol-dependent setting, we arrive at a sound tree rule.

Definition 5.18 (n-bisimulation for Protocol-Dependent Gossip Models)**.** *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 P an instance of the following two conditions.

2. ***(Forth)** For every $t \in W$ we have: if $s \sim^P_a t$ then there is a $t' \in W'$ such that $s' \sim'^P_a t'$ 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 $M, t \leftrightarrows_{n-1} M', t'$.*

The n-bisimulation is an equivalence relation between states in a model. We write $f(n)$ to denote the number of n-bisimilarity classes.

We define the degree of a formula as follows. As discussed in chapter 2, we treat protocol conditions as subformulas of the K_a^P modality and use the maximum degree of all strict subformulas similar to conjunction.

Definition 5.19 (Degree)**.** *Given a formula $\varphi \in \mathcal{L}^{P1}$, its* degree *$d(\varphi)$ is defined recursively by*

$$\begin{aligned} d(S_a b) &= 0 \\ d(\neg\varphi) &= d(\varphi) \\ d(\varphi \wedge \psi) &= max\{d(\varphi), d(\psi)\} \\ d(K_a^P \varphi) &= 1 + max\{d(\chi) \mid \chi \in sub(K_a^P \varphi) \setminus \{K_a^P \varphi\}\} \\ d([ab]\varphi) &= d(\varphi) \end{aligned}$$

where $sub(K_a^P \varphi) \setminus \{K_a^P \varphi\} = \{d(\varphi)\} \cup \{d(P_{ab}) \mid a \neq b \in Ag\}$ as in definition 4.6.

Accounting for the degree of protocol conditions is not necessary for showing modal equivalence: simply increasing the degree (by 1) for each K_a^P modality is sufficient. It is however important for the next result, which limits the degree of the equivalent call-free formula of any $\varphi \in \mathcal{L}^{P1}$.

Lemma 5.20. *For any formula $\varphi \in \mathcal{L}^{P1}$, its call-reduction $\psi \in \mathcal{L}^{P1}_{[-]}$ has degree $d(\psi) \leq d(\varphi)$.*

Proof. We use induction on the structure of φ in the same way as the proof of lemma 5.16. We do not repeat all steps, but observe that for each constructed call-free formula ψ in that proof, the statement holds and can be verified by computing its degree. We explicitly show the induction step for K_c^P with $c = a$.

Induction Hypothesis. Let $\chi \in \mathcal{L}^{P1}_{[-]}$ be a formula. For each strict subformula χ' of χ with call reduction ψ' we have $d(\psi') \leq d(\chi')$.

Induction Step. Suppose $\chi = K_c^P \chi'$. Thus $\varphi = [ab]K_c^P \chi'$ for $c \notin \{a, b\}$. Let $n = d(\varphi)$.

By lemma 5.16 we find that the call reduction of φ is $\psi = (P_{ab} \to \bigvee_{R \subseteq S}(O_b R \wedge K_a^P(P_{ab} \to (O_b R \to \psi'))))$, where ψ' is the call reduction of χ'. By definition, each of the strict subformulas of χ has a degree at most $n-1$. By IH on χ' we have thus $d(\psi') \leq d(\chi') < d(\chi) = n$ and as P_{ab} is a strict subformula of χ we have $d(P_{ab}) < d(\chi) = n$. By computing the degree of ψ we conclude that $d(\psi) \leq n$ as required. □

The following lemma is a well-known result in modal logic [BRV01]. While the lemma is usually an equivalence, we will only need the forward direction, which states that truth is preserved under n-bisimulation.

Lemma 5.21. *Let M, s and M', s' be pointed gossip models and $n \in \mathbb{N}$. If $M, s \leftrightarrows_n M', s'$ then for every formula $\varphi \in \mathcal{L}^{P1}_{[-]}$ with $d(\varphi) \leq n$ we have $M, s \models \varphi \iff M', s' \models \varphi$.*

Proof. Let M, s and M', s' be pointed gossip models. Suppose $M, s \leftrightarrows_n M', s'$ for some $n \in \mathbb{N}$. Let $\varphi \in \mathcal{L}^{P1}_{[-]}$ be arbitrary such that $d(\varphi) \leq n$. We use strong induction on n to show that $M, s \models \varphi \iff M', s' \models \varphi$.

Induction Hypothesis. Let n be arbitrary. Suppose that for all $m < n$ we have that if $M, s \leftrightarrows_m M', s'$ then for all ψ such that $d(\psi) \leq m$ we have $M, s \models \psi \iff M', s' \models \psi$.

Induction Step. Let $n = d(\varphi)$ be arbitrary and suppose the induction hypothesis holds for all $m < n$.

- The atomic and boolean cases follow immediately from **Atoms**.
- $\varphi = K_a^P \psi$.

  ($\Longrightarrow$) Suppose $M, s \models \varphi$. For contradiction, suppose that $M', s' \not\models \varphi$. Then there exists a world v' such that $s' \sim_a^P t'$ and $M', t' \not\models \psi$. By **Back** there exists a t in M such that $s \sim_a^P t$ and $M, t \leftrightarrows_{n-1} M', t'$. By

definition of $d(K_a^P\psi)$ we have $d(\psi) \leq n-1$, so by IH we have $M,t \not\models \psi$. However, this implies $M,s \not\models \varphi$ which contradicts the initial supposition. We conclude instead that $M',s' \models \varphi$.

($\Longleftarrow$) We suppose $M',s' \models \varphi$ and use **Forth** in the same manner. □

Lemma 5.21 only applies to call-free formulas. We could show that bisimulation is preserved under calls in order to extend this result to the full language [DHK20], but we will instead use the call reductions from lemma 5.16. With this approach we should be careful to use the degree of the call-reduction $\psi \in \mathcal{L}^{P1}_{[-]}$, rather than of the formula $\varphi \in \mathcal{L}^{P1}$ itself. However, it suffices to use $d(\varphi)$ because we have $d(\psi) \leq d(\varphi)$ by lemma 5.20.

Lemma 5.22. *Let M,s and M',s' be pointed gossip models and $n \in \mathbb{N}$ some natural number. For any $\varphi \in \mathcal{L}^{P1}$ with $d(\varphi) \leq n$ we have $M,s \leftrightarrows_n M',s'$ implies $M,s \models \varphi \iff M',s' \models \varphi$.*

Proof. Let $\varphi \in \mathcal{L}^{P1}$ be arbitrary and let $n = d(\varphi)$. By lemma 5.16 we have an equivalent formula $\psi \in \mathcal{L}^{P1}_{[-]}$ such that $\vdash \varphi \leftrightarrow \psi$ and moreover $d(\psi) \leq d(\varphi)$ by lemma 5.20. Let two models and states M,s and M',s' be arbitrary such that they are n-bisimilar. By lemma 5.21 we find that $M,s \models \psi$ if and only if $M',s' \models \psi$. However, $\varphi \leftrightarrow \psi$ is true in both states by soundness of $\mathbb{G}\mathbf{P1}^{\bullet}$. Therefore we conclude that $M,s \models \varphi$ if and only if $M',s' \models \varphi$. □

We have shown that n-bisimulation preserves truth for $\mathcal{L}^{P1}$. We can use this to bound the length of call sequences that satisfy a formula in a gossip model to the number of n-bisimilarity classes that its states have.

Lemma 5.23. *Let $\varphi \in \mathcal{L}^{P1}$ 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 $f(n)$ the number of n-bisimilarity classes.*

Proof. Let $(w,\sigma.\tau)$ and φ be arbitrary and suppose $(w,\sigma.\tau) \models \varphi$. Furthermore let $|\tau| > f(n)$, as otherwise the statement is already satisfied. As τ contains more calls than there are n-bisimilarity classes, there must be two different initial fragments τ_1 and τ_2 of τ such that $(w,\tau_1) \leftrightarrows_n (w,\tau_2)$. W.l.o.g. let τ_1 be the shortest of the two.

Let τ_3 be the remainder after τ_2, such that $\tau = \tau_2.\tau_3$. Then $(w,\sigma.\tau_2) \models [\tau_3]\varphi$ by assumption. Moreover we have $d([\tau_3]\varphi) = d(\varphi) = n$ by definition. Since $(w,\sigma.\tau_1) \leftrightarrows_n (w,\sigma.\tau_2)$, we find by lemma 5.22 that $(w,\sigma.\tau_1) \models [\tau_3]\varphi$. Hence we conclude that $(w,\sigma.\tau_1.\tau_3) \models \varphi$.

Observe that $|\sigma.\tau_1.\tau_3| < |\sigma.\tau_2.\tau_3| = |\sigma.\tau|$ because $\tau_1 \neq \tau_2$. If $|\sigma.\tau_1.\tau_3| \leq f(n)$, we are done. Else, we repeat the argument until we have found a sequence that is at most length $f(n)$. □

With the above bound, we can ensure that any truth in a model will be true within at most $f(n)$ many calls. The value of $f(n)$ is furthermore finite as there are only finitely many atoms [DHK20].

Next we will characterise the tree model. In the root of the synchronous tree model it is common knowledge that every agent only knows their own secret. We cannot express common knowledge in $\mathcal{L}^{P1}_{[-]}$, nor in any of the other gossip languages, but we can express an approximation thereof.

While in basic gossip models reflexivity was guaranteed by the model, protocol-dependent gossip models do not make this guarantee. An initial model with one world that has the correct valuation but an empty relation will induce a gossip model that vacuously satisfies any K_a^P and therefore also root_n for any n. To avoid this we must stipulate that each agent has a serial relation, i.e. there exists a relation at each step. We do so by including an existential modality $\widehat{K}_a^P\mathsf{root}_n^P$ to the recursive step.

Definition 5.24. *Recall that φ_ϵ was defined as the formula that is only true when every agent only knows their own secret, which is the valuation of the root of the tree model. We define the following formula* root *recursively to approximate common knowledge of this distribution.*

$$\mathsf{root}_0^P := \varphi_\epsilon$$
$$\mathsf{root}_{i+1}^P := \mathsf{root}_i^P \wedge \bigwedge_{a\in Ag} (K_a^P\mathsf{root}_i^P \wedge \widehat{K}_a^P\mathsf{root}_i^P)$$

The following lemma shows that satisfying root_n^P guarantees n-bisimilarity to the root of the tree model.

Lemma 5.25. *Let $n \in \mathbb{N}$. For any model M and gossip state (w, σ) we have that*

$$M, (w, \sigma) \models \mathsf{root}_n^P \text{ implies } M, (w, \sigma) \leftrightarrows_n M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \epsilon).$$

Proof. Let M and (w, σ) be an arbitrary model and state. Suppose that $M, (w, \sigma) \models \mathsf{root}_n^P$. We show that $M, (w, \sigma) \leftrightarrows_n M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \epsilon)$ by induction on n.

Base case. Let $n = 0$. We have $M, (w, \sigma) \models \mathsf{root}_0^P$ if and only if (w, σ) agrees on all atoms by definition of $\mathsf{root}_0^P = \varphi_\epsilon$. Hence $M, (w, \sigma) \leftrightarrows_0 M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \epsilon)$.

Induction Hypothesis. For arbitrary n we have $M, (w, \sigma) \models \mathsf{root}_n^P$ implies $M, (w, \sigma) \leftrightarrows_n M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \epsilon)$.

Induction Step. Suppose $M, (w, \sigma) \models \mathsf{root}_{n+1}^P$ for arbitrary an model and state. In particular then $M, (w, \sigma) \models \mathsf{root}_n^P$, so by IH $M, (w, \sigma) \leftrightarrows_n M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \epsilon)$. We show that $M, (w, \sigma) \leftrightarrows_{n+1} M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \epsilon)$ too. **Atoms** is immediate from the bisimulation.

Forth. Suppose there is some state v such that $(w, \sigma) \sim_a^P v$. Because $M, (w, \sigma) \models \mathsf{root}_{n+1}^P$ we have $M, v \models \mathsf{root}_n^P$ and in particular $M, v \models K_a^P \mathsf{root}_n^P$. By IH thus $M, v \leftrightarrows_n (w_{\mathsf{root}}, \epsilon)$. As $(w_{\mathsf{root}}, \epsilon) \sim_a^P (w_{\mathsf{root}}, \epsilon)$ by definition of $M^{\bullet}_{\mathsf{PDtree}}$, we are done.

Back. Suppose there is some state v' such that $(w_{\mathsf{root}}, \epsilon) \sim_a^P v'$. Then by definition of $M^{\bullet}_{\mathsf{PDtree}}$ we must have $v' = (w_{\mathsf{root}}, \epsilon)$. Because $(w, \sigma) \models \mathsf{root}_{n+1}$ and in particular it satisfies $\widehat{K}_a^P \mathsf{root}_n^P$, there exists a state v such that $(w, \sigma) \sim_a^P v$. As moreover $v \models \mathsf{root}_n^P$, we have by IH that $M, v \leftrightarrows_n M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \epsilon)$ and we are done. □

Satisfying root_n^P therefore provides us with n-bisimilarity to the root of the tree model. We can relate this to truths in all other states of the tree model. The following fact is immediate from the semantics of the call modality and states that for any formula φ that is satisfied at some call sequence σ in the tree model, the formula $[\sigma]\varphi$ is satisfied at its root.

Fact 5.26. *For all formulas $\varphi \in \mathcal{L}_{[-]}^{P1}$ we have $M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \sigma) \models \varphi$ if and only if $M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \epsilon) \models [\sigma]\varphi$.*

We can now use the preservation of truth under n-bisimulation to show that root_n^P is a sufficiently close approximation of the tree model.

Lemma 5.27. *Let $\varphi \in \mathcal{L}_{[-]}^{P1}$ such that $d(\varphi) \leq n$. For any model M and gossip state (w, τ) we have that*

$$M, (w, \tau) \models \mathsf{root}_n^P \text{ implies } M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \epsilon) \models [\sigma]\varphi \iff M, (w, \tau) \models [\sigma]\varphi.$$

Proof. Suppose $M, (w, \tau) \models \mathsf{root}_n^P$. By lemma 5.25 this implies that $M, (w, \tau) \leftrightarrows_n M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \epsilon)$. Modal equivalence up to degree n then follows from lemma 5.22. □

Combining fact 5.26 and lemma 5.27 we get the following corollary, which is the basis of the final rule of the proof system.

Corollary 5.28. *Let $\varphi \in \mathcal{L}_{[-]}^{P1}$ be a formula with degree at most n. Then $M^{\bullet}_{\mathsf{PDtree}}, (w_{\mathsf{root}}, \sigma) \models \varphi$ if and only if for all models M and sequences τ we have, we have $M, (w, \tau) \models \mathsf{root}_n^P \to [\sigma]\varphi$.*

Finally, we can apply the bound shown in lemma 5.23 to limit the number of call sequences we need to consider. We obtain the rule $\mathbf{Tree}^{\bullet}(P)$ in Table 5.3 and add this rule to the proof system $\mathbb{G}\mathbf{P1}^{\bullet}$ to obtain $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}}$.

Lemma 5.29. *φ is valid on the tree model $M^{\bullet}_{\mathsf{PDtree}}(P)$ if and only if it is provable in $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}}$. Formally, we have $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}} \vdash \varphi$ if and only if for all call sequences σ we have $M^{\bullet}_{\mathsf{PDtree}}(P), (w_{\mathsf{root}}, \sigma) \models \varphi$.*

Proof. Let $\varphi \in \mathcal{L}_{[-]}^{P1}$ be arbitrary.

($\Longrightarrow$) Suppose for all sequences σ we have $M^{\bullet}_{\mathsf{PDtree}}(P), (w_{\mathsf{root}}, \sigma) \models \varphi$. By corollary 5.28 this holds if and only if we have for all sequences σ that for all models M and all sequences τ that $M, (w, \tau) \models \mathsf{root}_n^P \to [\sigma]\varphi$, hence

Table 5.3. Tree rule for synchronous protocol-dependent gossip, where n is the degree $d(\varphi)$ and $f(n)$ is the number of n-bisimilarity classes.

Tree Rule for synchronous single-protocol gossip	
$\mathbf{Tree}^{\bullet}(P)$	If $\vdash \mathsf{root}_n^P \to [\sigma]\varphi$ for all σ s.t. $\lvert\sigma\rvert \leq f(n)$ then $\vdash \varphi$

$\mathsf{root}_n^P \to [\sigma]\varphi$ is a validity on the class of synchronous protocol-dependent gossip models. Following Theorem 5.17 we find by completeness of $\mathbb{G}\mathbf{P1}^{\bullet}$ that $\mathbb{G}\mathbf{P1}^{\bullet} \vdash \mathsf{root}_n^P \to [\sigma]\varphi$. As $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}}$ is stronger than $\mathbb{G}\mathbf{P1}^{\bullet}$, we then too obtain $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}} \vdash \mathsf{root}_n^P \to [\sigma]\varphi$. Finally by $\mathbf{Tree}^{\bullet}(P)$ we obtain $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}} \vdash \varphi$ as required.

($\Longleftarrow$) Suppose $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}} \vdash \varphi$. By soundness of all rules and axioms in $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}}$ we find that $M^{\bullet}_{\mathsf{PDtree}}, \sigma \models \varphi$ for all call sequences σ. □

Theorem 5.30. $\mathbf{Tree}^{\bullet}(P)$ *is sound and complete for* $M^{\bullet}_{\mathsf{PDtree}}(P)$.

Chapter 6

Towards Completeness for Multi-Protocol-Dependent Gossip

In chapter 5 we have introduced the proof systems $\mathbb{G}\mathbf{P1}^{\bullet}$ and $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}}$ and shown that these are complete on arbitrary models and the tree model in a single protocol setting. We now discuss how to extend this result to multiple protocols, working towards an axiomatisation for gossip settings that allow protocol-dependent knowledge modalities for all protocols. The main objective is to account for protocol interactions. When there are multiple protocols in a gossip model, their epistemic relations may interact with each other: knowledge based on protocol P may imply knowledge based on protocol Q.

We discuss how one can adapt the proof system $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}}$ to a system $\mathbb{GP}^{\bullet}_{\mathsf{tree}}$ that is sound for any number of protocols. We conjecture $\mathbb{GP}^{\bullet}_{\mathsf{tree}}$ is complete for the tree model in the multi-protocol setting.

6.1 Soundness of Existing Axioms

Recall from definition 4.8 that a protocol-dependent initial model has a separate relation for each protocol P. The soundness proofs of lemmas 5.1 and 5.2 for $\mathbb{G}\mathbf{P1}$ can be generalised to arbitrary protocols: each relation is still transitive and symmetric, and each satisfies the secret introspection properties. Similarly, the proof of lemma 5.14 for the the protocol-dependent observance axioms only depends on the protocol's own epistemic relation.

Let $\mathbb{GP}$ and $\mathbb{GP}^{\bullet}$ be the proof systems $\mathbb{G}\mathbf{P1}$ and $\mathbb{G}\mathbf{P1}^{\bullet}$ with each P-dependent axiom repeated for all $P \in \mathbb{P}$. The following corollaries now follow naturally from the earlier proofs.

Corollary 6.1. *The proof system $\mathbb{GP}$ is sound for initial models with any set of protocols $\mathcal{P} \subseteq \mathbb{P}$.*

Corollary 6.2. *The proof system $\mathbb{GP}$ is sound for gossip models with any set of protocols $\mathcal{P} \subseteq \mathbb{P}$.*

Corollary 6.3. *The proof system $\mathbb{GP}^{\bullet}$ is sound for synchronous gossip models with any set of protocols $\mathcal{P} \subseteq \mathbb{P}$.*

6.2 A General Tree Rule

The last part of the axiomatisation is the tree rule. The tree rule as defined in Table 5.3 for a single protocol is no longer sound in a multi-protocol setting, but we can make it sound by adapting the formula we use to approximate the root world. Recall that root_n intends to approximate the initial root world and in particular the common knowledge of this world up to degree n.

In chapter 5 we used a single protocol, but we can define the $\mathsf{root}^{\mathcal{P}}_n$ for any finite set of protocols $\mathcal{P}$ by including a conjunction over all knowledge modalities. This way we ensure that the modalities for all protocols in $\mathcal{P}$ approximate the common knowledge.

Definition 6.4. *Let $\mathcal{P}$ be a finite set of protocols. We define $\mathsf{root}_n^{\mathcal{P}}$ as follows.*

$$\mathsf{root}_0^{\mathcal{P}} := \varphi_\epsilon$$
$$\mathsf{root}_{i+1}^{\mathcal{P}} := \mathsf{root}_i^{\mathcal{P}} \wedge \bigwedge_{P\in\mathcal{P}, a\in Ag} (K_a^P \mathsf{root}_i^{\mathcal{P}} \wedge \widehat{K}_a^P \mathsf{root}_i^{\mathcal{P}})$$

Let $\mathbf{Tree}^{\bullet}(\mathcal{P})$ *then be the following rule, where $n = d(\varphi)$ and $f(n)$ is the number of n-bisimilarity classes.*

$$\textit{If } \vdash \mathsf{root}_n^{\mathcal{P}} \to [\sigma]\varphi \textit{ for all } \sigma \textit{ s.t. } |\sigma| \leq f(n) \textit{ then } \vdash \varphi.$$

It is clear that this approach does not work for infinite sets of protocols. We can however solve this by observing that any formula φ only uses a finite number of protocol-dependent knowledge modalities. Let $\mathsf{prot}(\varphi)$ be the set of protocols for which there are modalities in φ. It suffices to approximate the tree model for these modalities. The result is a generalised version of the tree rule in Table 6.1.

Table 6.1. The Synchronous Tree Rule for Multi-Protocol-Dependent Gossip for any number of protocols in $\mathbb{P}$. With $n = d(\varphi)$ and $f(n)$ the number of n-bisimilarity classes.

Tree Rule for All Protocols	
$\mathbf{Tree}^{\bullet}(\mathbb{P})$	If $\vdash \mathsf{root}_n^{\mathsf{prot}(\varphi)} \to [\sigma]\varphi$ for all σ s.t. $\lvert\sigma\rvert \leq f(n)$ then $\vdash \varphi$

While $\mathbf{Tree}^{\bullet}(\mathbb{P})$ is parameterized with $\mathbb{P}$, the rule would work just as well for any other set of protocols, and even in the single protocol setting. The set of protocols $\mathsf{prot}(\varphi)$ determines the relevant protocols in φ.

Conjecture 6.5. *The* $\mathbf{Tree}^{\bullet}(\mathbb{P})$ *rule is sound for the synchronous tree model with any set of protocols $\mathcal{P} \subseteq \mathbb{P}$.*

While this conjecture is positive, there are some considerations about the tree rule in a setting with infinitely many protocols. Recall from definition 4.8 that we have introduced separate epistemic relations for each protocol in the initial models. This means that the modal similarity type of our language is no longer finite.

There are two moments where this may come into play. Firstly we have used preservation of truth under n-bisimulation, recall lemma 5.21. This is a well-known property of n-bisimulation, as can be found in proposition 2.31 in [BRV01]. This source however requires additionally that the language has a finite modal similarity type.

While it appears that the proof of lemma 5.21 does not depend on this, we do run into another problem. We have used the number of n-bisimilarity classes to finitely bound the number of call sequences in the conditions of the tree rule. However, with infinitely many modalities, there may be infinitely many n-bisimilarity classes. This might not be fatal, but does limit the application of such a rule in theorem provers.

We could potentially solve this problem, but by doing so most likely will introduce another: we can obtain a finite modal similarity type if we instead view the modality K_a^P as a single modality K_a with different subformulas for each protocol. However, in doing so we remove the possibility to define separate protocol-dependent relations in the canonical initial model. This is most likely a necessary part in the canonical model construction as presented in chapter 5 when adapted to multiple protocols.

6.3 Protocol Interactions and Completeness

Finally, there is one more caveat. In a setting with more than one protocol, we should be careful about how protocols interact with each other. The following rule is an example of such interaction. Whenever the protocol conditions of P imply the conditions of Q, then knowledge relative to Q implies knowledge relative to P. The rule is sound on $M^{\bullet}_{\mathsf{PDtree}}$.

Table 6.2. The Global Protocol Inclusion Rule for protocol interactions.

Global Protocol Inclusion Rule	
GPI	If $\vdash P_{ab} \to Q_{ab}$ for all a, b then $\vdash K_i^Q \varphi \to K_i^P \varphi$ for all i

Before we show soundness, we prove an intermediate lemma that simplifies the inclusion of protocols to inclusion of their epistemic relations.

Lemma 6.6. *Let $M^{\bullet}_{\mathsf{PDtree}}$ be the synchronous tree model for protocol-dependent gossip. If $P_{ab} \to Q_{ab}$ is true globally, then for every two call sequences σ, τ we have that $\sigma \sim_i^P \tau$ implies $\sigma \sim_i^Q \tau$.*

Proof. Suppose that $P_{ab} \to Q_{ab}$ is true globally for all $a, b \in Ag$. Let σ and τ be arbitrary sequences such that $\sigma \sim_i^P \tau$. By synchronicity and definition of $\sim_i^P$ they have the same length. We use induction on the length n.

Base Case. Suppose $\sigma = \epsilon$. Then also $\tau = \epsilon$. Then from the definition of $\sim_i^Q$ we too have $\epsilon \sim_i^Q \epsilon$.

Induction Hypothesis. Let n be arbitrary $|\sigma| = n$ and suppose we have for all τ that $\sigma \sim_i^P \tau$ implies $\sigma \sim_i^Q \tau$.

Induction Step. We show the claim for $\sigma.ij$ with length $n+1$. We distinguish two cases where the last call in both sequences is equal and where it is not.

- Suppose $\sigma.ij \sim_i^P \tau.ij$. Then we have $\sigma \sim_i^P \tau$, $V_j(\sigma) = V_j(\tau)$, $\sigma \models P_{ij}$, and $\tau \models P_{ij}$. By the IH we have $\sigma \sim_i^Q \tau$ and by assumption we have $\sigma \models Q_{ij}$ and $\tau \models Q_{ij}$. By definition we find $\sigma.ij \sim_i^Q \tau.ij$.
- Suppose $\sigma.ab \sim_i^P \tau.cd$ for $a, b, c, d \neq i$. Then $\sigma \sim_i^P \tau$, and $\sigma \models P_{ab}$, and $\tau \models P_{cd}$. By assumption then also $\sigma, \models Q_{ab}$ and $\tau \models Q_{cd}$ and by IH $\sigma \sim_i^Q \tau$. Hence, by definition of $\sim_i^Q$ we have $\sigma.ab \sim_i^Q \tau.ab$. □

Lemma 6.7. *The* **GPI** *rule is sound for the synchronous protocol-dependent tree model $M^{\bullet}_{\mathsf{PDtree}}$.*

Proof. Suppose for all agents a, b we have that $P_{ab} \to Q_{ab}$ is true for all call sequences. Let agent i and formula φ be arbitrary. Consider an arbitrary call sequence σ such that $\sigma \models K_i^Q \varphi$.

Let $T = \{\tau \mid \sigma \sim_i^P \tau\}$ and $T' = \{\tau \mid \sigma \sim_i^Q \tau\}$. By definition of K_i^Q we have $\tau \models \varphi$ for all $\tau \in T'$. By lemma 6.6 we have $T \subseteq T'$. Hence we have $\tau \models \varphi$ for all $\tau \in T$ too, and by definition of K_i^P we conclude that $\sigma \models K_i^P \varphi$. As σ, i, and φ were chosen arbitrarily, this holds for all sequences, agents, and formulas. □

While this rule is sound, we conjecture that no such rules are needed if we restrict ourselves to the tree model, due to two properties that this model has. Firstly, it is reflexive. Therefore any protocol violation must be caused by a call, specifically by violating its protocol condition. Moreover, the epistemic relations $\sim_a^P$ for each protocol P in the initial root model are identical to each other. Hence the only effect of protocols on agent knowledge is via the protocol conditions.

We already have axioms that encapsulate this effect, namely the observance axioms. We therefore conjecture the proof system $\mathbb{GP}^{\bullet}_{\mathsf{tree}}$ to be complete for the tree model $M^{\bullet}_{\mathsf{PDtree}}$ with all protocols.

Conjecture 6.8. *$\mathbb{GP}^{\bullet}_{\mathsf{tree}}$ is sound and complete for the language $\mathcal{L}^P$ on the synchronous protocol-dependent tree model $M^{\bullet}_{\mathsf{PDtree}}$.*

Chapter 7

Expressivity of Multi-Protocol-Dependent Gossip

In this section we analyse the expressivity of the protocol-dependent knowledge modalities by comparing them to the basic language. We specifically look at the language $\mathcal{L}^P$, which contains a modality for each protocol in $\mathbb{P}$. It turns out that we can use combinations of protocols to define formulas that characterise many properties of the protocol-dependent tree model.

Throughout this section we assume the synchronous protocol-dependent tree model $M^{\bullet}_{\mathsf{PDtree}}$. Recall that we write $\sigma \models \varphi$ and generally only mention the call sequence τ when referring to a state $(w_{\mathsf{root}}, \tau)$.

7.1 Bisimulation for Basic Gossip Models

We start by finding states that satisfy the same formulas in the basic language $\mathcal{L}^G$ and will use bisimulation for this. We had already defined n-bisimulation for *protocol-dependent* gossip models, but we now define *unbounded* bisimulation for *basic* gossip models.

Definition 7.1 (Bisimulation)**.** *Let $M = \langle W, \sim, V\rangle$ and $M' = \langle W', \sim', V'\rangle$ be basic gossip models. A relation $Z \subseteq W \times W'$ is a* bisimulation *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' such that $s' \sim_a t'$ and tZt';*
3. *(**Back**) For each agent a, if $s' \sim_a t'$ then there is a t such that $s \sim_a t$ and tZt'.*

We call two states s, t bisimilar *and write $s \leftrightarrows t$ if there exists a bisimulation Z such that sZt.*

In particular we will now use bisimulation to relate two call sequences in the synchronous basic tree model: we use $\sigma \leftrightarrows \tau$ as a shorthand for $M^{\bullet}_{\mathsf{tree}}, (w_{\mathsf{root}}, \sigma) \leftrightarrows M^{\bullet}_{\mathsf{tree}}, (w_{\mathsf{root}}, \tau)$.

Lemma 7.2. *If two call sequences in $M^{\bullet}_{\mathsf{tree}}$ are bisimilar then they satisfy the same formulas in $\mathcal{L}^G$.*

Proof. Let σ, σ' be call sequences such that $\sigma \leftrightarrows \sigma'$. By lemma 3.4 it suffices to consider call-free formulas, so let $\varphi \in \mathcal{L}^G_{[-]}$. We show that $\sigma \models \varphi$ iff $\sigma' \models \varphi$ by induction on the structure of φ.

Base Case. Suppose $\varphi = S_a b$. By **Atoms** we have $\sigma \models S_a b \iff \sigma' \models S_a b$ as required.

Induction Hypothesis. Let $\varphi \in \mathcal{L}^G$ be arbitrary. Suppose that for all strict subformulas ψ of φ and for all call sequences σ, σ' we have $\sigma \leftrightarrows \sigma'$ implies $\sigma \models \psi \iff \sigma' \models \psi$.

Induction Step.

- Suppose $\varphi = \neg\psi$. We use the induction hypothesis on ψ definition of $\neg$ to find

$$\sigma \models \varphi \overset{\text{def}\neg}{\iff} \sigma \not\models \psi \overset{\text{IH}}{\iff} \sigma' \not\models \psi \overset{\text{def}\neg}{\iff} \sigma' \models \varphi.$$

- Suppose $\varphi = \psi \wedge \chi$. We use the induction hypothesis on ψ and χ and definition of $\wedge$ to find

$$\sigma \models \varphi \overset{\text{def}\wedge}{\iff} (\sigma \models \psi \text{ and } \sigma \models \chi) \overset{\text{IH}}{\iff} (\sigma' \models \psi \text{ and } \sigma' \models \chi) \overset{\text{def}\wedge}{\iff} \sigma' \models \varphi.$$

- Suppose $\varphi = K_a\psi$. For the forwards direction, suppose we have $\sigma \models K_a\psi$. Then for all τ such that $\sigma \sim_a \tau$ we have $\tau \models \psi$ by semantics of K_a. Let τ be an arbitrary sequence such that $\sigma' \sim'_a \tau$. We show that $\tau \models \psi$. By assumption $\sigma \underline{\leftrightarrow} \sigma'$ and so by **Back**, there exists a call sequence τ such that $\sigma \sim_a \tau$ and $\tau \underline{\leftrightarrow} \tau$. By the IH on ψ, this bisimilarity yields $\tau \models \psi$. As τ was arbitrary, this holds for all τ such that $\sigma' \sim'_a \tau$. Thus by semantics of K_a we conclude that $\sigma' \models K_a\psi$.

  For the backwards direction we use **Forth** and proceed analogously. □

In order to find bisimilar states in $M^{\bullet}_{\mathsf{tree}}$, we may be tempted to use disjoint call sequences. However, one might easily overlook epistemic relations that invalidate a bisimulation, as example 7.3 shows. Recall from definition 2.16 that the synchronous epistemic relation only relates sequences of the same length.

Example 7.3 (Disjoint Call Sequences)**.** *Let $Ag = \{a, b, c, d\}$. Consider the two call sequences $\sigma = ab.cd$ and $\tau = cd.ab$. The following is the partial call graph for these sequences.*

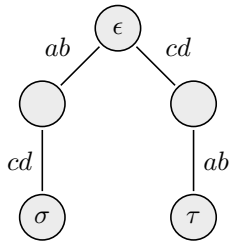

*On the surface it looks like these two sequences might be bisimilar: their atoms are the same and the epistemic relations between the above five states satisfy the **Forth** and **Back** conditions.*

However, at either σ or τ we have for each agent precisely two indistinguishable call sequences that all lie outside the picture, denoted below. Each of these sequences is a permutation of the 3 other agents for the call that the agent was not involved in. These sequence each satisfy different atoms, so there can be no bisimulation Z containing (σ, τ).

***Agent a**: $ab.cd \sim_a ab.bd$ **and** $ab.cd \sim_a ab.bc$.*

***Agent b**: $ab.cd \sim_b ab.ad$ **and** $ab.cd \sim_b ab.ac$.*

***Agent c**: $ab.cd \sim_c ad.cd$ **and** $ab.cd \sim_c bd.cd$.*

***Agent d**: $ab.cd \sim_d ac.cd$ **and** $ab.cd \sim_d bc.cd$.*

We do however have call sequences that are bisimilar, as example 7.4 shows.

Example 7.4. *Let $Ag = \{a, b\}$. Consider $\sigma = ab$ and $\tau = ba.ab$ with the following partial call graph.*

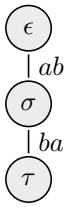

As there are only two agents, the only calls that can be executed are repeated calls back and forth. The model is therefore a binary tree, but the indistinguishability relations and atoms do not depend on the order of calls. We let Z be the bisimulation that contains 2 equivalence classes: one containing only the empty sequence ϵ and one containing all other call sequence.

Lemma 7.5. *In example 7.4, σ and τ are bisimilar.*

Proof. Let Z contain the total relation between all non-empty call sequences as well as the reflexive relation $\epsilon Z \epsilon$. Observe that thus $(\sigma, \tau) \in Z$ too. We claim that Z is a bisimulation.

First we observe that all atoms are satisfied after the first call, whether it is ab or ba. Next, both agents can exactly 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 $\sigma' Z \tau$, observe that they are non-empty call sequences because ϵ only occurs reflexively in Z, hence the secrets are already shared. Thus the atoms remain unchanged.

2. (**Forth**/**Back**) There are no epistemic relations except the reflexive ones, and each sequence has one. **Forth** and **Back** are satisfied by copying this reflexive relation. □

Because σ and τ are bisimilar, we know by lemma 7.2 that σ and τ model the same formulas in $\mathcal{L}^G$. Conversely, there is no formula $\varphi \in \mathcal{L}^G$ whose truth differs between σ and τ.

Corollary 7.6. *In example 7.4, σ and τ satisfy the same formulas in $\mathcal{L}^G$.*

7.2 Expressivity of the Protocol-Dependent Language

Next we look at the expressivity of the protocol-dependent language. Most importantly, we will show the existence of formulas in $\mathcal{L}^P$ that characterise properties in σ and τ from example 7.4 and even a formula that characterises the two call sequences uniquely.

We also show that adding a new operator Cab for *"Call ab has happened in the past"* [DGR23] does not increase expressivity of $\mathcal{L}^P$.

We will make heavy use of the global alarm: the property that all agents turn insane as soon as a P-illegal call takes place. Recall that we can characterise an illegal sequence in the tree model with the formula $K_a \bot$. Its dual $\widehat{K}_a^P \top$ constitutes the lack of a violation and therefore a P-permitted sequence.

7.3 Counting Formulas

We first construct a way to count the length of any call sequence using protocol-dependent knowledge. We do so by recursively defining a protocol for every natural number that depends on the previous protocol.

Definition 7.7 (Counting Protocols)**.** *We define the following protocols to determine the number of calls. For all natural numbers $k \geq 0$ and all agents $a \neq b$, and an arbitrary agent u we have:*

$$P_{ab}^0 := \bot \qquad \text{"Allow no calls"}$$
$$P_{ab}^{k+1} := \widehat{K}_u^{P^k} \top \qquad \text{"The previous protocol has not been violated"}$$

We can leverage the global alarm of these protocols. A sequence that violates protocol P^n corresponds to it exceeding length n. Recall however that a protocol can only be violated by an illegal call. Conversely, the lack of a call can never cause violation. The counting protocols therefore act like an upper bound, but not a lower bound.

Lemma 7.8. *For all agents u, call sequences σ and all $k \geq 0$ we have*

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

Proof. Let agent u and call sequence σ be arbitrary. We use induction on k.

Base case. Let $k = 0$.

($\Longrightarrow$) Suppose $\sigma \models \widehat{K}_u^{P^0} \top$. 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 we conclude that $|\sigma| \leq 0$.

($\Longleftarrow$) Suppose $|\sigma| \leq 0$. Then $\sigma = \epsilon$. As the empty call sequence cannot violate any protocol, it can also not violate P^0. Hence $\sigma \models \neg K_u^{P^0} \bot$ which implies $\sigma \models \widehat{K}_u^{P^0} \top$.

Induction Hypothesis. Let k be arbitrary and suppose we have $\sigma \models \widehat{K}_u^{P^k} \top$ if and only if $|\sigma| \leq k$

Induction Step. Suppose the induction hypothesis holds for arbitrary k. We show it holds for $k+1$. Let $\sigma = \tau.ab$ be an arbitrary call sequence such that $|\sigma| = k+1$. Using the semantics of $[ab]$ and $\widehat{K}_u$ and soundness of **Fnc** we get the following equivalences.

$$\tau.ab \models \widehat{K}_u^{P^{k+1}} \top \iff \tau \models [ab] \widehat{K}_u^{P^{k+1}} \top \iff \tau \models \neg [ab] K_u^{P^{k+1}} \bot$$

We then distinguish three cases for agent u to find that $\tau \models \neg [ab] K_u^{P^{k+1}} \bot \iff \tau \models \widehat{K}_u^{P^k} \top$. By applying the IH on $\tau \models \widehat{K}_u^{P^k} \top$ we find that $|\tau| \leq k$ and conclude that $|\sigma| \leq k+1$.

- Suppose $u \notin \{a, b\}$. We get the following equivalences. 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 that we have a reflexive relation $\tau \sim_u^{P^{k+1}} \tau$ which is sufficient to obtain $\tau \models \widehat{K}_u^{P^{k+1}} (P^{k+1}_{ab} \wedge [ab]\top)$. This satisfies the disjunct for $de = ab$.

$$\begin{aligned}
& \tau \models \neg [ab] K_u^{P^{k+1}} \bot \\
\iff & \tau \models \neg (P^{k+1}_{ab} \to \bigwedge_{d,e \neq u} K_u^{P^{k+1}} (P^{k+1}_{de} \to [de]\bot)) && (\mathbf{Pri}^{\bullet}_1(P^{k+1})) \\
\iff & \tau \models P^{k+1}_{ab} \wedge \neg (\bigwedge_{d,e \neq u} K_u^{P^{k+1}} (P^{k+1}_{de} \to [de]\bot)) && \text{(De Morgan)} \\
\iff & \tau \models P^{k+1}_{ab} \wedge \bigvee_{d,e \neq u} \widehat{K}_u^{P^{k+1}} (P^{k+1}_{de} \wedge \neg [de]\bot) && \text{(De Morgan, Sem. } \widehat{K}_u) \\
\iff & \tau \models P^{k+1}_{ab} \wedge \bigvee_{d,e \neq u} \widehat{K}_u^{P^{k+1}} (P^{k+1}_{de} \wedge [de]\top) && (\mathbf{Fnc}) \\
\iff & \tau \models \widehat{K}_u^{P^k} \top && (*) \\
\iff & |\tau| \leq k && \text{(By IH)} \\
\iff & |\sigma| \leq k+1
\end{aligned}$$

- Suppose $u = a$. We get the following equivalences, which follow the same structure as the previous case. For the backwards direction of the step at $(\dagger)$ observe that there is precisely one set $Q \subseteq \mathbb{S}$ such that $O_b Q$. For all other sets $T \neq Q$ the conjunct is satisfied with $\neg O_b T$. 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 $\tau \models \widehat{K}_u^{P^{k+1}} (P^{k+1}_{ab} \wedge O_b Q \wedge [ab]\top)$.

$$
\begin{aligned}
& \tau \models \neg[ab] K_u^{P^{k+1}} \bot \\
\iff & \tau \models \neg(P_{ab} \to \bigvee_{R \subseteq \mathbb{S}} (O_b R \wedge K_u^P(P_{ab} \to (O_b R \to [ab]\bot)))) && (\mathbf{Obs}_1^{\bullet}(P^{k+1})) \\
\iff & \tau \models P_{ab} \wedge \bigwedge_{R \subseteq \mathbb{S}} \neg(O_b R \wedge K_u^P(P_{ab} \to (O_b R \to [ab]\bot))) && \text{(De Morgan)} \\
\iff & \tau \models P_{ab} \wedge \bigwedge_{R \subseteq \mathbb{S}} \neg O_b R \vee \neg K_u^P(P_{ab} \to (O_b R \to [ab]\bot)) && \text{(De Morgan)} \\
\iff & \tau \models P_{ab} \wedge \bigwedge_{R \subseteq \mathbb{S}} \neg O_b R \vee \widehat{K}_u^P(P_{ab} \wedge O_b R \wedge \neg[ab]\bot) && \text{(De Morgan, Sem. } \widehat{K}_u) \\
\iff & \tau \models P_{ab} \wedge \bigwedge_{R \subseteq \mathbb{S}} \neg O_b R \vee \widehat{K}_u^P(P_{ab} \wedge O_b R \wedge [ab]\top) && (\mathbf{Fnc}) \\
\iff & \tau \models \widehat{K}_u^{P^k} \top && (\dagger) \\
\iff & |\tau| \leq k && \text{(By IH)} \\
\iff & |\sigma| \leq k+1
\end{aligned}
$$

- Suppose $u = b$. We repeat the steps for $u = a$ and instead apply $\mathbf{Obs}_2^{\bullet}(P^{k+1})$. □

The truth of these formulas does not depend on the choice of an agent, so we pick an arbitrary agent and call them u for "unlucky". After all, we will abuse their knowledge base by purposefully driving them insane with respect to our crafted protocols.

While the global alarm itself cannot be used to create a lower bound, we can easily negate the counting formula, which also negates the bound to form a strict lower bound. What we get is a characterisation of the length of a call sequence.

Definition 7.9 (Counting Formulas)**.** *We define $\varphi_0 := \varphi_\epsilon$ 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 7.10. *We have for every call sequence σ and all $k \in \mathbb{N}$ that*

$$\sigma \models \varphi_k \text{ if and only if } |\sigma| = k.$$

Proof. We prove the statement directly using the result of lemma 7.8.

$$
\begin{aligned}
& \sigma \models \varphi_k \\
\iff & \sigma \models \widehat{K}_u^{P^k} \top \wedge K_u^{P^{k-1}} \bot && \text{(def. } \varphi_k) \\
\iff & |\sigma| \leq k \text{ and } \neg(|\sigma| \leq k-1) && \text{(lemma 7.8)} \\
\iff & |\sigma| \leq k \text{ and } |\sigma| > k-1 \\
\iff & |\sigma| = k && \square
\end{aligned}
$$

This result is sufficient to conclude that $\mathcal{L}^P$ is more expressive than $\mathcal{L}^G$. Recall that the two sequences in example 7.4 were $\sigma = ab$ and $\tau = ab.ba$ and thus of different length, yet no formula in $\mathcal{L}^G$ can distinguish them.

Theorem 7.11. *The protocol-dependent language $\mathcal{L}^P$ is more expressive than the basic language $\mathcal{L}^G$.*

Proof. We show that there are two sequences σ, τ such that they agree on all formulas in $\mathcal{L}^G$ but there exists a formula $\varphi \in \mathcal{L}^P$ such that $\sigma \models \varphi$ and $\tau \not\models \varphi$.

Let $\sigma = ab$ and $\tau = ab.ba$ as per example 7.4. As shown in lemma 7.2 they agree on all formulas in $\mathcal{L}^G$. Consider $\varphi_1 \in \mathcal{L}^P$. By lemma 7.10 we have $\sigma \models \varphi_1$ and $\sigma \not\models \varphi_1$. □

7.4 Forcing Call Sequences

With the use of counting formulas we obtain a lot of expressivity. We can for instance define a protocol that only allows one specific call at the right moment. Combining multiple of these protocols recursively, we get a protocol that enforces an entire sequence.

Definition 7.12 (Sequence Protocols)**.** *Let $\sigma = a_1b_1.a_2b_2...a_mb_m$ be a non-empty call sequence of length m. We define the following protocols for $k \in \{0, ..., m\}$.*

$$
\begin{aligned}
P^{\sigma,0}_{ab} &:= \bot \\
P^{\sigma,k+1}_{ij} &:= \begin{cases} \widehat{K}_u^{P^{\sigma,\leq k}} \top \wedge \varphi_k & ij = a_{k+1}b_{k+1} \\ \bot & otherwise \end{cases} \\
P^{\sigma,\leq k}_{ab} &:= \bigvee_{j \leq k} P^{\sigma,j}_{ab}
\end{aligned}
$$

We call $P^{\sigma,\leq m}$ the sequence protocol *of σ and simply write P^σ.*

Recall that a protocol cannot enforce that a call sequence is actually completed. In order to ensure that we recognise the entire sequence and not just a prefix, we combine the global alarm of the sequence protocol with the counting formula for the length of the protocol.

Definition 7.13 (Sequence Formulas & Sequence Prefix Formulas)**.** *For any sequence σ of length m, we define two formulas $\varphi_{\leq\sigma}$ and φ_σ, which we call the* σ-prefix formula *and* σ-sequence formula *respectively.*

$$
\begin{aligned}
\varphi_{\leq\sigma} &:= \widehat{K}_u^{P^\sigma} \top && \sigma\text{-prefix formula} \\
\varphi_\sigma &:= \widehat{K}_u^{P^\sigma} \top \wedge \varphi_m && \sigma\text{-sequence formula}
\end{aligned}
$$

We will show that φ_σ exactly enforces σ, but first prove that $\widehat{K}_u^{P^{\sigma\leq,m}}$ enforces a prefix of σ of at most length m.

Lemma 7.14. *Let σ be a call sequence. We have for all sequences τ and all m such that $0 \leq m \leq |\sigma|$*

$$\tau \models \widehat{K}_u^{P^{\sigma,\leq m}} \text{ if and only if } \tau \text{ is a prefix of } \sigma \text{ and } |\tau| \leq m.$$

Proof. Let σ be an arbitrary call sequence. We denote σ by a series of calls as follows: $\sigma = a_1b_1.a_2b_2.a_3b_3....$ Let τ be an arbitrary call sequence. We use strong induction on m.

Induction Hypothesis. Let m be arbitrary and suppose $\tau \models \varphi_{\leq k}$ if and only if τ is a prefix of σ and $|\tau| \leq k$ holds for all $k < m$.

Induction Step. If $\tau = \epsilon$ then we are done vacuously as ϵ cannot violate any protocol and is a prefix no longer than 0 of every sequence σ. Let $\tau = \mu.ab$ therefore be non-empty. We have

$$\tau \models \widehat{K}_u^{P^{\sigma,\leq m}} \top \text{ if and only if } \mu \models P^{\sigma,\leq m}_{ab} \wedge [ab]\widehat{K}_a^{P^{\sigma,\leq m}} \top.$$

Since $P^{\sigma,\leq m}_{ab} = \bigvee_{j\leq m} P^{\sigma,j}_{ab}$, we have this if and only if $\mu \models P^{\sigma,k+1}_{ab}$ for some $k < m$. By definition of $P^{\sigma,k+1}_{ab}$

$$\mu \models \widehat{K}_u^{P^{\sigma,\leq k}} \wedge \varphi_k \wedge [ab]\widehat{K}_a^{P^{\sigma,\leq m}} \top.$$

We use the IH on k to find that μ is a prefix of σ of at most length k and moreover $|\mu| = k$ by lemma 7.10. Moreover, call ab can only have been $P^{\sigma,\leq m}$ permitted if $ab = a_{k+1}b_{k+1}$. Then $\tau = \mu.a_{k+1}b_{k+1}$, which is still a prefix of σ. Because finally $|\tau| = k + 1 \leq m$, we are done.

This concludes the induction on m. As τ and σ arbitrary, this holds for all call sequences. □

Lemma 7.15. *We have $\tau \models \varphi_\sigma$ iff $\tau = \sigma$.*

Proof. The proof is immediate from the result of lemma 7.10 and lemma 7.14. We have $\tau \models \widehat{K}_a^{P^{\sigma,\leq m}} \wedge \varphi_m$ if and only if τ is a prefix of σ and $|\tau| = m = |\sigma|$. Hence $\tau = \sigma$. □

7.5 Atom Formulas

We can also define protocols that are violated only when an atom becomes true. Clearly we can already express atoms: they are after all atoms in the language. We will however derive some other properties from these formulas.

Definition 7.16 (Atom Protocols)**.** *For each atom S_ab we define the protocol P^{S_ab} as follows. Let for all agents x, y the protocol condition be*

$$P_{xy}^{S_ab} := \neg[xy]S_ab.$$

That is, the protocol allows any sequence of calls that do not make the atom S_ab true. Or conversely: it does not allow any call sequence that achieves S_ab.

The definition of atom protocols is generally straightforward. We can apply again the use of the global alarm to express the truth of atoms in terms of protocol-dependent knowledge.

Definition 7.17 (Atom Formulas)**.** *Let $\varphi_{S_ab} := K_u^{P^{S_ab}} \bot$ for agents $a \neq b \in Ag$.*

Lemma 7.18. *Let $a \neq b \in Ag$. We have $\sigma \models \varphi_{S_ab}$ iff $\sigma \models S_ab$*

Proof. Let σ be some call sequence and $a, b \in Ag$ agents such that $a \neq b$.

($\Longrightarrow$) Suppose $\sigma \models \varphi_{S_ab}$. Then there must have been some call xy such that $\sigma = \tau.xy.\mu$ and xy violated the protocol, that is $\tau \models \neg P_{xy}^{S_ab} = \neg\neg[xy]S_ab$. We thus find $\tau.xy \models S_ab$ and as the truth of positive atoms is preserved under calls by definition 2.15, we conclude that $\sigma \models S_ab$.

($\Longleftarrow$) Suppose $\sigma \models S_ab$. Then $\sigma \neq \epsilon$ since $\epsilon \not\models S_ab$ by definition 4.15. So we can write $\sigma = \tau.xy.\mu$ where xy is the call after which S_ab first became true. Then $\tau \models [xy]S_ab$, which equals $\tau \models \neg P_{xy}$ and so $\tau.xy$ was P-illegal. Hence by definition 4.11 also σ was P-illegal and we conclude $\sigma \models K_u^{P^{S_ab}} \bot$. □

There are two benefits in using the atom formula over the atom itself. Firstly, the global alarm will ensure that all agents know (relative to the protocol) that S_ab is true. Secondly, we can generalise this approach for other formulas.

It may be tempting to try this approach with arbitrary formulas, but this will fail. The reason the atom formulas can use protocol conditions to force the atoms, is because the atoms are monotone in the sense that whenever they become true, their truth is preserved under calls. This aligns with the property that violated protocols remain violated after adding more calls.

7.6 Called operator

The language used by [DGR23] introduces a new operator Cab, which is true if a call ab has happened in the past. Formally: $\sigma \models Cab$ iff $ab \in \sigma$.

Whether a call has happened is not something that can be expressed in the basic language of gossip, so this addition increases the expressivity of $\mathcal{L}^G$. For the protocol-dependent language $\mathcal{L}^P$ it does not: we can define a protocol that is violated precisely when call ab happens. That this is possible, follows from the observation that Cab is a monotonic formula.

Definition 7.19 (Call Formula)**.** *Let $a \neq b$ be agents. We then define the* call formula $\varphi_{ab} := K_a^{P^{ab}} \bot$*, where P^{ab} is the protocol allowing any call to happen except call ab, defined by the following conditions.*

$$P_{ab} = \bot$$
$$P_{xy} = \top \quad \forall xy \neq ab$$

The following corollary is immediate from these protocol conditions and the semantics of protocol violation.

Corollary 7.20. *$\sigma \models \varphi_{ab}$ if and only if $ab \in \sigma$.*

Chapter 8

Discussion

Protocol-Dependent Validities and Reasoning about History

Call reductions have played a large role in the axiomatisation of the basic gossip problem by [DHK20]. It seems reasonable that calls are somewhat superfluous: the effects of a call, whether it is to the secret distribution or to agent knowledge, can be predicted before executing it. The role of calls in protocol-dependent gossip is much bigger. Calls can be permitted or not, and executing an illegal call has large consequences for the knowledge base of agents. It is therefore somewhat surprising that the axiomatisation for protocol-dependent gossip still has call reductions. While the effects have changed however, they remain clear before execution: it is clear that an agent will go insane after an illegal call.

Still, the axiomatisation of protocol-dependent gossip looks almost identical to the basic proof system. Besides the loss of reflexivity, the only significant change is the addition of protocol conditions to the observance axioms. It is at least remarkable that the only meaningful change is in axioms that are used as reduction axioms. This should mean that any additional knowledge that is gained from protocol-dependent knowledge, must be gained from doing calls.

The strength of protocol-dependency seems to lie in the agent's ability to unite their perception of the history – the initial world and any calls that have been made – with the constraints of the protocol – which calls are P-permitted. This ability is highest when the history is clearest, and it is clearest in the tree model. We therefore suspect that the most interesting validities for protocol-dependent knowledge are consequences of the tree rule. However, this rule does not have the most insightful structure and does not provide hints as to what sorts of validities these are.

Definition of Protocol-Dependent Initial Models

We have defined a more general class of models for protocol-dependent gossip using arbitrary initial models. They allow two properties that one might not expect from initial states of the protocol-dependent gossip problem: there are different epistemic relations for each protocol, and protocols may even be violated already. Usually, we assume that protocols can only be violated after a call, and up until violation, they agree with the basic epistemic relation that an agent would have had.

We justify this with two reasons. Firstly and purely technical: we need the flexibility to construct a canonical initial model. Secondly, we are mainly interested in an axiomatisation for the tree model, so there is no problem in utilising a larger class of models to achieve this. It might be possible however to show completeness by other means for these more well-behaved models and there are two reasons to attempt this.

Most importantly, it seems that these initial models – which the tree model is one of – have the property that all protocol interactions follow from the protocol conditions. There would in such a case be no need for extra axioms to relate different protocols. Secondly, these models satisfy the necessary conditions for the global alarm to take

effect, which has been the main source of expressivity. We may therefore also expect interesting behaviour of the protocol-dependent modalities on such models.

The Role of the Global Alarm

In chapter 7 we have seen that protocol-dependent knowledge modalities can express many things. Most notably, the global alarm helps characterise each individual gossip state in the tree model. The global alarm is however a strange semantic artefact and we claim that it should not be part of the semantics of protocol-dependent knowledge for two reasons.

1. **Exposure of calls**. External agents should not know which call took place in a non-transparent setting, but the global alarm informs them it was an illegal call (as opposed to a permitted call). This violates the privacy level: agents can suddenly infer more than they should be able to about which calls may have taken place.

2. **Knowledge of the conditions**. External agents should not necessarily know whether a call was permitted, because they might not know whether the protocol condition was violated. We do not usually require that a protocol is epistemic for external agents, only for the agent who executes the call (the caller).

We could solve the second problem extrinsically by restricting ourselves only to protocols that satisfy these requirements. rather than regular epistemic protocols[1] we could require *globally epistemic protocols* such that any agent knows whenever any condition holds. This puts a strong restriction on permissible protocols. Moreover, it does not solve the first problem.

A more elegant solution lies in redefining the semantics of K^P so that the alarm $K^P\bot$ only reaches the agent at some later point when they become sure that none of the call sequences that can have happened are P-permitted. We propose informally the following alternative semantics.

- After a call ab that was P-illegal, agent c does not know this call could have happened. In synchronous gossip, the agent instead only considers possible the set of P-permitted calls (that they were not involved in).

These semantics come with a downside: the agent will reach a state of false belief. This might be an acceptable trade-off, seeing that we already have this problem in P-illegal states. One could argue however that in the current semantics this is more acceptable because knowledge is still truthful when restricted to P-permitted states.

[1] Protocols where the caller knows the protocol condition holds when it does: $P_{ab} \to K_a P_{ab}$. See for instance [AGH15].

Chapter 9

Conclusion

We introduce the first proof systems for gossip including protocol-dependent knowledge modalities. Most importantly, we show that $\mathbb{G}\mathbf{P1}^{\bullet}_{\mathsf{tree}}$ is sound and complete on the synchronous protocol-dependent tree model restricted to one protocol-dependent modality. Additionally we show that the weaker system $\mathbb{G}\mathbf{P1}^{\bullet}$ is sound and strongly complete for arbitrary models in this setting, and demonstrate that it may be possible to extend the proof system to multiple protocols and in particular to the synchronous tree model $M^{\bullet}_{\mathsf{PDtree}}$ as defined by [Dit+19]. To this end, we propose a generalised system $\mathbb{GP}^{\bullet}_{\mathsf{tree}}$ that we conjecture is sound and complete for $M^{\bullet}_{\mathsf{PDtree}}$.

In doing so we unify the semantics for protocol-dependent gossip with semantics using arbitrary initial models. Meanwhile, we also identify a sound and complete proof system for other semantics of basic gossip in [DG24].

We show that using of protocol-dependent modalities in the language of gossip increases its expressivity significantly, but also raises questions about the desirability of its semantics, in particular the property of the global alarm that violates assumptions in the gossip problem.

Future Work

The axiomatisation for the multi-protocol setting remains an open question. Our method may be further extended to multiple protocols, but requires additional axioms for protocol interactions in arbitrary models. The system $\mathbb{GP}^{\bullet}_{\mathsf{tree}}$ might be a suitable candidate as a sound and complete axiomatisation for the synchronous protocol-dependent tree model, although a different approach is required to show completeness without the use of arbitrary initial models.

Future research into the expressivity of protocol-dependent knowledge modalities is needed to better understand the effects of the modality, both in different settings of gossip and settings outside gossip. This research may take interest in analysing the expressivity of languages restricted to epistemic protocols or defining other restricted classes of protocols. Within gossip, expressivity of the protocol-dependent language can also be tested on other models than the tree model, such as models including arbitrary initial states.

Besides work in the direction of axiomatisations, future work could consider modifying the semantics of protocol-dependent knowledge to avoid the global alarm property. A last direction is to find the syntactic derivation for both $\neg EEExp_{Ag}$ and $\neg K_a EExp_{Ag}$.

Bibliography

[AGH15] Krzysztof R. Apt, Davide Grossi, and Wiebe van der Hoek. "Epistemic Protocols for Distributed Gossiping". In: *Proceedings Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2015*. Volume 215. EPTCS. 2015, pages 51–66. https://doi.org/10.4204/EPTCS.215.5.

[Att+14] M. Attamah, Hans Ditmarsch, D. Grossi, and Wiebe Hoek. "Knowledge and Gossip". In: *Frontiers in Artificial Intelligence and Applications* 263 (Jan. 2014), pages 21–26. https://doi.org/10.3233/978-1-61499-419-0-21.

[AW17] Krzysztof R. Apt and Dominik Wojtczak. "Common Knowledge in a Logic of Gossips". In: *Proceedings Sixteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2017*. Volume 251. EPTCS. 2017, pages 10–27. https://doi.org/10.4204/EPTCS.251.2.

[BRV01] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science 53. Cambridge University Press, 2001. 554 pages. ISBN: 978-0-521-80200-0.

[BS72] Brenda Baker and Robert Shostak. "Gossips and Telephones". In: *Discrete Mathematics* 2.3 (1972), pages 191–193. https://doi.org/10.1016/0012-365X(72)90001-5.

[DG22] Hans van Ditmarsch and Malvin Gattinger. "The Limits to Gossip: Second-order Shared Knowledge of All Secrets Is Unsatisfiable". In: *WoLLIC 2022*. Edited by Agata Ciabattoni, Elaine Pimentel, and Ruy de Queiroz. 2022, pages 237–249. https://doi.org/10.1007/978-3-031-15298-6_15.

[DG24] Hans van Ditmarsch and Malvin Gattinger. "You Can Only Be Lucky Once: Optimal Gossip for Epistemic Goals". In: *Mathematical Structures in Computer Science* (2024). ISSN: 0960-1295, 1469-8072. https://doi.org/10.1017/S0960129524000082.

[DGR23] Hans van Ditmarsch, Malvin Gattinger, and Rahim Ramezanian. "Everyone Knows That Everyone Knows: Gossip Protocols for Super Experts". In: *Studia Logica* 111.3 (2023), pages 453–499. ISSN: 0039-3215, 1572-8730. https://doi.org/10.1007/s11225-022-10032-3.

[DHK20] Hans Van Ditmarsch, Wiebe Van Der Hoek, and Louwe B. Kuijer. "The Logic of Gossiping". In: *Artificial Intelligence* 286 (2020), page 103306. ISSN: 00043702. https://doi.org/10.1016/j.artint.2020.103306.

[Dit+14] Hans Van Ditmarsch, Sujata Ghosh, Rineke Verbrugge, and Yanjing Wang. "Hidden Protocols: Modifying Our Expectations in an Evolving World". In: *Artificial Intelligence* 208 (2014), pages 18–40.

[Dit+19] Hans van Ditmarsch, Malvin Gattinger, Louwe B. Kuijer, and Pere Pardo. "Strengthening Gossip Protocols Using Protocol-Dependent Knowledge". In: *Journal of Applied Logics - IfCoLog Journal of Logics and their Applications* 6.1 (2019), pages 157–203. http://arxiv.org/abs/1907.12321.

[Gat18] Malvin Gattinger. "New Directions in Model Checking Dynamic Epistemic Logic". PhD thesis. University of Amsterdam, 2018. ISBN: 978-94-028-1025-7. https://malv.in/phdthesis.

[HHI13] Wesley H Holliday, Tomohiro Hoshi, and Thomas F Icard III. "Information Dynamics and Uniform Substitution". In: *Synthese. An International Journal for Epistemology, Methodology and Philosophy of Science* 190 (Suppl 1 2013), pages 31–55.

[HLL99] Mor Harchol-Balter, Tom Leighton, and Daniel Lewin. "Resource Discovery in Distributed Networks". In: *Proceedings of the Eighteenth Annual ACM Symposium on Principles of Distributed Computing*. Podc '99. New York, NY, USA: Association for Computing Machinery, 1999, pages 229–237. ISBN: 1-58113-099-6. https://doi.org/10.1145/301308.301362.

[HM17] Andreas Herzig and Faustine Maffre. "How to Share Knowledge by Gossiping". In: *AI Communications* 30.1 (2017), pages 1–17. https://doi.org/10.3233/AIC-170723.

[HMP21] Andreas Herzig, Frédéric Maris, and Elise Perrotin. "A Dynamic Epistemic Logic with Finite Iteration and Parallel Composition". In: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning* 18.1 (2021), pages 676–680. ISSN: 2334-1033. https://doi.org/10.24963/kr.2021/68.

[Kar+00] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking. "Randomized Rumor Spreading". In: *Proceedings 41st Annual Symposium on Foundations of Computer Science*. 2000, pages 565–574. https://doi.org/10.1109/SFCS.2000.892324.

[Kle17] Rana Klein. "The Logical Dynamics of Gossip: An Analysis in Dynamic Epistemic Logic". MSc thesis. University of Amsterdam, 2017. https://eprints.illc.uva.nl/id/eprint/1567/1/MoL-2017-26.text.pdf.

[Tij71] Robert Tijdeman. "On a Telephone Problem". In: *Nieuw Archief voor Wiskunde* 3.19 (1971), pages 188–192.