Reasoning about Communication Graphs

Eric Pacuit* Rohit Parikh†

July 19, 2006

Abstract

Let us assume that some agents are connected by a *communication graph*. In the communication graph, an edge from agent i to agent j means that agent i can directly receive information from agent j. Agent i can then refine its own information by learning information that j has, including information acquired by j from another agent, k. We introduce a multi-agent modal logic with knowledge modalities and a modality representing communication among agents. Among other properties, we show that the logic is decidable, that it completely characterizes the communication graph, and that it satisfies the basic properties of the space logic of [18].

1 Introduction

The topic "who knew what and when" is not just of interest to epistemic logicians. Often it is the subject of political scandals (both real and imagined). For example, consider the much talked about Valerie Plame affair. A July 2003 column in the Washington Post reported that Plame was an undercover CIA operative. This column generated much controversy due to the fact that such information (the identity of CIA operatives) is restricted to the relevant government officials. Of course, in this situation, we know full well "Who knew what and when": in July of 2003, Robert Novak (the author of the article) knew that Plame was a CIA operative. What creates a scandal in this situation is *how* Novak came to know such information. Since the CIA goes to great lengths

*ILLC, University of Amsterdam `epacuit@staff.science.uva.nl`, `staff.science.uva.nl/~epacuit`

†CS, Math and Philosophy Departments, Brooklyn College and The Graduate Center of CUNY, 365 5th Avenue, New York City, NY, 10016, `rparikh@gc.cuny.edu`, `www.sci.brooklyn.cuny.edu/~rparikh` Research of both authors supported under the PSC-CUNY FRAP program.

to ensure that communication about sensitive information is contained within its own organization, the only way Novak could have known that Plame was a CIA operative was if a communication channel had been created between Novak and someone inside the CIA organization.

To put this a bit more formally, given a set of agents $\mathcal{A}$, call any graph $\mathcal{G} = (\mathcal{A}, E)$ a **communication graph** where the intended interpretation of an edge between agent i and agent j is that i and j *can communicate*. In this setting, the CIA can be represented as a connected component of $\mathcal{G}$. Given that the CIA is the only group of agents that (initially) knows the identity of CIA operatives, and Novak is not an element of the CIA component of $\mathcal{G}$ then we can conclude that Novak did not originally know the identity of CIA operatives and no amount of communication *that respects the graph* $\mathcal{G}$ can create a situation in which Novak does know the identity of a CIA operative. Thus Novak's report in the Washington Post implied that our original communication graph was incorrect[1]. That is, there must be an edge (or a chain) between Novak and *some* agent inside the CIA component. Since in principle, Novak could be connected to any member of the CIA component, much resources and time has been spent discussing the possible edges.

In this paper we develop[2] a multi-agent epistemic logic with a communication modality where agents are assumed to communicate according to some fixed communication graph. Agents are assumed to have some private information at the outset, but may refine their information by acquiring information possessed by other agents, possibly via yet other agents. That is, each agent is initially informed about the truth values of a finite set of propositional variables. Agents are assumed to be connected by a *communication graph*. In the communication graph, an edge from agent i to agent j means that agent i can directly receive information from agent j. Agent i can then refine its information by learning information that j has, including information acquired by j from another agent, k.

In keeping with the CIA-theme, we give an example from [19] of the type of situations that we have in mind. Let $K_i\phi$ mean that according to i's current information ϕ is true. Given a communication graph $\mathcal{G} = (\mathcal{A}, E)$, we say that a sequence of communications (i learns a fact from j who learns a fact from k, and so on) **respects the communication graph** if agents only communicate with their immediate neighbors in $\mathcal{G}$. Let $\Diamond\phi$ mean that ϕ becomes true after a sequence of communications that respects the communication graph. Suppose now that ϕ is a formula representing the exact whereabouts of Bin Laden, and that Bob, the CIA operative in charge of maintaining this information knows ϕ. In particular, $K_{\text{Bob}}\phi$, but suppose that at the moment, Bush does not know the exact whereabouts of Bin Laden ($\neg K_{\text{Bush}}\phi$). Presumably Bush *can* find out the exact whereabouts of Bin Laden ($\Diamond K_{\text{Bush}}\phi$) by going through Hayden, but of course, *we* cannot

[1] Of course, it could also mean that we were incorrect about the agents' initial information — Novak could have had previous knowledge about the identity of CIA agents. In this paper, we are interested in studying communication and so will not consider this case.

[2] This framework was first presented in [19].

find out such information ($\neg\Diamond K_{\mathrm{E}}\phi \wedge \neg\Diamond K_{\mathrm{R}}\phi$) since we do not have the appropriate security clearance. Clearly, then, as a *pre-requisite* for Bush learning ϕ, Hayden will also have to come to know ϕ. We can represent this situation by the following formula:

$$\neg K_{\mathrm{Bush}}\phi \wedge \Box(K_{\mathrm{Bush}}\phi \rightarrow K_{\mathrm{Hayden}}\phi)$$

where $\Box$ is the dual of diamond ($\Box\phi$ is true if ϕ is true after every sequence of communications that respect the communication graph).

Section 2 gives the details of our framework and the main results. Section 3 contains a discussion of some other relevant literature. We conclude the paper by discussing our underlying assumptions and provide pointers to future work.

2 The Logic of Communication Graphs

This section describes the logic of communication graphs, $\mathcal{K}(\mathcal{G})$, introduced in [19]. The intended application is to reason about the flow of information among a group of agents whose communication is restricted by some communication graph. We begin by making some simplifying assumptions about the nature of the communication. Thus the logic presented here should be viewed as a first step towards a general logic to reason about the situations described in the Introduction.

Let $\mathcal{A}$ be a set of agents. A **communication graph** is a directed graph $\mathcal{G}_{\mathcal{A}} = (\mathcal{A}, E)$ where $E \subseteq \mathcal{A}\times\mathcal{A} - \{(i,i) \mid i \in \mathcal{A}\}$. Intuitively $(i,j) \in E$ means that i can directly receive information from agent j, but *without* j knowing this fact. Thus an edge from i to j in the communication graph represents a one-sided relationship between i and j. Agent i has access to *any* piece of information that agent j knows (but in a restricted language). We have introduced this 'one sidedness' restriction in order to simplify our semantics, but also because such situations of one sided learning occur naturally. A common situation that is helpful to keep in mind is accessing a website. We can think of agent j as creating a website in which everything he *currently* knows is available, and if there is an edge between i and j then agent i can access this website without j being aware that the site is being accessed. Another important application is *spying* where one person accesses another's information without the latter being aware that information is being leaked. Naturally j may have been able to access some other agent k's website and had updated some of her own information. Therefore, it is important to stress that when i accesses j's website, he is accessing j's current information which may include part of what k knew initially.

In order for any communication to take place, we must assume that the agents understand a common language. Thus we assume a set At of propositional variables, understood by all the agents, but with only specific agents knowing their actual values at the start. Letters p, q, etc, will denote elements of At. The agents will have some information –

knowledge of the truth values of some elements of At, but refine that information by acquiring information possessed by other agents, possibly via yet other agents. This implies that if agents are restricted in whom they can communicate with, then this fact will restrict the knowledge they can acquire.

The assumption that i can access all of j's information is a significant idealization from common situations, but becomes more realistic if we think of this information as being confined to facts expressible as truth functional combinations of some small set of basic propositions. Thus our idealization rests on two assumptions:

1. All the agents share a common language, and
2. The agents make available all possible pieces of (purely propositional) information which they know *and which are expressible in this common language.*

The language is a multi-agent modal language with a communication modality. The formula $K_i\phi$ will be interpreted as "according to i's current information, i knows ϕ", and $\Diamond\phi$ will be interpreted as "after some communications (which respect the communication graph), ϕ becomes true". For example the formula

$$K_j\phi \to \Diamond K_i\phi$$

is intended to express the statement: "If agent j (currently) knows ϕ, then after some communication, agent i can come to know ϕ". Let At be a finite set of propositional variables. A well-formed formula of $\mathcal{K}(\mathcal{G})$ has the following syntactic form

$$\phi \ := p \mid \neg\psi \mid \phi \wedge \psi \mid K_i\phi \mid \Diamond\phi$$

where $p \in \mathsf{At}$. We abbreviate $\neg K_i\neg\phi$ and $\neg\Diamond\neg\phi$ by $L_i\phi$ and $\Box\phi$ respectively, and use the standard abbreviations for the propositional connectives ($\vee$, $\to$, and $\bot$). Let $\mathcal{L}_{\mathcal{K}(\mathcal{G})}$ denote the set of well-formed formulas of $\mathcal{K}(\mathcal{G})$. We also define $\mathcal{L}_0(\mathsf{At})$, (or simply $\mathcal{L}_0$ if At is fixed or understood), to be the set of ground formulas, i.e., the set of formulas constructed from At using $\neg, \wedge$ only.

2.1 Semantics

The semantics described here combines ideas both from the subset models of [18] and the history based models of Parikh and Ramanajum (see [21, 22]). We assume that agents are initially given some private information and then communicate according to some fixed communication graph $\mathcal{G}$. The semantics in this section is intended to formalize what agents know and may come to know after some communication.

Initially, each agent i knows or is informed (say by nature) of the truth values of a certain subset At_i of propositional variables, and the At_i *as well as this fact are common knowledge*. Thus the other agents know that i knows the truth values of elements of At_i, but, typically, not what these values actually are. We shall not assume that the At_i are disjoint, but for this paper we *will* assume that the At_i together add up to all of At. Thus if At_i and At_j intersect, then agents i, j will share information at the very beginning. Let W be the set of boolean valuations on At. An element $v \in W$ is called a *state*. We use 1 for the truth value *true* and 0 for the truth value *false*. Initially each agent i is given a boolean valuation $v_i : \mathsf{At}_i \to \{0, 1\}$. This initial distribution of information among the agents can be represented by a vector $\vec{v} = (v_1, \ldots, v_n)$. Of course, since we are modeling knowledge and not belief, these initial boolean valuations must be compatible. I.e., for each i, j, v_i and v_j agree on $\mathsf{At}_i \cap \mathsf{At}_j$. Call any vector of partial boolean valuations $\vec{v} = (v_1, \ldots, v_n)$ **consistent** if for each $p \in dom(v_i) \cap dom(v_j)$, $v_i(p) = v_j(p)$ for all $i, j = 1, \ldots, n$. Note that there is a 1-1 correspondence between consistent vectors and states w as we defined them earlier. We shall assume that only such consistent vectors arise as initial information. All this information is common knowledge and only the precise values of the v_i are private.

Definition 1 *Let* At *be a finite set of propositional variables and* $\mathcal{A} = \{1, \ldots, n\}$ *a finite set of agents. Given the distribution of sublanguages* $\vec{\mathsf{At}} = (\mathsf{At}_1, \ldots, \mathsf{At}_n)$*, an* **initial information vector for** $\vec{\mathsf{At}}$ *is any consistent vector* $\vec{v} = (v_1, \ldots, v_n)$ *of partial boolean valuations such that for each* $i \in \mathcal{A}$*,* $dom(v_i) = \mathsf{At}_i$*.*

We assume that the only communications that take place are about the physical world. But we do allow agents to learn objective facts which are not atomic, but may be complex, like $p \vee q$ where $p, q \in \mathsf{At}$. Now note that if agent i learned some *literal* from agent j, then there is a simple way to update i's valuation v_i with this new information by just adding the truth value of another propositional symbol. However, if i learns a more general ground formula from agent j, then the situation will be more complex. For instance if the agent knows p and learns $q \vee r$ then the agent now has three valuations on the set $\{p, q, r\}$ which cannot be described in terms of a partial valuation on a subset of At.

Fix a communication graph $\mathcal{G}$ and suppose that agent i learns some ground fact ϕ (directly) from agent j. Of course, there must be an edge from agent i to agent j in $\mathcal{G}$. This situation will be represented by the tuple (i, j, ϕ) and will be called a **communication event**. For technical reasons we assume that all formulas in a communication event are expressed in a canonical *disjunctive normal form* (DNF). That is, we assume ϕ is a set $\{C_1, \ldots, C_k\}$ where each C_i is a consistent finite set of elements of At and their negations. Let $\mathcal{L}_{DNF}(\mathsf{At})$ be the set of all such sets. Each set $\{C_1, \ldots, C_k\}$ represents the formula $\bigvee_{i=1,\ldots,k} \bigwedge C_i$. Recall that for each formula $\psi \in \mathcal{L}_0(\mathsf{At})$ there is a unique element $\{C_1, \ldots, C_k\} \in \mathcal{L}_{DNF}(\mathsf{At})$ such that $\bigvee_{i=1,\ldots,k} \bigwedge C_i$ is logically equivalent to ψ. In what follows, we will sometimes use ϕ to mean either a formula (an element of $\mathcal{L}(\mathsf{At})$ or $\mathcal{L}_0(\mathsf{At})$) or an element of $\mathcal{L}_{DNF}(\mathsf{At})$ and trust that this ambiguity of notation will not cause any confusion. Of course we could use a unique element of $\mathcal{L}_0(\mathsf{At})$ to express a member of $\mathcal{L}_{DNF}(\mathsf{At})$ and that too would work.

Definition 2 *Let $\mathcal{G} = (\mathcal{A}, E_{\mathcal{G}})$ be a communication graph. A tuple (i, j, ϕ), where $\phi \in \mathcal{L}_{DNF}(\mathsf{At})$ and $(i, j) \in E_{\mathcal{G}}$, is called a* **communication event**. *Then $\Sigma_{\mathcal{G}} = \{(i, j, \phi) \mid \phi \in \mathcal{L}_{DNF}, (i, j) \in E_{\mathcal{G}}\}$ is the set of all possible communication events (given the communication graph $\mathcal{G}$).*

The following fact will be needed in what follows. The proof is well-known and is left to the reader.

Lemma 3 *Suppose that there are k elements in At and n elements in $\mathcal{A}$. Then for any communication graph $\mathcal{G}$, there are at most $n \times n \times (2^{2^k})$ elements in $\Sigma_{\mathcal{G}}$.*

Given the set of events $\Sigma_{\mathcal{G}}$, a **history** is a finite sequence of events. I.e., $H \in \Sigma_{\mathcal{G}}^*$. The empty history will be denoted ϵ. The following notions are standard (see [21, 22] for more information). Given two histories H, H', say $H \preceq H'$ iff $H' = HH''$ for some history H'', i.e., H is an initial segment of H'. Obviously, $\preceq$ is a partial order. If H is a history, and (i, j, ϕ) is a communication event, then H followed by (i, j, ϕ) will be written $H; (i, j, \phi)$. Given a history H, let $\lambda_i(H)$ be i's local history corresponding to H. I.e., $\lambda_i(H)$ is a sequence of events that i can "see". Given our assumption of "one-sided communication", for this paper we use the following definition of λ_i: Map each event of the form (i, j, ϕ) to itself, and map other events (m, j, ψ) with $m \neq i$ to the null string while preserving the order among events.

Definition 4 *Fix a finite set of agents $\mathcal{A} = \{1, \ldots, n\}$ and a finite set of propositional variables At along with subsets $(\mathsf{At}_1, ..., \mathsf{At}_n)$. A* **communication graph frame** *is a pair $\langle \mathcal{G}, \vec{\mathsf{At}} \rangle$ where $\mathcal{G}$ is a communication graph, and $\vec{\mathsf{At}} = (\mathsf{At}_1, ..., \mathsf{At}_n)$ is an assignment of sub-languages to the agents. A* **communication graph model** *based on a frame $\langle \mathcal{G}, \vec{\mathsf{At}} \rangle$ is a triple $\langle \mathcal{G}, \vec{\mathsf{At}}, \vec{v} \rangle$, where $\vec{v}$ is an initial information vector for $\vec{\mathsf{At}}$.*

Now we address two issues. One is that not all histories are legal. For an event (i, j, ϕ) to take place after a history H, it must be the case that $(i, j) \in E_{\mathcal{G}}$, and that after H (and before (i, j, ϕ)), j already knew ϕ. Clearly i cannot learn from j something which j did not know. Whether a history is justified depends not only on the initial valuation, but also on the set of communications that have taken place prior to each communication in the history.

The second issue is that the information which an agent learns by "reading" a formula ϕ may be *more* than just the fact that ϕ is true. For suppose that i learns $p \vee q$ from j, but j is not connected, directly or indirectly, to anyone who might know the initial truth value of q. In this case i has learned *more* than $p \vee q$, i has learned p as well. For the only way that j could have known $p \vee q$ is if j knew p in which case p must be true. Our definition of the semantics below will address both these issues.

We first introduce the notion of i-equivalence among histories. Intuitively, two histories are i-equivalent if those communications which i takes active part in, are the same.

Definition 5 *Let w be a state and H a finite history. Define the relation $\sim_i$ as follows: $(w, H) \sim_i (v, H')$ iff $w_{|\mathsf{At}_i} = v_{|\mathsf{At}_i}$ and $\lambda_i(H) = \lambda_i(H')$.*

Formulas will be interpreted at pairs (w, H) where w is a state (boolean valuation) and H is a history (a finite sequence of communication events).

To deal with the notion of legal or justified history we introduce a propositional symbol L which is satisfied only by legal pairs (w, H). (We may also write $L(w, H)$ to indicate that the pair (w, H) is legal.) Since L can only be defined in terms of knowledge, and knowledge in turn requires quantification over legal histories, we shall need mutual recursion.

Definition 6 *Given a communication graph and the corresponding model $\mathcal{M} = \langle \mathcal{G}, \vec{\mathsf{At}}, \vec{v} \rangle$, and pair (w, H), we define the legality of (w, H) and the truth $\models_{\mathcal{M}}$ of a formula as follows:*

- $w, \epsilon \models_{\mathcal{M}} L$
- $w, H; (i, j, \phi) \models_{\mathcal{M}} L$ *iff* $w, H \models_{\mathcal{M}} L$, $(i, j) \in E$ *and* $w, H \models_{\mathcal{M}} K_j \phi$
- $w, H \models_{\mathcal{M}} p$ *iff* $w(p) = 1$, *where* $p \in \mathsf{At}$
- $w, H \models_{\mathcal{M}} \neg\phi$ *iff* $w, H \not\models_{\mathcal{M}} \phi$
- $w, H \models_{\mathcal{M}} \phi \wedge \psi$ *iff* $w, H \models_{\mathcal{M}} \phi$ *and* $w, H \models_{\mathcal{M}} \psi$
- $w, H \models_{\mathcal{M}} \Diamond\phi$ *iff* $\exists H'$, $H \preceq H'$, $L(w, H')$, *and* $w, H' \models_{\mathcal{M}} \phi$
- $w, H \models_{\mathcal{M}} K_i\phi$ *iff* $\forall (v, H')$ *if* $(w, H) \sim_i (v, H')$, *and* $L(v, H')$, *then* $v, H' \models_{\mathcal{M}} \phi$

Unless otherwise stated, we will only consider legal pairs (w, H), i.e., pairs (w, H) such that $w, H \models L$. We say ϕ is **valid in** $\mathcal{M}$, $\models_{\mathcal{M}} \phi$ if for all (w, H), $w, H \models_{\mathcal{M}} \phi$. Finally, we say ϕ is **valid in the communication graph frame** $\mathcal{F}$ if ϕ is valid in all models based on $\mathcal{F}$.

There are two notions of validity relevant for our study. The first is relative to a fixed communication graph. Let $\mathcal{G}$ be a fixed communication graph. We say that a formula

$\phi \in \mathcal{L}_{\mathcal{K}(\mathcal{G})}$ is **$\mathcal{G}$-valid** provided ϕ is valid in all communication graph frames $\mathcal{F}$ based on $\mathcal{G}$. A formula $\phi \in \mathcal{L}_{\mathcal{K}(\mathcal{G})}$ is **valid** if ϕ is $\mathcal{G}$-valid for all communication graphs $\mathcal{G}$. Of course validity implies $\mathcal{G}$-validity, but not vice versa. We write $\models_{\mathcal{G}} \phi$ if ϕ is $\mathcal{G}$-valid and $\models \phi$ if ϕ is valid.

Some comments are in order concerning the above definition of truth. L is defined in terms of the knowledge operator, and the knowledge operator uses L in its definition. This definition is of course fine since the definition uses recursion on the length of the formula. Still, is not clear that the process of determining whether a history is legal will terminate in a *finite* amount of time. For whether or not the history-state pair $(w, H; (i, j, \phi))$ is legal depends on whether (w, H) satisfies $K_j\phi$ which, in turn, depends on a set of histories which may be *longer* than H. We now show that the process of determining whether a history is legal will terminate.

We first need some notation. A **one-step compression** of H is a history H' which is obtained by deleting *one* second or subsequent occurrences of an event (i, j, ϕ) from H. I.e., if (i, j, ϕ) has occurred twice, then eliminate some later occurrence. Let $c(H)$ denote the maximally compressed history. $c(H)$ is generated by including each (i, j, ϕ) event of H exactly once according to the following order on events: e comes before e' iff the first occurrence of e in H came before the first occurrence of e' in H. The key observation is that the legality of a history-state pair (w, H) depends only on the legality of the pair $(w, c(H))$.

Lemma 7 *Let $\mathcal{G}$ be a communication graph, $\Sigma_{\mathcal{G}}$ a set of events and H be any history over $\Sigma_{\mathcal{G}}$. Suppose that w is a state. Then*

- *For all ϕ, and all j, $w, H \models K_j\phi$ iff $w, c(H) \models K_j\phi$*
- *(w, H) is legal iff $(w, c(H))$ is legal.*

Proof Clearly it is sufficient to prove the two conditions when $c(H)$ is replaced by an H' obtained from H by the elimination of *one* extra event. Therefore we shall make this assumption in the rest of the proof. Thus $H = H_1eH_2eH_3$ and $H' = H_1eH_2H_3$. Here e is some event (i, j, ϕ).

We show that $(w, H), (w, H')$ satisfy the same formulas ψ.

Clearly this is true if ψ is atomic and the argument also goes through for boolean combinations.

Suppose $\psi = \Diamond\theta$ and $w, H \models \psi$. Then there is H_4 such that $w, H; H_4 \models \theta$. By induction hypothesis $w, H'; H_4 \models \theta$ (it is not hard to see that it is legal) and hence $w, H' \models \Diamond\psi$. The converse is similar.

Suppose $\psi = K_r\theta$ and $w, H \models \psi$ where $r \neq i$. This case is easy as H, H' are r-equivalent.

What if $r = i$? $w, H \models K_i\theta$ iff for all v, H'' such that $(v, H'') \sim_i (w, H)$, $v, H'' \models \theta$. But since e has already occurred in both H, H', the possible v in question are the same for both. The H'' will have two e events and eliminating the second e will yield an H''' such that $v, H''' \models \theta$ iff $w, H'' \models \theta$. Thus the case $r = i$ also works.

The proof that compression preserves the legality of illegality is now immediate for it depends on some knowledge formulas being true. But that issue is not affected by the elimination of extra events e. □

With this Lemma we can show that the process of determining if a state-history pair (w, H) is legal terminates. More formally,

Proposition 8 *Let $\mathcal{M} = \langle \mathcal{G}, \vec{\mathsf{At}}, \vec{v} \rangle$ be a communication graph model. For any $H \in \Sigma^*_{\mathcal{G}}$ and state w, the question "Does $(w, H) \models L$?" terminates in a finite amount of time.*

Proof This is now immediate by the previous lemma. When asking whether $w, H \models K_i\phi$ we need to look at pairs (v, H') which are i-equivalent to w, H. But now we can confine ourselves to H' is which no (r, j, ψ) event with $r \neq i$ occurs twice, and these are bounded in length. □

2.2 Some Results

We now state the basic results about the logic of communication graphs.

We already defined $c(H)$ earlier. We say that two histories are c**-equivalent**, written $C(H, H')$, if $c(H) = c(H')$. Clearly H, H' have the same semantic properties as $c(H) = c(H')$ and hence as each other. Moreover, one is legal iff the other is legal. It follows also that for every H_1, $H; H_1$ is legal iff $H'; H_1$ is. Thus C is a bisimulation.

In the following, a history H is called w**-maximal** if (w, H) is legal and all possible (finitely many) communication events have taken place at least once. Another consequence of the above lemma is the existence of a **maximal history** (relative to some w).

Theorem 9 *1. If a formula ϕ is satisfiable in some graph model $(\mathcal{G}, \vec{\mathsf{At}})$ then it is satisfiable in a history in which no communication (i, j, ϕ) occurs twice.*

2. *If H is w-maximal, then for all formulas $\phi \in \mathcal{L}_{\mathcal{K}(\mathcal{G})}$, $\models \phi \to \Box\phi$.*

3. *If H is w-maximal and H' is any history compatible with w such that $H \preceq H'$, then for all formulas $\phi \in \mathcal{L}_{\mathcal{K}(\mathcal{G})}$, if $w, H' \models \phi$ then $w, H \models \phi$.*

4. *If H and H' are w-maximal, then for each formula $\phi \in \mathcal{L}_{\mathcal{K}(\mathcal{G})}$, $w, H \models \phi$ iff $w, H' \models \phi$.*

Proof The first three parts follow from lemma 7. We will prove the last statement: for any state w, if H and H' are w-maximal histories, then (w, H) and (w, H') satisfy the same formulas. The proof is by induction on ϕ.

The base case and boolean connectives are straightforward. Suppose that ϕ is of the form $\Diamond\psi$. Let w be an arbitrary state and suppose that H and H' are w-maximal histories. Suppose that $(w, H) \models \Diamond\psi$. Then there is some H'' such that $H \preceq H''$, (w, H'') is legal and $(w, H'') \models \phi$. By part 3 above, $w, H \models \phi$ and by the induction hypothesis, $w, H' \models \phi$. Hence, $w, H' \models \Diamond\phi$. Thus if $w, H \models \Diamond\psi$ then $w, H' \models \Diamond\psi$. The other direction is similar.

For the knowledge case we need the following claim:

Claim Let w be a state. Suppose H_1 and H_2 are w-maximal and (v, H_3) is legal. If $(w, H_1) \sim_i (v, H_3)$, then there is a history H_4 which is v-maximal such that $(w, H_2) \sim_i (v, H_4)$.

Proof of claim Let w be a states and suppose that H_1 and H_2 are w-maximal histories and (v, H_3) is a legal history-state pair such that $(w, H_1) \sim_i (v, H_3)$. Then v and w must agree on every atom which is not only known to i, but also on *every* atom known to some other agent whom i can read directly or indirectly. For at maximality, i already knows the truth values of all the sets At_j where j is directly or indirectly accessible from i. Thus we can find a legal history-state pair (v, H_4') such that $(w, H_2) \sim_i (v, H_4')$. It is not hard to see that H_4' can be extended to a v-maximal history.

Returning to the induction proof. Suppose that ϕ is of the form $K_i\psi$, w is an arbitrary state, and H and H' are w-maximal histories. Suppose that $(w, H) \models K_i\psi$ and $(w, H') \not\models K_i\psi$. Then there is a history state pair (v, H'') such that $(w, H') \sim_i (v, H'')$, (v, H'') is legal and $v, H'' \not\models \psi$. Note that without loss of generality we can assume that H'' is v-maximal (every legal history-state pair can be extended to a maximal history, furthermore this extension cannot differ on the truth value of ϕ since H'' already contains *all* events of the form (i, j, χ).). By the above claim (let $H_1 = H'$, $H_2 = H$ and $H_3 = H''$), there is a v-maximal history H''' with $(w, H) \sim_i (v, H''')$. By the induction hypothesis, (v, H'') and (v, H''') satisfy the same formulas. Hence $(w, H) \sim_i (v, H''')$ and $v, H''' \not\models \psi$. This contradicts the assumption that $w, H \models K_i\psi$. Hence, for any w-maximal histories H and H', $w, H \models K_i\psi$ iff $w, H' \models K_i\psi$. $\square$

This result immediately gives us a decision procedure as we can limit the length of the history which might satisfy some given formula ϕ.

Notice that if we restrict our attention to maximal histories, then the following property will be satisfied: for any two agents i and j if there is a path in the communication graph from i to j, then any ground fact that j knows, i will also know. In this case, we can say that j *dominates* i. This is Fitting's "dominance" relation discussed in section 3.

The following simple result demonstrates that given any sequence of communications H, the agents know at least the set of formulas that are implied by the set of formulas in H. That is, given a legal pair (w, H), let $X_i(w, H)$ be the set of states that agent i considers possible if the actual state is w and the communication between the agents evolved according to H. Given a formula $\phi \in \mathcal{L}_0(\mathsf{At})$, let $\widehat{\phi} = \{w \mid w \in W,\ w(\phi) = 1\}$. Now we formally define $X_i(w, H)$ recursively as follows

1. $X_i(w, \epsilon) = \{v \mid v_{|\mathsf{At}_i} = w_{|\mathsf{At}_i}\}$
2. $X_i(w, H; (i, j, \phi)) = X_i(w, H) \cap \hat{\phi}$
3. if $i \neq m$ then $X_i(w, H; (m, j, \phi)) = X_i(w, H)$

The following theorem shows that agents know at least the formulas implied by the set $X_i(w, H)$. The proof can be found in [19] and will not be repeated here.

Theorem 10 *Let $\mathcal{M} = \langle \mathcal{G}, \vec{\mathsf{At}}, \vec{v} \rangle$ be any communication graph model and ϕ a ground formula. If $X_i(w, H) \subseteq \widehat{\phi}$, then $(w, H) \models_{\mathcal{M}} K_i(\phi)$.*

As we saw above, the converse is not true. That is, there are formulas that an agent can come to know which are not implied by the set $X_i(w, H)$. These are the formulas that agents can deduce given their knowledge of the communication graph.

The following axioms and rules are known to be sound and complete with respect to the family of all subset spaces ([18]). Thus they represent the core set of axioms and rules for any **topologic.**

1. All propositional tautologies
2. $(p \to \Box p) \wedge (\neg p \to \Box \neg p)$, for $p \in \mathsf{At}$.
3. $\Box(\phi \to \psi) \to (\Box\phi \to \Box\psi)$
4. $\Box\phi \to \phi$

5. $\Box\phi \to \Box\Box\phi$

6. $K_i(\phi \to \psi) \to (K_i\phi \to K_i\psi)$

7. $K_i\phi \to \phi$

8. $K_i\phi \to K_iK_i\phi$

9. $\neg K_i\phi \to K_i\neg K_i\phi$

10. (Cross axiom) $K_i\Box\phi \to \Box K_i\phi$

We include the following rules: modus ponens, K_i necessitation and $\Box$ necessitation. It is easy to verify that the above axioms and rules are valid, i.e., valid in all frames based on any communication graphs. We only demonstrate that the cross axiom $K_i\Box\phi \to \Box K_i\phi$ is valid. It is easier to consider it in its contrapositive form: $\Diamond L_i\phi \to L_i\Diamond\phi$. This is interpreted as follows: if there is a sequence of updates that lead agent i to consider ϕ possible, then i already thinks it possible that there is a sequence of updates after which ϕ becomes true.

Proposition 11 *The axiom scheme $\Diamond L_i\phi \to L_i\Diamond\phi$ is valid.*

Proof Let $\mathcal{G}$ be an arbitrary communication graph and $\mathcal{M} = \langle \mathcal{G}, \vec{\mathsf{At}}, \vec{v} \rangle$ any communication graph model based on $\mathcal{G}$. Suppose that (w, H) is a legal state-history pair. Suppose that $w, H \models \Diamond L_i\phi$. Then there exists H' with $H \preceq H'$ such that $w, H' \models L_i\phi$. Hence there is a pair (v, H'') such that $(w, H') \sim_i (v, H'')$ and $v, H'' \models_{\mathcal{M}} \phi$. Let H''' be any sequence such that $\lambda_i(H) = \lambda_i(H''')$ and $H''' \preceq H''$. Such a history must exist since $H \preceq H'$ and $H' \sim_i H''$. Since $H \preceq H'$, $\lambda_i(H) \preceq \lambda_i(H') = \lambda_i(H'')$. Therefore, we need only let H''' be any initial segment of H'' containing $\lambda_i(H)$. By definition of L, all initial sequences of a legal history are legal. Therefore, since $v, H'' \models_{\mathcal{M}} \phi$, $v, H''' \models \Diamond\phi$; and since $H \sim_i H'''$, $w, H \models_{\mathcal{M}} L_i\Diamond\phi$. $\Box$

The next lemma follows from the existence of maximal histories.

Lemma 12 *The axiom $\Box\Diamond\phi \leftrightarrow \Diamond\Box\phi$ is valid.*

Proof Let $\mathcal{G}$ be an arbitrary communication graph, $\mathcal{M} = \langle \mathcal{G}, \vec{\mathsf{At}}, \vec{v} \rangle$ any communication graph model based on $\mathcal{G}$ and (w, H) a legal history-state pair. Suppose that $w, H \models_{\mathcal{M}} \Box\Diamond\phi$. Let H' be a w-maximal history extending H, then $w, H' \models_{\mathcal{M}} \Diamond\phi$ and hence there is a history H'' such that $H' \preceq H''$ and $w, H'' \models_{\mathcal{M}} \phi$. Since H' is maximal, by theorem 9 $w, H' \models \phi$. By theorem 9 again, $w, H' \models \Box\phi$. Hence $w, H \models \Diamond\Box\phi$.

Conversely, suppose that $w, H \models_{\mathcal{M}} \Diamond\Box\phi$. Then there is a history H' such that $H \preceq H'$ and $w, H' \models_{\mathcal{M}} \Box\phi$. Let H'_m be a w-maximal history that extends H'. Then $w, H'_m \models_{\mathcal{M}} \phi$. Let H'' be any history that extends H and H''_m be a w-maximal history that extends H''. By theorem 9, since $w, H'_m \models_{\mathcal{M}} \phi$, $w, H''_m \models_{\mathcal{M}} \phi$. Hence $w, H'' \models_{\mathcal{M}} \Diamond\phi$ and hence $w, H \models_{\mathcal{M}} \Box\Diamond\phi$. $\square$

If we fix a communication graph, then there are more formulas which are valid.

Lemma 13 *Let $\mathcal{G} = (\mathcal{A}, E)$ be a communication graph. Then for each $(i, j) \in E$, for all $l \in \mathcal{A}$ such that $l \neq i$ and $l \neq j$ and all ground formulas ϕ, the scheme*

$$K_j\phi \wedge \neg K_l\phi \to \Diamond(K_i\phi \wedge \neg K_l\phi)$$

is $\mathcal{G}$-valid.

Proof Let $\mathcal{G}$ be an arbitrary communication graph and $\mathcal{M}$ a communication graph model based on $\mathcal{G}$. Suppose that $w, H \models_{\mathcal{M}} K_j\phi \wedge \neg K_l\phi$. Then j knows ϕ and hence i can read ϕ directly from j's website. More formally, $H; (i, j, \phi)$ is a legal history (provided that H is legal). The agent l is none the wiser as $\lambda_l(H) = \lambda_l(H; (i, j, \phi))$. Therefore, $w, H; (i, j, \phi) \models_{\mathcal{M}} K_i\phi \wedge \neg K_l\phi$. $\square$

The converse is not quite true. For suppose that agent i is connected to agent j via (exactly) two other agents l_1, l_2. Then a fact known to j can be learned by i without l_1 finding out about it, ditto for l_2 and for any agent beside l_1, l_2. However, in this scenario, it is impossible for i to find out some ϕ from j with neither of l_1, l_2 knowing ϕ.

3 Related Literature

Communication graphs, or communication networks, and more generally social networks[3] have been studied by a number of different communities. Most notably, social networks are an important topic in sociology. But computer scientists have also had quite a lot to say about networks. The goal of this section is not to survey this vast amount of literature, but rather to give some details about a few papers most relevant to the framework which will be discussed in the rest of this paper. Needless to say, this list of topics is not meant to be complete.

[3] A social network is a graph on a set of agents where edges represent some social interaction such as acquaintances, coauthors on a mathematical paper, costars in a movie, and so on.

Coordination Problems Suppose there are two generals A and B with armies at the top of two mountains with a valley in between them. As the story goes, the generals attempt to coordinate their action by sending messages back and forth. However, since the communication channel is unreliable (the messengers must travel through dangerous territory), common knowledge of the time to attack cannot be achieved. This puzzle, called the generals problem, has been topic of much discussion. See [8] for a formal treatment and discussion of relevant literature and [17] for a discussion of common knowledge as it relates to coordination problems. Of course, if there was a reliable communication channel between the two generals, then they could easily coordinate their actions. Thus the existence of a communication graph (in this case just an edge between the two generals) facilitates coordination. This raises some interesting questions about the structure of the communication graph on a group of agents and its affect on coordination.

In [7, 6], Michael Chwe investigates the general question of when the structure of the communication graph can facilitate coordination among a group of agents. Chwe considers the following situation. There is a finite set of agents $\mathcal{A}$. Each agent must decide whether or not to revolt against the current government. That is, assume that agents choose between r (revolt) and s (stay at home). The agents are assumed to use the following decision rule: "I'll go if you go". That is, for a particular agent i, the greater the number of agents i believes will choose r, the higher the utility i assigns to r, provided agent i is willing to revolt. More formally, it is assumed that each agent is either willing to revolt (w) or not willing to revolt (nw). Then a utility function for agent i maps elements of $\{w, nw\} \times \{r, s\}^n$ to real numbers with the constraint that if an agent is not willing to revolt, then the utility of revolting is zero regardless the opinions of i's neighbors. It is assumed that the agents are connected by a communication graph $\mathcal{G} = (\mathcal{A}, E)$. Here $(i, j) \in E$ is intended to mean "agent i talks to agent j". That is i informs j as to whether or not he will revolt. Thus the set $B_i = \{j \mid (j, i) \in E\}$ is the set of agents for which i knows which action they will perform. Finally, it is assumed that the communication graph is common knowledge and that each agent has a prior belief[4] about which agents will revolt.

Chwe considers the question "which communication graphs enable the group to revolt?" To that end, a strategic game $\Gamma(\mathcal{G}, \{\pi\}_{i\in\mathcal{A}})$ is defined, where π_i is agent i's prior beliefs. In this game, an agent i's decision to revolt depends on its prior beliefs π_i and the set B_i of agents that i has communicated with. Of course if agent i's prior belief assigns high enough probability to a large enough group of agents revolting, then that agent will revolt *regardless of the communication graph.* Thus the interesting question is which communication graph will enable the group to revolt *regardless of the agents prior probabilities.* Chwe calls such communication graphs a **sufficient network**.

The main result of the paper [7] is a characterization of minimal sufficient networks. First of all, it should be obvious that if a communication graphs $\mathcal{G}$ enables a group to revolt, then so will any communication graph $\mathcal{G}'$ which is just like $\mathcal{G}$ except with additional

[4]Beliefs are represented by probability functions over the set $\{r, s\}^n$.

edges (it is straightforward to prove this in Chwe's framework). Chwe showed that any sufficient network has the following property: there is a finite set of cliques such that

1. Each agent is in at least one clique,
2. there is a relation $\rightarrow$ between the cliques that characterizes the edge relation in the graph. That is there is an edge from i to j iff there is some clique containing i and some clique containing j that are related by $\rightarrow$, and
3. the cliques are totally ordered by $\rightarrow$.

This result provides an interesting perspective on collective action. The communication graph facilitates the group's ability to share information and thus enabling group action. The logic presented in Section 2 is intended to make clear precisely how an agent's information can change in situations similar to the one described above.

Agreeing to Disagree In 1976, Aumann proved a fascinating result [1]. Suppose that two agents have the same prior probability and update their probability of an event E with some private information using Bayes' rule. Then Aumann showed that if the posterior probability of E is common knowledge, then they must assign the *same* posterior probability to the event E. In other words, if agents have the same prior probability and update using Bayes' rule, then the agents cannot "agree to disagree" about their posterior probabilities. See [4] for a nice discussion of this result and the literature that it generated. An immediate question that comes to mind is "*How* do the posterior probabilities become common knowledge?" Starting with Geanakoplos and Polemarchakis [11], a number of papers have addressed this issue [5, 2, 20, 16, 13].

The key idea is that common knowledge arises through communication. Suppose there are two agents who agree on a prior probability function. Suppose that each agent receives some private information concerning an event E and updates their probability function accordingly. Geanakoplos and Polemarchakis [11] show that if the agents each announce their posterior probabilities and update with this new information, then the probabilities will eventually become common knowledge and the probabilities will be equal. Similar to Chwe's analysis described above, the existence of a communication graph (with an edge between the two agents) enables consensus about the posterior probabilities.

Parikh and Krasucki [20] look at the general situation where there may be more than two agents[5] and communication is restricted by a communication graph. They show that under certain assumptions about the communication graph, consensus can be reached even though the posterior probabilities of the agents may not be common knowledge[6].

[5]Cave [5] also considers more than two agents, but assumes all communications are public announcements.

[6]This point was formally clarified by Heifetz in [13]. He demonstrates how to enrich the underlying partition space with time stamps in order to formalize precisely when events become common knowledge.

Before stating their result, some clarification is needed. Note that a communication graph tells us which agent *can* communicate with which agent, but not when two agents *do* communicate. To represent this information, Parikh and Krasucki introduce the notion of a **protocol**. A protocol is a pair of functions (r, s) where $r : \mathbb{N} \to \mathcal{A}$ and $s : \mathbb{N} \to \mathcal{A}$. Intuitively, $(r(t), s(t))$ means that $r(t)$ receives a message from $s(t)$ at time t. Say a protocol (r, s) respects a communication graph $\mathcal{G} = (\mathcal{A}, E)$ if for each $t \in \mathbb{N}$, $(r(t), s(t)) \in E$. A protocol is said to be **fair** provided every agent can send a message to any other agent, either directly or indirectly, infinitely often[7]. Parikh and Krasucki show that if the agents are assumed to have finitely many information sets, then for any protocol if the agents sends the current probability[8] (conditioned on the agent's current information set) of proposition A, then after a finite amount of time t for each agent i, the messages received after time t will not change i's information set. Furthermore, if the protocol is assumed to be fair (i.e., the communication graph is strongly connected) then all the agents will eventually assign the same probability to A. Krasucki takes the analysis further in [16] and provides conditions on the protocol (and implicitly on the underlying communication graph) which will guarantee consensus regardless of the agents' initial information.

Similar to Chwe's analysis, Parikh and Krasucki's analysis shows that the structure of the communication graph is a key aspect of consensus in a group. We end this section with a different interpretation of a communication graph that provides a very interesting perspective on many-valued modal logic.

Many-Valued Modal Logic In [9, 10], Fitting discusses various families of many-valued modal logics. In these frameworks, Fitting assumes that both the valuation of the formulas and the accessibility relation are many-valued. The frameworks are motivated in [10] with the help of a structure similar to a communication graph. As in this paper, it is assumed that there is a graph with the set of agents as nodes. However, Fitting interprets an edge in this graph as follows: "i is related to j" means that i *dominates* j, where 'dominate' means that "j says that something is true whenever i says it is". It is assumed that this relation is a partial order. Each agent is assumed to have its own valuation and accessibility relation on a set of possible worlds. Fitting is interested in which modal formulas the agents will agree on in a particular world. Two semantics are presented that solve this problem.

Deciding which agents agree on a particular formula in a common language naturally suggests a many-valued modal logic where formulas are assigned subsets of agents. Suppose that ϕ is assigned the set $\mathcal{B} \subseteq \mathcal{A}$ in a state w, then the intended interpretation

[7] Consult [20] for a formal definition of "fairness".

[8] Actually, Parikh and Krasucki consider a more general setting. They assume agents communicate the value of some function f that maps events to real numbers. Intuitively, f should be thought of as the probability of an event given some information set. The only condition imposed on f is a convexity condition: for any two disjoint close subsets X and Y, $f(X \cup Y)$ lies in the open interval between $f(X)$ and $f(Y)$. Here closed is defined with respect to the agents information sets. This generalizes a condition imposed by Cave [5] and Bacharach [2].

is that each agent in $\mathcal{B}$ agrees on the truth value of ϕ in w. The domination relation outlaws certain subsets of agents as truth values. For example, if i dominates j, then $\{i\}$ cannot be a possible truth value since j is assumed to *always* agree with i on the set of true formulas. The domination relation also provides an intuitionistic flavor to the underlying logic. For example, consider the formula $\neg\phi$ and suppose i dominates j. Now it is consistent with the notion of domination that i can consider ϕ false and j considers ϕ true. In this case, if we interpret $\neg$ classically, then i considers $\neg\phi$ true while j considers $\neg\phi$ false, which contradicts the fact that i dominates j. Thus, we are forced to say that i considers $\neg\phi$ true if *i and all agents that i dominates* consider that ϕ is false. Fitting offers two semantics which take these observations into account. One is a combination of Kripke intuitionistic models and Kripke muli-modal models and the second is a many-valued Kripke modal model. The two semantics are shown to be equivalent and a sound and complete axiomatization is offered.

Rasiowa and Marek offer a similar interpretation of a communication graph [25, 24]. In their framework an edge from i to j means that "j is more perceptive than i". If this is the case, then j's valuation of a proposition p is "better than" i's valuation of that same variable. Rasiowa and Marek provide a framework to reason about formulas on which there is consensus among the agents. The framework discussed in the rest of this paper can be seen as an attempt to explain *how* an agent i can come to dominate another agent j. That is, assuming the agents start with consistent (partial) theories, i can dominate j if there is a (possibly indirect) communication channel from i to j and j asks i about the truth value of all formulas in the language.

Dynamic Epistemic Semantics The study of *Dynamic Epistemic Logic* attempts to combine ideas from dynamic logics of actions and epistemic logic. The main idea is to start with a formal model that represents the uncertainty of an agent in a social situation. Then we can define an 'epistemic update' operation that represents the effect of a communicatory action, such as a public announcement, on the original model. For example, publicly announcing a true formula ϕ, converts the current model to a submodel in which ϕ is true at each state. Starting with [23] and more recently [3, 15, 29, 12, 27], logical systems have been developed with the intent to capture the dynamics of information in a social situation. Chapter 4 of Kooi's dissertation [15] and the forthcoming book [30] contain a thorough discussion of the current state of affairs.

These logics use **PDL** style operators to represent an epistemic update. For example, if $!\phi$ is intended to mean a public announcement of ϕ, then $\langle !\phi\rangle K_i\psi$ is intended to mean that after ϕ is publicly announced, agent i knows ψ. From this point of view, our communication modality $\Diamond$ can be understood as existentially quantifying over a sequence of *private* epistemic updates. However, there are some important differences between the semantics presented in this paper and the semantics found in the dynamic epistemic logic literature. First of all, in our semantics, communication is limited by the communication graph. Secondly, we do not consider general epistemic updates as is common in the literature, but rather study a specific type of epistemic update and its connection with a communication graph. Most important is the fact that the history of communications plays a key role in the definition of knowledge in this paper. The general approach of dynamic epistemic semantics is to define update operations mapping Kripke structures to other Kripke structures intended to represent the effect of an epistemic update on the first Kripke structure. For example, a public announcement of ϕ selects the submodel of a Kripke structure in which ϕ is true at every state. The definition of knowledge after an epistemic update is the usual definition, i.e., ϕ is known by i at state w if ϕ is true in all states that i considers possible from state w in the updated Kripke structure.

Floris Roelofson introduces communication graphs to the dynamic epistemic logic setting in [26]. The framework is more general than the one presented in this paper in three respects. First, the communication graph is a relation on the collection of subsets of $\mathcal{A}$, where an edge between $\mathcal{B}_1 \subseteq \mathcal{A}$ and $\mathcal{B}_2 \subseteq \mathcal{A}$ means that group $\mathcal{B}_1$ can communicate with group $\mathcal{B}_2$. Thus a communication event in this framework is a tuple $(\mathcal{B}_1, \mathcal{B}_2, \phi)$ intended to mean that group $\mathcal{B}_1$ sends a message whose content is ϕ to group $\mathcal{B}_2$; and a precondition for an event $(\mathcal{B}_1, \mathcal{B}_2, \phi)$ to take place is that ϕ is common knowledge among group $\mathcal{B}_1$[9]. Second, there is no assumption that messages between groups be restricted to ground formulas. Finally, Roelofson does not assume that the communication graph is common knowledge. Therefore, there may be messages that can "update" the agent's view of the communication graph. So, should our models be viewed as an interesting special case of these more general model? The answer to this question is not straightforward as the history of communication plays a crucial role in the semantics presented in this paper. A full comparison between these two approaches can be found in [14] (cf. [28] for a comparison between history based semantics and dynamic epistemic semantics).

4 Conclusions

In this paper we have introduced a logic of knowledge and communication. Communication among agents is restricted by a communication graph, and idealized in the sense that the agents are unaware when their knowledge base is being accessed. We have shown that the communication graph is characterized by the validities of formulas in models based on that communication graph, and that our logic is decidable.

Moving on to future work. Standard questions such as finding an elegant complete axiomatization will be studied. The semantics described in the Section 2 rests on some strong underlying assumptions. Below we briefly sketch how to remove some of these assumptions.

One-way communication As discussed in the introduction, an edge from i to j means that i can read j's website *without* j knowing that it's website is being read. Thus a

[9] We *could* of course consider other cases where some members of $\mathcal{B}_1$ communicate with some members of $\mathcal{B}_2$

communication event (i, j, ϕ) only changes i's knowledge. This can be formally verified by noting that if H and $H; (i, j, \phi)$ are w-legal histories, then by the definition of the λ_j function, $\lambda_j(H) = \lambda_j(H; (i, j, \phi))$. Thus $(w, H) \sim_j (v, H')$ iff $(w, H; (i, j, \phi)) \sim_j (v, H')$ and so j's knowledge is unchanged by the presence of the event (i, j, ϕ). We can model *conscious communication* by changing the definition of the local view function. Define the λ_i^* as follows: given a history H, let $\lambda_i^*(H)$ map events of the form (i, j, ϕ) and (j, i, ϕ) to themselves and all other events in which i does not occur in the first two components to the null event.

Starting with theories Another natural extension is to consider situations in which agents have a preference over which information they will read from another agent's website. Thus for example, if one hears that an English Ph.D. student and his adviser recently had a meeting, then one is justified in assuming that they probably did not discuss the existence of non-recursive sets, even though the adviser may conceivably know this fact. Given that this preference over the formulas under discussion among different groups of agents is common knowledge, each agent can regard some (legal) histories as being more or less likely than other (legal) histories. From this ordering over histories, we can define a *defeasible* knowledge operator for each agent. The operator is defeasible in the sense that agents may be wrong, i.e., it *is* after all possible that the English student and his adviser actually spent the meeting discussing the fact that there must be a non-recursive set.

References

[1] Aumann, R. Agreeing to disagree. *Ann. Statist. 4*, 6 (1976), 1236 – 1239.

[2] Bacharach, M. Some extensions of a claim of Aumann in an axiomatic model of knowledge. *Journal of Economic Theory 37* (1985), 167 – 190.

[3] Baltag, A., and Moss, L. Logics for epistemic programs. *Synthese: Knowledge, Rationality, and Action 2* (2004), 165 – 224.

[4] Bonanno, G., and Nehring, K. Agreeing to disagree: A survey. Available at `http://www.econ.ucdavis.edu/faculty/bonanno/wpapers.htm`, 1997.

[5] Cave, L. Learning to agree. *Economics Letters 12* (1983), 147 – 152.

[6] Chwe, M. S.-Y. Structure and strategy in collective action. *American Journal of Sociology 105*, 1 (1999), 128 – 156.

[7] Chwe, M. S.-Y. Communication and coordination in social networks. *Review of Economic Studies 67* (2000), 1 – 16.

[8] Fagin, R., Halpern, J., Moses, Y., and Vardi, M. *Reasoning about Knowledge.* The MIT Press, Boston, 1995.

[9] Fitting, M. Many-valued modal logics. *Fundam. Inform. 15*, 3-4 (1991), 235–254.

[10] Fitting, M. Many-valued modal logics II. *Fundam. Inform. 17*, 1-2 (1992), 55 – 73.

[11] Geanakoplos, J., and Polemarchakis, H. We can't disagree forever. *Journal of Economic Theory 28*, 1 (1982), 192 – 200.

[12] Gerbrandy, J. *Bisimulations on Planet Kripke*. PhD thesis, University of Amsterdam, 1999.

[13] Heifetz, A. Comment on consensus without common knowledge. *Journal of Economic Theory 70* (1996), 273 — 277.

[14] Hoshi, T. The logic of communication graphs for group communication and the dyname epistemic logic with a future operator. Unpublished manuscript, 2006.

[15] Kooi, B. *Knowledge, Chance and Change*. PhD thesis, University of Groningen, 2003.

[16] Krasucki, P. Protocols forcing consensus. *Journal of Economic Theory 70* (1996), 266 – 272.

[17] Lewis, D. *Convention*. Harvard University Press, 1969.

[18] Moss, L., and Parikh, R. Topological reasoning and the logic of knowledge. In *Proceedings of TARK IV* (1992), Y. Moses, Ed., Morgan Kaufmann.

[19] Pacuit, E., and Parikh, R. The logic of communication graphs. In *Proceedings of DALT 2004* (2005), J. Leita, A. Omicini, P. Torroni, and P. Yolum, Eds., no. 3476 in Lecture Notes in AI, Springer, pp. 256 – 269.

[20] Parikh, R., and Krasucki, P. Communication, consensus, and knowledge. *Journal of Economic Theory 52*, 1 (1990), 178 – 189.

[21] Parikh, R., and Ramanujam, R. Distributed processes and the logic of knowledge. In *Logic of Programs* (1985), vol. 193 of *Lecture Notes in Computer Science*, Springer, pp. 256 – 268.

[22] Parikh, R., and Ramanujam, R. A knowledge based semantics of messages. *Journal of Logic, Language and Information 12* (2003), 453 – 467.

[23] Plaza, J. Logics of public communications. In *Proceedings, 4th International Symposium on Methodologies for Intelligent Systems* (1989).

[24] Rasiowa, H., and Marek, V. W. Mechanical proof systems for logic II, consensus programs and their processing. *J. Intell. Inf. Syst 2*, 2 (1993), 149 – 164.

[25] Rasiowa, H., and Marek, W. On reaching consensus by groups of intelligent agents. In *Methodologies for Intelligent Systems* (1989), Z. W. Ras, Ed., North-Holland, pp. 234 – 243.

[26] Roelofsen, F. Exploring logical perspectives on distributed information and its dynamics. Master's thesis, ILLC, University of Amsterdam, 2005.

[27] van Benthem, J. 'One is a lonely number': on the logic of communication. Tech. Rep. Technical Report PP-2002-27, ILLC, University of Amsterdam, 2002. Available at http://staff.science.uva.nl/∼johan/Muenster.pdf.

[28] van Benthem, J., and Pacuit, E. Epistemic temporal logic and its neighbors. Working Paper.

[29] van Ditmarsch, H. *Knowledge Games*. PhD thesis, University of Groningen, 2000.

[30] van Ditmarsch, H., van der Hoek, W., and Kooi, B. *Dynamic Epistemic Logic*. Springer-Verlag, Forthcoming, 2006.