The Epistemic Logic of Full Communication and Social Networks: An Analysis of Mediation and Network Formation

MSc Thesis *(Afstudeerscriptie)*

written by

Luca van der Kamp
(born June 24th, 1995 in Amsterdam, The Netherlands)

under the supervision of **Prof. Dr. Sonja Smets** and **Dr. Fernando R. Velázquez Quesada**, and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:
December 7th, 2021

Members of the Thesis Committee:
Dr. Malvin Gattinger
Dr. Ronald de Haan
Dr. Christian Schaffner (*chair*)
Prof. Dr. Sonja Smets
Dr. Fernando R. Velázquez Quesada

Institute for Logic, Language and Computation

Abstract

Recent developments in Dynamic Epistemic Logics of social networks formalise the relation between epistemics and communication restricted by the existence of connections between people situated in a social network. There are two key aspects to social networks: communication over a social networks is often mediated, the ability of a groups to communicate to one another is determined by the existence, or non-existence of crucial positions; and social networks are dynamic, they change form and shape as people form and lose friendships and relations. Logical studies on epistemics and networks almost never treat these two aspects of social networks to its fullest. We set out to analyse these two aspects and their relation to epistemics from a logical perspective. We formalise a dynamic epistemic logic of full and semi-public communication *over* a social network, Communication Logic, and provide its sound and complete axiomatisation. Through Communication Logic, we identify and study crucial positions in social networks that either enable or block the flow of knowledge between groups. Moreover, we construct a game-theoretic framework of network formation and change that treats network formation as something driven by the people inside the social network itself. We formulate axioms that capture properties of such network formation games, and identify properties of network formation implicit in most socio-economic studies of network formation. Finally, we sketch how to use coalition logic to talk about coalitional ability in these network formation games, and hint towards a unified logic of full communication over, and coalitional ability of a social network.

Acknowledgements

I would like to thank my supervisors Fernando and Sonja for their astounding help, support, feedback, and insights. Writing a thesis during a pandemic and its lock-downs was a challenge, and I could not have hoped for a better and more kindhearted supervision.

Contents

Chapter 1

Introduction

Today, communication typically does not occur at a public forum, where someone talks and everybody listens; rather it takes place in and between groups. The first is a mode of (public) announcement, the second of private conversation. Private conversation is distinct from announcement, not only by the temporal and spatial separation of conversation, but also by its social separation. The social (dis)connection of people determines the scope and the reach of private conversation. Once the general mode of communication is dominated by private conversation, communication, in turn, is dominated by the social network. The ability for people to share their knowledge with each-other is limited by the structure of the social network — not only by who they are in direct relations with, but also who their relations are related with and so on.

There are two fundamental, and equally important aspects of social networks that determine the extent of communication: crucial positions, concepts such as bridges, connectors, hubs; and network dynamics, its formation and its change. The former determines, in large, the propagation of information over the network. The latter directly determines who can communicate to whom, and controls the formation, existence, and placement of the crucial positions. We will embark on a study of the logic of these two aspects of social networks.

We attempt to formalise the logic of communication over a social network. In its simple form, such communication is the exchange of knowledge between individuals, restricted by the existence of social relations or communication channels. In its developed form, this simple exchange implies a dynamic of knowledge exchange across a network. The separation between these two forms is exactly the social network: a mediation of the direct exchange, where, instead of two persons sharing with each other what they know, their knowledge is exchanged via other people. We analyse this process of mediated communication — the diffusion of knowledge *over* a social network.

With mediation comes dependency. Groups depend on external parties for communication. We want to analyse these mediators of different epistemic processes of communication. What are their network-structural preconditions, and in what ways do they relate to different aspects of knowledge. Finally, we set out to analyse network formation, ultimately to analyse how mediation behaves in a dynamic network.

The interplay between communication, social networks, and epistemics has been the subject of studies of epistemic social network logics [17; 24; 27; 56; 58]. Crucial positions have been studied from a more graph-theoretic, often quantitative, perspective [28; 38; 43], or a sociological perspective, such as Gatekeeping Theory [18; 19]. Network dynamics are often treated economically [21; 32; 38; 39; 51] or socially [23]. More relevant to us, logics exist that combine network dynamics with epistemics and communication [52; 53; 55].

A study on epistemics, communication, and mediation requires a theory with certain qualities: it must have a conception of knowledge and communication; it must treat the social network as a "first-order" object of the theory; and its logic must be able to conceptualise intricate network formation dynamics. To our knowledge, there is no logic (or combination of logics) that meets all these requirements. If studies of mediating positions in social networks treat some form of diffusion over a network, they treat the diffusion of information, properties, or opinions, not of knowledge. With respects to studies of dynamic networks, dynamic epistemic logics with a dynamic social network exist, but they focus on a different side of such dynamics: they give a *descriptive* study of network dynamic, with logics that indicates what would (possibly or necessarily) happen after the network is changed in a certain, often simple, ways.

In this thesis, we depart to construct a logic, partly from the aforementioned studies and their logics, that pertains to all the required qualities: we study mediating positions of *epistemic* processes of diffusion, and we formulate a network dynamic that treats network formation as something driven by the people inside the social network itself. Instead of possibility we focus on *ability*. However, formalising a total theory of epistemics, communication, and network change is a large project — too large for a single thesis. Network formation, for example, is worthy of a thesis of its own. We therefore won't construct this theory in its entirety. Instead, we develop the groundwork to establish such a theory.

The topics of this thesis have a conceptual separation, that of communication and network formation, whose parts have a formalisation inherent to two different theories. The first is the material of Dynamic Epistemic Logic. The second is much more related to game theory. Logic and game theory are by no means separated — modern dynamic logics of game theory can capture many game theoretical concepts, and game theories of logic sometimes give a better understanding of logical concepts (see van Benthem [60]). Still, the two parts of this thesis are separated in their relation to logic: one, the epistemic, works *from* a logic; the another, the game-theoretic, works *towards* a logic. Hence, we divide this thesis in two parts that each develop their own theory. If so desired, these parts can be read separately. But they are meant to be the groundwork for a unified theory, and are therefore fundamentally interlinked. The communication part consists of Chapter 2 and 3; the network formation part of Chapter 4.

The outline of this thesis follows the conceptual development of the topics sketched earlier. In Chapter 2, we construct a logic of full communication: *Communication Logic*. This logic is like resolution logic [3] and the logic of semi-public events [16], but it contains a social network that dictates the reach of the knowledge resolution. We provide a sound and complete axiomatisation of Communication Logic based on the reduction technique often employed in dynamic epistemic logics. We develop Communication Logic alongside an analysis of iterated communication and knowledge diffusion. Finally, we treat knowledge resolution and its relation to full communication over a social network, and analyse some of its more non-intuitive properties with respects to the communication of formulas.

In Chapter 3, we explore the relation between mediating positions, the network structure, and knowledge resolution. In parts, this analysis runs parallel to the analysis of such positions from a more sociological perspective: Gatekeeping Theory, and its formalisation in the Master's Thesis of Belardinelli [19]. On the other hand, our analysis diverges from the social sciences in key aspects. The social sciences almost never treat communication as an epistemic process in its full sense. Instead, we depart our analysis from epistemic logic, and its form of communication as distributed knowledge resolution, constructed in Chapter 2. Furthermore, many social studies conceptualise crucial network positions as categories of the network in itself, whereas we treat these positions as mediators of communication and the epistemic process of knowledge realisation.

In Chapter 4, we develop network dynamics from the perspective of game theory, also using ideas from social choice theory. We conceptualise single-shot network formation games, and formulate "axioms" that capture properties of such formation. Then, we expand the single-shot games to an extensive form able to treat dynamics of social influence. In the last section, we sketch a coalition logic that formalises ability of network formation games.

We conclude our work in Chapter 5, where we tentatively hint towards a unification of the two parts of this thesis, and the many questions that arise from it.

1.1 Preliminaries

First, a preliminary section, where we treat epistemic logics. We assume that the reader is familiar with Kripke semantics for modal logic. See Blackburn, de Rijke, and Venema [20] otherwise.

The most popular approach to epistemic logics is a modal logic with *possible world semantics*: uncertainty about the state of the world is represented via indistinguishability relations (or similarity relations). Developed most prominently in Hintikka [36] among others.

In Epistemic Logic, knowledge of each agent $a \in A$ is expressed as a modality K_a over the indistinguishability relation $\sim_a$. Fix a set of atomic propositions Prop, and a non-empty and finite set of agents A.

Definition 1.1.1 (Syntax of $\mathcal{L}_{EL}$)**.** The language of epistemic logic $\mathcal{L}_{EL}$ is generated from the following Backus-Naur Form (BNF):

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi$$

with $p \in \text{Prop}$, and $a \in A$.

Material implication ($\rightarrow$), material equivalence ($\leftrightarrow$) disjunction ($\vee$), $\bot$, and $\top$ are defined in the standard way. $K_a\varphi$ is read as "a knows that φ". The language is interpreted on *epistemic models*.

Definition 1.1.2 (Epistemic Model)**.** An *epistemic frame* is a tuple $\mathfrak{F} := (W, (\sim_a)_{a\in A})$. where W is a set of worlds, and for any agent $a \in A$ $\sim_a$ is the indistinguishability relation over W. An *epistemic model* is an epistemic frame equipped with a valuation function: $\mathfrak{M} := (\mathfrak{F}, V)$ where $V : \text{Prop} \to \mathscr{P}(W)$ is the valuation function assigning a set of worlds to each proposition.

The knowledge modality K_a is interpreted as a standard $\Box$-modality over $\sim_a$.

Definition 1.1.3 (Semantics of Epistemic Logic)**.**

$$\begin{array}{lll}
\mathfrak{M}, w \Vdash p & \text{iff} & w \in V(p) \\
\mathfrak{M}, w \Vdash \neg\varphi & \text{iff} & \mathfrak{M}, w \nVdash \varphi \\
\mathfrak{M}, w \Vdash \varphi \wedge \psi & \text{iff} & \mathfrak{M}, w \Vdash \varphi \text{ and } \mathfrak{M}, w \Vdash \psi \\
\mathfrak{M}, w \Vdash K_a\varphi & \text{iff} & \text{if } w \sim_a v \text{ then } \mathfrak{M}, v \Vdash \varphi
\end{array}$$

There are many modalities for group knowledge. The one we will use in this thesis is *distributed knowledge*. For $G \subseteq A$, $D_G\varphi$ reads as "it is distributively known among G that φ". For $G \subseteq A$ let $\sim_G := \bigcap_{a\in G} \sim_a$. D_G is the $\Box$-modality over the relation $\sim_G$.

$$\mathfrak{M}, w \Vdash D_G\varphi \quad \text{iff} \quad \text{if } w \sim_G v \text{ then } \mathfrak{M}, v \Vdash \varphi$$

Distributed knowledge among G represents the knowledge that follows from the combined information of all agents in G.

Chapter 2

Communication Logic

In this chapter we will introduce a logic for reasoning about communication and the propagation of knowledge in a social network. We construct a dynamic epistemic logic with a modality for full communication. We provide an axiomatisation for the logic and investigate some of its properties.

2.1 Social Network

The premise of this thesis is communication in a social setting. A social setting must be understood as a setting dominated by a myriad of relations: friendship, like-mindedness, kinship, common interest, frequency of interaction (in virology for example), etc. Despite the myriad of relations, we will only consider a single relationship, albeit one that is influenced by the many other relationships: the relationship of knowledge acquisition. We understand the social network as relations of knowledge sources and flows: a has a social tie to b when a "listens" to b and takes what b says as knowledge. It is important to be particular about what is meant with a social relation from the start, since the type of relation dictates its relational properties.

2.1.1 Symmetry

In many logical studies on social networks [25; 49; 55–58], the social network is assumed to be symmetric. This makes sense; friendships (at least healthy ones) come from both sides, the interaction between both parties of the relation is symmetric, and both parties face each others as equals. Relations of like-mindedness, kinship, and common interest are also symmetric because these are comparative notions. However, in the setting of knowledge propagation and communication, relations are not necessarily symmetric. Sources such as the newspaper, television, and personal websites lack most forms of mutual contact. And even in social relations that are bidirectional, information flow sometimes is not: even if a student and a teacher are engaged in a reciprocal friendship, the flow of information between a student and a teacher in the setting of a lecture is most often from teacher to student. Lastly, even in conversation, the information flow could be unidirectional: a party of the conversation could simply not be listening, or someone might be sceptical of whatever someone else says. This is why we don't require the social network to be symmetric, like many other logics that study social relation in a communicational or epistemic setting [17; 24; 45; 52–54; 71].

2.1.2 Reflexivity

Although it does make sense to state that agents epistemically follow or listen to themselves, we do not assume reflexivity of the social network. The method of obtaining knowledge from oneself is distinct from that of obtaining knowledge from others (the relation regarded here). Even though we won't touch upon the former, this distinction is significant enough to demand a separate and particular treatment of the latter: as a relation that is not necessarily reflexive. Such a treatment is of importance in, for example, a setting of unreliable information channels, where information that is communicated over a channel is not guaranteed to reach its destination. It separates information whose source is another agent, from information whose source is oneself. The former would be susceptible to information loss, while the latter would not (at least for the ideal agents that epistemic logic typically deals with).

We also will not assume irreflexivity. Agents are allowed to follow themselves just as they would any other agent. Exemplary to the use of reflexivity is the opposite of above example, where communication is reliable, but an agent's recall is not. Then, agents have all the reason to follow themselves (e.g. read their own notes).

2.1.3 Formal Definitions

A social network (A, F) is a graph consisting of a finite set of agents (the vertices), A, and a relation over these agents (the edges), F. Throughout this thesis, we fix a finite and non-empty set of agents A. For ease of notation, define the *follow function* of F, $\mathcal{F} : A \to \mathscr{P}(A)$, such that $\mathcal{F}(a)$ denotes the set of agents that a follows: $\mathcal{F}(a) := \{x \mid a\ F\ x\}$. Also define $\mathcal{F}^+(a) := \mathcal{F}(a) \cup \{a\}$, the set of agents that a follows, including a itself. We extend these function to groups: for any $G \subseteq A$, $\mathcal{F}(G) := \bigcup_{a \in G} \mathcal{F}(a)$, and $\mathcal{F}^+(G) := \bigcup_{a \in G} \mathcal{F}^+(a)$.[1] Note that, throughout this thesis, G is used to denote a subset of A instead of the customary usage of G in graph theory as denoting a graph.

For $G \subseteq A$, we define the *restriction* of the follow function to a group $G \subseteq A$ as $\mathcal{F}|_G(H) := \mathcal{F}(H) \cap G$ (for any $H \subseteq A$); in this way, $\mathcal{F}|_G(H)$ contains the elements of $\mathcal{F}(H)$ that are also in G. Similarly, we define the inclusive restriction of the follow function to a group $G \subseteq A$ as $\mathcal{F}|^+_G(H) := \mathcal{F}|_G(H) \cup H$; in this way, $\mathcal{F}|^+_G(H)$ contains the elements of $\mathcal{F}(H)$ that are also in G *and* all the elements of H itself.

When the social network relation F is clear from the context we will not specify to what social relation the follow function belongs.

Definition 2.1.1 (Walks and Paths)**.** For $n \geq 1$, we call a sequence $P = (p_i)_{i=0}^n$ in A a *walk* from p_0 to p_n iff $p_0\ F \cdots F\ p_n$. For $G \subseteq A$, if $p_1, p_2, \ldots, p_{n-1} \in G$ then we say that P is a walk *in* G, G-walk for short, from p_0 to p_n. Note that p_0 and p_n do not have to be elements of G for P to be a G-walk. If furthermore $p_0, p_1, \ldots, p_n$ are all distinct, then we call P a path or G-path. The *length* of a walk or path is the number of edges in it, n for $P = (p_i)_{i=0}^n$. We will denote that there is a G-path from x to y by $x \to_G y$, and denote that there is a G-path of length n from x to y by $x \to_G^n y$. Finally, we will use the notation $x \to_G^{\leqslant n} y$ to denote that there is a G-path of length at most n from x to y, and $x \not\to_G^{\leqslant n}$ if there is no G-path from x to y of length at most n.

Note that we do not consider sequences of length 0, e.g. (a), as walks or paths. This is a stylistic choice, that is of use in the next chapter. As a result, there only is a path from a to itself when $a\ F\ a$. It allows us to avoid requiring identity in our language to express path-existence between two agents.

[1] Note that $\mathcal{F}^+(G) = \mathcal{F}(G) \cup G$.

2.2 Communication: Sharing All You Know

With the social network established, we will now define a logic for reasoning about knowledge and communication. In any setting about communication of knowledge between agents, at least three things have to be discussed: the object, the extent, and the reach, of communication.

With the *object of communication* we mean *what* is communicated between the agents. In most epistemic logics with communication, the objects of communication are propositions or formulas of a specific language. However, along with languages such as that described in Ågotnes and Wáng [3] and Baltag and Smets [16], we will take indistinguishability itself as the object of communication: agents communicate their knowledge about the (in)distinguishability of worlds.

With the *extent of communication* we mean *how much* an agent discloses with a communication act. Here two approaches are often taken: either an agent discloses their knowledge about a certain proposition or formula, as the $[a : \theta]$-modality in Xiong et al. [71], the $[i : \varphi]$-modality in Ruan and Thielscher [53], the $[F!\varphi]$-modality in Seligman, Liu, and Girard [56], and the *PDL*-modality $[\mathrm{send}_{Fa}(\psi)]$ in Seligman, Liu, and Girard [55]; or an agent discloses all they know [3; 11, the $!a$-system; 16; 24, ch. 4].[2] The interpretation of "all they know" depends on the object of communication. Because we discuss the communication of similarity relations, we will interpret "everything an agent knows" as meaning the information contained in the entire similarity relation of that agent. This as opposed to the common interpretation of "every formula the agent knows". This distinction is important because, as discussed in van der Hoek, van Linder, and Meyer [63], if communication is understood as communicating all formulas one knows, then distributed knowledge of a proposition does not always result in group knowledge of that proposition after full communication.[3]

With the *reach of communication* we mean both the reach of the object of communication and the reach of the knowledge of whether communication has taken place.[4] We will restrict the former by the social network structure, and assume that the latter reaches the entire social network — whether communication has taken place is publicly known. As such, lending the term from Baltag and Smets [16], our approach to communication is *semi-public*: all agents know that communication has taken place, they know who communicated to whom, they know that these agents communicated everything they knew, but whether they get to know the contents of the communication depends on the social network; only agents that follow the communicating agent get to know what is communicated. What all of this entails will become clear in some later examples.

2.2.1 Internal Communication

In line with *Dynamic Epistemic Logic* (*DEL*), we will model communication as a modality that transforms the epistemic model to one that reflects that communication has taken place. It is common to take an *external* approach to such communication modalities: logics such as Public Announcement Logic [14] have modalities for announcements by an external force. We will take an *internal* approach: communication is an act between agents inside the social network. Who communicates to who is determined by the social relationship between agents.

[2] Other approach: sharing all about certain issues, in [12; 24, ch. 5].

[3] An example of a situation in which group knowledge and knowledge after communication of propositions does not align is given in Example 2.3.4.

[4] Of course, this is not all there is to communication. One could go into much more detail about the reach of communication, such as what is allowed by the general *action models* (see Baltag, Moss, and Solecki [14] and Baltag and Moss [13]) of Dynamic Epistemic Logic. These action models make it possible to specify in much more detail the preconditions of sharing events, as well as the uncertainty that the agents have about them. Such events allow for, for example, the formalization of higher order properties of communication such as who should know that who knows that communication has taken place.

Such an approach is also taken in Baltag and Smets [16]. However, whereas they explicitly specify who communicates to whom in each communication update, this is implicit in our communication modality: who communicates to whom is determined by the underlying social network.

2.3 Communication Logic

Now we will construct a simple Dynamic Epistemic Logic for full communication: Communication Logic (CL). This logic has a distributed knowledge modality as well as a modality for *full communication* that follows the interpretation discussed in the last section. It allows us to reason about the situation after a group of agents communicate everything they know. We thereby abstract away from the exact details of the communication — from what exactly is said — and only model exchange of information in a perfect[5] setting. After exhaustive communication, what one is left with is what is called the *communication core* in van Benthem [61, p. 249], limited to the bounds of communication by the social network.

This has many interpretations; for example as agents that read some database or entity containing all information an agent possesses, as in Baltag and Smets [16], or as the knowledge *potential* of agents after allowing exhaustive (unlimited) communication to take place between other agents without going into detail about the exact conversations or sentences uttered. Potential is an important detail here, since, as proven in van Benthem [62], when communication is thought of as announcement of formulas (sentences) in a certain order, this order matters for the resulting knowledge state, and limits its ability to reach the “communication core”. Full communication abstracts away from such details.

2.3.1 Syntax

The language of Communication Logic $\mathcal{L}_{CL}$ consists of propositions, Boolean connectives, and distributed knowledge modalities. Additionally, it contains communication modalities for each $G \subseteq A$, $[!G]$, to be read as “if all agents in G communicate all they know to the agents that follow them”.

The syntax of Communication Logic also includes a representation of the social relation (this representation will be required for the axiomatisation). There are many approaches to syntactically embedding such information about a network structure in a logic. The simplest approach, taken in Baltag et al. [17], Carrington [24], Smets and Velázquez-Quesada [58], Roelofsen [52, ch. 3][6], Christoff and Rendsvig [27], and Ruan and Thielscher [53], is to embed the social network structure in the language as *propositions* — one for each possible edge in the network. The truth value of these propositions reflect the existence of their respective social relation.

A more complex (albeit more expressive, in the case of infinite networks) approach is taken in Seligman, Liu, and Girard [56], Seligman, Liu, and Girard [55], and Sano and Tojo [54]. Here they introduce a Kripke modality over the social relation reading ‘some of my friends’ or ‘all of my friends’. This requires formulas to be evaluated on agents-world pairs, as well as elements of hybrid logic [5] for expressing the existence of relations between specific agents in a network.

A first-order approach to embedding the social network in the syntax is taken in Liberman and Rendsvig [40], where the social relation is represented by a first-order predicate.

And finally, there are logics that do not embed the social relation in the syntax at all. Rather they specify the existence of a social relation by other means. For example, in the logic in Xiong et al. [71], the social following relation is expressed via belief and a

[5] Perfect in the sense that agents are able to communicate all they know, regardless of the imperfections of communication through a (formal) language.

[6] Here, communication channels are considered between sets of agents. As such, the propositions state the existence of relations between sets of agents.

communication modality (tweeting modality). A problem with such systems is that, most often, no formula that exactly defines the social relation exists. This is related to the effect that Xiong et al. [71] calls *ghost followers*: an agent might believe or know exactly what another agent does, no matter how many agents communicate how many times, even though that agent other does not socially follow the other agent.

For simplicity's sake, we will take the propositional approach. The language has propositions, $F_{a,b}$ for $a, b \in A$, that denote whether a follows b. $F_{a,b}$ is to be read as "a (epistemically) follows b".

We define the language of Communication Logic, $\mathcal{L}_{CL}$, inductively. Let Prop be a countable set of propositionals.

Definition 2.3.1 (Syntax of $\mathcal{L}_{CL}$)**.** The language $\mathcal{L}_{CL}$ is generated from the following Backus-Naur Form (BNF):

$$\varphi ::= p \mid \varphi \wedge \varphi \mid \neg\varphi \mid [!G]\varphi \mid D_G\varphi \mid F_{a,b}$$

with $p \in \text{Prop}$, $a, b \in A$, and $G \subseteq A$. Material implication ($\rightarrow$), material equivalence ($\leftrightarrow$) disjunction ($\vee$), $\bot$, and $\top$ are defined in the standard way.

D_G is the distributed knowledge modality, to be read in the standard way. Note that our language does not contain the knowledge modality K. Knowledge of an individual agent i is represented by $D_{\{i\}}$. For brevity's sake, let K_i be shorthand for $D_{\{i\}}$. There are other modalities for group knowledge, of which common knowledge is probably the one whose absence the reader has noticed. For our purposes however, distributed knowledge suffices. We do however discuss one other notion of group knowledge, that we define in terms of individual knowledge. Let $E_G\varphi$ denote that everyone in G knows φ:

Definition 2.3.2 (Everyone knows)**.** For any $G \subseteq A$, let:

$$E_G := \bigwedge_{a \in G} K_a.$$

2.3.2 Semantics

Formulas of the language are evaluated on *communication frames and models*: epistemic frames and models augmented with a social network.

Definition 2.3.3 (Communication frames and models)**.** A *communication frame* is a tuple $\mathfrak{F} = (W, F, (\sim_a)_{a \in A})$, where W is a non-empty set of possible worlds. $F \subseteq A \times A$ is the social *follow* relation. The follow relation F is not assumed to be reflexive or symmetric. $a\ F\ b$ intuitively reads as "a *follows* b" or "a listens to (receives) what b says". For each $a \in A$, $\sim_a \subseteq (W \times W)$ is the indistinguishability relation for a over all possible worlds W. We take the epistemic part of the model to be $S5$, this means that the indistinguishability relation is *symmetric*, *reflexive*, and *transitive*; i.e. $\sim_a$ is an equivalence relation. For any $G \subseteq A$, define $\sim_G := \bigcap_{a \in G} \sim_a$.

A *communication model* $\mathfrak{M} = (\mathfrak{F}, V)$ is a communication frame $\mathfrak{F}$ with a valuation function $V : \text{Prop} \rightarrow \mathscr{P}(W)$.

We assume that there is only one social network, i.e. that the social network is the same in all possible worlds, and thereby that the structure of the network is common knowledge.

We define the semantics of the communication modality using a *model update* that represents full communication of a set of agents $G \subseteq A$, using the interpretation we discussed in the previous section. In the updated model, all agents acquire the knowledge of the agents in G that they follow. After the communication update it is public knowledge that communication took place. From the perspective of a single agent a, such semi-public events can be represented by a restriction of their indistinguishability relation $\sim_a$ to the indistinguishability relations of the agents they follow which are part of G: $\mathcal{F}|_G(a)$. Moreover, we assume that in a communication update by G, when the agents in G communicate, all the other agents do nothing but receive information. We also assume that agents have *perfect memory*: any certainty about the state of the world they had before communication remains after a communication update. Since our social network is not assumed to be reflexive, we must therefore restrict the indistinguishability relation $\sim_a$ to that of $\sim_{\mathcal{F}|^+_G(a)} = \sim_{\mathcal{F}|_G(a)} \cap \sim_a$.

Definition 2.3.4 (Communication model updates)**.** For any $G \subseteq A$, a *communication (model) update* $!G$ is a function that transforms a model into one that reflects the situation after which all agents in G shared all they know with their followers, and all agents know that G did this. We denote the application of $!G$ to $\mathfrak{M}$ with $\mathfrak{M}^{!G}$. Define $!G$ as follows: $\mathfrak{M}^{!G} := (W, F, (\sim_a^{!G})_{a \in A}, V)$, where for any $i \in A$:[7] $\sim_i^{!G} := \bigcap_{j \in \mathcal{F}|^+_G(i)} \sim_j = \sim_{\mathcal{F}|^+_G(i)}$. For convenience, we write $!a$ for $!\{a\}$, for any $a \in A$.

Note that communication model updates are well-behaved, in that for any $a \in A$, $G \subseteq A$, and indistinguishability relation $\sim_a$: $\sim_a^{!G}$ is also an indistinguishability relation. This holds trivially since $\sim_a^{!G}$ is an intersection of equivalence relations over W, and equivalence relations are closed under intersection.

Formulas of communication logic are evaluated on worlds.

Definition 2.3.5 (Semantics)**.** The semantics of *CL* are as follows;

$\mathfrak{M}, w \Vdash p$	iff	$w \in V(p)$
$\mathfrak{M}, w \Vdash \neg\varphi$	iff	$\mathfrak{M}, w \nVdash \varphi$
$\mathfrak{M}, w \Vdash \varphi \wedge \psi$	iff	$\mathfrak{M}, w \Vdash \varphi$ and $\mathfrak{M}, w \Vdash \psi$
$\mathfrak{M}, w \Vdash D_G\varphi$	iff	if $w \sim_G v$ then $\mathfrak{M}, v \Vdash \varphi$
$\mathfrak{M}, w \Vdash [!G]\varphi$	iff	$\mathfrak{M}^{!G}, w \Vdash \varphi$
$\mathfrak{M}, w \Vdash F_{a,b}$	iff	$a\ F\ b$

We write $\mathfrak{M} \Vdash \varphi$ iff for all $w \in W$: $\mathfrak{M}, w \Vdash \varphi$; $\mathfrak{F}, w \Vdash \varphi$ iff for all valuations V: $(\mathfrak{F}, V), w \Vdash \varphi$; $\mathfrak{F} \Vdash \varphi$ iff for all $w \in W$: $\mathfrak{F}, w \Vdash \varphi$; and $\Vdash \varphi$ iff for all frames $\mathfrak{F}$: $\mathfrak{F} \Vdash \varphi$, i.e. iff φ is a validity.

For a set of formulas $\Gamma \subseteq \mathcal{L}_{CL}$, we write $\mathfrak{M}, w \Vdash \Gamma$ iff all formulas in Γ are true in w: $\forall \psi \in \Gamma$, $\mathfrak{M}, w \Vdash \psi$. The notation $\mathfrak{F}, w \Vdash \Gamma$, $\mathfrak{F} \Vdash \Gamma$, and $\Vdash \Gamma$ are defined like their single formula equivalent.

Finally, for a set of formulas $\Gamma \subseteq \mathcal{L}_{CL}$, and a formula $\varphi \in \mathcal{L}_{CL}$, we say that φ is a *semantic consequence* of Γ, notation $\Gamma \Vdash \varphi$, iff for all models $\mathfrak{M}$ and all worlds w: if $\mathfrak{M}, w \Vdash \Gamma$, then $\mathfrak{M}, w \Vdash \varphi$.

Now we will regard some examples to clarify the exact implications of full communication with (in)distinguishability as the object of communication. In all these examples, the circles represent possible worlds, with their names displayed above. Inside the circle we will write whether a proposition is true (p) or false ($\neg p$) in the world. The real world (the world we evaluate in) is denoted by a double circle. The indistinguishability relation is represented by a squiggly line. We will omit the reflexive relations and display the relation without any indication of direction due to its symmetry.

Social networks (and fragments thereof) have a slightly different presentation. We denote agent (vertices) by squares. The names of the agents are displayed inside the square. An arrow from a to b denotes that a follows b. All agents not displayed in such figures are assumed to have *no* social connections.

[7] Note the use of $\mathcal{F}|^+_G$ instead of $\mathcal{F}|_G$. This reflects the perfect memory assumption made earlier: agents always have access to their own knowledge when forming a new epistemic relation from what they have learned of the communicating agents. This is akin to the assumption that agents have access to their own database in Baltag and Smets [16].

Example 2.3.1. Consider three agents a, b, and c, and their social relations depicted below.

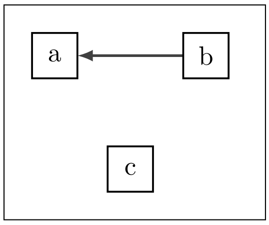

Figure 2.1: The social environment of a, b, and c.

Let $W = \{w, v\}$ and $V(p) = \{w\}$. Let the similarity relations be as depicted below in Figure 2.2 on the left.

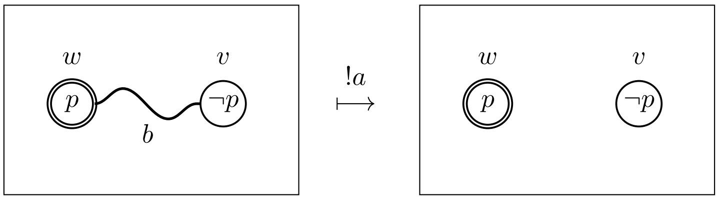

Figure 2.2: $(\sim_a)_{a\in A}$ and $(\sim_a^{!a})_{a\in A}$.

a and c know p and know from each other that they know p. Furthermore, b does not know p, but does know that a and c know *whether* p. After a communicates (i.e. in $\mathfrak{M}^{!a}$) the similarity relation is as shown on the right of Figure 2.2:

Before the communication update, c's similarity relation was such that c knew that a knew that p, hence after b's similarity relation is restricted to that of a in $\mathfrak{M}^{!G}$, c knows that b knows p. This is one example of how, under full communication updates, the *event* of communication is public — everybody knows that communication has taken place, and that everybody communicated everything they know.

Example 2.3.2. Consider a similar situation as before, but let the similarity relation be as shown in Figure 2.3 below on the left.

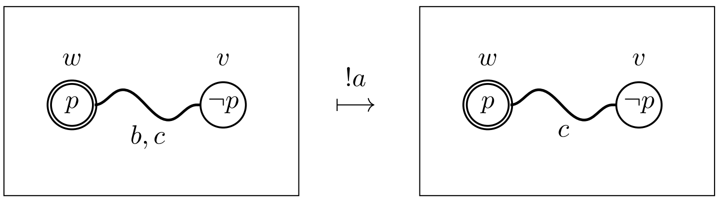

Figure 2.3: $(\sim_a)_{a\in A}$ and $(\sim_a^{!a})_{a\in A}$.

Now, c, like b, only knows that a knows *whether* p — $\mathfrak{M}, w \Vdash K_c(K_a p \vee K_a \neg p)$ but $\mathfrak{M}, w \nVdash K_c(K_a p) \vee K_c(K_a \neg p)$. After a communicates, that is in $\mathfrak{M}^{!a}$, the similarity relation is as shown on the right of Figure 2.3. As in the example before, b knows that p. However, now c only knows that b knows whether p. The content of the communication is *not* public. It is only known to the agents that follow a. The other agents only know that a shared everything they knew to the agents that follow a.

Example 2.3.3. A similar phenomenon occurs when c does not know that a and b know whether p. Consider the situation with the same social network as before, and a similarity relation as depicted in Figure 2.4 below on the left.

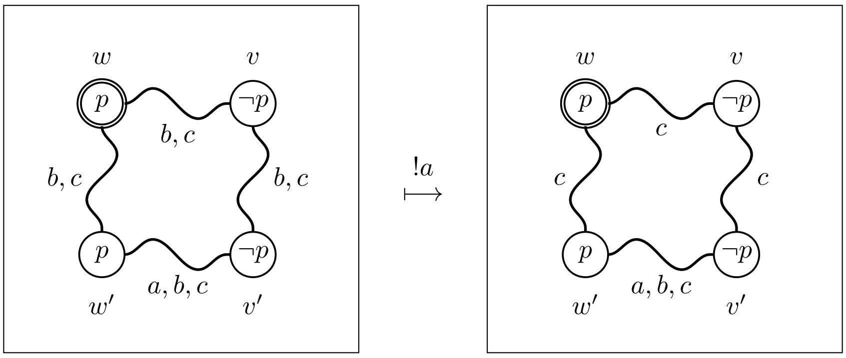

Figure 2.4: $(\sim_a)_{a\in A}$ and $(\sim_a^{!a})_{a\in A}$.

a knows that p, but both b and c don't know that a knows that p. After a communicates, that is in $\mathfrak{M}^{!a}$, the similarity relation is as shown on the right of Figure 2.4. Again, b knows that p. But c does not know that a or b know whether p.

Example 2.3.4. Finally, we consider an example that clarifies the difference between indistinguishability of worlds and knowledge about propositions as object of communication, from the perspective of distributed knowledge resolution. This example is based on an example in van der Hoek, van Linder, and Meyer [63]. Consider two agents, a and b, and take a model $\mathfrak{M}$ where $a\ F\ b$ and $b\ F\ a$.

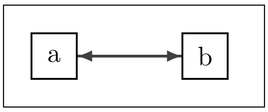

Figure 2.5: The social relations of a and b.

Let $(\sim_a)_{a\in A}$ of $\mathfrak{M}$ be as depicted below.

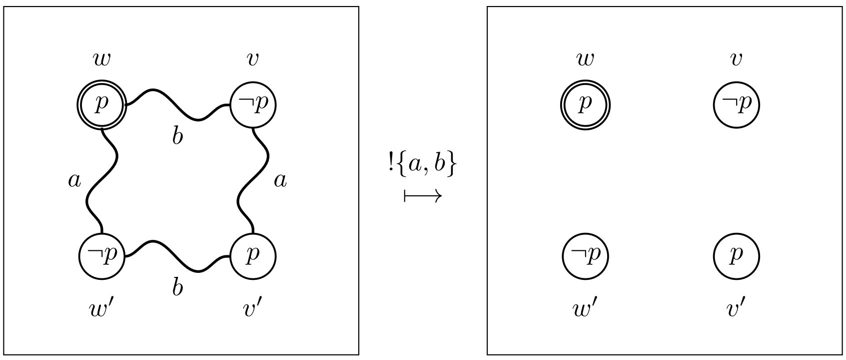

Figure 2.6: $(\sim_a)_{a\in A}$ and $(\sim_a^{!\{a,b\}})_{a\in A}$.

In $\mathfrak{M}$, $\{a,b\}$ distributively knows p. However, neither a nor b knows that p. As such, if propositions are the object of communication, and therefore the prerequisite for communicating about p is knowledge of p, then distributed knowledge of p is not obtainable through communication by either a or b. However, as we take similarity relations as the object of communication, we do have that $\mathfrak{M}^{!\{a,b\}}, w \Vdash K_a p \wedge K_b p$.

	$\mathbf{S5_n^D} + \texttt{KF}$		
(A1)	All tautologies of propositional calculus		
($\texttt{K}^\texttt{D}$)	$D_G\varphi \wedge D_G(\varphi \to \psi) \to D_G\psi$	$G \subseteq A$	(Distribution Axiom)
($\texttt{T}^\texttt{D}$)	$D_G\varphi \to \varphi$	$G \subseteq A$	(Knowledge Axiom/Veridicality)
($\texttt{4}^\texttt{D}$)	$D_G\varphi \to D_G D_G\varphi$	$G \subseteq A$	(Positive Introspection)
($\texttt{5}^\texttt{D}$)	$\neg D_G\varphi \to D_G\neg D_G\varphi$	$G \subseteq A$	(Negative Introspection)
(D2)	$D_G\varphi \to D_{G'}\varphi$ if $G \subseteq G'$	$G \subseteq A$	(Monotonicity)
(R1)	$\dfrac{\varphi, \varphi \to \psi}{\psi}$		(Modus Ponens)
(R2)	$\dfrac{\varphi}{D_G\varphi}$	$G \subseteq A$	(Knowledge Generalisation)
(KF)	$F_{a,b} \to E_A F_{a,b}$	$a, b \in A$	(Commonly Known Network)

Table 2.1: Axiom system for of CL^-

2.3.3 Axiomatisation

We will provide a sound and strongly complete axiomatisation of *CL*. We will prove its completeness using the *reduction technique*[8] from the field of Dynamic Epistemic Logic: we define truth-preserving reduction axioms that push the communication modalities inwards, finally eliminating them altogether at the atomic level of the formulas. This reduces the completeness proof of our axiomatisation to the completeness proof of epistemic logic with distributed knowledge and social network propositions.

The base logic to which we reduce Communication Logic is its *non-dynamic fragment*: CL^-. The language of this fragment $\mathcal{L}_{CL^-}$ contains everything but communication modalities: propositions, boolean connectives, the distributed knowledge modalities, and the network propositions.

Communication models have an indistinguishability relation that is reflexive, transitive and symmetric. Thus, the axiomatisation of CL^-, and consequently *CL*, is built upon the standard axiomatisation of *S5* extended to include distributed knowledge (i.e. an intersection modality): $\mathbf{S5^D}$. Such an axiomatisation is given in Fagin et al. [29] and Halpern and Moses [35] (here called $\mathbf{S_{5_n}^D}$).

These systems axiomatise the normal modal logic *S5*, extended to include the intersection modality (of distributed knowledge) up to n agents, $S5^D$. For this, they include an axiom schema connecting the (normal) knowledge modality to the distributed knowledge (intersection) modality ($K_i \leftrightarrow D_{\{i\}}$), as well as all standard axioms for *S5*: `K`, `D2`, `T`, `4`, `5`, and `N`. Since we don't bother ourselves with K_i, we can leave these out. The remaining axiom schemata of this system are included in Table 2.1.

We extend the logic of $S5^D$ with the set of network propositions. Recall that these propositions are distinct from Prop in that they are common knowledge in all models. For this, we employ the axiom `KF`: $F_{a,b} \to E_A F_{a,b}$. If we add `KF` to $\mathbf{S5_n^D}$ we come to a complete axiomatisation of CL^-: $\mathbf{S5_n^D} + \texttt{KF}$. This system is shown in Table 2.1. Note that $\texttt{4}^\texttt{D}$ is not necessary, as it is derivable from `A1`,$\texttt{K}^\texttt{D}$, $\texttt{T}^\texttt{D}$, $\texttt{5}^\texttt{D}$, `R1`, and `R2`.[9]

Strong completeness of this system for Communication Models can be proven using the standard technique of canonical model construction [20, ch. 4.2] with a "unraveling-folding" method [68], and can be based entirely on the canonical model construction in the completeness proof of $\mathbf{S5_n^D}$ given in Fagin et al. [29] and Halpern and Moses [35], or the

[8] See Wang and Cao [67] for a thorough overview of this technique and its uses.

[9] From `A1`, $\texttt{T}^\texttt{D}$, $\texttt{5}^\texttt{D}$, and `R1`: $\vdash D_G\varphi \to D_G\neg D_G\neg D_G\varphi$. Moreover, from $\texttt{5}^\texttt{D}$ and some propositional reasoning: $\vdash \neg D_G\neg D_G\varphi \to D_G\varphi$. The rest is derivable from $\texttt{K}^\texttt{D}$, `R1`, `R2` and some propositional reasoning.

simpler and more recent model construction in the strong completeness proof of Wáng and Ågotnes [69] (here the axiom system is called **Int**(**S5**)). Because such a construction is quite standard, we won't go into its details.

Do note that, with the addition of the `KF` axiom, it is ensured that if $F_{a,b}$ is true in a world, then it is common knowledge in that world (maximally consistent set) of the canonical model, as all instances of the `KF` axiom are included in all maximally CL^{-}-consistent sets,

The reduction of CL to the base logic CL^{-} requires an analysis of the interplay between the distributed knowledge modality and the communication modalities. We do this in much the same way as in Ågotnes and Wáng [3] and Baltag and Smets [16]. It will become clear that these dynamics come down to a change in the set of agents that the distributed knowledge modality operates on, like in Ågotnes and Wáng [3] and Baltag and Smets [16]. Only here, the particular change of this set depends on the social network, as what information the agents get depends on their follow relation. To show this, we first explore the relation between the pre-updated and post-updated similarity relation:

Lemma 2.3.1. *For any similarity relation $(\sim_a)_{a\in A}$, and any $G, H \subseteq A$: $w \sim_H^{!G} v$ iff $w \sim_{\mathcal{F}|_G^+(H)} v$.*

Proof.

$$\sim_H^{!G} = \bigcap_{i\in H} \sim_i^{!G} = \bigcap_{i\in H} \bigcap_{j\in\mathcal{F}|_G^+(i)} \sim_j = \bigcap_{j\in\bigcup_{i\in H}\mathcal{F}|_G^+(i)} \sim_j = \bigcap_{j\in\mathcal{F}|_G^+(H)} \sim_j = \sim_{\mathcal{F}|_G^+(H)}$$

□

The semantic result of Lemma 2.3.1 implies the following syntactic equivalences:

Theorem 2.3.2. *For any model $\mathfrak{M}$, any $G, H \subseteq A$, $\varphi \in \mathcal{L}_{CL}$, and $w \in W$:*

$$\mathfrak{M}, w \Vdash [!G]D_H\varphi \;\text{ iff }\; \mathfrak{M}, w \Vdash D_{\mathcal{F}|_G^+(H)}[!G]\varphi$$

Proof.

$$\begin{aligned}
\mathfrak{M}, w \Vdash [!G]D_H\varphi &\iff \mathfrak{M}^{!G}, w \Vdash D_H\varphi \\
&\iff \text{if } \; w \sim_H^{!G} v \; \text{ then } \; \mathfrak{M}^{!G}, v \Vdash \varphi \\
&\overset{2.3.1}{\iff} \text{if } \; w \sim_{\mathcal{F}|_G^+(H)} v \; \text{ then } \; \mathfrak{M}^{!G}, v \Vdash \varphi \\
&\iff \text{if } \; w \sim_{\mathcal{F}|_G^+(H)} v \; \text{ then } \; \mathfrak{M}, v \Vdash [!G]\varphi \\
&\iff \mathfrak{M}, w \Vdash D_{\mathcal{F}|_G^+(H)}[!G]\varphi
\end{aligned}$$

□

The equivalence of Theorem 2.3.2 provides us with a basis for the reduction of CL to CL^{-}. However, this equivalence isn't suitable as it is currently stated: the contents of $\mathcal{F}|_G^+(H)$ depends on the state of the social network. An axiom that works for the entire class of Communication Models, regardless of the social network of these models, must be stated independent of the social network. We will construct such an axiom from Theorem 2.3.2 in steps.

Given that we are dealing with **S5**, axiom `D2` holds: if a set of agents distributively knows φ, then any superset of that set will also distributively know φ. Therefore, $[!G]D_H\varphi$ holds iff there is a subset of G which is contained in H's follow sets, that, together with H, distributively knows that $[!G]\varphi$. Therefore, if we can describe "H is contained in the follow set of G" in terms of Communication Logic, then we can describe the validity we are after.

$\mathbf{C_n}$		
(!G $\mathtt{D_H}$)	$[!G]D_H\varphi \leftrightarrow \bigvee_{G'\subseteq G} (\mathrm{Fol}(H, G') \wedge D_{H\cup G'}[!G]\varphi)$	$G, H \subseteq A$
(!G NEG)	$[!G]\neg\varphi \leftrightarrow \neg[!G]\varphi$	$G \subseteq A$
(!G CON)	$[!G](\varphi \wedge \psi) \leftrightarrow [!G]\varphi \wedge [!G]\psi$	$G \subseteq A$
(!G ATOM)	$[!G]p \leftrightarrow p$	$G \subseteq A$
(RE)	$\dfrac{\varphi \leftrightarrow \chi}{\psi \leftrightarrow \psi[\varphi/\chi]}$	

Table 2.2: Axiom schemas for communication in $\mathcal{L}_{CL}$. Here, $p \in \mathrm{Prop} \cup \{F_{a,b} \mid a, b \in A\}$.

A set H is contained in the follow set of G, $H \subseteq \mathcal{F}(G)$, when every agent in H is followed by an agent in G; i.e. for all $h \in H$ there is a $g \in G$ such that $g\ F\ h$. Since A is finite, a conjunction of disjunctions over $F_{a,b}$ propositions suffices.

$$\mathrm{Fol}(G, H) := \begin{cases} \top & G = \emptyset \\ \bigwedge_{h\in H} \bigvee_{g\in G} F_{g,h} & \text{otherwise} \end{cases}$$

As we did with $\mathcal{F}(a)$, this formula defines an expression that extends the notion of following to groups. Such an extension entails that all the information G possesses is accessible through communication by H. We will discuss variations of these extensions in the next chapter, when we discuss connectors and directly connected sets.

Now, we have sufficient tools to formulate the validity.

Proposition 2.3.1. *For any $G \subseteq A$:*

$$\Vdash [!G]D_H\varphi \leftrightarrow \bigvee_{G'\subseteq G} (\mathrm{Fol}(H, G') \wedge D_{H\cup G'}[!G]\varphi)$$

Proof. Let M be an arbitrary model, and $w \in W$.

($\Rightarrow$) Note that for arbitrary $G, H \subseteq A$, by definition of $\mathcal{F}|_G(H)$ and Fol, $\mathrm{Fol}(H, \mathcal{F}|_G(H))$ is a validity. Recall that $\mathcal{F}|_G(H) \subseteq G$. Then:

$$\begin{aligned} & M, w \Vdash [!G]D_H\varphi \\ \iff & M, w \Vdash D_{\mathcal{F}|^+_G(H)}[!G]\varphi && \text{(Theorem 2.3.2)} \\ \iff & M, w \Vdash D_{H\cup\mathcal{F}|_G(H)}[!G]\varphi && \text{(Def. of } \mathcal{F}|^+_G) \\ \iff & M, w \Vdash \mathrm{Fol}(H, \mathcal{F}|_G(H)) \wedge D_{H\cup\mathcal{F}|_G(H)}[!G]\varphi \\ \implies & M, w \Vdash \bigvee_{G'\subseteq G} (\mathrm{Fol}(H, G') \wedge D_{H\cup G'}[!G]\varphi) \end{aligned}$$

($\Leftarrow$) Assume $M, w \Vdash \bigvee_{G'\subseteq G}(\mathrm{Fol}(H, G') \wedge D_{H\cup G'}[!G]\varphi)$. If there is an $G' \subseteq G$ s.t. $M, w \Vdash \mathrm{Fol}(H, G')$ and $M, w \Vdash D_{H\cup G'}[!G]\varphi$, then for any $g \in G'$ there exists a $h \in H$ such that $h\ F\ g$. This means that $g \in \mathcal{F}(H)$. Thus, $G' \subseteq \mathcal{F}(H)$. Since $G' \subseteq G$, $G' \subseteq \mathcal{F}|_G(H)$ and thus $H \cup G' \subseteq \mathcal{F}|^+_G(H)$. By soundness of D2, and since $M, w \Vdash D_{H\cup F}[!G]\varphi$, $M, w \Vdash D_{\mathcal{F}|^+_G(H)}[!G]\varphi$. By Theorem 2.3.2 $M, w \Vdash [!G]D_G\varphi$.

□

We can now phrase the reduction schemata for the communication modalities. These are shown in Table 2.2. As we will show, when read from left to right, Table 2.2 forms a reduction system for $\mathcal{L}_{CL}$ whose normal forms do not contain any communication update modalities.

Therefore, these reduction rules bring about a translation from formulas of $\mathcal{L}_{CL}$ to formulas of $\mathcal{L}_{CL^-}$. This translation is a function $\mathcal{T} : \mathcal{L}_{CL} \to \mathcal{L}_{CL^-}$. We recursively define this translation as follows:

Definition 2.3.6 ($\mathcal{L}_{CL^-}$ translation)**.** Let $p \in \text{Prop} \cup \{F_{a,b} \mid a, b \in A\}$

$$\begin{aligned}
\mathcal{T}(p) &= p & \mathcal{T}([!G]p) &= \mathcal{T}(p)\\
\mathcal{T}(\neg\varphi) &= \neg\mathcal{T}(\varphi) & \mathcal{T}([!G]\neg\varphi) &= \mathcal{T}(\neg[!G]\varphi)\\
\mathcal{T}(\varphi \wedge \psi) &= \mathcal{T}(\varphi) \wedge \mathcal{T}(\psi) & \mathcal{T}([!G](\varphi \wedge \psi)) &= \mathcal{T}([!G]\varphi \wedge [!G]\psi)\\
\mathcal{T}(D_G\varphi) &= D_G\mathcal{T}(\varphi) & \mathcal{T}([!G]D_H\varphi) &= \mathcal{T}(\bigvee_{F \subseteq G}(\text{Fol}(H, F) \wedge D_{F\cup H}[!G]\varphi))\\
& & \mathcal{T}([!G][!H]\varphi) &= \mathcal{T}([!G]\mathcal{T}([!H]\varphi))
\end{aligned}$$

This translation recursively defines a function from formulas of $\mathcal{L}_{CL}$ to their normal form in the rewriting system displayed in Table 2.2. In the recursive definition of this translation, the first (propositional) case in the left column is the terminating one; the next three are *pass-through* rules, continuing the rewriting on a deeper level (on less complex subformulas). The formulas in the right column are the actual rewriting rules. The last rule in the right column "carries" the translation over the first communication modality of any occurrence of at least two consecutive communication modalities. This postpones the rewriting and elimination of the communication modalities to a deeper (or more inside) level. Hence, our reduction is of what is called the "inside-out" style in Wang and Cao [67].

We can show that this rewriting system is terminating by assigning a fitting complexity measure to formulas, such that the translation always reduces the complexity. Such a complexity measure can easily be defined, since each rewriting rule always either eliminates a communication modality or moves a communication modality inwards.

This reduction system has two crucial properties. First, the normal forms of this system are part of the non-dynamic fragment of $\mathcal{L}_{CL}$, $\mathcal{L}_{CL^-}$. This follows from the observation that the rewriting system pushes the communication modalities inwards over all possible symbols of the language, finally eliminating the modality all-together when it occurs just above the propositional level (in a form $[!G]p$ for $G \subseteq A, p \in \text{Prop} \cup \{F_{a,b} \mid a, b \in A\}$). Second, the translation the reduction system brings about is such that for any $\varphi \in \mathcal{L}_{CL}$: $\mathfrak{M}, w \Vdash \varphi$ iff $\mathfrak{M}, w \Vdash \mathcal{T}(\varphi)$.

Proposition 2.3.2. *For any $\varphi \in \mathcal{L}_{CL}$, $\vdash_{CL} \varphi \leftrightarrow \mathcal{T}(\varphi)$ and $\mathcal{T}(\varphi) \in \mathcal{L}_{CL^-}$.*

Proof. We prove this by induction on the complexity of formulas. The base case of atomic formulas as well as the induction steps of $\varphi = \neg\psi$, $\varphi = \psi \wedge \chi$, and $\varphi = D_G\psi$ are trivial.

The cases of $\varphi = [!G]p$, $\varphi = [!G](\psi \wedge \chi)$, $\varphi = [!G]\neg\psi$, and $\varphi = [!G]D_H\psi$ are easily proven using `!G ATOM`, `!G CON`, `!G NEG`, and `!G` $\mathtt{D_H}$ respectively.

Finally, for the case $\varphi = [!G][!H]\psi$, we must show that $\vdash_{CL} [!G][!H]\psi \leftrightarrow \mathcal{T}([!G]\mathcal{T}([!H]\psi))$. By the induction hypothesis, $\vdash_{CL} [!H]\psi \leftrightarrow \mathcal{T}([!H]\psi)$ and $\mathcal{T}([!H]\psi) \in \mathcal{L}_{CL^-}$. Using `RE`: $\vdash_{CL} [!G][!H]\psi \leftrightarrow [!G]\mathcal{T}([!H]\psi)$. Hence, $\mathcal{T}([!H]\psi)$ does not contain any communication modalities. Therefore, $\vdash_{CL} [!G]\mathcal{T}([!H]\psi) \leftrightarrow \mathcal{T}([!G]\mathcal{T}([!H]\psi))$ and $\mathcal{T}([!G]\mathcal{T}(!H]\psi)) \in \mathcal{L}_{CL^-}$ by one of the previous cases of the induction. By repeated application of Modus Ponens therefore $\vdash_{CL} [!G][!H]\psi \leftrightarrow \mathcal{T}([!G]\mathcal{T}([!H]\psi))$. The rest follows by Modus Ponens. □

By this translation, *CL* is as expressive as CL^-. Soundness of $\mathbf{S5_n^D} + \mathtt{KF} + \mathbf{C_n}$ follows from soundness of the individual axioms of this system with respect to communication models.

Theorem 2.3.3 (Soundness)**.** *The axioms of $\mathbf{S5_n^D} + \mathtt{KF} + \mathbf{C_n}$ are sound with respect to communication models.*

Proof. Soundness of $\mathbf{S5_n^D}$ is provided in Fagin et al. [29]. Soundness of the `KF` axiom follows from the observation that $F_{a,b}$ is true in a world w iff it is true in all worlds. Soundness of $\mathbf{C_n}$ amounts to showing the validity of its axioms. `!G` $\mathtt{D_H}$ is valid by Proposition 2.3.1. For the other axioms, let $\mathfrak{M}$ be an arbitrary communication model, let $w \in W$ be an arbitrary world. Then:

(i) (`!G ATOM`) $\mathfrak{M}, w \Vdash [!G]p \leftrightarrow p$ since $\mathfrak{M}^{!G}$ differs from $\mathfrak{M}$ only in $(\sim_a)_{a \in A}$.

(ii) (`!G NEG`) $\mathfrak{M}, w \Vdash [!G]\neg\varphi \iff \mathfrak{M}^{!G}, w \Vdash \neg\varphi \iff \mathfrak{M}^{!G}, w \nVdash \varphi \iff \mathfrak{M}, w \nVdash [!G]\varphi \iff \mathfrak{M}, w \Vdash \neg[!G]\varphi$.

(iii) (`!G CON`) $\mathfrak{M}, w \Vdash [!G](\varphi \wedge \psi) \iff \mathfrak{M}^{!G}, w \Vdash \varphi \wedge \psi \iff \mathfrak{M}^{!G}, w \Vdash \varphi$ and $\mathfrak{M}^{!G}, w \Vdash \psi \iff \mathfrak{M}, w \Vdash [!G]\varphi$ and $\mathfrak{M}, w \Vdash [!G]\psi \iff \mathfrak{M}, w \Vdash [!G]\varphi \wedge [!G]\psi$.

(iv) (`RE`) Assume $\Vdash \varphi \leftrightarrow \chi$. We will prove that $\Vdash \psi \leftrightarrow \psi[\varphi/\chi]$ by induction on the complexity of ψ. The base case of $\psi = p$ as well as the cases of $\psi = \psi' \wedge \psi''$, and $\psi = \neg\psi'$ are trivial.

For the case of $\psi = D_G\psi'$ we must show that for any $\mathfrak{M}$ and $w \in W$: $\mathfrak{M}, w \Vdash D_G\psi' \iff \mathfrak{M}, w \Vdash (D_G\psi')[\varphi/\chi]$. If $\varphi = D_G\psi'$ then this holds trivially, else we have that $(D_G\psi')[\varphi/\chi] = D_G(\psi'[\varphi/\chi])$. Furthermore, $\mathfrak{M}, w \Vdash D_G\psi' \iff$ if $\forall v\ w \sim_G v$ then $\mathfrak{M}, v \Vdash \psi \overset{I.H.}{\iff}$ if $\forall v\ w \sim_G v$ then $\mathfrak{M}, v \Vdash \psi'[\varphi/\chi] \iff \mathfrak{M}, w \Vdash D_G(\psi[\varphi/\chi])$.

Finally, for the case of $\psi = [!G]\psi'$ we must show that for any $\mathfrak{M}$ and $w \in W$: $\mathfrak{M}, w \Vdash [!G]\psi' \iff \mathfrak{M}, w \Vdash ([!G]\psi')[\varphi/\chi]$. If $\varphi = [!G]\psi'$ then this holds trivially, else we have that $([!G]\psi')[\varphi/\chi] = [!G](\psi'[\varphi/\chi])$. Furthermore, $\mathfrak{M}, w \Vdash [!G]\psi' \iff \mathfrak{M}^{!G}, w \Vdash \psi' \overset{I.H.}{\iff} \mathfrak{M}^{!G}, w \Vdash \psi'[\varphi/\chi] \iff \mathfrak{M}, w \Vdash [!G](\psi[\varphi/\chi])$. □

Theorem 2.3.4 (Strong completeness). $\mathbf{S5^D_n}$ + `KF` + $\mathbf{C_n}$ *is strongly complete with respect to communication frames: for any set of formulas* $\Gamma \subseteq \mathcal{L}_{CL}$ *and formula* $\varphi \in \mathcal{L}_{CL}$*, if* $\Gamma \Vdash \varphi$ *then* $\Gamma \vdash \varphi$.

Proof. Strong completeness is obtained from the strong completeness proof of $\mathbf{S5^D_n}$ + `KF` w.r.t. communication frames for the logic CL^- sketched above, and because the axioms of $\mathbf{C_n}$ reduce $\mathcal{L}_{CL}$ formulas to formulas of $\mathcal{L}_{CL^-}$ via the translation in Definition 2.3.6. Assume $\Gamma \Vdash_{CL} \varphi$, we have to show that $\Gamma \vdash_{CL} \varphi$. Let $\Gamma^{\mathcal{T}} = \{\mathcal{T}(\varphi) \mid \varphi \in \Gamma\}$. By Proposition 2.3.2 and soundness of $\mathbf{S5^D_n} + \mathbf{C_n}$ we can show that $\Vdash_{CL} \varphi \leftrightarrow \mathcal{T}(\varphi)$, and that $\Vdash_{CL} \Gamma$ iff $\Vdash_{CL} \Gamma^{\mathcal{T}}$. Hence, $\Gamma^{\mathcal{T}} \Vdash_{CL} \mathcal{T}(\varphi)$. Since $\mathcal{T}(\varphi) \in \mathcal{L}_{CL^-}$, $\Gamma^{\mathcal{T}} \subseteq \mathcal{L}_{CL^-}$, and the semantics of CL^- coincides with the semantics of CL: $\Gamma^{\mathcal{T}} \Vdash_{CL^-} \mathcal{T}(\varphi)$. By completeness of $\mathbf{CL}^-$ for communication frames: $\Gamma^{\mathcal{T}} \vdash_{CL^-} \mathcal{T}(\varphi)$. Since the axiom system of CL includes the system $\mathbf{CL}^-$: $\Gamma^{\mathcal{T}} \vdash_{CL} \mathcal{T}(\varphi)$. Thus, by Proposition 2.3.2 and Modus Ponens ($R1$) $\Gamma^{\mathcal{T}} \vdash_{CL} \varphi$. Finally, by definition of $\Gamma^{\mathcal{T}}$ Proposition 2.3.2, and Modus Ponens (`R1`): $\Gamma \vdash_{CL} \Gamma^{\mathcal{T}}$. Hence, $\Gamma \vdash_{CL} \varphi$. □

2.4 Iterated Communication and Fixed Points

In this section we explore properties of iterated communication updates. The most general form of such an iteration is that of subsequent communication updates by n possibly distinct sets of agents $(G_i)_{i=0}^n$. Because such an iteration is quite involved, we will instead treat a simpler variant of iterated communication that sufficiently serves the purposes of this thesis: iterated communication by the same group $G \subseteq A$.

2.4.1 Iterated Communication by All

First, consider two consecutive communication updates by all the agents in the social network A. On the first communication update of A, each agent in the network will have updated their epistemic state with the knowledge of the agents that they follow (for agent i this is $\mathcal{F}(i)$). This happens globally, in the whole network. Thus, on the second update, all agents will have updated their epistemic state not only with the knowledge of the agents they follow, but also with the knowledge of the agents that are followed by the agents that they follow. This is clear from the following example.

Example 2.4.1. Consider a setting of four agents $A = \{1, 2, 3, 4\}$. Let the network be as depicted in Figure 2.7. Here the vertices represent agents, the text of the vertices represent the similarity relation of the respective agent, and edges represent the follow relation F. After the first $!A$-update, $\sim_2^{!A} = \sim_2 \cap \sim_3 \cap \sim_4$, $\sim_1^{!A} = \sim_1 \cap \sim_2$, $\sim_3^{!A} = \sim_3$, and $\sim_4^{!A} = \sim_3 \cap \sim_4$. After a second update $\sim_1^{!A!A} = \sim_1^{!A} \cap \sim_2^{!A} = \sim_1 \cap \sim_2 \cap \sim_3 \cap \sim_4$.

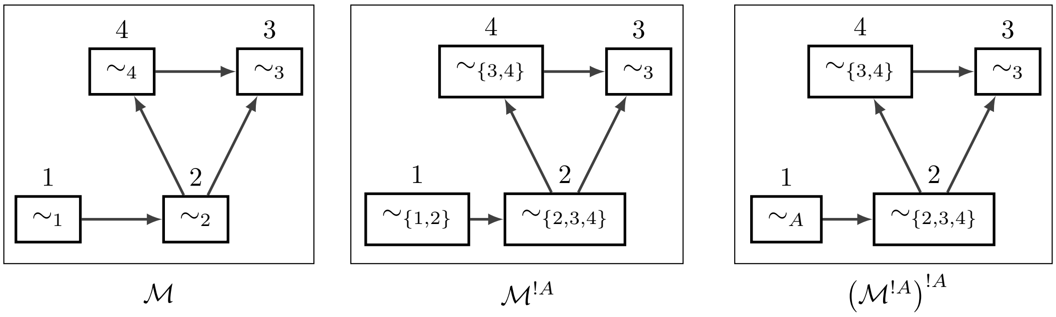

Figure 2.7

What is finally reached, when all agents communicate enough times, is the "best" epistemic state that the social network allows. It is the communication core restricted by the social network. This is generally different from the "best epistemic state", the communication core proper, as described in van Benthem [61, p. 249], which is only ensured to be reached after everybody communicates a sufficient number of times in a fully connected network.

2.4.2 Iterated Communication by Some

If we regard a more general setting of iterated communication by a subset of the agents in the network $G \subseteq A$, something similar plays out. On the first communication update by G, each agent in the network will have updated their epistemic state with the knowledge of the agents that they follow who are part of G, for agent i this is $\mathcal{F}|_G(i)$. On the second update, all agents will have updated their epistemic state not only with these agents, but also with knowledge of the agents in G that are followed by agents in G that they follow. This is clear from the following example.

Example 2.4.2. We start from the model in Example 2.4.1. But now we update the model with a communication update of $\{3, 4\}$. The situation after the first and second update is shown in Figure 2.8. Again, the vertices represent agents, the text in the vertices represents the similarity relation of the respective agent, and edges represent the follow relation F.

After the first $!\{3, 4\}$-update, $\sim_1^{!\{3,4\}} = \sim_1$, $\sim_2^{!\{3,4\}} = \sim_2 \cap \sim_3 \cap \sim_4$, $\sim_3^{!\{3,4\}} = \sim_3$, and $\sim_4^{!\{3,4\}} = \sim_3 \cap \sim_4$. Now, after a second update, still $\sim_1^{!\{3,4\}!\{3,4\}} = \sim_1$, since 1 does not follow anyone in $\{3, 4\}$.

As we can see, if only agents in G communicate, information won't reach the "best epistemic state possible", like it did with a communication update by all agents. Still propagation takes place, only this time restricted to the social network we get when we remove all edges that are not to an agent in G — the social network $F|_G = \{(i, j) \in F \mid j \in G\}$. This will be shown in a later proposition (when we present the characterisation of $\mathcal{F}|_\cdot^{+n}$).

But first, we consider iterated communication by a group G more formally. Let $\circ$ be functional composition. Then in general:

Proposition 2.4.1. *For any $i \in A$ and $G \subseteq A$,* $\left(\sim_i^{!G}\right)^{!G} = \bigcap_{j \in (\mathcal{F}|_G^+ \circ \mathcal{F}|_G^+)(i)} \sim_j$

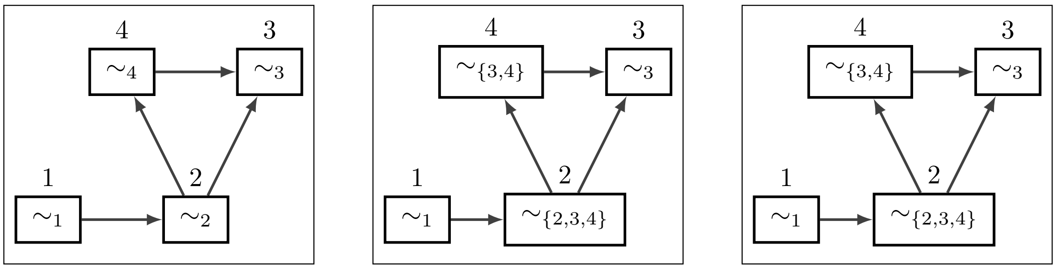

$\mathcal{M}$ $\quad$ $\mathcal{M}^{!\{3,4\}}$ $\quad$ $\left(\mathcal{M}^{!\{3,4\}}\right)^{!\{3,4\}}$

Figure 2.8

Proof.

$$\left(\sim_i^{!G}\right)^{!G} = \bigcap_{j\in\mathcal{F}|_G^+(i)} \sim_k^{!G} = \bigcap_{j\in\mathcal{F}|_G^+(i)}\bigcap_{k\in\mathcal{F}|_G^+(j)} \sim_k = \bigcap_{k\in\bigcup_{j\in\mathcal{F}|_G^+(i)}\mathcal{F}|_G^+(j)} \sim_k = \bigcap_{k\in\mathcal{F}|_G^+(\mathcal{F}|_G^+(i))} \sim_i$$

□

We define the iterated application of a communication update $!G$, $!G\otimes n$, as follows.

Definition 2.4.1 (Iterated communication updates & notation)**.** Let $G\subseteq A$, $i\in A$, and $\mathfrak{F}$ be a frame with similarity relation $\sim$ and $n\in\mathbb{N}$. We write $!G\otimes n$ as a shorthand for n applications of $!G$ to $\sim_i$. Let:

$$\sim_i^{!G\otimes n} := \begin{cases} \sim_i & \text{if } n=0 \\ (\sim_i^{!G})^{!G\otimes n-1} & \text{if } n>0 \end{cases}$$

Similarly, for n applications of $!G$ to $\mathfrak{F}$ we write $\mathfrak{F}^{!G\otimes n}$. Let $\mathfrak{F}^{!G\otimes n} := \mathfrak{F}$, and for $n>0$ $\mathfrak{F}^{!G\otimes n} := (\mathfrak{F}^{!G})^{!G\otimes n-1}$. For $\underbrace{\mathcal{F}|_G^+\circ\mathcal{F}|_G^+\circ\cdots\circ\mathcal{F}|_G^+}_{n}$ we write $\mathcal{F}|_G^{+n}(i)$, where $\mathcal{F}|_G^{+0}$ is the identity function (over $\mathscr{P}(A)$). We will use the same notation for $\mathcal{F}$ and $\mathcal{F}|_G$.

Syntactically, we denote $n\in\mathbb{N}$ repeated communication modalities as $[!G]^n$. To be thorough, we will also use this notation for 0 applications of a communication modality.

The following proposition identifies the effect of repeated communication updates on the indistinguishability relation of agents.

Proposition 2.4.2. *For any $i\in A$, $G\subseteq A$, and $n\in\mathbb{N}$:* $\sim_i^{!G\otimes n} = \sim_{\mathcal{F}|_G^{+n}(i)}$

Proof. By induction on n. Recall that $\sim_{\mathcal{F}|_G^{+n}(i)} = \bigcap_{j\in\mathcal{F}|_G^{+n}(i)}\sim_j$. Of course, $\sim_i^{!G\otimes 1} = \bigcap_{j\in\mathcal{F}|_G^{+1}(i)}\sim_j$. Furthermore: $\sim_i^{!G\otimes n} = \left(\sim_i^{!G\otimes n-1}\right)^{!G} = \bigcap_{j\in\mathcal{F}|_G^+(i)}\sim_j^{!G\otimes n-1} \overset{IH}{=} \bigcap_{j\in\mathcal{F}|_G^+(i)}\bigcap_{k\in\mathcal{F}|_G^{+n-1}(j)}\sim_k = \bigcap_{k\in\left(\bigcup_{j\in\mathcal{F}|_G^+(i)}\mathcal{F}|_G^{+n-1}(j)\right)}\sim_k = \bigcap_{k\in\mathcal{F}|_G^{+n-1}(\mathcal{F}|_G^+(j))}\sim_k$ □

This proposition shows that $!G\otimes n$ transforms the indistinguishability relation of each agent $a\in A$ in a model to that of $\mathcal{F}|_G^{+n}(a)$. Thus, to analyse the properties of repeated communication updates, we must analyse the properties of $\mathcal{F}|_G^{+n}$.

First note that $\mathcal{F}|_G^+(H)$ is distributive over $\cup$:

Lemma 2.4.1 ($\mathcal{F}|_G^+$ is distributive over $\cup$)**.** *For any $G\subseteq A$ and for any two sets $X,Y\subseteq A$:*

$$\mathcal{F}|_G^+(X\cup Y) = \mathcal{F}|_G^+(X)\cup\mathcal{F}|_G^+(Y)$$

Proof. Follows from associativity of $\cup$: $\mathcal{F}|_G^+(X \cup Y) = \bigcup_{x \in X \cup Y} \mathcal{F}|_G^+(x) = \bigcup_{x \in X} \mathcal{F}|_G^+(x) \cup \bigcup_{y \in Y} \mathcal{F}|_G^+(y) = \mathcal{F}|_G^+(X) \cup \mathcal{F}|_G^+(Y)$ ☐

As a direct result of this, repetition of $\mathcal{F}|_G^{+^n}$ on sets of agents ($\mathcal{F}|_G^{+^n}(H)$ for $H \subseteq A$) behaves as a proper extension of repetition on individual agents ($\mathcal{F}|_G^{+^n}(h)$ for $h \in A$) in that $\mathcal{F}|_G^{+^n}(H)$ is equal to $\bigcup_{h \in H} \mathcal{F}|_G^{+^n}(h)$.

Proposition 2.4.3. $\mathcal{F}|_G^{+^n}(H) = \bigcup_{h \in H} \mathcal{F}|_G^{+^n}(h)$

Proof. By a trivial induction over n using Proposition 2.4.1. ☐

Also, as we would expect, $\mathcal{F}|_G^{+^n}(H)$ as a function of n is monotone:

Proposition 2.4.4 (Monotonicity of $\mathcal{F}|_G^+$ in n)**.** *For any $G, H \subseteq A$ and any $1 < n$: $\mathcal{F}|_G^{+^{n+1}}(H) \supseteq \mathcal{F}|_G^{+^n}(H)$*

Proof. By definition of $\mathcal{F}|_G^+(H)$, it holds that $\mathcal{F}|_G^+(H) \subseteq \mathcal{F}|_G^+(H')$ for $H \subseteq H'$. The rest follows by induction over n. ☐

Recall that $\mathcal{F}|_G^+(a)$ is the set of agents in G that a follows, together with a themselves. $\mathcal{F}|_G^{+^n}(a)$ also has an intuitive graph-theoretic interpretation. To characterise $\mathcal{F}|_G^{+^n}$, we will first show that for any $x \in A$, $\mathcal{F}|_G^+$ applied to the set of agents reachable from x with a G-path of at most length $n-1$, extends it to the set of agents reachable from x by a G-path with a length of at most n:

Lemma 2.4.2. $\mathcal{F}|_G^+(\{y \mid x \to_G^{\leqslant n-1} y \;\; and \;\; y \in G\}) = \{y \mid x \to_G^{\leqslant n} y \;\; and \;\; y \in G\}$

Proof. $\mathcal{F}|_G^+(\{y \mid x \to_G^{\leqslant n-1} y \text{ and } y \in G\}) = \mathcal{F}|_G(\{y \mid x \to_G^{\leqslant n-1} y \text{ and } y \in G\}) \cup \{y \mid x \to_G^{\leqslant n-1} y \text{ and } y \in G\}$. $\mathcal{F}|_G(\{y \mid x \to_G^{\leqslant n-1} y \text{ and } y \in G\})$ is the set of agents in G that are followed by agents in G for which x has a G-path with a length of at most $n-1$. This is the set of agents in G for which x has a G-path with a length between 2 and n. ☐

The characterisation of $\mathcal{F}|_\cdot^{+^n}$ is rather intuitive: for $G \subseteq A$, $\mathcal{F}|_G^{+^n}(i)$ contains exactly the subset of G that is reachable from i by a path *in G* of length at most n:

Proposition 2.4.5 (Characterisation of $\mathcal{F}|_\cdot^{+^n}$)**.** *For any $G \subseteq A$ and $n > 0$: $\mathcal{F}|_G^{+^n}(x) = \{y \mid x \to_G^{\leqslant n} y \text{ and } y \in G\} \cup \{x\}$*

Proof. By induction on n. The base case of $n = 1$ follows from the definition of $\mathcal{F}|_G^+$: $\mathcal{F}|_G^{+^1}(x) = \mathcal{F}|_G^+(x) = (\mathcal{F}(x) \cap G) \cup \{x\}$. For the induction step, note the following:

$$
\begin{aligned}
\mathcal{F}|_G^{+^n}(x) &= \mathcal{F}|_G^+(\mathcal{F}|_G^{+^{n-1}}(x)) \\
&= \mathcal{F}|_G^+(\{y \mid x \to_G^{\leqslant n-1} y \text{ and } y \in G\} \cup \{x\}) && \text{(I.H.)} \\
&= \mathcal{F}|_G^+(\{y \mid x \to_G^{\leqslant n-1} y \text{ and } y \in G\}) \cup \mathcal{F}|_G^+(x) && \text{(Proposition 2.4.1)} \\
&= \{y \mid x \to_G^{\leqslant n} y \text{ and } y \in G\} \cup \{x\} && \text{(Lemma 2.4.2)}
\end{aligned}
$$
☐

As a corollary, a similar result is obtained for $\mathcal{F}|_G^{+^n}$ extended to sets.

Corollary 2.4.3. $\mathcal{F}|_G^{+^n}(H) = \{y \mid x \in H \;\; and \;\; x \to_G^{\leq n} y \;\; and \;\; y \in G\} \cup H$

Therefore, the result of n communication updates by G relative to the social network F is exactly the result of n full communication updates by A relative to the network containing only the follow relations to members of G.

2.4.3 Fixed Points

It remains to be shown that after some number of communication updates by all agents A, the "best epistemic state", or communicational core, relative to a social network is reached. To show this, we will show that communication updates have a *fixed point*, that there is a model $\mathfrak{M}_f$ such that $\mathfrak{M}_f^{!G} = \mathfrak{M}_f$. And moreover, that this model will be reachable by iterated communication. I.e. that such a fixed point $\mathfrak{M}_f$ can be calculated from any model $\mathfrak{M}$ by a repeated number of applications of $!G$. By Proposition 2.4.2, we can show this by looking at the fixed points of $\mathcal{F}|_G^{+n}$.

The process of iterated communication updates by G is a process of continually removing more connections from the similarity relations of agents. At the nth iteration of communication, for each agent $i \in A$, starting from $X_0 := \{i\}$, i's similarity relation becomes such that it includes the information of (excludes the relations in) the similarity relations of all the agents in $X_n := \mathcal{F}|_G^+(X_{n-1})$, so that $\sim_i^{!G \otimes n} = \sim_{X_n}$. This process of extending the set X_{n-1} to $\mathcal{F}|_G^+(X_{n-1})$ stops exactly when X_{n-1} is a *fixed point* of $\mathcal{F}|_G^+$, i.e. when $\mathcal{F}|_G^+(X_{n-1}) = X_{n-1}$.

As $\mathcal{F}|_G^+(i) \subseteq A$, A is a trivial such fixed point of $\mathcal{F}|_G^+$. Also, recall that $\mathcal{F}|_G^+(X) = (\mathcal{F}(X) \cap G) \cup X \subseteq X \cup G$. Therefore, any set X such that $G \subseteq X$ is also a fixed point of $\mathcal{F}|_G^+$. In general, for any agent i and no matter the shape of the social network, we reach such a fixed point at the latest in stage $|G|$ — after $|G|$ communication updates by G: $X_{|G|} = \mathcal{F}|_G^{+|G|}(i)$ is a fixed point of $\mathcal{F}|_G^+$.

Proposition 2.4.6 (Fixed points and iteration)**.** *For any $G \subseteq A$ and any $i \in A$, $\mathcal{F}|_G^{+|G|}(i)$ is a fixed point of $\mathcal{F}|_G^+$:*

$$\mathcal{F}|_G^{+|G|+1}(i) = \mathcal{F}|_G^{+|G|}(i)$$

Proof. ($\supseteq$) follows trivially from Proposition 2.4.4

($\subseteq$) Take any $x \in \mathcal{F}|_G^{+|G|}(i)$. By Proposition 2.4.5, this holds iff either $x = i$, or $i \to_G^{\leqslant |G|+1} x$ and $x \in G$. In the former case, $x \in \mathcal{F}|_G^{+|G|+1}(i)$ by definition of $\mathcal{F}|_G^+$. Since any element of a path cannot be equal to any other, the maximum length of the path in the latter case must be $|G|$. Hence, $i \to_G^{\leqslant |G|} x$. Therefore, $x \in \mathcal{F}|_G^{+|G|}(i)$. ☐

Of course, for a fixed point X of $\mathcal{F}|_G^+$ it holds that $\mathcal{F}|_G^{+n}(X) = X$ for any $n > 0$. Hence, as a corollary to the previous identification of a fixed point for $\mathcal{F}|_G^+$, we get that:

Corollary 2.4.4 (Fixed points and iteration)**.** *For any $G \subseteq A$ and any $n \geq |G| > 0$:*

$$\mathcal{F}|_G^{+n}(i) = \mathcal{F}|_G^{+|G|}(i)$$

In terms of Communication Logic, this leads to the following:

Proposition 2.4.7. *For any $\varphi \in \mathcal{L}_{CL}$, $G \subseteq A$, and $n > |G|$:*

$$\Vdash [!G]^n\varphi \leftrightarrow [!G]^{n+1}\varphi \qquad\qquad \Vdash [!G]^n\varphi \leftrightarrow [!G]^{|G|}\varphi$$

Proof. By induction on the complexity of φ. Let $\mathfrak{F}$ be an arbitrary communication frame. Cases $\varphi = p, \varphi = \neg\psi, \varphi = \psi \vee \psi$ are trivial.

For case $\varphi = D_H\psi$: $\mathfrak{F} \Vdash [!G]^n D_H\psi \iff \mathfrak{F} \Vdash D_{\mathcal{F}|_G^{+n}(H)}\psi$ by repeated application of Theorem 2.3.2. By Corollary 2.4.4, this is equivalent to $\mathfrak{F} \Vdash D_{\mathcal{F}|_G^{+|G|}(H)}[!G]^n\psi$. By the I.H. this is equivalent to $\mathfrak{F} \Vdash D_{\mathcal{F}|_G^{+|G|}(H)}\psi$. ☐

2.5 Communication and Knowledge

Communication as formalised in Communication Logic is a direct exchange of all knowledge of an individual to their followers. Iterated communication, then, formalises the indirect effects of such direct communication when agents repeatedly share all they know. In this section, we will look at the effects such exchange has on the knowledge of agents. We will keep this analysis to the former, more simple, direct communication. In the next chapter, we will further develop this analysis towards iterated communication.

The result of directed communication from person to person is, possibly, that the latter gets to know some formulas. For this, some knowledge among the two persons is required. Distributed knowledge can help us formulate this requirement. Naively, if the person talking (the sender, s) together with the person listening (the receiver, r) distributively know that φ, then the listener gets to know φ after the person talking shares all they know. In this way, the distributed knowledge among s and r of φ is *realised* in r. This follows the intuitive interpretation of distributed knowledge: the knowledge that every individual in a group will have if they somehow combine their knowledge — that distributed knowledge somehow "pre-encodes" what a group gets to know after they share all they know with one another. In this section we will analyse why this intuition is, for the most part, wrong. We will formulate what the actual epistemic and network-structural preconditions are for an agent to get to know a formula after communication, and identify a class of formulas for which the above intuition does hold. But we start by reviewing *distributed knowledge resolution* as discussed in Ågotnes and Wáng [3]. Firstly, because distributed knowledge resolution is intrinsically linked with the full communication modality of Communication Logic, and secondly, because of its relation to the movement from distributed knowledge to individual knowledge.

2.5.1 Resolution Operators

Ågotnes and Wáng [3] describe a logic of distributed knowledge *resolution*. In their logic, a resolution operator R_G models the act of "resolving" the distributed knowledge of G, of the members of a group G sharing all their information with one another. There is an immediate connection between this distributed knowledge resolution and full communication in Communication Logic: if all agents in G follow one other, and nobody else follows any of them, then $[!G]$ is exactly this resolution operator R_G. The full communication modality of Communication Logic models the resolution of distributed knowledge *restricted by the network*.

Let us be precise about the resolution of distributed knowledge. What is resolved by full communication is the *similarity of worlds*. The worlds that are distributively regarded as distinguishable by a group, become distinguishable by all its members. In essence, resolution removes the separation between the semantics of distributed knowledge of a group and of (individual) knowledge of its members — letting the latter be equal to the former. In doing so, the epistemic state of the world and the knowledge of the agents is changed. Because formulas can refer to the epistemic state of the world, the resolution of distributed knowledge does not imply that the *formulas* that are distributively known by a group G become known by its members, This difference between the semantics of resolution and of distributed knowledge is why the standard semantics of distributed knowledge do not align with the intuition of distributed knowledge as that what is knowable by its members through infinite and full communication [3].

2.5.2 Realisation of Distributed Knowledge

Now we will analyse the effects full communication has on knowledge in Communication Logic. These effects are different from the resolution of distributed knowledge (which considers the effects of a group sharing all they know with one another, and only one another) because in Communication Logic, the extent to which agents share all they know with one another is restricted by the network. Therefore, we consider a weaker but more general process of the aggregation of distributed knowledge. As to not confuse the resolution of distributed knowledge and the effects of communication in Communication Logic, we call this process the *realisation* of distributed knowledge.

Definition 2.5.1 (Distributed knowledge realisation)**.** By *distributed knowledge realisation* we describe the process (movement, concentration, or aggregation) through which a set of agents share their individual knowledge, which intuitively (but not necessarily effectively, see the discussion above and below) makes what is distributively known by them, individually known by (some of) its members.

Full communication in Communication Logic only results in the realisation of distributed knowledge when the network allows it. To what members the distributed knowledge is realised also depends on the network structure. These network-structural requirements, and the different kinds of distributed knowledge realisation that result from it will be discussed in more detail in the next chapter. The "syntactic" side of this distributed knowledge realisation, its effect, is that agents might "realise" formulas that were distributively known by the communicating group. We will now discuss this effect of distributed knowledge realisation more thoroughly.

First, we regard the simplest setting of distributed knowledge realisation and its effects on knowledge: the realisation of distributed knowledge and its effect on the knowledge about proposition $\gamma \in \text{Prop}$. Regard some group $G \subseteq A$ that distributively knows γ. To resolve the distributed knowledge of the group to an individual $g \in G$ by direct communication, that individual must follow all agents in $G \setminus \{g\}$.[10] If G distributively knows γ , then after G communicates all they know, g will know γ. This is clear from Theorem 2.3.2 and the `!G ATOM` axiom. Expressed in Communication Logic we have that:

$$\Vdash \left(D_G\gamma \wedge \bigwedge_{g' \in G\setminus\{g\}} F_{g,g'} \right) \to [!G]K_g\gamma$$

Naively, one might think that the same formula holds if we move to the effects of realisation with respects to distributed knowledge about any formula $\varphi \in \mathcal{L}_{CL}$. However, if we replace the proposition γ with any formula of Communication Logic, we stumble upon a problem: the above will not hold in general. If, for example, we consider the formula $\varphi = p \wedge \neg K_g p$ (shaped somewhat like a Moore's sentence), then:

$$\nVdash \left(D_G\varphi \wedge \bigwedge_{g' \in G\setminus\{g\}} F_{g,g'} \right) \to [!G]K_g\varphi. \tag{2.1}$$

Such sentences that don't follow the intuition about communication, exist not because of some issue with the framework, but because the basic epistemic language is "too expressive". The non-dynamic fragment of Communication Logic can articulate the epistemic state of agents, and can therefore distinguish between a model pre-communication ($\mathfrak{M}$) and one post-communication ($\mathfrak{M}^{!G}$). It is exactly such formulas that touch the divide between the semantics of distributed knowledge resolution operators and distributed knowledge itself: distributed knowledge of a group is the combined knowledge they implicitly have *now*; the knowledge that a member of the group has after distributed knowledge resolution is that what they know when they have actually combined all their knowledge. Therefore, as we will see in the next chapter, the effects of distributed knowledge realisation are that individuals

[10]Of course, it could well be that the individual g already knows γ, or needs less information than the information of all agents in G to deduce that γ. However, in the general case g must follow all other agents in G. Recall that an agent does not need to follow themselves because we assume that agents do not forget.

get to know formulas that the distributively group knew to be true *after communication*: if G distributively knows that φ will hold *after G communicates*, and g follows all agents in G, then g will know that φ after communication by G. For any $\varphi \in \mathcal{L}_{CL}$, $G \subseteq A$, and $g \in G$:

$$\Vdash D_G[!G]\varphi \wedge \bigwedge_{g' \in G} F_{g,g'} \rightarrow [!G]K_g\varphi. \tag{2.2}$$

2.5.3 Successful Formulas and Model Update Invariance

The phenomenon of (2.1) is not particular to Communication Logic, but can be encountered in all extensions of epistemic logic that include some form of modality which changes the similarity relation. In Public Announcement Logic for example, the formula $[!\varphi]\varphi$ is not generally valid, and neither is $[!\varphi]K_a\varphi$. Regard the formula discussed in van Ditmarsch and Kooi [64]: $\psi = K_a(p \wedge \neg K_b p)$. This formula is true in some model $\mathfrak{M}$ and world w such that a knows p and knows that b does not know p. However, $\mathfrak{M}, w \nVdash_{PAL} [\psi]\psi$, as after the announcement of ψ, b knows that p. Therefore, after the announcement of ψ, the truth value of ψ changes. The same phenomenon occurs in a communication setting of full disclosure that does not rely on social networks. For example, in Baltag and Smets [16], for every $a \in A$, the formula $D_G\varphi \rightarrow [!G]K_a\varphi$, where $[!G]$ is read as "all agents read the information possessed by G", is not valid. A counterexample is the Moore-like sentence similar to the counterexample in Communication Logic, $\varphi = (p \wedge \neg K_a p)$. $D_G p$ implies that a will know p after all agents read the information possessed by G. Therefore, $[!G]K_a\varphi$ is false.

For Public Announcement Logic, if we restrict ourselves to formulas whose truth-value is not affected by the public announcement operation — formulas φ such that $\Vdash \varphi \leftrightarrow [\varphi]\varphi$ — then we get validities equivalent to 2.1. Such formulas are called *successful* in Public Announcement Logic: a formula φ is successful iff $[\varphi]\varphi$ is a validity — iff the formula stays true after it is truthfully announced. For a study on successful formulas see van Ditmarsch and Kooi [64]. An analogous restriction to successful formulas can be made in the logic of Baltag and Smets [16]: taking only φ such that $\varphi \leftrightarrow [!G]\varphi$ is a validity.

We can formulate a similar class of successful formulas in Communication Logic — formulas whose truth value is invariant over a communication update. A formula $\varphi \in \mathcal{L}_{CL}$ is *model update invariant* over $!G$, or $!G$-invariant, in a frame iff its truth-value does not change after an $!G$ update. A formula is $!G$-invariant (in general) when it is $!G$-invariant in all frames.

Definition 2.5.2 (Model Update Invariance)**.** A $\varphi \in \mathcal{L}_{CL}$ is $!G$-invariant in a frame $\mathfrak{F}$ iff $\mathfrak{F} \Vdash \varphi \leftrightarrow [!G]\varphi$. φ is $!G$-invariant (in general) iff it is $!G$-invariant in all frames: $\Vdash \varphi \leftrightarrow [!G]\varphi$.

Returning to the setting of communication, distributed knowledge realisation, and its effects on knowledge: when knowledge is realised from a group to an individual, that individual gets to know all $!G$-invariant formulas that the group distributively knew. For $!G$-invariant formulas φ, (2.1) *will* hold:

$$\Vdash D_G\varphi \wedge \bigwedge_{g' \in G} F_{g,g'} \rightarrow [!G]K_g\varphi. \tag{2.3}$$

Note that as this will hold for each $g \in G$, distributed knowledge is realised in a group G when G shares all they know, and every agents in G follows one another.

We already discussed a class of model update invariant formulas when we analysed fixed points of iterated communication. Recall that for $n > |G|$, $\Vdash [!G]^n\varphi \leftrightarrow [!G]^{|G|}\varphi$ (Proposition 2.4.7). Hereby, for $n > |G|$, the class of formulas of the form $[!G]^n\psi$ is $!G$-invariant.

Corollary 2.5.1. *For $n \geq |G|$ formulas of the form $[!G]^n\varphi$ are $!G$-invariant.*

Recall that the reason that the formulas in Proposition 2.4.7 are valid is due to fixed points of $\mathcal{F}|^+_G$. In all frames, and for all $H \subseteq A$, $\mathcal{F}|^{+|G|}_G(H)$ is a fixed point of $\mathcal{F}|^+_G$. For this reason also, formulas of the form $[!G]^n\varphi$ are $!G$-invariant. Using these fixed points, we can identify a more general class of formulas that is $!G$-invariant in a frame. For $G \subseteq A$, and for formulas of the non-dynamic fragment of Communication Logic $\varphi \in \mathcal{L}_{CL^-}$, if for all occurrences of the distributed knowledge modality D_H in φ, H is a fixed point of the inclusive restricted follow function $\mathcal{F}|^+_G$ of the social relation in the frame $\mathfrak{F}$, then φ is $!G$-invariant in $\mathfrak{F}$.

Proposition 2.5.1. *A formula φ of the non-dynamic fragment of $\mathcal{L}_{CL}$, $\mathcal{L}_{CL^-}$, is $!G$-invariant in a frame $\mathfrak{F}$ if for every modality of the form D_H that occurs in φ, we have that H is a fixed point of $\mathcal{F}|^+_G$.*

Proof. By induction on the complexity of $\varphi \in \mathcal{L}_{CL^-}$. The base case follows by validity of `!G ATOM`: if $\varphi = p$ for $p \in \text{Prop} \cup \{F_{a,b} \mid a, b \in A\}$, then $\vdash p \leftrightarrow [!G]p$. For the induction step, distinguish the following cases:

1. $\varphi = \psi \wedge \chi$. By validity of `!G CON`, $\Vdash [!G](\psi \wedge \chi) \leftrightarrow ([!G]\psi \wedge [!G]\chi)$. By the I.H. $\mathfrak{F} \Vdash \psi \leftrightarrow [!G]\psi$ and $\mathfrak{F} \Vdash \chi \leftrightarrow [!G]\chi$. Therefore, $\mathfrak{F} \Vdash \varphi \leftrightarrow [!G]\varphi$.

2. $\varphi = \neg\psi$. By the I.H. and by validity of `!G NEG`, $\mathfrak{F} \Vdash \neg\psi \leftrightarrow [!G]\neg\psi$.

3. $\varphi = D_H\psi$: By the I.H. $\mathfrak{F} \Vdash \psi \leftrightarrow [!G]\psi$. Therefore, $\mathfrak{F} \Vdash D_H(\psi \leftrightarrow [!G]\psi)$ and hence $\mathfrak{F} \Vdash D_H\psi \leftrightarrow D_H[!G]\psi$. By Theorem 2.3.2: $\mathfrak{F} \Vdash D_H[!G]\psi$ iff $\mathfrak{F} \Vdash [!G]D_{H'}\varphi$ where H' is a set such that $H = \mathcal{F}|^+_G(H')$. By assumption $\mathcal{F}|^+_G(H) = H$. Hence, $\mathfrak{F} \Vdash D_H\varphi \leftrightarrow [!G]D_H\varphi$. ☐

Corollary 2.5.2. *Any φ of the propositional fragment of $\mathcal{L}_{CL}$ (the fragment consisting of propositions, network propositions, and Boolean connectives) is $!G$-invariant for any $G \subseteq A$.*

By interpretation of $\mathcal{F}|^+_G$, this means that all the agents in G that are followed by agents in H are a subset of H itself. Note that this is not a strict identification of successful $\mathcal{L}_{CL^-}$ formulas.[11] By the translation of $\mathcal{L}_{CL}$ to $\mathcal{L}_{CL^-}$ (Definition 2.3.6), this result can be transferred to all formulas of Communication Logic, as a formula φ is $!G$-invariant if its $\mathcal{L}_{CL^-}$-translation is.

Proposition 2.5.2. *For $\varphi \in \mathcal{L}_{CL}$ and $G \subseteq A$: if $\mathcal{T}(\varphi)$ is $!G$-invariant, then so is φ.*

Proof. Assume $\Vdash \mathcal{T}(\varphi) \leftrightarrow [!G]\mathcal{T}(\varphi)$. By completeness, $\vdash \mathcal{T}(\varphi) \leftrightarrow [!G]\mathcal{T}(\varphi)$. By Proposition 2.3.2, $\vdash \varphi \leftrightarrow \mathcal{T}(\varphi)$. Therefore, by `RE`: $\vdash \varphi \leftrightarrow [!G]\varphi$. Hence, by soundness, $\Vdash \varphi \leftrightarrow [!G]\varphi$. ☐

Consequently, we can define a more general family of $!G$-invariant formulas in $\mathcal{L}_{CL}$.

Corollary 2.5.3. *A formula $\varphi \in \mathcal{L}_{CL}$ is $!G$-invariant in a frame $\mathfrak{F}$ if for every modality of the form D_H that occurs in $\mathcal{T}(\varphi)$, H is a fixed point of $\mathcal{F}|^+_G$.*

We can further extend the notion of invariance over iterated communication updates.

Definition 2.5.3 ($!G \otimes n$ invariance)**.** A formula φ is $!G\otimes n$-invariant in $\mathfrak{F}$ iff $\mathfrak{F} \Vdash [!G]^n\varphi \iff \mathfrak{F} \Vdash \varphi$.

A similar family of $!G \otimes n$ invariant formulas can be identified by fixed points.

Proposition 2.5.3. *For any $n > 0$, a formula $\varphi \in \mathcal{L}_{CL^-}$ is $!G \otimes n$-invariant in a frame $\mathfrak{F}$ if for any occurrence of D_H, H is a fixed point of $\mathcal{F}|^{+n}_G$.*

[11] A larger class of $!G$-invariant formulas can be specified by only requiring H to be a fixed point of $\mathcal{F}|^+_G$ for only the D_H that matter for the truth value of φ. This is beyond the scope of this thesis.

Proof. This proof is similar to that of Proposition 2.5.1. $\square$

We can similarly extend this result to general formulas of Communication Logic by the translation of $\mathcal{L}_{CL^-}$ to $\mathcal{L}_{CL}$.

Corollary 2.5.4. *For any $\varphi \in \mathcal{L}_{CL}$, φ is $!G \otimes n$-invariant in a model $\mathfrak{M}$ if any occurrence of D_H in $\mathcal{T}(\varphi)$: H is a fixed point for $\mathcal{F}|_G^{+n}$.*

2.6 Summary

In this chapter we constructed an epistemic logic, Communication Logic, with distributed knowledge and a modality for full communication — the act of agents communicating all they know to their followers in a directed social network. We provided a sound and complete axiomatisation of Communication Logic based on the reduction technique often used for Dynamic Epistemic Logics. Then, we analysed the effects of repeated acts of communication by a single group. Such repeated communication reaches a fixed point. The epistemic state that is reached is the communication core restricted by the social network. Then, we looked at the effects of full communication on the knowledge of agents. We investigated *distributed knowledge realisation*, the process by which a group, through communication, makes what is distributively known to them, individually known by some of its members. We found that, given that an agent follows all members of a group G, that agent gets to know, not all formulas that were distributively known by G, but all formulas that were distributively known by G to be true *after communication*. We showed that this is not a phenomenon particular to Communication Logic, but a common trait of dynamic extensions of epistemic logic. We then introduced $!G$-invariant formulas, a class of formulas alike successful formulas as discussed in van Ditmarsch and Kooi [64]. This is the class of formulas that, when distributively known by a group, and given the right network, do become individually known by members of the group through distributed knowledge realisation.

In the next chapter we will analyse communication and its relation to distributed knowledge realisation in more detail. Specifically, we will explore the network-structural and epistemic preconditions for distributed knowledge realisation driven by iterated full communication, and the roles that the agents play in it.

Chapter 3

Crucial Positions in Communication

In the last chapter we constructed a logic of full communication in a social network. We concluded the chapter with an analysis of the knowledge realisation, the concentration of that what is distributively known by a group to some of its members through full communication over a social network. We conjectured that such a realisation can come about under certain structural network conditions, and that its effects with respects to knowledge about formulas follow the intuition only for a specific class of formulas: model update invariant formulas.

In this chapter we will analyse the relation between distributed knowledge realisation, the network structure, and its effects on knowledge about formulas more thoroughly. We will define and examine network-structural notions crucial to processes of knowledge propagation through the network, and examine their interplay via Communication Logic.

Social networks act as the structure for the propagation of many things, such as diseases, beliefs, behaviour, or information. Often it is implicitly assumed that the propagation concerns properties that are solely determined by the transmission of atomic units in possession of the agents — units such as viruses, bits, opinions, or preferences. The axiom of transmission in such propagation is "a will have p if x amount of a's neighbours already has p". We call an account of propagation *atomic* when it abides by this axiom of atomic transmission. In a way, such a perspective is also applicable here: through communication, agents propagate their similarity relation through the network. And, from the perspective of knowledge, agents propagate the property of "knowing φ" through the network. However, the propagation of the latter does not behave like the propagation of a single unit. Knowing (the truth value of) φ is not dependent on a single unit or object in possession of agents. Rather it depends on the specific combination of units, the similarity relations, that they possess. As such, in full communication, where the object of communication is similarity, the axiom of atomic transmission is not applicable to the propagation of knowledge; it could well be that a gets to know φ from its neighbours even if none of the neighbours know φ themselves, simply because a has some combination of similarity relations crucial for determining φ's truth value that the neighbours of a do not. Furthermore, knowledge about φ is not always possessed by a single agent, because it can be distributed between a set of agents. And finally, even if agents have knowledge about a formula φ, they might not know φ anymore after communication, as φ can become false.

We will analyse the processes of knowledge propagation while upholding an *epistemic account of communication*, where the driving force of knowledge propagation through a network is the realisation of distributed knowledge. We will iteratively develop structural notions that are related to distributed knowledge realisation — specific types of realisation of a more generic form than we have addressed in the previous chapter.

In the following section, we discuss the network-structural requirements for the realisation of distributed knowledge of a set of agents to all its members: *groups.* In Section 3.2, we work towards structures that capture a more general, directed, analysis of distributed knowledge realisation — one that is more suited for directed social networks, where information is propagated *from* senders to receivers through a third party, *connectors.* Subsequently, we investigate the negation of connectors, *blocking sets*, sets that can block the realisation of distributed knowledge by *not* communicating. Finally, in Section 3.5 we discuss the contributions that our work makes to the field of the social sciences. In particular, all concepts defined in the this chapter are closely related to concepts defined in Belardinelli [19]. The definitions that will follow, function as a generalisation of these concepts to a directed setting.

3.1 Realising Distributed Knowledge

Recall the setting of communication discussed in Section 2.5, where, through knowledge realisation, the distributed knowledge about a formula among a group of agents becomes known by a member of that group. As conjectured in (2.3), this is possible only for model update invariant formulas, and under certain network-structural conditions (that g follows all other agents in G). We will prove a variant of this conjecture later in this chapter.

To keep things simple, we will not bother with model update invariance of formulas in this chapter. Rather, we will work with formulas of distributed knowledge realisation like those of conjecture (2.2). This conjecture is as follows:

$$\Vdash \left(D_G[!G]\varphi \wedge \bigwedge_{g' \in G \setminus \{g\}} F_{g,g'} \right) \to [!G]K_g\varphi.$$

As this conjecture will hold for all members of G, we get that direct communication can bring about knowledge realisation in a social network where all the agents of G follow one another:

$$\Vdash \left(D_G[!G]\varphi \wedge \bigwedge_{g \in G} \bigwedge_{g' \in G \setminus \{g\}} F_{g,g'} \right) \to [!G]E_G\varphi.$$

The conditions for knowledge realisation above are not at all minimal. Knowledge realisation can just as well be brought about when agents are further apart in the social network by iterated communication. This is where the connection between knowledge realisation and knowledge propagation becomes apparent. We conjecture the following: if for every agent $g' \in G$, there is a G-path of length n from g to g', and G distributively knows that after G communicates n times, φ holds, then after G communicates n times, g will know φ. To express this in Communication Logic requires a formula that states the existence of a P-path from a set of agents to another, of a certain maximum length. Throughout this chapter let $\bigvee^0 \varphi := \bot$, $\bigvee_{i \in \emptyset} \varphi := \bot$, $\bigwedge^0 \varphi := \top$, and $\bigwedge_{i \in \emptyset} \varphi := \top$, in accordance with $\vee$ and $\wedge$ being the join and meet respectively. Recall that the length of a path is equal to the number of edges in it. First define "Walk", stating the existence of a P-walk from $a \in A$ to $b \in A$, of length n. Recall that a P-walk from x to y of length n is a sequence $(P_i)_{i=0}^n$, such that $n \geq 1$, $p_0\ F \cdots F\ p_n$, $p_0 = x$, $p_n = y$, and $p_1, \dots p_{n-1} \in P$. We do not consider sequences of length 1, e.g. (a), as walks or paths. This avoids the requirement of identity in our language to express that there exists a path between two agents.[1]

[1] If we were to take (a) as a walk, then walks, and subsequently max-paths, (and connectors and blocking sets) are only "expressible" in communication logic in the sense that there is a *schema* that expresses that there is a walk from a to b in a model.

Definition 3.1.1 (Walk formula)**.** Let $P \subseteq A$, $a, b \in A$, and $n \in \mathbb{N}$. Then:

$$\text{Walk}(P,a,b,n) := \begin{cases} \bot & \text{if } n = 0 \\ F_{a,b} \vee \bigvee_{(p_1,\dots,p_{n-1}) \in P^{n-1}} \left(F_{a,p_1} \wedge \bigwedge_{i=1}^{n-2} F_{p_i,p_{i+1}} \wedge F_{p_{n-1},b} \right) & \text{otherwise} \end{cases}$$

Let the formula $\leq$Path(P, a, b, n) denote that there is a P-path of length $\leq n$ from a to b. As this formula is a statement about maximal paths, and there is a P-path of maximal length n iff there is a P-walk of maximal length n, maximal P-paths are definable in terms of the existence of P-walks:

Definition 3.1.2 (Maximum path formula)**.** Let $P \subseteq A$, $a, b \in A$, and $n \in \mathbb{N}$, then:

$$\leq\text{Path}(P,a,b,n) = \bigvee_{i=0}^{n} \text{Walk}(P,a,b,n).$$

As a shorthand, let $\leq$Path$(P, a, b) := \leq$Path$(P, a, b, |A| - 1)$.

Proposition 3.1.1 (Correctness of Walk and $\leq$Path)**.** *For any $a, b \in A$, $P \subseteq A$, and frame $\mathfrak{F}$: there is a P-walk from a to b of length n in $\mathfrak{F}$ iff $\mathfrak{F} \Vdash$ Walk(P, a, b, n). Furthermore, there is a P-path from a to b of length $\leq n$ in $\mathfrak{F}$ iff $\mathfrak{F} \Vdash \leq$Path(P, a, b, n).*

Proof. There is a P-walk from a to b of length 1 iff $a\ F\ b$. There is a P-walk from a to b of length n iff there are $p_1, \dots, p_{n-1} \in P$ such that $a\ F\ p_1\ F \cdots F\ p_{n-1}\ F\ b$. This is exactly what Walk(P, a, b, n) states. Furthermore, as any P-walk is always at least as long as a P-path and any P-path is also a P-walk, there is a P-path of length at most n iff there is a p-walk of length at most n. This is what $\leq$Path(P, a, b, n) states. ☐

Using $\leq$Path, we can conjecture a formula that expresses the distributed knowledge realisation of φ among a set G to *one* of its members through iterated communication:

$$\Vdash D_G[!G]^n\varphi \wedge \bigwedge_{g' \in G \setminus \{g\}} \leq\text{Path}(G, g, g', n) \to [!G]^n K_g\varphi. \tag{3.1}$$

Finally, as this conjecture should hold for all agents in G, it has a collective variant that represents distributed knowledge realisation in G to *all* its members:

$$\Vdash D_G[!G]^n\varphi \wedge \bigwedge_{g \in G} \bigwedge_{g' \in G \setminus \{g\}} \leq\text{Path}(G, g, g', n) \to [!G]^n E_G\varphi. \tag{3.2}$$

We will prove conjecture (3.2) in this section, and prove forms (3.1) and (2.2) in the following section. Conjecture (3.2) has a relation to knowledge propagation though the network as a result of communication. The set G described above has a particular ability to bring about this propagation themselves: by communication, G can make whatever they distributively know after communication, individually known by its members. Indirect, iterated, communication by a group G is the general device for exactly this process: under minimally allowing conditions for knowledge realisation, iterated communication by G realises distributed knowledge of G to its members. In the remainder of this section we will analyse such sets G, their structural requirements, and their epistemic-communicational implications.

3.1.1 Groups

The set of agents G in the formula (3.2) is a set such that there is a G-path from any agent in G to any other agent in G. We call such sets *groups*.[2] We quantify groups according to

[2]We will discuss variations of groups in Section 3.2.7.

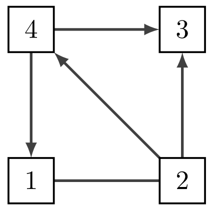

Figure 3.1

their *width* or *diameter*: the length of the longest shortest path from any two agents in the group. We call groups whose longest shortest path between any two distinct members has a length of at most n: n-groups.To avoid ambiguity, from now on we will call any $G \subseteq A$ a *set* of agents, and only call G a group when it is a k-group for some $k \in \mathbb{N}$.[3] For completeness' sake, we call singleton sets $\{a\} \subseteq A$: 0-groups.

Definition 3.1.3 (n-group)**.** For $G \subseteq A$ and frame $\mathfrak{F}$: any $G \subseteq A$ such that $|G| = 1$ is a 0-group in $\mathfrak{F}$. For $G \subseteq A$ such that $|G| > 1$, G is an n-group iff for any $g, g' \in G$ such that $g \neq g'$: $g \rightarrow_G^{\leqslant n} g'$, We call G a group in $\mathfrak{F}$ iff it is an n-group in $\mathfrak{F}$ from some $n \in \mathbb{N}$.

We will omit in which frame G is a (n-)group when it is clear from the context, or the particular frame is irrelevant.

Let the following example illustrate the definition and structure of groups.

Example 3.1.1. Consider the social network depicted in Figure 3.1, where an arrow-less line indicates a bidirectional social connection. In this four-agent setting, there is no path from 3 to 4 or from 3 to 2. As such, $\{1,2,3,4\}$ (or any other set S such that $\{1,3\} \subseteq S$, $\{2,3\} \subseteq S$, or $\{3,4\} \subseteq S$) is not a group. The shortest path from 1 to 4 has a length of 2, and all other possible shortest paths between the agents in $\{1,2,4\}$ have a length of 1. Therefore, $\{1,2,4\}$ forms a 2-group, and $\{1,2\}$ forms a 1-group. Finally, the singleton sets $\{1\}$, $\{2\}$, $\{3\}$, and $\{4\}$ form 0-groups.

As n-groups are not required to be minimal over n, for $m \geq n$ any n-group is also an m-group. Thus, $\{1,2\}$ is also a 2-group, and a 3-group etc.

As groups are defined in terms of the existence of maximal path, n-groups can be expressed in Communication Logic using the previously constructed $\leq$Path formulas.

Definition 3.1.4 (Group formula)**.** For any non-empty $G \subseteq A$, define Group as follows.

$$\text{Group}(G, n) := \bigwedge_{g \in G} \bigwedge_{g' \in G \setminus \{g\}} \leq\text{Path}(G, g, g', n)$$

As an abbreviation, we write Group(G) for Group($G, |A| - 1$), i.e. to state that G is a group in general.

Correctness of the group formulas follow from the correctness of $\leq$Path.

Proposition 3.1.2 (Correctness of Group-formula)**.** *For non-empty $G \subseteq A$, $n \in \mathbb{N}$, G is an n-group in $\mathfrak{F}$ iff $\mathfrak{F} \Vdash \text{Group}(G, n)$.*

Proof. This follows from the correctness of $\leq$Path (Proposition 3.1.1). $\square$

3.1.2 Groups and Knowledge Realisation

As conjectured in (3.2), k-groups are sets of agents that can realise distributed knowledge to all its members by communicating k times. To show this, we will first consider the

[3]In the literature a k-group is also called a k-clan [41].

semantic implications of communication by a group. If $G = \{g_1, \ldots, g_m\}$ is a k-group in a frame $\mathfrak{F}$, then in the frame that results from n acts of communication $\mathfrak{F}^{!G\otimes n}$, each resulting epistemic relation $\sim_{g_i}^{!G\otimes n}$ is the intersection of all relations in $\{\sim_{g_1}, \ldots, \sim_{g_m}\}$. Therefore, after communication, all agents in a group can distinguish all the worlds that the entire group could before communication. This entails that all agents know exactly as much as what was distributively known by the entire group before communication, modulo the issue with non-invariant formulas discussed in Section 2.5.

Proposition 3.1.3 (Semantic k-group knowledge realisation)**.** *For $G \subseteq A$, $n \in \mathbb{N}$, and a frame $\mathfrak{F}$, if G is an n-group in $\mathfrak{F}$, then:*

$$\forall g \in G \ \ \sim_g^{!G\otimes n} = \sim_G$$

Proof. For $g, g' \in G$ such that $g \neq g'$, by Proposition 2.4.5 $g \to_G^{\leqslant n} g' \Rightarrow g' \in \mathcal{F}|_G^{+n}(g)$, Therefore if G is an n-group, then for all $g, g' \in G$ $g' \in \mathcal{F}|_G^{+n}(g)$ i.e. $\mathcal{F}|_G^{+n}(g) = G$. Thus, by Proposition 2.4.2 $\sim_g^{!G\otimes n} = \sim_{\mathcal{F}|_G^{+n}(g)} = \sim_G$. □

This semantic result for groups implies the validity of (3.2). That is, if an n-group G distributively knows that after they communicate n times, φ holds, then after G communicates n times, everyone in G individually knows that φ.

Proposition 3.1.4 (Syntactic k-group knowledge realisation)**.** *Let $G \subseteq A$, then:*

$$\Vdash (\mathrm{Group}(G, n) \wedge D_G[!G]^n \varphi) \to [!G]^n E_G \varphi$$

Proof. Fix an arbitrary communication model $\mathfrak{M}$ and $w \in W$. By Proposition 3.1.3, if G is an n-group then $\forall g \in G \ \ \sim_g^{!G\otimes n} = \sim_G$. Therefore, $D_G[!G]^n$ implies $[!G]^n \bigwedge_{g\in G} K_g \varphi$. □

Groups, therefore, are a sufficient network-structural property to realise distributed knowledge to individual knowledge among all its members.

3.1.3 Summary

In this section we have identified the topological prerequisites for the realisation of distributed knowledge of G to individual knowledge among its members — namely for the set of agents to form a group. A k-group can do this by communicating k times.

3.2 Directional Distributed Knowledge Realisation

The particular form of knowledge realisation discussed in the previous section, one that realises the distributed knowledge of a group to all its members, does not cover all interesting settings of knowledge propagation through a social network. Consider the following example.

Example 3.2.1. A set of construction workers is tasked with building a new building. The specifics of how the building must be constructed are distributively known by the project management team: the architects, engineers, and city planners. The knowledge of the project managers, together with the actual knowledge of how to build buildings, is enough to know exactly how to build the building.

In such a setting, it is distributively known by the management team together with the construction workers how to build the building. The only thing in the way of the builders to start construction is the realisation of this knowledge to every builder. However, the interpretation of knowledge realisation we used in the previous section does not apply here: it is not required, and often not desired, for the project managers to know exactly how the building should be constructed. Only the construction workers must know.

A more conventional interpretation of communication through a network, especially in a directional network, is the transmission of information from some *sending* party to some *receiving* party. In terms of knowledge realisation, this entails that distributed knowledge of the senders and receivers becomes individual knowledge only among the receivers.[4] In a way, we have already regarded such a setting in conjecture (2.2)and (3.1), where G is the set of senders, and g is the receiver.

The simplest settings of directional knowledge realisation is one where a single agent $s \in A$, the sender, communicates with another agent $r \in A$, the receiver. This realises the distributed knowledge of $\{s, r\}$ to r. Assume that s and r together distributively know that after s communicates, φ is true. Furthermore, assume the most elementary topological precondition for communication from the sender to the receiver: that r follows s. Then after s communicates all they know, r will know φ. Expressed in Communication Logic:

$$\Vdash D_{\{s,r\}}[!s]\varphi \wedge F_{r,s} \rightarrow [!s]K_r\varphi.$$

In a more general setting, the receiving party is a *set* of agents $R \subseteq A$. Then, any receiving agent $r \in R$ together with the sending agent s must distributively know that φ is true after s communicates. s must also follow every agent in R. If this is the case, then after s communicates, every agent in R will know φ:

$$\Vdash \bigwedge_{r\in R} D_{\{s,r\}}[!s]\varphi \wedge \bigwedge_{r\in R} F_{r,s} \rightarrow [!s]E_R\varphi.$$

We can make the same generalisation with the sending party, taking a set of senders S rather than an individual agent. Then the sets S together with any r must distributively know that after S communicates, φ is true. Moreover, all the information of S must be accessible by all agents in R. Ergo, all agents of S must be followed by all agents in R. If this true, then after the agents in S communicate all they know, the agents in R will individually know φ.

$$\Vdash \bigwedge_{r\in R} D_{S\cup\{r\}}[!S]\varphi \wedge \bigwedge_{s\in S}\bigwedge_{r\in R} F_{r,s} \rightarrow [!S]E_R\varphi. \tag{3.3}$$

Knowledge can also be realised if the receivers do not directly follow the senders, by repeated communication. Then, the structural requirements are that there exists an $S \cup R$-path from all agents in R to all agents in S. Now, information is propagated by communication of both S and R, and n such communications must take place, where n is the length of the longest shortest $S \cup R$-path from any agent in R to any agent in S:

$$\Vdash \bigwedge_{r\in R} D_{S\cup\{r\}}[!(S\cup R)]^n\varphi \wedge \bigwedge_{s\in S}\bigwedge_{r\in R} \leq\mathrm{Path}(S\cup R, s, r, n) \rightarrow [!(S\cup R)]^n E_R\varphi. \tag{3.4}$$

3.2.1 The Third Party in Communication

Knowledge can even be realised when no $S \cup R$-paths exist from R to S, if we include agents outside S and R in the communication. Expanding on the previous example, the project managers S might not directly inform the construction workers, but inform people tasked with briefing the construction workers. Then, information is propagated through a *third party* $P \subseteq A$. The formula for such a propagation is equivalent to (3.4), with all $S \cup R$ replaced by $S \cup P \cup R$:

$$\Vdash \bigwedge_{r\in R} D_{S\cup\{r\}}[!(S\cup P\cup R)]^n\varphi \wedge \bigwedge_{s\in S}\bigwedge_{r\in R} \leq\mathrm{Path}(S\cup P\cup R, s, r, n) \rightarrow [!(S\cup P\cup R)]^n E_R\varphi. \tag{3.5}$$

[4]Note that the previous setting is a specific case of this directional knowledge realisation where the receiving party and the sending party are the same set. As such, groups can be formulated in terms of the notions we will define in this section. We will show this in Section 3.2.7.

Note that this formula assumes that the receivers communicate, and the senders communicate more than once. This is not always necessary, but keeps the formulas simple. We expand on this in Section 3.2.6.

We will prove (3.5) and some of its variations later in this section. Validity of (3.3), (3.4), and the other formulas in this subsection follow from the validity of (3.5).

3.2.2 Different Forms of Knowledge Realisation

The directional form of knowledge realisation brings other shortcomings of the definition of knowledge realisation to light. Consider the following example.

Example 3.2.2. A set of detectives R is investigating a crime. They have investigated the crime scene and collected evidence. As this evidence was inconclusive, they have collected a set of witnesses S to solve the case. What these witnesses know, together with the evidence that the detectives collected is enough to know who the killer is.

In such a setting, it is distributively known by the witnesses and the detectives who the killer is. The only thing in the way of the detectives to know who the killer is, is the realisation of this distributed knowledge. Again, the interpretation of knowledge realisation we used in the previous section does not apply here: it is not required, and often not desired, for the witnesses to know who the killer is. Only the detectives must know. Unlike the previous example however, the detectives are not required to all know who the killer is. It is enough for some detective to know. Sometimes, a sufficient *required* knowledge distribution in R for the knowledge realisation to be regarded as *successful* is something weaker. We can express this alternative result of knowledge realisation by the dual of “everybody knows that” (E_R): “somebody knows that” (S_R), whose definition is as follows.

Definition 3.2.1 (S_H)**.** For any $H \subseteq A$ and $\varphi \in \mathcal{L}_{CL}$, let: $S_H\varphi := \bigvee_{h \in H} K_h\varphi$.

Moreover, if, for example, the detectives report their knowledge to their supervisors, it is sufficient for the detectives to distributively know who the killer is. In this way, it is only required for the distributed knowledge among the senders and receivers to be *partially* realised in the receivers. This partially realised distributed knowledge can then be fully realised to a form of individual knowledge among the supervisors by another process of directional knowledge realisation developed before. It is important to treat this double, split, process of knowledge realisation separately, because the third party that allows the distributed knowledge to be realised could differ. For example, while police officers might act as the third party for the realisation of distributed knowledge among the witnesses and detectives to distributed knowledge among the detectives, a secretary might act as a third party for the realisation of the distributed knowledge among the detectives to individual knowledge among the heads of investigation.

We will consider precisely these three *results* of directed knowledge realisation: the partial realisation of distributed knowledge among the receivers, and full realisation of distributed knowledge to the “first-order” extensions of individual knowledge among the receivers: individual knowledge among some or all receivers. We come to the following definition of directed knowledge realisation.

Definition 3.2.2 (Directed knowledge realisation)**.** By *(directed) distributed knowledge realisation* we describe the process (movement, concentration, or aggregation) through which agents in a group share their individual knowledge, which intuitively (but not necessarily effectively) makes what is distributively known by the senders and receivers, some form of knowledge by the receivers. We call the former the *precondition* of knowledge realisation, and call the latter its *result.* We distinguish three specific forms of directed knowledge realisation: directed knowledge realisation towards (1) distributed knowledge, (2) individual knowledge by some, and (3) individual knowledge by all receivers. To keep things simple, from here on we won’t emphasise that (1), the distributed knowledge realisation towards distributed knowledge, is a form of *partial* knowledge realisation

3.2.3 Connectors

We have seen that knowledge realisation is achieved when three requirements are met: a set of sending agents together with a set of receiving agents distributively know some formula, the structure of the network allows for knowledge flow from the sending set to a receiving set through a propagator, and the sending set and propagator communicate a sufficient number of times. We will call such propagators *connectors*.

A connector is a set that enables the realisation of distributed knowledge among the senders and receivers to knowledge among the receivers. Furthermore, we have seen that there are multiple types of knowledge realisation. As will become clear in this section, different types of connector enable different types of knowledge realisation, from different kinds of assumed knowledge distributions in the sending and receiving set, towards multiple kinds of knowledge distributions in the receiving set.

There are two approaches to finding appropriate definitions of connectors: we can work from the different forms of knowledge realisation towards topological properties that characterize connectors which make these forms of knowledge realisation possible; or conversely, we can conjecture topological properties for connectors and find their corresponding epistemic implications with respect to knowledge realisation. We will take the latter approach, as there is an intuitive common dividing characteristic on the network-structural side of connectors: connectors, regardless of their particular definition, allow the information of agents in S to reach agents in R. Otherwise communication from S to R would be impossible. Hence, any connector must satisfy that: for $s \in S$ and $r \in R$, $r \to_{S \cup C \cup R} s$. In extending this relation to a stronger one, six qualitatively distinct definitions of connectors arise, derived from the six[5] possible first-order quantifications of the path relation $\to_G$ over elements of senders and receivers. We assign names to these variants of the form $Q_1^x Q_2^y$, according to their quantification over the sending and receiving set, where Q_1 and Q_2 are one of $\exists, \forall$, and x, y are one of s, r. We call $Q_1^x Q_2^y$ the connector's *type*, and denote the set of types by $\mathcal{T} := \{\forall^r\exists^s, \exists^s\exists^r, \forall^s\exists^r, \forall^s\forall^r, \exists^r\forall^s, \exists^s\forall^r\}$.

Definition 3.2.3 (Connector)**.** Let $S, R, C \subseteq A$. Define:[6]

$$
\begin{aligned}
\mathbf{C}^{S,R,C}_{\forall^r\exists^s} &:= \forall r \in R \exists s \in S\ r \to_{S\cup C\cup R} s &\qquad \mathbf{C}^{S,R,C}_{\exists^s\exists^r} &:= \exists s \in S \exists r \in R\ r \to_{S\cup C\cup R} s \\
\mathbf{C}^{S,R,C}_{\forall^s\exists^r} &:= \forall s \in S \exists r \in R\ r \to_{S\cup C\cup R} s &\qquad \mathbf{C}^{S,R,C}_{\forall^s\forall^r} &:= \forall s \in S \forall r \in R\ r \to_{S\cup C\cup R} s \\
\mathbf{C}^{S,R,C}_{\exists^r\forall^s} &:= \exists r \in R \forall s \in S\ r \to_{S\cup C\cup R} s &\qquad \mathbf{C}^{S,R,C}_{\exists^s\forall^r} &:= \exists s \in S \forall r \in R\ r \to_{S\cup C\cup R} s
\end{aligned}
$$

For $t \in \mathcal{T}$, the set C is a *t-connector from S to R* in $\mathfrak{F}$ iff $\mathbf{C}^{S,R,C}_t$ holds in $\mathfrak{F}$. We will omit in which frame C is a connector when it is clear from the context, or if the particular frame is irrelevant.

Recall that information travels in the reverse direction of the F relation: when $a\ F\ b$, the information travels from b to a. Therefore, the direction of the connector is the reverse of the direction of the F path it provides, so that the direction for the connector reflects the direction of information flow.

It will become clear that the six types of connectors all correspond to particular forms of knowledge realisation, with different assumed knowledge distributions pre-communication, and different resulting forms of knowledge post-communication. We already developed one form in (3.4), for $\forall^s\forall^r$-connectors.

To clarify the network-structural shapes of the different connector types, consider the following example.

[5] Six instead of eight because the order of quantification for a double $\exists$ or double $\forall$ does not matter.

[6] For overlapping S and R, the requirement of a path from agents in R to agents in S is a bit too strong as no path is needed for $s \in S$, $r \in R$ s.t. $s = r$. However, as Communication Logic cannot express identity of agents, we won't concern ourselves with such cases.

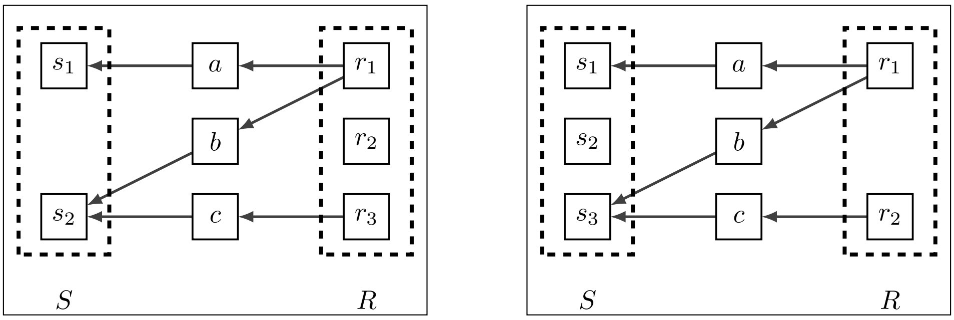

(a) (b)

Figure 3.2

Example 3.2.3. In Figure 3.2a, $\{a\}$ is a $\exists^s\exists^r$ connector from S to R, $\{a, c\}$ is a $\exists^s\exists^r$ and a $\forall^s\exists^r$ connector from S to R, and $\{a, b\}$ is a $\exists^s\exists^r$, $\forall^s\exists^r$, and $\exists^r\forall^s$ connector from S to R.

In Figure 3.2b, $\{a\}$ is a $\exists^s\exists^r$ connector from S to R, $\{a, c\}$ is a $\exists^s\exists^r$ and a $\forall^r\exists^s$ connector from S to R, and $\{a, b\}$ is a $\exists^s\exists^r$, $\forall^r\exists^s$, and $\exists^s\forall^r$ connector from S to R,

A connector is not required to be minimal, in that not necessarily all its elements are essential for the set to function as a connector. We will discuss such minimal connectors later. And, although it is the most intuitive setting, we don't require connectors to be disjoint from the sending and receiving set. However, a set is a connector iff its part disjoint from the sending and receiving set is.

Proposition 3.2.1 (Disjoint connectors)**.** *For sets $S, R \subseteq A$ and $t \in \mathcal{T}$, the set $C \subseteq A$ is a t-connector from S to R iff $C \setminus (S \cup R)$ is.*

Proof. Follows trivially by the definition and since $R \cup C \cup S = R \cup (C \setminus (S \cup R)) \cup S$. $\square$

The existence of a connector depends on the existence of paths from receiving agents to sending agents. Consequently, as the set of all agents A is finite, we can express connectors in Communication Logic by conjunctions and disjunctions of $\leq$Path formulas.

Definition 3.2.4 (Connector formulas)**.** Let $C, S, R \subseteq A$. Recall that $\leq\text{Path}(C, S, R) := \leq\text{Path}(C, S, R, |A| - 1)$. Define:

$$\forall^r\exists^s(C, S, R) := \bigwedge_{r\in R}\bigvee_{s\in S} \leq\text{Path}(G, r, s) \qquad \exists^s\exists^r(C, S, R) := \bigvee_{s\in S}\bigvee_{r\in R} \leq\text{Path}(G, r, s)$$

$$\forall^s\exists^r(C, S, R) := \bigwedge_{s\in S}\bigvee_{r\in R} \leq\text{Path}(G, r, s) \qquad \forall^s\forall^r(C, S, R) := \bigwedge_{s\in S}\bigwedge_{r\in R} \leq\text{Path}(G, r, s)$$

$$\exists^r\forall^s(C, S, R) := \bigvee_{r\in R}\bigwedge_{s\in S} \leq\text{Path}(G, r, s) \qquad \exists^r\forall^s(C, S, R) := \bigvee_{s\in S}\bigwedge_{r\in R} \leq\text{Path}(G, r, s)$$

Correctness of these formulas follows from the correctness of $\leq$Path.

Proposition 3.2.2 (Connector formula correctness)**.** *For a frame $\mathfrak{F}$, sets $S, R, C \subseteq A$, and $t \in \mathcal{T}$: C is a t-connector in $\mathfrak{F}$ iff $\mathfrak{F} \Vdash t(C, S, R)$.*

Proof. This follows from the correctness of $\leq$Path (Proposition 3.1.1). $\square$

As shown in Example 3.2.3, a connector can be of multiple types. The way the types of the particular connectors in this example overlap is no coincidence. Given that the sending and receiving set are non-empty, a $\forall^s\exists^r$-connector always also is a $\exists^r\forall^s$-connector, and a $\exists^r\forall^s$-connector also always is a $\exists^s\exists^r$-connector. This because, for connectors between non-empty sets, the connector types form a hierarchy that converges at the two extreme connector types: $\forall^s\forall^r$ and $\exists^s\exists^r$.

Proposition 3.2.3 (Connector hierarchy). *For non-empty sets $S, R \subseteq A$, and a set $C \subseteq A$:*

$$\Vdash \forall^s\forall^r(C,S,R) \to \exists^r\forall^s(C,S,R) \qquad \Vdash \exists^r\forall^s(C,S,R) \to \forall^s\exists^r(C,S,R)$$
$$\Vdash \forall^s\exists^r(C,S,R) \to \exists^s\exists^r(C,S,R)$$

$$\Vdash \forall^s\forall^r(C,S,R) \to \exists^s\forall^r(C,S,R) \qquad \Vdash \exists^s\forall^r(C,S,R) \to \forall^r\exists^s(C,S,R)$$
$$\Vdash \forall^r\exists^s(C,S,R) \to \exists^s\exists^r(C,S,R).$$

Proof. This follows from the definition of connectors, as for non-empty sets $S, R \subseteq A$ and $C \subseteq A$ $\forall s \in S \forall r \in R\ r \to_G s \implies \exists r \in R \forall s \in S\ r \to_G s \implies \forall s \in S \exists r \in R\ r \to_G s \implies \exists r \in R \exists s \in S\ r \to_G s$ and $\forall s \in S \forall r \in R\ r \to_G s \implies \exists s \in s \forall r \in R\ r \to_G s \implies \forall r \in R \exists s \in S\ r \to_G s \implies \exists r \in R \exists s \in S\ r \to_G s$. ☐

If we regard connectors from a group to another group, the six types of connectors become equivalent, in that the right direction of the above implications also holds.

Proposition 3.2.4 (Connector hierarchy over groups). *For non-empty sets $S, R \subseteq A$, set $C \subseteq A$, and $t, t' \in \mathcal{T}$:*

$$\Vdash \text{Group}(S) \wedge \text{Group}(R) \to t(C,S,R) \leftrightarrow t'(C,S,R).$$

Proof. As groups are sets such that all agents have a path to each other, having a path *to* some agent in a group is equivalent to having a path to all agents in that group. Similarly, having a path *from* an agent in some group is equivalent to having a path from all agents in that group. ☐

3.2.4 Connector Latency

As with groups, we can identify how many times a sender, receiver and connector must communicate before knowledge is realised. We call this the *latency* of a connector. Latency is not only an interesting property of a connector, but as a sending and receiving set can be connected by multiple connectors with the same type, it is also a quantitative way to distinguish such connectors from each other. And although, in the case of connectors of the same type, latency is only a quantitative difference between connectors that are qualitatively equal, latency will play a dominant role in social networks where communication happens scarcely.

For a connector from S to R, latency, of course, is a function of the length of the minimal path from agents in R to agents in S. Hence, we can define latency in terms of $r \to_G^{\leqslant n} s$, and consequently by quantification of the $\mathbf{C}_t^{S,R,C}$ conditionals in Definition 3.2.3.

Definition 3.2.5 (Connector latency). Let $S, R \subseteq A$, $n \in \mathbb{N}^+$, and:

$$\mathbf{C}^{n,S,R,C}_{\forall^r\exists^s} := \forall r \in R \exists s \in S\ r \to^{\leqslant n}_{S \cup C \cup R} s \qquad \mathbf{C}^{n,S,R,C}_{\exists^s\exists^r} := \exists s \in S \exists r \in R\ r \to^{\leqslant n}_{S \cup C \cup R} s$$
$$\mathbf{C}^{n,S,R,C}_{\forall^s\exists^r} := \forall s \in S \exists r \in R\ r \to^{\leqslant n}_{S \cup C \cup R} s \qquad \mathbf{C}^{n,S,R,C}_{\forall^s\forall^r} := \forall s \in S \forall r \in R\ r \to^{\leqslant n}_{S \cup C \cup R} s$$
$$\mathbf{C}^{n,S,R,C}_{\exists^r\forall^s} := \exists r \in R \forall s \in S\ r \to^{\leqslant n}_{S \cup C \cup R} s \qquad \mathbf{C}^{n,S,R,C}_{\exists^s\forall^r} := \exists s \in S \forall r \in R\ r \to^{\leqslant n}_{S \cup C \cup R} s$$

For $t \in \mathcal{T}$ and a frame $\mathfrak{F}$, we say that a connector $C \subseteq A$ from S to R is a t-n-connector from S to R (in $\mathfrak{F}$) iff $\mathbf{C}_t^{n,S,R,C}$ holds (in $\mathfrak{F}$).

The following example illustrates connector latency, and how it can quantitatively distinguish two connectors of the same type.

Example 3.2.4. Consider the network depicted in Figure 3.3. $\{c\}$ and $\{d\}$ both provide a path from an agent in G_r to an agent in G_s of length 2. Furthermore, both G_s and G_r are 2-groups. As G_s and G_r are groups, both $\{c\}$ and $\{d\}$ are t-connectors from G_s to G_r for

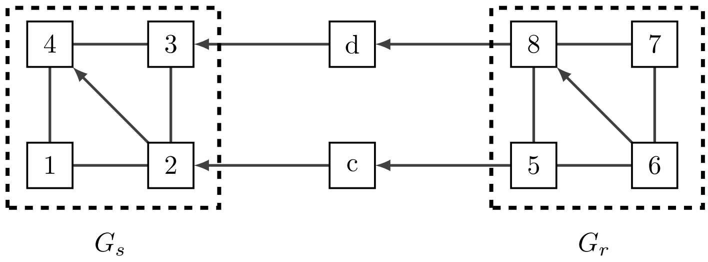

Figure 3.3

all $t \in \mathcal{T}$. Yet, there are situations in which G_s or G_r could prefer c over d and vice versa. c follows an agent that has a path of length 1 to every other agent in G_s, whereas d follows an agent that does not. c therefore is quicker at collecting all the information of the agents in G_s. d however is followed by an agent in G_r that is the only agent that is followed by every other agent in G_r. Hence, d will *distribute* the information it receives from G_s quicker through G_r.

This is reflected by the latencies of $\{c\}$ and $\{d\}$. $\{d\}$ is a $\exists^s\forall^r$-n-connector for $n \geq 3$, as the longest shortest path from any agent in G_r to 3 through $G_r \cup \{d\} \cup G_s$ has a length of 3 (of course d therefore is a $\exists^s\forall^r$-n-connector for all $n \geq 3$). $\{c\}$ however is not, as the longest shortest $G_s \cup \{c\} \cup G_r$-path from 7 to any agent in G_s has a length of 4. Instead, c is a $\exists^s\forall^r$-n-connector for $n \geq 4$. Conversely, $\{c\}$ is a $\forall^s\exists^r$-n-connector for $n \geq 3$, whereas $\{d\}$ is not a $\forall^s\exists^r$-3-connector. Instead, d is a $\forall^s\exists^r$-n-connector for $n \geq 4$.

For any t-connector, we can find a number n such that the connector has a latency of n.

Proposition 3.2.5 (Quantitative and qualitative connectors)**.** *C is a t-connector from S to R iff there is an $n \leq |A| - 1$ such that C is a t-n-connector.*

Proof. Because A is finite, the length of all minimal paths are finite as well. Therefore, for an $n \leq |A| - 1$ $s \to_G r \Leftrightarrow s \to_G^n$. ☐

Recall that the part of a connector disjoint from the sending and receiving set forms a connector of the same type (Proposition 3.2.1). The part of the connector disjoint from the sending and receiving sets also forms a connector of the same latency.

Proposition 3.2.6 (Disjoint connector latency)**.** *For $S, R \subseteq A$, and $t \in \mathcal{T}$, $C \subseteq A$ is a t-n-connector from S to R iff $C \setminus (S \cap R)$ is.*

Proof. For a connector C the definitions of latency all quantify over the binary relation $\to_{S \cup C \cup R}^{\leq n}$, and $S \cup C \cup R = S \cup (C \setminus (R \cup S)) \cup R$. ☐

As with the qualitative notion of connectors, we can express that a set is a connector with a certain latency in Communication Logic using $\leq$Path formulas.

Definition 3.2.6 (Connector latency formulas)**.** For $C, S, R \subseteq A$, and $n \in \mathbb{N}^+$:

$$\forall^r\exists^s(C,n,S,R) := \bigwedge_{r \in R} \bigvee_{s \in S} \leq\text{Path}(G,r,s,n) \quad \exists^s\exists^r(C,n,S,R) := \bigvee_{r \in R} \bigvee_{s \in S} \leq\text{Path}(G,r,s,n)$$
$$\forall^s\exists^r(C,n,S,R) := \bigwedge_{s \in S} \bigvee_{r \in R} \leq\text{Path}(G,r,s,n) \quad \forall^s\forall^r(C,n,S,R) := \bigwedge_{s \in S} \bigwedge_{r \in R} \leq\text{Path}(G,r,s,n)$$
$$\exists^r\forall^s(C,n,S,R) := \bigvee_{r \in R} \bigwedge_{s \in S} \leq\text{Path}(G,r,s,n) \quad \exists^r\forall^s(C,n,S,R) := \bigvee_{s \in S} \bigwedge_{r \in R} \leq\text{Path}(G,r,s,n)$$

The connector latency formulas are correct by the correctness of the $\leq$Path formulas.

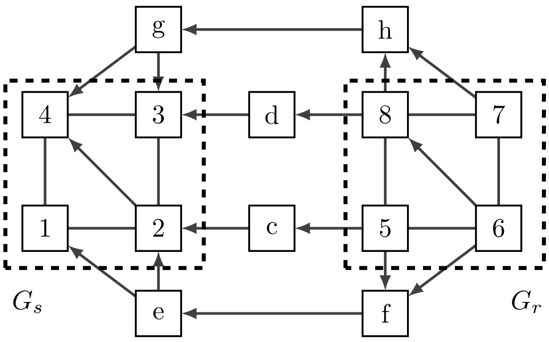

(a)

	$\{c\}$	$\{d\}$	$\{e, f, g, h\}$	A
$\forall^s\forall^r$	5	5	4	4
$\exists^r\forall^s$	3	4	4	3
$\forall^s\exists^r$	3	4	3	3
$\exists^s\forall^r$	4	3	4	3
$\forall^r\exists^s$	4	3	3	3
$\exists^s\exists^r$	2	2	3	2

(b)

Figure 3.4: A table (on the right) of the latencies of connectors from G_s to G_r of the network displayed on the left.

Proposition 3.2.7 (Connector latency formula correctness). *For a frame $\mathfrak{F}$, non-empty $S, R \subseteq A$, and for $t \in \mathcal{T}$: C is a t-n-connector from S to R in $\mathfrak{F}$ iff $\mathfrak{F} \Vdash Q_1^x Q_2^y(C, n, S, R)$.*

Proof. C is a t-n-connector from S to R in $\mathfrak{F}$ iff $\forall s \in S \exists r \in R r \to_G^{\leqslant n} s$. This is exactly what $\forall^s\exists^r(C, n, S, R)$ states. Similarly for the other types of connectors. $\square$

The connector hierarchy also applies to quantitatively delineated connectors.

Proposition 3.2.8 (Connector latency hierarchy). *Regard the non-empty sets $S, R \subseteq A$, and the set $B \subseteq A$. Then:*

$$\Vdash \forall^s\forall^r(C, n, S, R) \to \exists^r\forall^s(C, n, S, R) \qquad \Vdash \exists^r\forall^s(C, n, S, R) \to \forall^s\exists^r(C, n, S, R)$$
$$\Vdash \forall^s\exists^r(C, n, S, R) \to \exists^s\exists^r(C, n, S, R)$$

$$\Vdash \forall^s\forall^r(C, n, S, R) \to \exists^s\forall^r(C, n, S, R) \qquad \Vdash \exists^s\forall^r(C, n, S, R) \to \forall^r\exists^s(C, n, S, R)$$
$$\Vdash \forall^r\exists^s(C, n, S, R) \to \exists^s\exists^r(C, n, S, R).$$

Proof. This follows trivially because Proposition 3.2.2 extends to implications of the quantified versions of $\mathbf{C}_t^{S,R,C}$. $\square$

If the sending and receiving sets are groups, then the connector types become equivalent in that any t-connector between that sender and receiver will also be a t'-connector of any other type, as shown in Proposition 3.2.4. Even so, the typed *latencies* of the connectors could be distinct from each other, in that if the sending and receiving sets are groups, then a t-n-connector will, in general, not also be a t'-n-connector for another type t'.[7] That is, not all the right-to-left directions of the implications in Proposition 3.2.8 follow when S and R groups. This holds for *all* six types of latency. We will show this by a more elaborate example of connector latency based on Example 3.2.4.

Example 3.2.5. Consider the network shown in Figure 3.4a. In this network, $G_s = \{1, 2, 3, 4\}$ and $G_r = \{5, 6, 7, 8\}$ form 2-groups Furthermore $\{c\}$, $\{d\}$. $\{c, d, e, f\}$, and $\{c, d, e, f, g, h\}$ form connectors (among others) from G_s to G_r. In Table 3.4b their latencies are shown.

$\{c\}$ and $\{d\}$ are equal with respect to $\exists^s\exists^r$ and $\forall^s\forall^r$-latency. However, as noted in Example 3.2.4 $\{c\}$ is quicker at collecting knowledge from all of G_s whereas $\{d\}$ is better in distributing the knowledge among G_s. This is reflected in the difference between $\{c\}$ and $\{d\}$ with respect to $\forall^s\exists^r$ and $\forall^r\exists^s$-latency — with respect to the former $\{c\}$ is quicker than

[7]For any type t', a t-n-connector will however be a t'-m-connector for some m possibly distinct from n.

$\{d\}$ and with respect to latter $\{d\}$ is quicker. $\{e, f, g, h\}$ connects all agents in G_r to an agent in G_s by a path with a length of at most 3. Therefore, it has a $\forall^r\exists^s$ of 3. Similarly, it has a $\forall^s\exists^r$-latency of 3 and a $\exists^s\exists^r$-latency of 3. The $\forall^s\forall^r$-latency of $\{e, f, g, h\}$ is 4 since the path from 7 to 1 through $G_s \cup \{e, f, g, h\} \cup G_r$ has length of 4. Similarly, it has a $\exists^r\forall^s$ and $\exists^s\forall^r$-latency of 4 by for example the shortest path from 7 to 1.

Finally, $\{c, d, e, f, g, h\}$ has the minimum latency in each respective category out of all previous discussed connectors as it contains all these connectors.

If the sending and receiving sets are groups (as in example 3.2.4), the "widths" of the sending and receiving groups form an upper bound on the differences between the latencies of the different connector types.

Proposition 3.2.9 (Connector latency and groups)**.** *If R is an l-group in $\mathfrak{F}$ then for $m > l$:*

$$\mathfrak{F} \Vdash \exists^s\exists^r(C, n, S, R) \to \forall^r\exists^s(C, n+l, S, R) \quad \mathfrak{F} \Vdash \exists^s\exists^r(C, n, S, R) \to \exists^s\forall^r(C, n+l, S, R)$$
$$\mathfrak{F} \Vdash \exists^r\forall^s(C, n, S, R) \to \forall^s\forall^r(C, n+l, S, R) \quad \mathfrak{F} \Vdash \forall^s\exists^r(C, n, S, R) \to \forall^s\forall^r(C, n+l, S, R)$$

And, if S is a k-group in $\mathfrak{F}$ then for $n > k$:

$$\mathfrak{F} \Vdash \exists^s\exists^r(C, n, S, R) \to \forall^s\exists^r(C, n+k, S, R) \quad \mathfrak{F} \Vdash \exists^s\exists^r(C, n, S, R) \to \exists^r\forall^s(C, n+k, S, R)$$
$$\mathfrak{F} \Vdash \exists^s\forall^r(C, n, S, R) \to \forall^s\forall^r(C, n+k, S, R) \quad \mathfrak{F} \Vdash \forall^r\exists^s(C, n, S, R) \to \forall^s\forall^r(C, n+k, S, R)$$

Proof. Let $m = n + k$ and $o = n + l$. If S and R are a k-group and l-group respectively, then by the defining property of groups we have that:

$$\mathbf{C}^{n,S,R,C}_{\exists^r\forall^s} \Rightarrow \mathbf{C}^{o,S,R,C}_{\forall^s\forall^r} \quad \mathbf{C}^{n,S,R,C}_{\forall^s\exists^r} \Rightarrow \mathbf{C}^{o,S,R,C}_{\forall^s\forall^r} \quad \mathbf{C}^{n,S,R,C}_{\exists^s\forall^r} \Rightarrow \mathbf{C}^{m,S,R,C}_{\forall^s\forall^r} \quad \mathbf{C}^{n,S,R,C}_{\forall^r\exists^s} \Rightarrow \mathbf{C}^{m,S,R,C}_{\forall^s\forall^r}$$
$$\mathbf{C}^{n,S,R,C}_{\exists^s\exists^r} \Rightarrow \mathbf{C}^{o,S,R,C}_{\forall^r\exists^s} \quad \mathbf{C}^{n,S,R,C}_{\exists^s\exists^r} \Rightarrow \mathbf{C}^{o,S,R,C}_{\exists^s\forall^r} \quad \mathbf{C}^{n,S,R,C}_{\exists^s\exists^r} \Rightarrow \mathbf{C}^{m,S,R,C}_{\exists^r\forall^s} \quad \mathbf{C}^{n,S,R,C}_{\exists^s\exists^r} \Rightarrow \mathbf{C}^{m,S,R,C}_{\forall^s\exists^r}$$

□

By Proposition 3.2.9, it could be tempting to assume that for connectors, some of the six types of latency can be calculated from others by adding the "width" of the groups they connect. This is not the case however, as there are examples where the right-to-left directions of the implications in Proposition 3.2.9 do not hold. Again, Example 3.2.4 adequately shows this: here G_s and G_r are both 2-groups, the connector $\{d\}$ has a $\forall^r\exists^s$-latency of 4 and a $\exists^s\exists^r$-latency of 3. For the right-to-left direction between the other types, different connectors in Example 3.2.4 serve as counter-examples.

3.2.5 Connectors and Knowledge Realisation

As conjectured in formula (3.5), connectors are sets of agents such that the senders, receivers and the connector can realise distributed knowledge among the senders and receivers to individual knowledge among the receivers. We will show that the specific type of connector that allows exactly this is $\forall^s\forall^r$. The other types correspond to different forms of knowledge realisation, where the distribution of knowledge among the senders and receivers pre-communication and the receivers post communication differ slightly. The resulting knowledge distribution of these types of knowledge realisation are of a stronger form than distributed knowledge. To show this conjecture, we first analyse the semantic implications of connectors.

Without assuming a specific frame (or network topology), if C is a connector from S to R with a latency of n, then after S, R, and C communicate n times, the similarity relations of agents in R are restricted with the similarity relations of agents in S. The exact agents this concerns depend on the type of connector. More formally, define the following propositions:

Definition 3.2.7 (Type similarity inclusion)**.** For non-empty sets $S, R \subseteq A$, a set $C \subseteq A$, $n \in \mathbb{N}^+$, let $G = S \cup C \cup R$, and define

$$
\begin{aligned}
&\subseteq^{n,S,R,C}_{\forall^r\exists^s} := \forall r \in R \exists s \in S \ \sim_r^{!G\otimes n} \subseteq \sim_s &\quad& \subseteq^{n,S,R,C}_{\exists^s\exists^r} := \exists s \in S \exists r \in R \ \sim_r^{!G\otimes n} \subseteq \sim_s \\
&\subseteq^{n,S,R,C}_{\forall^s\exists^r} := \forall s \in S \exists r \in R \ \sim_r^{!G\otimes n} \subseteq \sim_s &\quad& \subseteq^{n,S,R,C}_{\forall^s\forall^r} := \forall s \in S \forall r \in R \ \sim_r^{!G\otimes n} \subseteq \sim_s \\
&\subseteq^{n,S,R,C}_{\exists^r\forall^s} := \exists r \in R \forall s \in S \ \sim_r^{!G\otimes n} \subseteq \sim_s &\quad& \subseteq^{n,S,R,C}_{\exists^s\forall^r} := \exists s \in S \forall r \in R \ \sim_r^{!G\otimes n} \subseteq \sim_s
\end{aligned}
$$

Proposition 3.2.10 (Connector Semantic)**.** *For non-empty sets $S, R \subseteq A$ a frame $\mathfrak{F}$, a set $C \subseteq A$, $n \in \mathbb{N}^+$, and $G = S \cup C \cup R$: if C is a t-connector from S to R with a latency of n then $\subseteq^{n,S,R,C}_t$ holds.*

Proof. Let $G = S \cup R \cup C$. Recall that C has an t-latency of n iff $\mathbf{C}^{n,S,R,C}_t$ holds. Let $m \in \mathbb{N}^+$, $r \in R$, and $s \in S$. By the definition of paths and $\sim_r^{!G\otimes m}$, and the characterisation of $\mathcal{F}|^{+m}_G$ (Proposition 2.4.5): $r \to^{\leqslant m}_G s \Longrightarrow s \in \mathcal{F}|^{+m}_G(r) \Longrightarrow \sim_r^{!G\otimes n} \subseteq \sim_s$. This proves that $\mathbf{C}^{n,S,R,C}_t \Longrightarrow \subseteq^{n,S,R,C}_t$. ☐

Like the semantic results for groups, these semantic results for connectors have a syntactic counterpart related to knowledge realisation. We will present these in the form of formulas of Communication Logic. As conjectured in the beginning of this chapter, the formulas take on the form $(\xi \wedge \chi) \to \mu$, where ξ is the topological precondition, χ is the epistemic precondition, and μ is the epistemic result of knowledge realisation.

The topological preconditions state the existence of a connector with a latency of n. For knowledge to be realised, a combination of senders and receivers together must distributively know some formula $[!(S \cup C \cup R)]^n \varphi$. This is the epistemic precondition χ. The epistemic results are of the forms defined before, either knowledge realisation towards individual knowledge by all in R, E_R; individual knowledge by some in R, "somebody in R knows", S_R; or distributed knowledge among R, D_R. Consequently, μ in $(\xi \wedge \chi) \to \mu$ takes on the form $[!(C \cup R \cup S)]^m \Box_R \varphi$, where $\Box$ is a knowledge modality, one of D, E, or S.

We can deduce the knowledge realisation of formulas from the $\subseteq^{n,S,R,C}_t$ conditions in Proposition 3.2.10. This results from the following: if the similarity relation of a set X is contained in the similarity relation of a set Y after Z iteratively communicates n times, then if $X \cup Y$ distributively know that $[!Z]^n \varphi$, then X will distributively know that φ after Z communicates n times.

Lemma 3.2.1. *For any frame $\mathfrak{F}$ and $X, Y, Z \subseteq A$ such that $\sim_X^{!Z\otimes n} \subseteq \sim_Y$:*

$$\mathfrak{F} \Vdash D_{X\cup Y}[!Z]^n \varphi \to [!Z]^n D_X \varphi$$

Proof. This follows from the semantics of communication and distributed knowledge, worked out in the previous chapter. ☐

Recall that the six connector types form a hierarchy. We will go through the six types of connectors by order of type, listing their specific epistemic precondition ξ, and epistemic post-condition μ. The hierarchy is shaped as two diverging branches from the weakest type $\exists^s\exists^r$, converging again at the strongest type $\forall^s\forall^r$. We start from $\exists^s\exists^r$, first expanding the $\exists^s\exists^r$–$\forall^s\exists^r$–$\exists^r\forall^s$–$\forall^s\forall^r$ branch, and then the leftover types $\forall^r\exists^s$ and $\exists^s\forall^r$ from the $\exists^s\exists^r$–$\forall^r\exists^s$–$\exists^s\forall^r$–$\forall^s\forall^r$ branch. At each level of the hierarchy, from weakest to strongest type, the connector types correspond to an increasing number of $\mathcal{L}_{CL}$-formulas. Formulas for $\forall^s\exists^r$ also holds for $\exists^r\forall^s$ and $\forall^s\forall^r$-connectors, but not for $\exists^s\exists^r$ connectors etc. At each stage of this hierarchy, we therefore only develop the formulas that belong to this stage up. We do not show that these formulas are not valid for the lower stages, but counterexamples do arise out of Figure 3.2, when paired with appropriate similarity relations.

Connectors of the weakest type ($\exists^s\exists^r$) correspond to two forms of knowledge realisation: towards distributed knowledge or towards individual knowledge by some. To come to distributed knowledge of φ in R, a particular s, together with all receivers R must distributively know that after s, R, and C communicate n times, φ holds. Because $\exists^s\exists^r$-connectors abstract away form the particular s they connect from or to, it must hold that for all $s \in S$, $D_{\{s\}\cup R}[!(S \cup R \cup C)]^n\varphi$.[8] To come to individual knowledge by some receiver, by similar reasoning, it must hold that for all $s \in S$ and $r \in R$, s and r together know that after S, C, and R communicate n times, φ must hold.

Proposition 3.2.11. *For non-empty sets $S, R \subseteq A$, a set $C \subseteq A$, $n \in \mathbb{N}^+$, and $\varphi \in \mathcal{L}_{CL}$:*

$$\Vdash \left(\exists^s\exists^r(C, n, S, R) \wedge \bigwedge_{s\in S} D_{\{s\}\cup R}[!(S \cup C \cup R)]^n\varphi\right) \to [!(S \cup C \cup R)]^n D_R\varphi$$

$$\Vdash \left(\exists^s\exists^r(C, n, S, R) \wedge \bigwedge_{s\in S}\bigwedge_{r\in R} D_{\{s,r\}}[!(S \cup C \cup R)]^n\varphi\right) \to [!(S \cup C \cup R)]^n S_R\varphi$$

Proof. If C is a $\exists^s\exists^r$-n-connector from S to R, then $\subseteq_{\exists^s\exists^r}$ holds by Proposition 3.2.10. Therefore, $\exists s \in S \exists r \in R \ \sim_r^{!(S\cup C\cup R)\otimes n} \subseteq \sim_s$.

(1) $\exists s \in S \exists r \in R \ \sim_r^{!(S\cup C\cup R)\otimes n} \subseteq \sim_s$ implies $\exists s \in S \ \sim_R^{!(S\cup C\cup R)\otimes n} \subseteq \sim_s$. Call this s: $s^\star$. Assume that $\mathfrak{F} \Vdash \bigwedge_{s\in S} D_{\{s\}\cup R}[!(S \cup C \cup R)]^n\varphi$. We get that $\mathfrak{F} \Vdash [!(S \cup C \cup R)]D_R\varphi$, by taking $X = R$, $Y = \{s^\star\}$, and $Z = S \cup C \cup R$ in Lemma 3.2.1.

(2) Call these $s^\star$ and $r^\star$ respectively. Assume that $\mathfrak{F} \Vdash \bigwedge_{s\in S}\bigwedge_{r\in R} D_{\{s,r\}}[!(S \cup C \cup R)]^n\varphi$. Therefore, in particular, $\mathfrak{F} \Vdash D_{\{s^\star,r^\star\}}[!G]^n\varphi$. We get that $\mathfrak{F} \Vdash [!(S \cup C \cup R)]K_{r^\star}\varphi$, by taking $X = \{r^\star\}$, $Y = \{s^\star\}$, and $Z = S \cup C \cup R$ in Lemma 3.2.1. ☐

Like $\exists^s\exists^r$-connectors, $\forall^s\exists^r$-connectors facilitate knowledge realisation towards D_R and S_R. As $\forall^s\exists^r$-connectors are stronger than $\exists^s\exists^r$-connectors, the epistemic preconditions for these forms of knowledge realisation are weaker than the preconditions above: $\forall^s\exists^r$-connectors are able to transfer the knowledge of all senders to some receiver. Hence, to reach distributed knowledge in R, it is enough for S and R to distributively know that $[!(S \cup C \cup R)]^n\varphi$. Likewise, to reach S_R, *some* sender together with any receiver must distributively know that $[!(S \cup C \cup R)]^n\varphi$.

Proposition 3.2.12. *For non-empty sets $S, R \subseteq A$, a set $C \subseteq A$, $n \in \mathbb{N}^+$, and $\varphi \in \mathcal{L}_{CL}$:*

$$\Vdash (\forall^s\exists^r(C, n, S, R) \wedge D_{S\cup R}[!(S \cup C \cup R)]^n\varphi) \to [!(S \cup C \cup R)]^n D_R\varphi$$

$$\Vdash \left(\forall^s\exists^r(C, n, S, R) \wedge \bigvee_{s\in S}\bigwedge_{r\in R} D_{\{s,r\}}[!(S \cup C \cup R)]^n\varphi\right) \to [!(S \cup C \cup R)]^n S_R\varphi$$

Proof. If C is a $\forall^s\exists^r$-n-connector from S to R, then $\subseteq_{\forall^s\exists^r}$ holds by Proposition 3.2.10. Therefore, $\forall s \in S \exists r \in R \ \sim_r^{!(S\cup C\cup R)\otimes n} \subseteq \sim_s$.

(1) Hence $\sim_R^{!(S\cup C\cup R)\otimes n} \subseteq \sim_S$. Assume that $\mathfrak{F} \Vdash D_{S\cup R}[!(S \cup C \cup R)]^n\varphi$. We get that $\mathfrak{F} \Vdash [!(S \cup C \cup R)]D_R\varphi$, by taking $X = R$, $Y = S$, and $Z = S \cup C \cup R$ in Lemma 3.2.1.

(2) For each $s \in S$, call this r: r_s. Assume that $\mathfrak{F} \Vdash \bigvee_{s\in S}\bigwedge_{r\in R} D_{\{s,r\}}[!(S \cup C \cup R)]^n\varphi$. Then, in particular, $\mathfrak{F} \Vdash \bigvee_{s\in S} D_{\{s,r_s\}}[!(S \cup C \cup R)]^n\varphi$. Let $s^\star$ be an s that satisfies this. We get that $\mathfrak{F} \Vdash [!(S \cup C \cup R)]S_R\varphi$, by taking $X = \{r_{s^\star}\}$, $Y = \{s^\star\}$, and $Z = S \cup C \cup R$ in Lemma 3.2.1. ☐

Connectors of the stronger $\exists^r\forall^s$ type are able to transfer all knowledge of S to one particular receiver (as opposed to possibly distinct individuals for $\forall^s\exists^r$-connectors). For $\exists^r\forall^s$-connectors, therefore, the epistemic precondition for knowledge realisation towards S_R is even weaker: all receivers r together with S must distributively know that $[!(S \cup C \cup R)]^n\varphi$.

[8] Of course, we could "pull" this s inwards in the formula below, so that we get $(\bigvee_{s\in S}(\bigvee_{r\in R} \leq\mathrm{Path}(S \cup C \cup R, r, s, n)) \wedge D_{\{s\}\cup R}[!(S \cup C \cup R)]^n\varphi) \to [!(S \cup C \cup R)]^n D_R\varphi$.

Proposition 3.2.13. *For non-empty sets $S, R \subseteq A$, a set $C \subseteq A$, $n \in \mathbb{N}^+$, and $\varphi \in \mathcal{L}_{CL}$:*

$$\Vdash \left(\exists^r \forall^s (C, n, S, R) \wedge \bigwedge_{r \in R} D_{S \cup \{r\}} [!(S \cup C \cup R)]^n \varphi \right) \to [!(S \cup C \cup R)]^n S_R \varphi$$

Proof. If C is a $\exists^r\forall^s$-n-connector from S to R, then $\subseteq_{\exists^r\forall^s}$ holds by Proposition 3.2.10. Therefore, $\exists r \in R \forall s \in S \sim_r^{!(S \cup C \cup R) \otimes n} \subseteq \sim_s$. Hence, $\exists r \in R \sim_r^{!(S \cup C \cup R) \otimes n} \subseteq \sim_S$. Call this r $r^\star$. Assume that $\mathfrak{F} \Vdash \bigwedge_{r \in R} D_{S \cup \{r\}} [!(S \cup C \cup R)]^n \varphi$. We get that $\mathfrak{F} \Vdash [!(S \cup C \cup R)] K_{r^\star} \varphi$, by taking $X = \{r^\star\}$, $Y = S$, and $Z = S \cup C \cup R$ in Lemma 3.2.1. ☐

Connectors of the strongest type, $\forall^s\forall^r$, correspond to knowledge realisation towards individual knowledge by some, or by all receivers. The epistemic preconditions for these are respectively that some or all receiver together with all senders distributively know that $[!(S \cup C \cup R)]^n \varphi$.

Proposition 3.2.14. *For non-empty sets $S, R \subseteq A$, a set $C \subseteq A$, $n \in \mathbb{N}^+$, and $\varphi \in \mathcal{L}_{CL}$:*

$$\Vdash \left(\forall^s \forall^r (C, n, S, R) \wedge \bigwedge_{r \in R} D_{S \cup \{r\}} [!(S \cup C \cup R)]^n \varphi \right) \to [!(S \cup C \cup R)]^n E_R \varphi$$

$$\Vdash \left(\forall^s \forall^r (C, n, S, R) \wedge \bigvee_{r \in R} D_{S \cup \{r\}} [!(S \cup C \cup R)]^n \varphi \right) \to [!(S \cup C \cup R)]^n S_R \varphi$$

Proof. If C is a $\forall^s\forall^r$-n-connector from S to R, then $\subseteq_{\forall^s\forall^r}$ holds by Proposition 3.2.10. Therefore, $\forall s \in S \forall r \in R \sim_r^{!(S \cup C \cup R) \otimes n} \subseteq \sim_s$. Hence, $\sim_r^{!(S \cup C \cup R) \otimes n} \subseteq \sim_S$.

(1) Assume that $\mathfrak{F} \Vdash \bigwedge_{r \in R} D_{S \cup \{r\}} [!(S \cup C \cup R)]^n \varphi$. For all $r \in R$ we get that $\mathfrak{F} \Vdash [!(S \cup C \cup R)] K_r \varphi$, by taking $X = \{r\}$, $Y = S$, and $Z = S \cup C \cup R$ in Lemma 3.2.1.

(2) Assume that $\mathfrak{F} \Vdash \bigvee_{r \in R} D_{S \cup \{r\}} [!(S \cup C \cup R)]^n \varphi$. Call an r that satisfies this $r^\star$. We get that $\mathfrak{F} \Vdash [!(S \cup C \cup R)] K_{r^\star} \varphi$, by taking $X = \{r^\star\}$, $Y = S$, and $Z = S \cup C \cup R$ in Lemma 3.2.1. ☐

This concludes the $\exists^s\exists^r$–$\forall^s\exists^r$–$\exists^r\forall^s$–$\forall^s\forall^r$ branch of connector types. For the other branch, the remaining types to discuss are $\forall^r\exists^s$ and $\exists^s\forall^r$; of which the former is the weakest.

$\forall^r\exists^s$-connectors correspond to knowledge realisation towards S_R and E_R. For the former, it is required that for all senders, there is some receiver such that together they distributively know that $[!(S \cup C \cup R)]^n \varphi$. For the latter, it is required that any sender and receiver together distributively know that $[!(S \cup C \cup R)]^n \varphi$.

Proposition 3.2.15. *For non-empty sets $S, R \subseteq A$, a set $C \subseteq A$, $n \in \mathbb{N}^+$, and $\varphi \in \mathcal{L}_{CL}$:*

$$\Vdash (\forall^r \exists^s (C, n, S, R) \wedge \bigwedge_{s \in S} \bigwedge_{r \in R} D_{\{s,r\}} [!(S \cup C \cup R)]^n \varphi) \to [!(S \cup C \cup R)]^n E_R \varphi$$

$$\Vdash (\forall^r \exists^s (C, n, S, R) \wedge \bigwedge_{s \in S} \bigvee_{r \in R} D_{\{s,r\}} [!(S \cup C \cup R)]^n \varphi) \to [!(S \cup C \cup R)]^n S_R \varphi$$

Proof. If C is a $\forall^r\exists^s$-n-connector from S to R, then $\subseteq_{\forall^r\exists^s}$ holds by Proposition 3.2.10. Therefore, $\forall r \in R \exists s \in S \sim_r^{!(S \cup C \cup R) \otimes n} \subseteq \sim_s$. Call the element of S that makes this statement true for $r \in R$: s_r.

(1) Assume that $\mathfrak{F} \Vdash \bigwedge_{s \in S} \bigwedge_{r \in R} D_{\{s,r\}} [!(S \cup C \cup R)]^n \varphi$, then in particular we have that $\bigwedge_{r \in R} D_{\{s_r,r\}} [!(S \cup C \cup R)]^n \varphi$. For all $r \in R$ we get that $\mathfrak{F} \Vdash [!(S \cup C \cup R)] K_r \varphi$, by taking $X = \{r\}$, $Y = \{s_r\}$, and $Z = S \cup C \cup R$ in Lemma 3.2.1.

(weakest) type	epistemic precondition (ξ)	result (μ)	Proposition
$\forall^s\forall^r$	$\bigwedge_{r\in R} D_{S\cup\{r\}}[!G]^n\varphi$	$[!G]^n E_R\varphi$	3.2.14
	$\bigvee_{r\in R} D_{S\cup\{r\}}[!G]^n\varphi$	$[!G]^n S_R\varphi$	
$\exists^r\forall^s$	$\bigwedge_{r\in R} D_{S\cup\{r\}}[!G]^n\varphi$	$[!G]^n S_R\varphi$	3.2.13
$\forall^s\exists^r$	$\bigvee_{s\in S}\bigwedge_{r\in R} D_{\{s,r\}}[!G]^n\varphi$	$[!G]^n S_R\varphi$	3.2.12
	$D_{S\cup R}[!G]^n\varphi$	$[!G]^n D_R\varphi$	
$\exists^s\exists^r$	$\bigwedge_{s\in S}\bigwedge_{r\in R} D_{\{s,r\}}[!G]^n\varphi$	$[!G]^n S_R\varphi$	3.2.11
	$\bigwedge_{s\in S} D_{\{s\}\cup R}[!G]^n\varphi$	$[!G]^n D_R\varphi$	
$\forall^r\exists^s$	$\bigwedge_{s\in S}\bigwedge_{r\in R} D_{\{s,r\}}[!G]^n\varphi$	$[!G]^n E_R\varphi$	3.2.15
	$\bigwedge_{s\in S}\bigvee_{r\in R} D_{\{s,r\}}[!G]^n\varphi$	$[!G]^n S_R\varphi$	

Table 3.1: Summary of connector formulas

(2) Assume that $\mathfrak{F} \Vdash \bigvee_{r\in R}\bigwedge_{s\in S} D_{\{s,r\}}[!(S\cup C\cup R)]^n\varphi$, then in particular we have that $\bigvee_{r\in R} D_{\{s_r,r\}}[!(S\cup C\cup R)]^n\varphi$. Let $r^\star$ be an $r \in R$ that satisfies this. We get that $\mathfrak{F} \Vdash [!(S\cup C\cup R)]K_{r^\star}\varphi$, by taking $X = \{r^\star\}$, $Y = \{s_{r^\star}\}$, and $Z = S\cup C\cup R$ in Lemma 3.2.1. □

Finally, there are no formulas representing knowledge realisation that hold for $\exists^s\forall^r$-connectors, that we have not already discussed.[9] This because of the particular (first-order) abstraction we consider in the precondition of knowledge realisation: that either *all* or *some* senders together with all or some receivers distributively know something. For $\exists^s\forall^r$-connectors, the knowledge realisation formulas with these preconditions are already implied by the formulas for the weaker type $\forall^r\exists^s$. Moreover, the variant of knowledge realisation for $\exists^s\forall^r$-connectors that, for the senders, is not abstracted to sets, i.e. the variant whose precondition states that a specific sender s together with any receiver know something, has already (implicitly) been discussed: $\exists^s\forall^r(C,n,S,R)$ implies that for some particular s, $\forall^s\forall^r(C,n,\{s\},R)$ holds, and hence that the formulas in Proposition 3.2.14 hold when $S = \{s\}$. This form of knowledge realisation *is* unique to connectors of the $\exists^s\forall^r$ type (and the stronger $\forall^s\forall^r$ type), as it does not hold for $\forall^r\exists^s$-connectors.

For a summary of the kinds of knowledge realisation made possible each connector type, see Table 3.1. When moving down this table, the connector type necessary for the knowledge realisation (of ξ to μ) becomes weaker. At the same time, either the epistemic result μ gets weaker, or the epistemic precondition ξ gets stronger.

3.2.6 Less Chatty Communication

The particular forms of knowledge realisation we presented for connectors are very "chatty": the senders, receivers, and the connector all communicate n times. Most often however, less communication is required. Only n small subsets of these agents have to communicate subsequently for knowledge realisation to take place. As stated in Section 2.4, this has been a matter of simplification for the sake of clarity and presentation. We will not make an attempt to distill the actual, minimally required, communication from the formulas of the last section, even though this should be an easy task given the simple structure of connectors. We will, however, make two remarks on connector types that require iterative communication by a smaller subset of agents, as this does not complicate things too much.

The formulas of connectors whose type has an existential quantification over the receivers ($\exists^s\exists^r$, $\forall^s\exists^r$, and $\exists^r\forall^s$) will also hold if we replace the communication modalities $[S\cup C\cup R]$ with $[S\cup C]$. For such connector types, the receivers do not have to communicate for the respective knowledge realisation of these types to happen.

[9] As we will see in the next section, there *are* formulas that hold for the negation of $\exists^s\forall^r$-connectors, that do not hold for the negation of $\forall^r\exists^s$-connectors.

Proposition 3.2.16 (Silent receivers)**.** *For a connector of type* $t \in \{\exists^s\exists^r, \forall^s\exists^r, \exists^r\forall^s\}$*, Proposition 3.2.10, and subsequently the formulas of Proposition 3.2.11, 3.2.12, and 3.2.13 also hold if we replace* $[!(S \cup C \cup R)]^n$ *with* $[!(S \cup C)]^n$*.*

Proof. This follows from the proof of Proposition 3.2.10 and since if there is a path of length n from some $r \in R$ to some s through $S \cup R \cup C$, then there also is a path from some $r' \in R$ to that s through $S \cup C$ of length $\leq n$. Hence, $\exists r \in R\ r \rightarrow^{\leqslant n}_{S \cup C \cup R} s \Rightarrow \exists r \in R\ r \rightarrow^{\leqslant n}_{S \cup C} s \Rightarrow \exists r \in R\ s \in \mathcal{F}|^{+}_{S \cup C}{}^{n}(r)$, by the characterisation of $\mathcal{F}|^{+}_{S \cup C}{}^{n}$ (Proposition 2.4.5). ☐

Something similar also holds for the senders under other connector types: the formulas of connectors whose type has an existential quantification over the senders ($\exists^s\exists^r$, $\exists^s\forall^r$, $\forall^r\exists^s$) will also hold if we replace the iterated communication modalities $[S \cup C \cup R]^n$ with $[S][C \cup R]^{n-1}$. For such connector types, the senders only have to communicate once, simultaneously, for the respective knowledge realisation of these types to happen.

Proposition 3.2.17 (Reserved senders)**.** *For a connector of a type* $t \in \{\exists^s\exists^r, \exists^s\forall^r, \forall^r\exists^s\}$ *Proposition 3.2.10, and subsequently the formulas of Proposition 3.2.11 and 3.2.15, also hold if we replace* $[!(S \cup C \cup R)]^n$ *with* $[!S][!(C \cup R)]^{n-1}$*.*

Proof. This follows from the proof of Proposition 3.2.10 and since if there is a path of length n from some $r \in R$ to some s through $S \cup R \cup C$, then there also is a path from that r to some $s' \in S$ through $R \cup C$ of length $\leq n$. Hence, $\exists s \in S\ r \rightarrow^{\leqslant n}_{S \cup C \cup R} s \Rightarrow \exists x \in R \cup C \exists s \in S\ r \rightarrow^{\leqslant n-1}_{R \cup C} x$ and $x\ F\ s \Rightarrow \exists s \in S\ s \in \mathcal{F}|^{+}_{S}(\mathcal{F}|^{+}_{R \cup C}{}^{n-1}(r))$ by Proposition 2.4.5. ☐

Finally, there is a specific case of connectors that is special, when knowledge can be realised by communication of only the sending and receiving sets. In such cases, $\emptyset$ is a t-connector from S to R. We call the sets R and S *directly t-connected* in these cases.

Definition 3.2.8 (Direct connectedness)**.** For sets $S, R \subseteq A$, $t \in \mathcal{T}$, and $n \in \mathbb{N}^+$ S is directly t-connected to R iff $\emptyset$ is a t-connector from S to R. Furthermore, S is directly t-n-connected to R iff $\emptyset$ is a t-n-connector from S to R.

Senders S and receivers R that are directly connected are able to realise distributed knowledge of $S \cup R$ to knowledge of R by communication of S and R alone; because if $C = \emptyset$, then $S \cup C \cup R = S \cup R$ in the realisation formulas. Direct connectedness is related to the formula used in the axiomatisation of Communication Logic in the previous chapter. $\mathrm{Fol}(G, H)$ states that H is directly $\forall^s\exists^r$-1-connected to G.

3.2.7 Groups and Connectors

We have presented knowledge realisation two folded. In the previous section, we looked at the most common interpretation, the realisation of distributed knowledge among a set of agents to individual knowledge among its members. In this section, we presented a variant of knowledge realisation that is more in line with the common interpretation of communication over a directed network: the realisation of distributed knowledge among a set of senders and receivers to knowledge among the receivers. We claimed that the latter is a generalisation of the former. Now we will prove this. We will show that groups are also definable in terms of connectors.

Proposition 3.2.18 (Equivalence of definitions)**.** *For* $n \in \mathbb{N}$*,* $t \in \mathcal{T}$*, and a frame* $\mathfrak{F}$*:* $G \subseteq A$ *is an n-group in* $\mathfrak{F}$ *iff for all* $g \in G$*,* G *is a* $\forall^s\forall^r$*-n-connector from* G *to* $\{g\}$ *in* $\mathfrak{F}$*.*

Moreover, for $n \in \mathbb{N}$*,* $t \in \mathcal{T}$*, and a frame* $\mathfrak{F}$ *whose social relation is reflexive:* $G \subseteq A$ *is an n-group in* $\mathfrak{F}$ *iff* G *is a* $\forall^s\forall^r$*-n-connector from* G *to* G *in* $\mathfrak{F}$*.*

Proof. $\forall g \in G : \mathbf{C}^{n, G \setminus \{g\}, \{g\}, G}_{\forall^s\forall^r} \iff \forall g \in G\ \forall g' \in G \setminus \{g\}\ g \rightarrow^{\leq n}_{G} g'$. This is the case iff G is an n-group by the definition of n-groups. $\mathbf{C}^{n, G, G \setminus \{g\}, G}_{\forall^s\forall^r} = \forall g, g' \in G\ g \rightarrow^{\leq n}_{G} g'$. If $\mathfrak{F}$ is reflexive, then this is the case iff G is an n-group. ☐

The reductions of groups to connectors are reflected in the semantic and syntactic results of groups, as these results are specific cases of results for connectors. If for all $g \in G$, G is a connector from $S_g = G \setminus \{g\}$ to $R_g = \{g\}$, then $\bigwedge_{g \in G} \varphi_g$ implies the result in Proposition 3.1.4, where φ_g is the knowledge realisation formula for G as a connector from S_g to R_g obtained from Proposition 3.2.14. Similarly and less complicated: if we assume that the network is reflexive, and assign $S = R = C = G$, then the formula in Proposition 3.2.14 is equivalent to the formula in Proposition 3.1.4.

This correspondence between groups and connectors not only shows that connectors and their format of knowledge realisation are more general than groups, it also hints towards extensions to the notion of groups and undirected knowledge realisation that we have not discussed: notions of groups equivalent to sets that form a connector from themselves to themselves of the other connector types, and their forms of knowledge realisation. Note that we then get three extra (first-order) group types $\exists\exists$, $\exists\forall$, and $\forall\exists$, besides the type $\forall\forall$ which we discussed in the previous section (under the generic name "group").

Assuming reflexivity of the social network for simplicity's sake, we can then understand, for example, an *$\exists\forall$-group* as a set G such that G is a $\exists^r\forall^s$-connector (or equivalently a $\exists^s\forall^r$-connector) from G to G. This gives us a notion of a set that can realise distributed knowledge among G to "someone in G knows", $S_G\varphi$. The formula for this knowledge realisation is obtained by assigning $S = R = C = G$ (or $S = G \setminus \{g\}$, $R = \{g\}$) in Proposition 3.2.13. The results of other notions of groups, and their respective knowledge realisation, can similarly be obtained form the other propositions: 3.2.11, 3.2.15, and 3.2.12.[10]

3.2.8 Summary

We have now developed directed knowledge realisation that realises distributed knowledge of a set of senders and receivers to knowledge among only the receivers — either distributively, individually by some, or by all receivers. We have explored six structural requirements for such realisation: the six types of connectors, $\forall^r\exists^s$, $\exists^s\exists^r$, $\forall^s\exists^r$, $\forall^s\forall^r$, $\exists^r\forall^s$, and $\exists^s\forall^r$-connectors. First, we developed the connector types qualitatively, distinguishing between connectors by which receivers they connect to which senders. We have shown that the types, when regarded qualitatively, become equivalent under the assumption that the sending and receiving sets are groups. Then, we developed the types quantitatively, distinguishing between connectors also by how long their connections are from the receivers to the senders, using the notion of latency. We have shown that the quantitative and qualitative notions are related by the notion of fixed points discussed in the previous section. We provided formulas for the connector types, and have show that a connector of a certain type (and latency n) can realise a certain distribution of knowledge among the senders and receivers pre-communication to a certain distribution among the receivers post-communication, through (at most n) iterated communications by the sender, connector, and receiver.

3.3 Blocking Information Flow

In this section we will study the negation of what we discussed in the previous section: *blocking* communication from a sender to a receiver. We will investigate sets $B \subseteq A$, such that if B does not communicate, information cannot be realised from $S \cup R$ to R, *no matter what the other agents do.* We call such B *blocking sets.* As our framework only formalised the action of communicating, and not the action of not-communicating, we must investigate such sets by what their *complement* cannot achieve through communication. That is, we will investigate them as sets $B \subseteq A$, such that no matter how many times its complement

[10] Do note that some knowledge realisation formulas in the propositions for connectors will become equivalent to others, or become tautologies, when we assume that $S = R = C = G$ or $S = G \setminus \{g\}$ and $R = \{g\}$.

$\overline{B}$ communicates, distributed knowledge of $S \cup R$ does not become some form of group knowledge within R.

3.3.1 Blocking Sets

A set whose complement cannot bring about knowledge realisation through communication is a set whose complement is not a connector. As we consider six connector types, we will define blocking sets of the six types $\mathcal{T} = \{\forall^r\exists^s, \exists^s\exists^r, \forall^s\exists^r, \forall^s\forall^r, \exists^r\forall^s, \exists^s\forall^r\}$, such that for $t \in \mathcal{T}$, a set is a t-blocking set only when its complement is not a t-connector.

Definition 3.3.1 (Blocking sets)**.** For non-empty sets $S, R \subseteq A$ and $t \in \mathcal{T}$, $B \subseteq A$ is a t-blocking set from S to R (in $\mathfrak{F}$) iff $\overline{B}$ is not a t-connector from S to R (in $\mathfrak{F}$).

We don't require blocking sets to be disjoint from the sending or receiving set. However, a set B is a blocking set from S to R iff its part disjoint from S and R is.

Proposition 3.3.1 (Disjoint blocking sets)**.** *For non-empty sets $S, R \subseteq A$ and $t \in \mathcal{T}$, a set $B \subseteq A$ is a t-blocking set iff $B \setminus (S \cup R)$ is.*

Proof. $\overline{B \setminus (S \cup R)} = \overline{B} \cup S \cup R$. The rest follows from Proposition 3.2.1. ☐

When the sending and receiving sets are such that the senders directly follow the receivers, no set B is a blocking set.

Proposition 3.3.2 (Blocking sets & direct connectedness)**.** *For all $t \in \mathcal{T}$ and non-empty sets $S, R \subseteq A$ such that S is directly t-connected to R,* no *$B \subseteq A$ is a t-blocking set.*

Proof. Let $B \subseteq A$ be an arbitrary set. S is directly t-connector to R iff $\emptyset$ is a t-connector from S to R. As being a connector is closed under supersets, $\overline{B}$ also is a t-connector. Therefore, B is not a t-blocking set. ☐

If there are no connectors from a set S to a set R, then all sets are blocking sets.

Proposition 3.3.3 (Blocking sets & no connectors)**.** *For non-empty sets $S, R \subseteq A$, if there is no t-connector $C \subseteq A$ from S to R,* every *$B \subseteq A$ is a t-blocking set.*

Proof. This follows from the definition of blocking sets. ☐

Because connectors are expressible in Communication Logic, and the set of agents A is finite, blocking sets are also expressible in Communication Logic.

Proposition 3.3.4 (Blocking set formula)**.** *For non-empty $S, R \subseteq A$, frame $\mathfrak{F}$, and $Q_1^x Q_2^y \in \mathcal{T}$: $B \subseteq A$ is a t-blocking set from S to R in $\mathfrak{F}$ iff $\mathfrak{F} \Vdash \neg t(\overline{B}, S, R)$.*

Proof. This follows from Proposition 3.3.1 and the correctness of the formula for $Q_1^x Q_2^y$-connectors. ☐

As with connectors, blocking sets form a hierarchy. This hierarchy is inverse to the connector hierarchy.

Corollary 3.3.1 (Blocking set hierarchy)**.** *For non-empty sets $S, R \subseteq A$, and a set $B \subseteq A$:*

$$\Vdash \neg\exists^s\exists^r(\overline{B}, S, R) \to \neg\forall^r\exists^s(\overline{B}, S, R) \qquad \Vdash \neg\forall^r\exists^s(\overline{B}, S, R) \to \neg\exists^s\forall^r(\overline{B}, S, R)$$
$$\Vdash \neg\exists^s\forall^r(\overline{B}, S, R) \to \neg\forall^s\forall^r(\overline{B}, S, R)$$

$$\Vdash \neg\exists^s\exists^r(\overline{B}, S, R) \to \neg\forall^s\exists^r(\overline{B}, S, R) \qquad \Vdash \neg\forall^s\exists^r(\overline{B}, S, R) \to \neg\exists^r\forall^s(\overline{B}, S, R)$$
$$\Vdash \neg\exists^r\forall^s(\overline{B}, S, R) \to \neg\forall^s\forall^r(\overline{B}, S, R).$$

Furthermore, for blocking sets from a *group* S to a *group* R, the six types are equivalent. Proposition 3.2.4 sufficiently shows this.

3.3.2 Delaying Sets

As blocking sets are defined in terms of connectors, and we quantitatively distinguished connectors using the notion of latency, we can also quantitatively distinguish blocking sets from each other. Where a blocking set B should block knowledge realisation by not communicating, given that the agents outside B iteratively communicate an unlimited number of times, delaying sets have to block knowledge realisation if the agents outside B iteratively communicate a *limited* number of times. Thereby, n-delaying sets are sets that, when they don't communicate, make sure that no other set can realise distributed knowledge of $S \cup R$ to R by communicating n times or less.

Definition 3.3.2 (Delaying sets)**.** For non-empty sets $S, R \subseteq A$, $n \in \mathbb{N}^+$, and $t \in \mathcal{T}$, the set $B \subseteq A$ is a t-n-delaying set from S to R iff $\overline{B}$ is not a t-n-connector.

Delaying sets are a proper quantification of blocking sets in that t-blocking sets are equivalent to sets that are t-n-delaying for all $n \in \mathbb{N}^+$.

Proposition 3.3.5 (Delaying sets and blocking sets)**.** *B is a t-blocking set from S to R iff B is a t-n-delaying set for all $n \in \mathbb{N}^+$.*

Proof. This follows from Proposition 3.2.5 and the definition of blocking and delaying sets. ☐

By the fixed-point of communication discussed in the last chapter, B is a blocking set iff it is a $|A| - 1$-delaying set.

Corollary 3.3.2. *B is a t-blocking set from S to R iff B is a t-$|A| - 1$-delaying set for all.*

The equivalence between blocking sets and blocking sets disjoint from the sending and receiving set also holds for delaying sets.

Proposition 3.3.6 (Disjoint delaying sets)**.** *For non-empty sets $S, R \subseteq A$, $t \in \mathcal{T}$, and $n \in \mathbb{N}^+$ a set $B \subseteq A$ is a t-n-delaying set iff $B \setminus (S \cup R)$ is.*

Proof. Similar to the proof of Proposition 3.3.1. ☐

If there are no connectors from a set S to a set R with a latency of n, then all sets are n-delaying sets.

Proposition 3.3.7 (Delaying sets & no connectors)**.** *For non-empty sets $S, R \subseteq A$, if there is no t-n-connector $C \subseteq A$ from S to R (i.e. if A is not a t-n-connector from S to R)* every *$B \subseteq A$ is a t-n-delaying set.*

Proof. This follows from the definition of delaying sets. ☐

In terms of network structure, an n-delaying set B is a set such that there is no $\overline{B}$-path of length $\leq n$ from agents in R to agents in S. As such, every path of length $\leq n$ from agents in R to agents in S contains some agent in B, which can make these paths useless for communication by not communicating. What particular agents in R and S this concerns depends on the type of delaying set. Recall the path conditions defined in Definition 3.2.5, then:

Proposition 3.3.8. *For non-empty sets $S, R \subseteq A$, and $n \in \mathbb{N}^+$, B is a t-n-delaying set from S to R iff $\mathbf{C}_t^{n,S,R,\overline{B}}$ does not hold*

As quantitatively distinguished connectors form a hierarchy, delaying sets also form a hierarchy. This hierarchy is immediate from the connector latency hierarchy in Proposition 3.2.8.

Corollary 3.3.3 (Delaying set hierarchy)**.**

$\Vdash \neg\exists^s\exists^r(\overline{B},n,S,R) \to \neg\forall^r\exists^s(\overline{B},n,S,R)$ $\quad\Vdash \neg\forall^r\exists^s(\overline{B},n,S,R) \to \neg\exists^s\forall^r(\overline{B},n,S,R)$

$\Vdash \neg\exists^s\forall^r(\overline{B},n,S,R) \to \neg\forall^s\forall^r(\overline{B},n,S,R)$

$\Vdash \neg\exists^s\exists^r(\overline{B},n,,R) \to \neg\forall^s\exists^r(\overline{B},n,S,R)$ $\quad\Vdash \neg\forall^s\exists^r(\overline{B},n,S,R) \to \neg\exists^r\forall^s(\overline{B},n,S,R)$

$\Vdash \neg\exists^r\forall^s(\overline{B},n,S,R) \to \neg\forall^s\forall^r(\overline{B},n,S,R).$

3.3.3 Blocking Sets, Delaying Sets, and Knowledge Realisation

Now we will explore the communicational properties of delaying sets, and consequently blocking sets, with respect to knowledge realisation. Again, we will first discuss the semantic implications and thereafter the syntactic. By the definition of blocking sets and delaying sets, an analysis of blocking and delaying sets is congruent to an analysis of the negation of connectors.

If C is not a connector with a latency of n, then after the senders, receivers, and connector communicate n times, the uncertainty of agents in R is at least as much as the agents in $S' \cup R \cup C$ for an $S' \subset S$.

Definition 3.3.3 (Type similarity exclusion)**.** For non-empty sets $S, R \subseteq A$, a set $C \subseteq A$, and $n \in \mathbb{N}^+$, let $G = S \cup C \cup R$. Define:

$$\begin{array}{lrl}
\supseteq^{n,S,R,C}_{\forall^r\exists^s} := & \exists r \in R & \sim^{!G\otimes n}_r \supseteq \sim_{\{r\}\cup(G\setminus S)} \\
\supseteq^{n,S,R,C}_{\exists^s\exists^r} := & \forall r \in R & \sim^{!G\otimes n}_r \supseteq \sim_{\{r\}\cup(G\setminus S)} \\
\supseteq^{n,S,R,C}_{\forall^s\exists^r} := & \exists s \in S \forall r \in R & \sim^{!G\otimes n}_r \supseteq \sim_{\{r\}\cup(G\setminus\{s\})} \\
\supseteq^{n,S,R,C}_{\forall^s\forall^r} := & \exists s \in S \exists r \in R & \sim^{!G\otimes n}_r \supseteq \sim_{\{r\}\cup(G\setminus\{s\})} \\
\supseteq^{n,S,R,C}_{\exists^r\forall^s} := & \forall r \in R \exists s \in S & \sim^{!G\otimes n}_r \supseteq \sim_{\{r\}\cup(G\setminus\{s\})} \\
\supseteq^{n,S,R,C}_{\exists^s\forall^r} := & \forall s \in S \exists r \in R & \sim^{!G\otimes n}_r \supseteq \sim_{\{r\}\cup(G\setminus\{s\})}
\end{array}$$

Proposition 3.3.9 (Non-connector semantics)**.** *For non-empty sets $S, R \subseteq A$ a frame $\mathfrak{F}$, a set $C \subseteq A$, $n \in \mathbb{N}^+$, and $G = S \cup C \cup R$: if C is* not *a t-connector from S to R with a latency of n then $\supseteq^{n,S,R,C}_t$ holds.*

Proof. Let "not $\mathbf{C}^{n,S,R,C}_t$" be denoted by $\overline{\mathbf{C}}^{n,S,R,C}_t$. $r \not\rightarrow^{\leqslant n}_G s \implies (s \notin \mathcal{F}|^{+n}_G(r)$ or $s = r) \implies \mathcal{F}|^{+n}_G(r) \subseteq \{r\} \cup (G \setminus \{s\}) \implies \sim_{\mathcal{F}|^{+n}_G(r)} \supseteq \sim_{\{r\}\cup(G\setminus\{s\})}$. Therefore, $\overline{\mathbf{C}}^{n,S,R,C}_t \implies \supseteq^{n,S,R,C}_t$ in the cases of $t \in \{\forall^s\exists^r, \forall^s\forall^r, \exists^r\forall^s, \exists^s\forall^r\}$.

To prove this is the cases for $t \in \{\forall^r\exists^s, \exists^s\exists^r\}$, note that $(\exists r \in R \forall s \in S\ s \notin \mathcal{F}|^{+n}_G(r)$ or $s = r) \implies \exists r \in R\ \mathcal{F}|^{+n}_G(r) \subseteq \{r\} \cup (G \setminus S)$. Moreover, $(\forall s \in S \forall r \in R\ s \notin \mathcal{F}|^{+n}_G(r)$ or $s = r) \implies \forall r \in R\ \mathcal{F}|^{+n}_G(r) \subseteq \{r\} \cup (G \setminus S)$. □

As an n-delaying set is such that $\overline{B}$ is not a connector with a latency of n, we obtain the semantic implications of being a delaying set.

Corollary 3.3.4 (Delaying set semantic)**.** *For non-empty sets $S, R \subseteq A$, a set $B \subseteq A$, $n \in \mathbb{N}^+$, and $t \in \mathcal{T}$: if B is a t-n-delaying set from S to R in $\mathfrak{F}$ then $\supseteq^{n,S,R,\overline{B}}_t$ holds.*

Proof. This follows from Definition 3.3.2 and Proposition 3.3.9. □

The $\supseteq^{n,S,R,\overline{B}}_t$ conditions imply syntactic results for a connector's negation (sets that are not a connector, i.e. non-connectors), blocking sets, and delaying sets. These results are of the form $(\xi \wedge \chi) \to \mu$, Where ξ is the topological preconditions, χ is the epistemic preconditions, and μ is the epistemic result.

The topological precondition ξ states the existence of an n-delaying set (the non-existence of a connector with a latency of n).

A blocking set B from S to R is not necessarily also a blocking set from, for example, $\overline{B} \setminus S$ to R. Even when B is a blocking set from S to R, knowledge about a formula can still flow from $\overline{B} \setminus S$ to R when $S \cup \overline{B} \cup R$ communicates. Therefore, for R to not know that φ after $S \cup \overline{B} \cup R$ communicate n times, we must state that $\overline{B} \setminus S$ does not know $[!S \cup \overline{B} \cup R]\varphi$. This is the epistemic precondition χ for blocking sets to act as blockers of knowledge reolution.

As blocking sets prevent communication from senders to receives, one might expect that the epistemic preconditions in the knowledge realisation formulas of delaying sets also include a *positive* statement regarding knowledge: the epistemic preconditions of the formulas for connectors presented in Section 3.2.5, formulas that state that some sending agents together with some receiving agents distributively know that $[!(S \cup \overline{B} \cup R)]^n\varphi$. In this way, the knowledge realisation formulas of delaying sets would explicitly state the actual relation of delaying sets to knowledge realisation: that, even though some senders and receivers distributively know that $[!(S \cup \overline{B} \cup R)]^n\varphi$, the receivers will not know it after n communications by $S \cup \overline{B} \cup R$. However, for the knowledge realisation formulas of delaying set to imply this, it is not necessary to *explicitly* state that $S \cup R$ know $[!(S \cup \overline{B} \cup R)]^n\varphi$ — a simple consequence of *validities*: the knowledge realisation formulas below hold in *all* models, and therefore also in the models where some senders together with some receivers distributively know that $[!(S \cup R \cup \overline{B})]^n\varphi$.

As with connectors, the epistemic results μ in $(\xi \wedge \chi) \to \mu$ take on three forms: $\neg[!(\overline{B} \cup R \cup S)]^n \Box_R \varphi$ where $\Box$ is a knowledge modality, one of D, S, or E.

The formulas of the epistemic preconditions require some explanation, as their readings are not immediately apparent. Essentially, all epistemic preconditions are statements "inverse" to their positive (connector) counterpart. For connectors, the epistemic preconditions state that certain pairs of senders and receivers have distributed knowledge about some formula. The upcoming non-connector (delaying set) preconditions state that *everybody involved in the communication but* certain senders have *no* knowledge about a formula. With "everybody involved" we mean the senders, receivers, and the non-connector itself (in case of blocking and delaying sets, the non-connector is $\overline{B}$.) The formula can still be distributively known by some senders together with some receivers, but only if certain sending agents are included. However, this reading of the preconditions only comes about if we assume that $S \cap R = \emptyset$ (which we have not done for generalities' sake). Note the following: for $S, R, \overline{B} \in A$ such that $S \cap R = \emptyset$, $G = S \cup R \cup \overline{B}$, and $\varphi \in \mathcal{L}_{CL}$

$$\Vdash \bigwedge_{r \in R} \neg D_{(G \setminus S) \cup \{r\}}\varphi \leftrightarrow \neg D_{G \setminus S}\varphi \qquad \Vdash \neg D_{(G \setminus S) \cup R}\varphi \leftrightarrow \neg D_{G \setminus S}\varphi$$

$$\Vdash \bigwedge_{s \in S} \neg D_{(G \setminus \{s\}) \cup R}\varphi \leftrightarrow \bigwedge_{s \in S} \neg D_{G \setminus \{s\}} \qquad \Vdash \bigwedge_{s \in S} \bigwedge_{r \in R} \neg D_{(G \setminus \{s\}) \cup \{r\}}\varphi \leftrightarrow \bigwedge_{s \in S} \neg D_{G \setminus \{s\}}$$

$$\Vdash \bigvee_{r \in R} \bigwedge_{s \in S} \neg D_{(G \setminus \{s\}) \cup \{r\}} \leftrightarrow \bigwedge_{s \in S} \neg D_{G \setminus \{s\}} \qquad \Vdash \bigvee_{s \in S} \bigwedge_{r \in R} \neg D_{(G \setminus \{s\}) \cup \{r\}}\varphi \leftrightarrow \bigvee_{s \in S} \neg D_{G \setminus \{s\}}\varphi$$

Keep these validities in mind in the upcoming propositions. We will use the reading of the simplified forms of epistemic preconditions, as these are more intuitive: "everybody except all senders does not know that φ", "for every sender, everybody without that sender does not know φ", and "for some sender, everybody without that sender does not know that φ".

The syntactic results for delaying sets can be derived from the following syntactic implications of similarity relation inclusion: if the similarity relation of a set Y is contained in the similarity relation of a set X after a set Z communicates n times, and if Y does not distributively know that $[!Z]^n\varphi$, then X will not know φ after Z communicates n times.

Lemma 3.3.5. *For a frame $\mathfrak{F}$ such that $\sim_X^{!Z\otimes n}\supseteq\sim_Y$: $\mathfrak{F}\Vdash\neg D_Y[!Z]^n\varphi\to[!Z]^n\neg D_X\varphi$.*

Proof. Assume that $\sim_X^{!Z\otimes n}\supseteq\sim_Y$. Then, $\sim_{\mathcal{F}|_Z^{+n}(X)}\supseteq\sim_Y$. Hence, $\mathfrak{F}\Vdash\neg D_Y[!Z]^n\varphi\to\neg D^n_{\mathcal{F}|_Z^+(X)}[!Z]^n\varphi$. Therefore, by Theorem 2.3.2, $\mathfrak{F}\Vdash\neg D_Y[!Z]^n\varphi\to\neg[!Z]^nD_X\varphi$. ☐

Recall that the six delaying set types form a hierarchy shaped as two diverging branches that converge again at the strongest type $\exists^s\exists^r$. Because of this, the delaying set types correspond to an increasing number of $\mathcal{L}_{CL}$-formulas that represent the ability to block knowledge realisation. We will go through the six types of delaying sets by order of type, listing their specific formulas, epistemic preconditions ξ, and epistemic post-conditions μ. We start from $\forall^s\forall^r$, first expanding the $\forall^s\forall^r$–$\exists^s\forall^r$–$\forall^r\exists^s$–$\exists^s\exists^r$ branch, and then the leftover types $\exists^r\forall^s$ and $\forall^s\exists^r$ of the $\forall^s\forall^r$–$\exists^r\forall^s$–$\forall^s\exists^r$–$\exists^s\exists^r$ branch. At each step through this hierarchy, we will develop the formulas belonging to the level of this hierarchy, so that a developed formula for $\forall^s\exists^r$ also holds for $\exists^r\forall^s$ and $\exists^s\exists^r$-connectors, but not for $\forall^s\forall^r$ connectors etc. Note that we do not show that these formulas are not valid for delaying sets of a lower type. But counterexamples do arise out of Figure 3.2, when paired with a proper similarity relation.

The weakest type, $\forall^s\forall^r$, blocks all agents in R of individually knowing a formula φ: even if everybody involved ($S\cup\overline{B}\cup R$) communicates n times, somebody still will not know that φ. The epistemic precondition for this is that for all senders s, everybody involved in communication except for s does not know that $[!(S\cup\overline{B}\cup R)]^n\varphi$. The full formula therefore is as follows.

Proposition 3.3.10. *For non-empty sets $S,R\subseteq A$, set $B\subseteq A$, let $G=S\cup\overline{B}\cup R$. Then, for $\varphi\in\mathcal{L}_{CL}$ and $n\in\mathbb{N}^+$:*

$$\Vdash\left(\neg\forall^s\forall^r(\overline{B},n,S,R)\wedge\bigwedge_{s\in S}\bigwedge_{r\in R}\neg D_{(G\setminus\{s\})\cup\{r\}}[!S\cup\overline{B}\cup R]^n\varphi\right)\to\neg[!S\cup\overline{B}\cup R]^nE_R\varphi$$

Proof. If $\overline{B}$ is a $\forall^s\forall^r$-n-delaying set, then by Corollary 3.3.4, $\supseteq^{n,S,R,\overline{B}}_{\forall^s\forall^r}$ holds: $\exists s\in S\exists r\in R\ \sim_r^{!G\otimes n}\supseteq\sim_{(G\setminus\{s\})\cup\{r\}}$. Call these r and s, $r^\star$ and $s^\star$ respectively. If $\mathfrak{F}\Vdash\bigwedge_{s\in S}\bigwedge_{r\in R}\neg D_{(G\setminus\{s\})\cup\{r\}}[!G]^n\varphi$, then in particular $\mathfrak{F}\Vdash\neg D_{(G\setminus\{s^\star\})\cup\{r^\star\}}[!G]^n\varphi$. By Lemma 3.3.5, taking $X=\{r^\star\}$, $Y=(G\setminus\{s^\star\})\cup\{r^\star\}$, and $Z=S\cup\overline{B}\cup R$, $\mathfrak{F}\Vdash\neg[!S\cup\overline{B}\cup R]^nK_{r^\star}\varphi$. Therefore, $\mathfrak{F}\Vdash\neg[!S\cup\overline{B}\cup R]^nE_R\varphi$. ☐

$\exists^r\forall^s$-delaying sets can block two forms of knowledge: individual knowledge by all, and by some receivers. Even if everybody involved ($S\cup\overline{B}\cup R$) communicate n times then, respectively, somebody or everybody in R still does not know that φ. For the former, everybody involved in communication except for some sender s must not know that $[!(S\cup\overline{B}\cup R)]^n\varphi$. For the latter, it must hold that for all senders s, everybody involved in communication except for s does not know that $[!(S\cup\overline{B}\cup R)]^n\varphi$.

Proposition 3.3.11. *For non-empty sets $S,R\subseteq A$, set $B\subseteq A$, let $G=S\cup\overline{B}\cup R$. Then, for $\varphi\in\mathcal{L}_{CL}$ and $n\in\mathbb{N}^+$:*

$$\Vdash\left(\neg\exists^s\forall^r(\overline{B},n,S,R)\wedge\bigvee_{s\in S}\bigwedge_{r\in R}\neg D_{(G\setminus\{s\})\cup\{r\}}[!S\cup\overline{B}\cup R]^n\varphi\right)\to\neg[!S\cup\overline{B}\cup R]^nE_R\varphi$$

$$\Vdash\left(\neg\exists^s\forall^r(\overline{B},n,S,R)\wedge\bigwedge_{s\in S}\bigwedge_{r\in R}\neg D_{(G\setminus\{s\})\cup\{r\}}[!S\cup\overline{B}\cup R]^n\varphi\right)\to\neg[!S\cup\overline{B}\cup R]^nS_R\varphi$$

Proof. If $\overline{B}$ is a $\exists^s\forall^r$-n-delaying set, then by Corollary 3.3.4, $\supseteq^{n,S,R,\overline{B}}_{\exists^s\forall^r}$ holds: $\forall s\in S\exists r\in R\ \sim_r^{!G\otimes n}\supseteq\sim_{(G\setminus\{s\})\cup\{r\}}$. For each $s\in S$, call the $r\in R$ that satisfies this: r_s.

(1) If $\mathfrak{F} \Vdash \bigvee_{s\in S} \bigwedge_{r\in R} \neg D_{(G\setminus\{s\})\cup\{r\}}[!G]^n\varphi$, then for some $s \in S$: $\mathfrak{F} \Vdash \neg D_{(G\setminus\{s\})\cup\{r_s\}}[!G]^n\varphi$. By Lemma 3.3.5, taking $X = \{r_s\}$, $Y = (G \setminus \{s\}) \cup \{r_s\}$, and $Z = S \cup \overline{B} \cup R$, $\mathfrak{F} \Vdash \neg[!S \cup \overline{B} \cup R]^n K_{r_s}\varphi$. Therefore, $\mathfrak{F} \Vdash \neg[!S \cup \overline{B} \cup R]^n E_R\varphi$.

(2) If $\mathfrak{F} \Vdash \bigwedge_{s\in S} \bigwedge_{r\in R} \neg D_{(G\setminus\{s\})\cup\{r\}}[!G]^n\varphi$, then for all $s \in S$, by Lemma 3.3.5, taking $X = \{r_s\}$, $Y = (G \setminus \{s\}) \cup \{r_s\}$, and $Z = S \cup \overline{B} \cup R$, $\mathfrak{F} \Vdash \neg[!S \cup \overline{B} \cup R]^n K_{r_s}\varphi$. Therefore, $\mathfrak{F} \Vdash \neg[!S \cup \overline{B} \cup R]^n S_R\varphi$. □

The stronger $\forall^r\exists^s$-delaying sets, can already block individual knowledge by all receivers when everybody involved except all senders S don't distributively know that $[!(S \cup \overline{B} \cup R)]^n\varphi$.

Proposition 3.3.12. *For non-empty sets $S, R \subseteq A$, set $B \subseteq A$, let $G = S \cup \overline{B} \cup R$. Then, for $\varphi \in \mathcal{L}_{CL}$ and $n \in \mathbb{N}^+$:*

$$\Vdash \left(\neg\forall^r\exists^s(\overline{B}, n, S, R) \wedge \bigwedge_{r\in R} \neg D_{(G\setminus S)\cup\{r\}}[!S \cup \overline{B} \cup R]^n\varphi\right) \to \neg[!S \cup \overline{B} \cup R]^n E_R\varphi$$

Proof. If $\overline{B}$ is a $\forall^r\exists^s$-n-delaying set, then by Corollary 3.3.4, $\supseteq^{n,S,R,\overline{B}}_{\forall^r\exists^s}$ holds: $\exists r \in R$ $\sim_r^{!S\cup\overline{B}\cup R\otimes n}\supseteq\sim_{(G\setminus S)\cup\{r\}}$. Call this r: $r^\star$. If $\bigwedge_{r\in R} \neg D_{(G\setminus S)\cup\{r\}}[!S \cup \overline{B} \cup R]^n\varphi$, then in particular $\mathfrak{F} \Vdash D_{(G\setminus S)\cup\{r^\star\}}[!S \cup \overline{B} \cup R]^n\varphi$. By Lemma 3.3.5, taking $X = \{r^\star\}$, $Y = (G\setminus S)\cup \{r^\star\}$, and $Z = S\cup\overline{B}\cup R$, $\mathfrak{F} \Vdash \neg[!S \cup \overline{B} \cup R]^n K_{r^\star}\varphi$. Therefore, $\mathfrak{F} \Vdash \neg[!S \cup \overline{B} \cup R]^n E_R\varphi$. □

The strongest type of delaying set, $\exists^s\exists^r$, can stop distributed knowledge among the receivers from happening. For this, everybody involved in communication except the senders S must not know that $[!(S \cup \overline{B} \cup R)]^n\varphi$.

Proposition 3.3.13. *For non-empty sets $S, R \subseteq A$, set $B \subseteq A$, let $G = S \cup \overline{B} \cup R$. Then, for $\varphi \in \mathcal{L}_{CL}$ and $n \in \mathbb{N}^+$:*

$$\Vdash \left(\neg\exists^s\exists^r(\overline{B}, n, S, R) \wedge \neg D_{(G\setminus S)\cup R}[!S \cup \overline{B} \cup R]^n\varphi\right) \to \neg[!S \cup \overline{B} \cup R]^n D_R\varphi$$

Proof. If $\overline{B}$ is an $\exists^s\exists^r$-n-delaying set, then by Corollary 3.3.4, $\supseteq^{n,S,R,\overline{B}}_{\exists^s\exists^r}$ holds. Hence, $\forall s \in S \forall r \in R, \sim_r^{!G\otimes n}\supseteq\sim_{(G\setminus\{s\})\cup\{r\}}$. Therefore, $\forall r \in R, \sim_r^{!G\otimes n}\supseteq\sim_{(G\setminus S)\cup\{r\}}$. Because $\sim_{(G\setminus S)\cup\{r\}}=\sim_r \cap \sim_{G\setminus S}$, we have $\sim_R^{!G\otimes n}= \bigcap_{r\in R} \sim_r^{!G\otimes n}\supseteq \bigcap_{r\in R} \sim_r \cap \sim_{G\setminus S}=\sim_{(G\setminus S)\cup R}$. Thus, if $\mathfrak{F} \Vdash \neg D_{(G\setminus S)\cup R}[!S \cup \overline{B} \cup R]^n\varphi$, then by Lemma 3.3.5, taking $X = R$, $Y = (G\setminus S)\cup R$, and $Z = S \cup \overline{B} \cup R$, $\mathfrak{F} \Vdash \neg[!S \cup \overline{B} \cup R]^n\neg D_R\varphi$. □

$\exists^r\forall^s$-delaying sets can block some R from knowing something, or all R from knowing something. For these, everybody except any sender must not know that $[!(S \cup \overline{B} \cup R)]^n\varphi$.

Proposition 3.3.14. *For non-empty sets $S, R \subseteq A$, set $B \subseteq A$, let $G = S \cup \overline{B} \cup R$. Then, for $\varphi \in \mathcal{L}_{CL}$ and $n \in \mathbb{N}^+$:*

$$\Vdash \left(\neg\exists^r\forall^s(\overline{B}, n, S, R) \wedge \bigwedge_{s\in S} \bigwedge_{r\in R} \neg D_{(G\setminus\{s\})\cup\{r\}}[!S \cup \overline{B} \cup R]^n\varphi\right) \to \neg[!S \cup \overline{B} \cup R]^n S_R\varphi$$

$$\Vdash \left(\neg\exists^r\forall^s(\overline{B}, n, S, R) \wedge \bigvee_{r\in R} \bigwedge_{s\in S} \neg D_{(G\setminus\{s\})\cup\{r\}}[!S \cup \overline{B} \cup R]^n\varphi\right) \to \neg[!S \cup \overline{B} \cup R]^n E_R\varphi$$

Proof. If $\overline{B}$ is a $\exists^r\forall^s$-n-delaying set, then by Corollary 3.3.4, $\supseteq^{n,S,R,\overline{B}}_{\exists^r\forall^s}$ holds: $\forall r \in R \exists s \in S$ $\sim_r^{!G\otimes n}\supseteq\sim_{(G\setminus\{s\})\cup\{r\}}$. For each $r \in R$, call the $s \in S$ that satisfies this: s_r.

(1) If $\mathfrak{F} \Vdash \bigwedge_{s\in S} \bigwedge_{r\in R} \neg D_{(G\setminus\{s\})\cup\{r\}}[!G]^n\varphi$, then also $\mathfrak{F} \Vdash \bigwedge_{r\in R} \neg D_{(G\setminus\{s_r\})\cup\{r\}}[!G]^n\varphi$. For any $r \in R$, by Lemma 3.3.5, taking $X = \{r\}$, $Y = (G \setminus \{s_r\}) \cup \{r\}$, and $Z = S \cup \overline{B} \cup R$, $\mathfrak{F} \Vdash \neg[!S \cup \overline{B} \cup R]^n K_r\varphi$. Therefore, $\mathfrak{F} \Vdash \neg[!S \cup \overline{B} \cup R]^n S_R\varphi$.

(weakest) type	epistemic precondition (ξ)	result (μ)	Proposition
$\neg\forall^s\forall^r$	$\bigwedge_{s\in S}\bigwedge_{r\in R}\neg D_{(G\setminus\{s\})\cup\{r\}}[!G]^n\varphi$	$\neg[!G]^n E_R\varphi$	3.3.10
$\neg\exists^r\forall^s$	$\bigvee_{r\in R}\bigwedge_{s\in S}\neg D_{(G\setminus\{s\})\cup\{r\}}[!G]^n\varphi$	$\neg[!G]^n E_R\varphi$	3.3.14
	$\bigwedge_{s\in S}\bigwedge_{r\in R}\neg D_{(G\setminus\{s\})\cup\{r\}}[!G]^n\varphi$	$\neg[!G]^n S_R\varphi$	
$\neg\forall^s\exists^r$	$\bigwedge_{s\in S}\neg D_{(G\setminus\{s\})\cup R}[!G]^n\varphi$	$\neg[!G]^n D_R\varphi$	3.3.15
$\neg\exists^s\exists^r$	$\neg D_{(G\setminus S)\cup R}[!G]^n\varphi$	$\neg[!G]^n D_R\varphi$	3.3.13
$\neg\exists^s\forall^r$	$\bigwedge_{s\in S}\bigwedge_{r\in R}\neg D_{(G\setminus\{s\})\cup\{r\}}[!G]^n\varphi$	$\neg[!G]^n S_R\varphi$	3.3.11
	$\bigvee_{s\in S}\bigwedge_{r\in R}\neg D_{(G\setminus\{s\})\cup\{r\}}[!G]^n\varphi$	$\neg[!G]^n E_R\varphi$	
$\neg\forall^r\exists^s$	$\bigwedge_{r\in R}\neg D_{(G\setminus S)\cup\{r\}}[!G]^n\varphi$	$\neg[!G]^n E_R\varphi$	3.3.12

Table 3.2: Summary of delaying set formulas

(2) If $\mathfrak{F}\Vdash\bigvee_{r\in R}\bigwedge_{s\in S}\neg D_{(G\setminus\{s\})\cup\{r\}}[!G]^n\varphi$, then there is some r such that in particular $\mathfrak{F}\Vdash\neg D_{(G\setminus\{s_r\})\cup\{r\}}[!G]^n\varphi$. By Lemma 3.3.5, taking $X=\{r\}$, $Y=(G\setminus\{s_r\})\cup\{r\}$, and $Z=S\cup\overline{B}\cup R$, $\mathfrak{F}\Vdash\neg[!S\cup\overline{B}\cup R]^n K_r\varphi$. Therefore, $\mathfrak{F}\Vdash\neg[!S\cup\overline{B}\cup R]^n E_R\varphi$. $\square$

And finally, $\forall^s\exists^r$-delaying sets can block R from distributively knowing a formula. For this, it must hold that for all senders s, everybody involved in communication except s do not distributively know that $[!(S\cup\overline{B}\cup R)]^n\varphi$.

Proposition 3.3.15. *For non-empty sets $S,R\subseteq A$, set $B\subseteq A$, let $G=S\cup\overline{B}\cup R$. Then, for $\varphi\in\mathcal{L}_{CL}$ and $n\in\mathbb{N}^+$:*

$$\Vdash\left(\neg\forall^s\exists^r(\overline{B},n,S,R)\wedge\bigwedge_{s\in S}\neg D_{(G\setminus\{s\})\cup R}[!S\cup\overline{B}\cup R]^n\varphi\right)\to\neg[!S\cup\overline{B}\cup R]^n D_R\varphi$$

Proof. If $\overline{B}$ is a $\forall^s\exists^r$-n-delaying set, then by Corollary 3.3.4, $\supseteq^{n,S,R,\overline{B}}_{\forall^s\exists^r}$ holds. Hence, $\exists s\in S\forall r\in R\ \sim_r^{!G\otimes n}\supseteq\sim_{(G\setminus\{s\})\cup\{r\}}$. As $\sim_{(G\setminus\{s\})\cup\{r\}}=\sim_r\cap\sim_{G\setminus\{s\}}$, $\exists s\in S\ \ \sim_R^{!G\otimes n}=\bigcap_{r\in R}\sim_r^{!G\otimes n}\supseteq\bigcap_{r\in R}\sim_r\cap\sim_{G\setminus\{s\}}=\sim_{(G\setminus\{s\})\cup R}$. Call this s: $s^\star$.

If $\mathfrak{F}\Vdash\bigwedge_{s\in S}\neg D_{(G\setminus\{s\})\cup R}[!S\cup\overline{B}\cup R]^n\varphi$, then also $\mathfrak{F}\Vdash\neg D_{(G\setminus\{s^\star\})\cup R}[!S\cup\overline{B}\cup R]^n\varphi$. Thus, by Lemma 3.3.5, taking $X=R$, $Y=(G\setminus\{s^\star\})\cup R$, and $Z=S\cup\overline{B}\cup R$, $\mathfrak{F}\Vdash\neg[!S\cup\overline{B}\cup R]^n\neg D_R\varphi$. $\square$

Recall that by Proposition 3.3.5, B is a t-blocking set iff it is a t-$(|A|-1)$-delaying set. Therefore, the delaying set results also give us results for blocking sets.

A rephrasing of all the $\supseteq^{n,S,R,\overline{B}}_t$ conditions in terms of $\sim_R$ is possible. However, such a rephrasing is only possible for two latency notions: $\exists^s\exists^r$ and $\forall^s\exists^r$. In the case of $\exists^r\forall^s$, $\exists s\in S\forall r\in R\ \ \sim_r^{!G\otimes n}\supseteq\sim_{\{r\}\cup G\setminus\{s\}}$ does not imply $\exists s\in S\ \sim_R\supseteq\sim_{R\cup G\setminus\{s\}}$ as the s in question are relative to each r and could therefore possibly be distinct. For the other types, the quantification over R in the $\supseteq^{n,S,R,\overline{B}}_t$ condition is existential, and these conditions can therefore also not be extended to results for $\sim_R$. As a result, we get fewer results for non-connectors, delaying sets, and blocking sets then we did for their positive counterpart, connectors. The other four latency notions do correspond to weaker[11] version of knowledge realisation: either blocking knowledge realisation towards “everybody in R knows”, E_R, or knowledge realisation towards “somebody in R knows”, S_R.

For a summary of the kinds of knowledge realisation made possible each connector type, see Table 3.2. This table is sorted to correspond to the order of Table 3.1. Moving down the table corresponds to a *stronger* delaying set type. Parallel, either the epistemic result gets stronger, or the epistemic precondition gets weaker.

[11]These are weaker because both $\not\Vdash\neg S_R\varphi\to\neg D_R\varphi$ and $\not\Vdash\neg E_R\varphi\to\neg D_R\varphi$.

3.3.4 Summary

We now have a description of the negative counterpart to connectors: blocking sets and delaying sets. Whereas connectors can realise distributed knowledge of a sending and a receiving set to the receiving set by communication, blocking sets can stop this realisation by *not communicating*. Equally, delaying sets can *delay* the realisation of knowledge to the receiving set by a certain amount of communication updates. Parallel to the six connector types, we have defined blocking and delaying sets of six types. We have shown that blocking sets (and delaying sets) are sets such that their complements are not connectors of a specific type (and latency). Therefore, we have analysed the capabilities of blocking and delaying sets by showing what not being a connector of a specific type and latency implies both in semantic and syntactic terms.

3.4 Minimality

Up until this point, we have discussed definitions of groups, connectors, and blocking sets, that are sufficient for certain epistemic-communicational results. However, these definitions are far from necessary for these epistemic-communicational results, . In this section we will work towards necessary conditions. We will do this by discussing multiple minimality notions of groups, connectors, and blocking sets. As connectors are the more general of the three (groups and blocking sets are definable in terms of connectors), we will discuss these minimality notions for connectors only. Minimality notions for blocking sets, delaying sets, and groups naturally arise out of the minimality notions for connectors. First we will discuss minimality with respect to latency, then we will discuss minimality with respect to membership.

3.4.1 Exact Connector Latency

The semantic and syntactic results of connectors and non-connectors allow for a *minimal* definition of t-n-connectors with respect to their latency n: sets of agents $C \subseteq A$ such that C is a t-n-connector and not a t-$(n-1)$-connector. Such connectors are special in that their latency n specifies the exact number of communication steps needed to realise knowledge corresponding to the type t. Any n-connector is such an exact m-connector for some $m \leq n$. We call this number m its *exact* t*-latency*.

Definition 3.4.1 (Latency-minimality)**.** We call a set C that is an t-n-connectors and not a t-$(n-1)$-connector *latency-minimal (for n)*.

Note that any t-connector is latency-minimal for some n. We call this n the minimal t-latency of a connector.

We can express that C is a latency-minimal t-n-connector as follows:

$$\mathrm{lmin}t(C, n, S, R) := t(C, n, S, R) \wedge \neg t(C, n-1, S, R).$$

The syntactic implications of latency-minimal connectors can be derived from the syntactic results of connectors and non-connectors by combining the two. We get exactly what we'd expect: under certain epistemic preconditions, latency-minimal t-n-connectors C are exactly those sets such that n communication acts by $S \cup C \cup R$ are necessary and sufficient for R to know a formula, either distributively or individually by some or all receivers. If we, for example, work out the formula for a latency-minimal $\forall^s\exists^r$-n-connector C from S to R, we get the following. Let $G = S \cup C \cup R$ be everybody involved in communication. If $G \setminus S$ do not distributively know that φ after G communicates $n-1$ or n times, any receiver together with any sender together distributively know that φ after G communicates $n-1$ or n times. Then, C is a latency-minimal $\forall^r\exists^s$-n-connector from S to R *iff* some receiver does not yet know φ after G communicates $n-1$ times, but after G communicates n times, all receivers know that φ.

$$\begin{aligned}\Vdash &\bigwedge_{s\in S}\bigwedge_{r\in R}\left(\neg D_{(G\setminus S)\cup\{r\}}[!G]^{n-1}\varphi \wedge \neg D_{(G\setminus S)\cup\{r\}}[!G]^{n}\varphi \wedge D_{\{s,r\}}([!G]^{n-1}\varphi \wedge [!G]^{n}\varphi)\right)\\ &\to (\mathrm{lmin}\forall^r\exists^s(C,n,S,R) \leftrightarrow ([!G]^{n-1}\neg E_R\varphi \wedge [!G]^n E_R\varphi)).\end{aligned}$$

As such, n is the exact number of times that a latency-minimal n-connector has to communicate to make sure that distributed knowledge is realised.

3.4.2 Minimal Connectors

We can define another notion of minimality with respect to the members of a connector. We define minimality of a connector relative to both its type, and to its type *and* latency. That is, we call C a *minimal* t-connector from S to R when all its members are essential for C to from a t-connector from S to R, and call C a *minimal* t-n-connector from S to R when all its members are essential to form a t-n-connector.

Definition 3.4.2 (Minimal connectors)**.** For $S, R, C \subseteq A$, $t \in \mathcal{T}$, and $n \in \mathbb{N}^+$ we call C a minimal t-connector from S to R iff for every $c \in C$, $C \setminus \{c\}$ is not t-connector from S to R.

Furthermore, we call a C a minimal t-n-connector from S to R iff for every $c \in C$, $C \setminus \{c\}$ is not t-n-connector from S to R.

We can express that C is a minimal t-n-connector as follows:

$$\mathrm{min}t(C,n,S,R) := t(C,n,S,R) \wedge \bigwedge_{c\in C}\neg t(C\setminus\{c\},n,S,R).$$

As we have provided formulas, and syntactic and semantic results for both forming and not forming a connector, minimal connectors and the syntactic and semantic implications of a minimal connector are obtainable using the tools provided in the previous subsection. These results are of a similar form as those of latency-minimal connectors. But instead of stating the minimally required number of iterated communications, these formulas state the minimally required members of the connector that must be included in the iterated communication updates for it to bring about knowledge of φ among the receivers. If we, for example, work out the formula for a $\forall^r\exists^s$-connector $C \subseteq A$, we get the following. For any $c \in C$ let $G^{-c} = S \cup C \setminus \{c\} \cup R$. If for any$c \in C$, $G \setminus S$ and $G^{-c} \setminus S$ respectively do not know that $[!G]^n\varphi$ and $[!G^{-c}]^n\varphi$, and any s and r together distributively know that φ after G or G^{-c} communicate n times, then C is a minimal $\forall^r\exists^s$-n-connector from S to R *iff* for every $c \in C$, some receiver does not know φ after G^{-c} communicate n times, and all receivers know that φ after G communicates n times.

$$\begin{aligned}\Vdash &\bigwedge_{c\in C}\bigwedge_{s\in S}\bigwedge_{r\in R}\left(\neg D_{(G\setminus S)\cup\{r\}}[!G]^{n}\varphi \wedge \neg D_{(G^{-c}\setminus S)\cup\{r\}}[!G^{-c}]^{n}\varphi \wedge D_{\{s,r\}}([!G]^{n}\varphi \wedge [!G^{-c}]^{n}\varphi)\right)\\ &\to \left(\mathrm{min}\forall^r\exists^s(C,n,S,R) \leftrightarrow \left(\bigwedge_{c\in C}[!G^{-c}]^n\neg E_R\varphi \wedge [!G]^n E_R\varphi\right)\right)\end{aligned}$$

As such, minimal connectors are connectors, such that any of its members is essential to realise distributed knowledge.

Besides these results, there is a particular relation between the minimal latency of $\exists^s\exists^r$-connectors and minimal $\exists^s\exists^r$-connectors.

Proposition 3.4.1. *For $S, R, C \subseteq A$, C is a minimal $\exists^s\exists^r$-connector from S to R iff C has a minimal $\exists^s\exists^r$-latency of $|C|+1$.*

Proof. For a sequence $P = (p_i)_{i=1}^{n}$ such that P forms an $S \cup C \cup R$-path from r to s, if there is a p_i for $1 < i < n$ such that $p_i \in S$ or $p_i \in R$, then $(p_j)_{j=0}^{i}$ or $(p_j)_{j=i}^{n}$ would form a path from an agent in R to an agent in S. Therefore, we can always find a sequence $P' = (p_i)_{i=m}^{n'}$ where m is the highest number such that $p_x \notin S \cup R$ for $0 < x < m$, $p_m \in R$, n' is the lowest number such that $p_x \notin S \cup R$ for $n' < x < n$, and $p_{n'} \in S$, i.e. P' forms a C-path from an agent in R to an agent in S. Hereby, $\exists s \in S \exists r \in R\ r \rightarrow_{S \cup C \cup R} s \iff \exists s \in S \exists r \in R\ r \rightarrow_C s$.

Moreover, as paths do not contain loops (and therefore do not visit an agent twice), $\exists s \in S \exists r \in R\ r \rightarrow_C s \iff \exists s \in S \exists r \in R\ r \rightarrow_C^{\leq |C|+1} s$; and $\exists s \in S \exists r \in R\ s \rightarrow_C^{\leq |C|} r \iff \exists c \in C \exists s \in S \exists r \in R\ s \rightarrow_{C \setminus \{c\}}^{\leq |C|} r$. Therefore, C is not a $\exists^s\exists^r$-connector from S to R iff $\forall c \in C \forall s \in S \forall r \in R\ s \not\rightarrow_{C \setminus \{c\}}^{|C|} r$.

C is a minimal $\exists^s\exists^r$-connector from S to R iff (i) C is a $\exists^s\exists^r$-connector, and (ii) for all $c \in C$, $C \setminus \{c\}$ is not a $\exists^s\exists^r$-connector from S to R. (i) $\iff \exists s \in S \exists r \in R\ r \rightarrow_C^{\leq |C|+1} \iff$ C has a $\exists^s\exists^r$-latency of $|C| + 1$. (ii) $\iff \forall c \in C \forall s \in S \forall r \in R\ r \not\rightarrow_{C \setminus \{c\}}^{\leq |C|} \iff \forall s \in S \forall r \in R\ r \not\rightarrow_C^{\leq |C|} \iff C$ does not have a $\exists^s\exists^r$-latency of $|C|$. Therefore, C is a minimal $\exists^s\exists^r$-connector from S to R iff C has a minimal $\exists^s\exists^r$-latency of $|C| + 1$. ☐

Corollary 3.4.1. *For two groups $G_s, G_r \subseteq A$ and $t \in \mathcal{T}$, $C \subseteq A$ is a minimal t-n-connector from G_s to G_r iff $C \setminus (S \cup R) = \emptyset$ and C has a $\exists^s\exists^r$-latency of $|C| + 1$.*

Proof. This follows from Proposition 3.2.4 and 3.4.1. ☐

3.5 Relations To Other Theories

Structural importance and centrality within a network is the subject of many theories and studies. In this section we will discuss our contributions to this field. We will discuss the relation between the concepts defined in this chapter and studies on diffusion in social networks, structural holes, and Gatekeeping Theory, and its formalisation in Belardinelli [19]. As the last study is closest to our setting, we will give it a more in-depth discussion. In particular, we will show that connectors and blocking sets are generalisations of structural notions in Belardinelli's work, adapted to a directed setting.

3.5.1 Diffusion in Social Network

There is a connection between the special abilities of blocking sets and what is discussed and formalised in Christoff and Naumov [26]. They consider the logic of *diffusion* in social network based on threshold models, an atomic perspective on the propagation of belief. Their logic is based on the work of Azimipour and Naumov [9], a logical study on the *lighthouse principle* of diffusion. Each agent has a threshold, a value denoting what percentage of neighbours must have adopted something for that agent to adopt it themselves. Given a network structure, these thresholds, and some *early adopters*, agents that have adopted behaviour without any peer-pressure, what are the relations of influence and propagation through peer-pressure in these threshold models from these early adopters to the rest of the group. Azimipour and Naumov [9] formalise the logic of this relation, denoted by $\rhd$, such that $S \rhd R$ iff if S were to adopt something without peer-pressure then R will eventually adopt it as well through peer-pressure.

Christoff and Naumov [26] extend this logic to incorporate an external set of *recalcitrant* agents: agents that are immune to influence. Recalcitrant agents are related to blocking sets, as blocking sets can be interpreted as sets of agents that, when taken to be recalcitrant to communication, make it impossible for one group to influence the knowledge of the other. Though an important distinction between recalcitrant agents and blocking sets is that, where recalcitrant agents do not adopt a given behaviour, agents in a blocking set *could* still adopt the knowledge about a formula in communication, as they still receive knowledge through communication and update their epistemic state accordingly. Instead, blocking sets stop the propagation of information by not communicating.

These studies abide by the axiom of atomic transmission: the propagation of belief is assumed to be solely determined by the transmission of atomic units. Our study paves a way for such behaviours to be studied in a non-atomic and epistemic setting. Connectors specify the agents required for "early adopters" of some knowledge φ to spread to other agents. Blocking sets, instead, specify agents that, when taken to be "recalcitrant" in communication, can block the spread of knowledge about φ to other agents.

Moreover, our study bring to light the opposite character of information spread in epistemics. When information spreads through a network, knowledge, rather, is concentrated in a smaller set. Connectors make it possible for information to spread from a sending set to a receiving set. And in doing so, they make it possible for information that is spread among the senders and receivers to be concentrated in the receivers. Whereas Christoff and Naumov [26] and Azimipour and Naumov [9] develop the logic of the former, we develop the logic of the latter.

3.5.2 Structural Holes & Redundancy

Burt [23] gives an account of relations that provide "network benefit" to each other by filling in a "structural hole" in the social network. Central to the theory of structural holes is the notion of (non-)redundancy of connections: structural holes are "separations between nonredundant contact"[23] in the social network. Here non-redundancy is indicated by connectedness in the network. Two agents are *redundant by cohesion* when they are connected by "a strong relationship".[12] Such agents are likely to share the same network benefits, as information will be shared between the two frequently. Two agents are *redundant by structural equivalence* when they both provide the same connections, and therefore have access to the same benefit.

In this chapter we regarded a setting where benefit is epistemic. This benefit flows not only through direct connections between agents, but also indirectly, through third parties by iterated communication. The latter is really distinct from the former in exactly the setting of directed knowledge realisation: when a set S together with a set R distributively know that φ, and a connector C allows for this knowledge to be realised, the connector together with the sender does not have to know that φ. The connector might not even get to know anything new when realising the distributed knowledge of φ to R by iterated communication. This chapter therefore can be a starting point for a formal account of structural holes in an epistemic setting where benefit, and therefore redundancy, is epistemic; and where connections between agents are not necessarily direct. Such an extension to Burt's work must start with the introduction of notions of *redundant connectors*. Here we will only hint towards definitions of such notions.

We can define a *cohesively redundant* connector from S to R as a set C such that C is a t-connector and S is directly t-connected to R (when $\emptyset$ is a connector from S to R). As we regard connections from sets to sets through a third party as well, it makes sense to define a stronger variant of cohesive redundancy: two sets are cohesively redundant when there is a connector from the one to the other. A t-connector from S to R is a *strong cohesively redundant* connector (up to n) iff it is not the only t-connector (with a latency $\leq n$) from S to R; i.e., the connector is not also a t-blocking set (t-n-delaying set).

Recall that minimal connectors are connectors such that all their members are essential for it to form a connector. Cohesively non-redundant connectors, instead, are connectors that are essential for S and R to be connected at all. Because of this, a strong cohesively

[12] Contrary to our social networks, the network in Burt [23] is bi-relational and bidirectional, it consists of weak and strong ties, so comparing Burt's work to ours only goes so far. One way around this difference is to assume that all connections are strong, as it is most likely that someone who someone shares all they know with is in a strong relationship with them.

non-redundant connector is necessarily also a minimal connector, but a minimal connector is not necessarily strong-cohesively non-redundant: the other connector C' from S to R that makes a connector C strong-cohesively redundant could contain agents outside C, $C' \nsubseteq C$; it could even be entirely disjoint from C.

Structural redundancy is not about whether S and R are connected, but about the overlap of the connections of S and R. We can define structural redundancy as follows: two sets of agents are *structurally redundant* when the sending set S is connected to agents that R is already connected to. A t-connector C is structurally redundant when there already is a connector C' from R to the agents followed by S, $\mathcal{F}^n(S)$, and this C' is disjoint from C. Or, when considering iterated communication and paths through the social network, a connector is structurally redundant when there already is a connector from R to the agents followed by S up to n social ties away ($\mathcal{F}^{+^n}(S)$), for some $n \in \mathbb{N}^+$. Further restrictions could be made to only regard overlapping agents whose information is reachable by a particular kind communication. Then structural redundancy is about whether there already is a t-connector from R to $\mathcal{F}|_G^{+^n}(S)$, for a particular $n \in \mathbb{N}^+$, $G \subseteq A$, and $t \in \mathcal{T}$.

Finally, Burt relates redundancy to mutual *benefit*. Therefore, a natural variant of structural and cohesive redundancy is *epistemical redundancy*. Senders S and receivers R are epistemically redundant when $S \cup R$ knows no more than R. Connectors from S to R fill in a structural hole with respect to epistemics when S and R are epistemically non-redundant. Of relevance to such a definition of redundancy are the comparative knowledge relations defined in van Ditmarsch, van der Hoek, and Kooi [65], and formalised in relation to communication in Baltag and Smets [16].

3.5.3 Gatekeeping Theory

Related to the concept of structural holes is Gatekeeping theory. Gatekeeping theory is concerned with social network positions, or agents, that have full control over the flow of information between two groups. These agents are called *gatekeepers*, as they "gatekeep" information between these groups. Gatekeeping as a concept has many interpretations. Barzilai-Nahon [18] makes an attempt to formulate these into a single theory. A more formal, and logical, attempt at this is undertaken in Belardinelli [19]. In it, several structurally important positions in a bidirectional and connected social network are identified and defined, characterisations of agents that enable information flow, block information flow, and have full control of information flow. These concepts are defined under the assumption that the social network is undirected (symmetric), irreflexive, and connected.

Gatekeepers are sets that have full control over the information flow between two groups. Along with gatekeepers, several other structural notions are discussed and formalised, of which we will focus on connectors and blocking sets.[13] For connectors, Belardinelli gives the following definitions. As the names of the concepts of Belardinelli and ours overlap, we will prefix Belardinelli's concepts with "bidirectional".

> Let G, G' be two disconnected groups [$G \cup G'$ is not a group]. We say that [$B \subseteq A$] is a [bidirectional] *connector* between G, G' iff $G \cup G' \cup B$ is a group. ([19, p. 12])

The definition of bidirectional connectors closely resembles the definition of a redundant connector provided in this chapter. Their differences are as follows: (1) a connector is defined bidirectionally: a bidrectional connector is a connector *between* G and G'; (2) a bidirectional connector is only defined between groups; and (3) a bidirectional connector must connect two sets that are not directly connected to each-other. The first two reflect the two major

[13] The other notions defined in Belardinelli's work are: bridges, bridging sets, C-local gatekeepers, gatekeeping sets, gatekeeping bridges, and the grand gatekeeper. Bridges are sets C such that C is a minimal $\exists^s\exists^r$-connector from S to R and from R to S; bridging sets are defined in terms of bridges; C-local gatekeepers are agents $c \in C$, such that $c \in B$ for some bridge $B \subseteq A$; gatekeeping sets are bridging sets that are also blocking sets; and the grand gatekeeper is a maximal bridging set.

differences between Belardinelli's social networks and ours: (1) reflects the assumption of a symmetric network made in Belardinelli [19], while (2) reflects the assumption of connectedness. Note that because of this difference, the definition of connectors is defined in terms of groups, whereas our definitions are all reducible to (definable in terms of) connectors (see Proposition 3.2.18, and 3.3.1).

For blocking sets, Belardinelli gives the following definition. We will take some liberty in the exact phrasing and (mostly stylistic) details of this definition, as some of these details do not match our definition.[14]

> Let G, G' be two disconnected groups [$G \cup G'$ is not a group] and consider some [$B \subseteq A$] [...]. We say that [B] is a [bidirectional] *blocking set* between G, G' iff every connector C between G, G', contains an element of [B], i.e. for all [$C \subseteq A$], if $G \cup G' \cup C$ is a group, then $C \cap [B] \neq \emptyset$. ([19, p. 28])

To make the connection between bidirectional blocking sets and the definition of blocking and delaying sets apparent, consider the following proposition.

Proposition 3.5.1. *For non-empty sets $S, R \subseteq A$ and $t \in \mathcal{T}$, a set $B \subseteq A$ is a t-blocking set (t-n-delaying set) $B \subseteq A$ in $\mathfrak{F}$ iff for every t-connector (t-n-connector) C from S to R in $\mathfrak{F}$, $C \cap B \neq \emptyset$.*

Proof. The right-to-left direction follows trivially as $\overline{B} \cap B = \emptyset$. The left-to-right direction follows from the fact that not being a t-connector, and t-n-connector is closed under supersets. □

The formal connections between these definitions become apparent when we consider frames whose social network relation is symmetric. For $G, G' \subseteq A$ such that $G \cup G'$ does not form a group, our definition of a connector coincide with that of Belardinelli [19]. The same holds for blocking sets.

Proposition 3.5.2 (Equivalence of definitions)**.** *Let $\mathfrak{F}$ be a frame whose social relation is symmetric and let $G, G' \subseteq A$ be groups such that $G \cup G'$ is not a group, then:*

(i) C is a bidirectional connector between G and G' in $\mathfrak{F}$ iff C is a t-connector from G to G' and from G' to G in $\mathfrak{F}$, for any $t \in \mathcal{T}$.

(ii) B is bidirectional blocking set between G and G' in $\mathfrak{F}$ iff B is a t-blocking set from G to G' and from G' to G in $\mathfrak{F}$, for any $t \in \mathcal{T}$.

Proof. Let $\mathfrak{F}$ be a frame whose social relation is symmetric.

(i) As G and G' are bidirectional groups, $G \cup C \cup G'$ is a bidirectional group iff there is a bidirectional path in $G \cup C \cup G'$ from some agent in G and some agent in G'. Furthermore, $G \cup G'$ is not a group iff there is no edge between some agents distinct agents in $G \cup G'$. As connector types are equivalent when the sending and receiving sets are groups (Proposition 3.2.4), $G \cup C \cup G'$ is a bidirectional group and $G \cup G'$ is not iff C is a t-connector from G to G' and from G' to G for any $t \in \mathcal{T}$.

(ii) Follows from (i), the definition of bidirectional blocking sets given above, and Proposition 3.5.1. □

[14] Belardinelli does not consider the entire set of agents A (or $\mathcal{A}$ in the case of Belardinelli) a blocking set, while these are always blocking sets in our case. Likewise, Belardinelli does not consider blocking sets that overlap with the two groups as blocking sets. In our case, we do. We have shown a correspondence between overlapping and non-overlapping blocking sets in Proposition 3.3.1. As these differences therefore seem mostly stylistic, we will ignore them.

Belardinelli formalises their concepts in *Network Logic*. Network logic is an agent-evaluated Propositional Dynamic Logic with elements of hybrid logic. The logic lacks any competent epistemic notions. Instead, it employs an atomic account of information and communication: information is represented as "bits" of data from a set $\mathcal{D}$. For $d \subseteq \mathcal{D}$, "d" is to be read as "I am informed of d" or "I have the data bits d". The logic formalises communication by *posting actions*. Agents can post data to their friends only if they are informed of said data. Sequences of posting actions model iterated communication actions. The language has a modality $\langle\langle A \rangle\rangle$ related to Coalition Logic [47] and STIT ("Seeing To It That") Logic [70]. $\langle\langle A \rangle\rangle\varphi$ denotes that A can bring about φ using some *sequence* of posting actions by agents in A. Thereby, the logic abstracts away from the number of times agents must communicate.

Characterisations in Network Logic are provided for all structural notions defined in Belardinelli [19]. Further formulas are provided for the communicational properties of these structural notions. We will only discuss the communicational formulas for connectors, as the others are of a similar form. The communicational formula for connectors is:

$$(G_1 \to d) \to \langle\langle C \rangle\rangle \exists (G_2 \wedge \Diamond d). \tag{3.6}$$

G_1 and G_2 are group nominals, reading "I am part of G_1/G_2"; $\exists$ is the universal quantification over all agents, to be read as "there is an agent such that"; $\Diamond$ is the diamond modality over the social relation, $\Diamond\varphi$ is read as "I have a neighbour such that φ". Hence, the formula must be read as follows: "if G_1 has the information d ($G_1 \to d$); then, through iterated communication, C can bring about ($\langle\langle C \rangle\rangle$) that there is an agent ($\exists$) who is part of G_2, and who has a neighbour who is informed about d ($G_2 \wedge \Diamond d$)." It is proven that this formula is true in a model only if, in that model, C forms a connector between the groups named G_1 and G_2, and these groups don't form a group together.

The meaning and form of (3.6) is similar to our communicational formulas for connectors in Proposition 3.2.11 to 3.2.15. It differs in that in (3.6): (i) it is not syntactically specified that $b_1, \dots, b_n$ form a bidirectional connector (even though formulas are provided for it); (ii) after communication, an agent in G_2 is not informed of b themselves, but they are socially related to someone who is; (iii) because the object that is communicated is propositional, the "epistemic" precondition ($G_1 \to d$) in (3.6) is simpler than the epistemic preconditions in Proposition 3.2.11 to 3.2.15, that state the distribution of the knowledge about φ among S and R; and (iv) the number of communication actions required is not stated, whereas this is stated in Proposition 3.2.11 to 3.2.15; (i) and (ii) are mostly stylistic differences; (iii) stems from the different approaches to information and epistemics in both logics; and (iv) is a result of the abstraction of the number of communication actions required that in Network Logic does.

To conclude, our work is related to Belardinelli's, and contributed to the logical analysis of Gatekeeping Theory, in two ways. Firstly, it functions as a generalisation of many of the definitions of Belardinelli to a setting of *directional* social networks. Secondly, it is a development of structural notions and their implications towards an epistemic account of communication. Moreover, we introduced two properties of connectors not discussed in the work of Belardinelli. We have developed a quantitative account of communication, taking into account the number of communication actions required for connectors to be effective, and we introduced quantitative distinctions between different types of communication and their respective knowledge realisation. These extensions allow us to distinguish the communicational implications of structural notions that are not distinguishable in Belardinelli's work. For example, the communicational formulas of bridges and bidirectional connectors, as defined by Belardinelli, are identical, whereas they are not in our case. Expressed in connectors, a bridge C is a minimal $\exists^s\exists^r$-connector. Hence, by Proposition 3.4.1, its communicational properties differ from non-minimal connectors in that, for bridges, it only takes $|C| + 1$ communication updates before distributed knowledge is realised.

Chapter 4

Network formation and games

In the previous chapter we have discussed distributed knowledge resolution and its relation to the structure of the social network: given a certain network structure N, a set of senders S, and a set of receivers R that together distributively know something, we identified what sets are able to bring about and block communication between S and R. In this chapter, we will build a framework to reason about how these networks are formed and changed, ultimately to reason about how the structural requirements of connectors and blocking sets come about.

Most network logics formalise social network dynamics *explicitly* in their syntax. In such logics, the network is formed and changed by syntactically specifying the edges to be removed or the ones that survive, depending on the perspective. This as opposed to some logics that only syntactically specify that an update takes place, but not what the update entails. In such logics, the edges to be removed or added is determined by the semantics of this update and the current model. Examples of the latter, implicit, approach are the logics of network formation through threshold models in Smets and Velázquez-Quesada [57] and Smets and Velázquez-Quesada [58], which employ modalities of threshold updates, where edges are formed depending on how similar agents are w.r.t. certain properties. Examples of logics of explicit social network dynamic are the logic with a basic edge deletion and addition modalities in Seligman, Liu, and Girard [55], and the logic with follow and unfollow modalities in Ruan and Thielscher [53]. These logics both modify the network by adding or deleting edges by name. In Roelofsen [52], connections between groups of agents are added or deleted based on the truth value of a precondition in a "reconfiguration event".

Logics of explicit network dynamics are closely related to relational update logics such as Sabotage Logic [7; 59], that introduces a modality for deleting "any edge" in a graph.[1] For a general study and formalization of relation update logics, see Areces, Fervari, and Hoffmann [6].

The common denominator of these logics is their *outside perspective* on the dynamics of the network structure. Rather than being manipulated by agents, the network is changed from the outside by modalities that induce a network change. The usual interpretation of such modalities is that of *possibility and necessity*: it is possibly or necessarily so that, when a social relation between a and b is added, a can know what b knows after communication, and so on. As such, these logics are *descriptive*, only indicating what would (possibly or necessarily) happen after an action takes place.

Instead, we take an *inside* perspective, where agents themselves form and change the network structure. Therefore, instead of possibility we focus on *ability*. Hereby, we bring

[1]Most dynamic epistemic logics such as Public Announcement Logic, but also our Communication Logic modify a relational structure in exactly such a way. These are not "designed to work" on social networks. Rather they present implicit modalities that modify a relational (Kripke) structure in such a way that aligns with their specific modal characteristics (often **S45** or **S5** knowledge).

network dynamics to closer resemble what will *actually* happen. A natural mathematical setting for this is *strategic games*.

In the upcoming section we develop a generic game-theoretic framework for network formation, where agents actively decide on and shape their network structure in a single-shot game. Then, in Section 4.2, we pose some reasonable restrictions on the class of games we consider. In Section 4.3, we give an overview of the games that underlie most socio-economic studies of network formation, and link these to the presented restrictions. Finally, in Section 4.4, we move to a setting of *network change*, and model games of network formation in a given social environment. We discuss variants of the restrictions for single-shot games, for this extensive setting. In the next chapter, we will develop a formal language to reason about ability in such games.

4.1 Games

In this section we will work towards a generic game-theoretic framework to reason about network formation and change. We will do this from the perspective of *strategic* games rather than cooperative[2] games, as we are not concerned with what coalitions will form under which "value distribution functions/rules", but rather on the coalitional ability of agents to shape the social network.

Game theory presents a dichotomy between simultaneous and non-simultaneous play. In the former, players choose their actions all at once, independent of each other, without regard of time; and thereby while not being informed of the actions of the others. In the latter players take turns choosing actions, allowing them to revise their strategy based on the moves played by the others before them. [44] In most social settings, relations are made and broken coincidentally, with no regards of any particular temporal order over the agents manipulating the network topology. For this reason, we take strategic games with simultaneous play as the starting point of our investigation of the dynamics of network formation.

4.1.1 Single-shot Games With Simultaneous Moves

We commence from single-shot games with simultaneous moves, typically called *strategic games*.[3] Strategic games are played by a set of players A. According to the standard definition, each player $a \in A$ is assigned a set of available actions Σ_a. A tuple containing a single action for each player (each from their respective action set) is called an *action profile*.

Usually, the outcome of a game is represented by *score*: given an action profile, each player is assigned a score by an *utility function*. It is assumed that players try to maximise their utility. In this way the utility function formalises player *incentive*. However, as we are interested in the process of network formation via game theory, and we are particularly interested in whether agents are able to shape the network in order to reach certain network positions, we let the outcome of a game to be a network topology. We will not assign agents any preference order over these networks. Therefore, rather than games, what we discuss is in actuality a game *form*. These game forms become a game when paired with a preference order. We will refer to these game forms as games for the sake of convenience.

Our games are played by a non-empty and finite set of players A. Let $n = |A|$ denote the number of players of the game. Let S be a non-empty set of possible outcome states. All players are asked to submit an action from their respective non-empty action or choice set Σ_i.

[2] For an overview of cooperative games and network formation see Jackson [37].

[3] For an overview of strategic games see Peters [50], in particular Chapter 6, or Osborne et al. [44] Section 2.1.

Definition 4.1.1 (Actions)**.** For each player $a \in A$, associate a non-empty choice set Σ_a. Let σ_A denote an n-tuple of choices, one for each player: $\sigma_A = (\sigma_1, \sigma_2, \ldots, \sigma_n)$ where $\sigma_1 \in \Sigma_1, \sigma_2 \in \Sigma_2, \ldots, \sigma_n \in \Sigma_n$. We call σ_A an action profile. For brevity's sake, when the set of all players A is clear from the context we will denote σ_A by σ.

For any $C \subseteq A$, let $\Sigma_C = \prod_{i \in C} \Sigma_i$ denote the set of all action profiles for C. We call Σ_C an action set or choice set for C. For any $C \subseteq A$ we denote an action profile for C by $\sigma_C \in \Sigma_C$. If $C = \{x\}$, we write σ_{-x} for $\sigma_{\overline{C}}$.

Players submit their actions σ_a simultaneously. Accordingly, players cannot act on the actions that the other players submit. Actions are aggregated into an outcome by the *outcome function o*. Let S be a non-empty set of outcome states. The function o maps action profiles for A to outcome states.

Definition 4.1.2 (Outcome function)**.** For a set of players A, and an action set Σ_A, an outcome function $o : \Sigma_A \to S$ maps each action profile to a state S.

Given a $C \subseteq A$, let $o(\sigma_C, \sigma_{\overline{C}}) = o(\sigma_A)$, where σ_A is the action profile on A induced by σ_C an $\sigma_{\overline{C}}$. For readability, we will omit the brackets in case of $C = \{x\}$; writing $o(\sigma_x, \sigma_{-x})$ for $o((\sigma_x), \sigma_{-x})$.

A strategic game consists of a set of action profiles for A, a set of outcome states S, and an outcome function that maps elements of Σ_A to S.

Definition 4.1.3 (Strategic game)**.** A strategic game $\mathcal{G}$ is a tuple

$$\mathcal{G} = (A, \Sigma_A, o, S)$$

where A is a non-empty set of agents, Σ_A is the action set, o is the outcome function, and S is a non-empty set of outcome states.

We model network formation as a strategic game as follows. Let $\mathbf{F}_A = \mathscr{P}(A \times A)$ denote the set of possible network structures on A. Each agent a is asked to submit a social network σ_a from their respective choice set of social network $\Sigma_a \subseteq \mathbf{F}_A$. The outcome function then produces an outcome network from the action profile of chosen social networks $o(\sigma_A) \in \mathbf{F}_A$ We interpret an agent's choice σ_a as the agent making an attempt to shape the network according to σ_a. We call these games *network games.*

Definition 4.1.4 (Network game)**.** A strategic game $\mathcal{G} = (A, \Sigma_A, o, S)$ is a network game iff for all $a \in A$: $\Sigma_a \subseteq \mathbf{F}_A$, and $S = \mathbf{F}_A$.

4.1.2 Social Choice Theory

There is a clear connection to our game-theoretic framework and *social choice theory*[4]. Social choice theory is concerned with the definition and properties of social choice functions: functions that aggregate the preference orders of agents over a decision to a social preference, much like how outcome functions "aggregate" the social networks proposed by agents to an outcome social network.

There are two key differences between classic social choice theory and our setting. Firstly, social choice theory is primarily concerned with *preference*, agents submit a preference order over their choices. In our framework, agents only submit a single choice. This could be interpreted as them submitting their most-preferred network. But this is rather deceiving as we are not concerned with the category of preference to begin with. Secondly, within classical social choice, it is assumed that the preference orders have a *universal domain*: there might be restrictions over the type of preference orders allowed (e.g. total orders, total preorders, single-peaked preferences), but these ordering are over all possible world states.

[4]For an overview of social choice theory see Fishburn [30] or Brandt et al. [22].

Hereby, unlike our setting, no predetermined restriction can be made on what each agent is allowed to change in the world.

There exists work on social choice theory that considers restrictions on the voting rights (the choice sets) of agents, for example social choice theory of rights. Originally, social choice theory of rights used similar models to classical social choice theory. However, as proposed and popularised by Gärdenfors [33], social choice theory of rights is also studied in game form, a framework similar to the strategic games we use.

Besides social choice theory of rights, our definition of a strategic game is also equal to those used in social choice theory from the non-normative perspective [1], and to the games typically used in Coalition Logic [47].

4.2 Axioms for Network Games

As it stands, network games are too general for them to represent network formation processes. Not all network games are interesting or make sense. In particular, the interpretation we gave to network games — that when an agent x submits a network σ_x, it means that x makes an attempt to shape the social network topology according to σ_x — does not directly follow from the definition of network games. Elements of Σ_A represent abstract choices, and only affect the possible games by limiting the number of choices an agent has. Their interpretation as attempts to shape a social network to a particular topology only come about by pairing them with particular outcome functions. We must therefore restrict the network games that we consider.

Like axiomatic social choice theory [22, ch. 2], we will take an axiomatic approach to identify these "sensible" network games. However, whereas the axiomatic approach in social choice theory is mostly normative [22], searching for axioms that express various normative properties of voting, our axioms are not to be interpreted in such a way. Instead, we take the axioms as simple propositions about the process of network formation. We will list some such axioms in the upcoming section. We do not intend for this list to be complete.

There are two crucial components of network games that we must discuss: properties of choice sets, indicating which proposals each agent is allowed to submit, and properties of outcome functions, indicating how these proposals lead to outcome networks.

Note that any restriction of the choice set is optional, in that such restrictions can be "embedded" in the outcome function by sending all choice sets not included in a restricted Σ_A to some dummy outcome. However, laying restrictions on the choice set is an intuitive way to come to certain network games, as we will see in the upcoming sections. Therefore, we will discuss such restrictions.

4.2.1 Properties of the Choice Set

We postpone the discussion on the exact contents of the choice sets to the next section. First we will regard their closure properties.

An important property of aggregation functions in social choice theory is *anonymity*. Anonymity requires that the name of the person that submits a preference order over the candidates does not matter. We can define a similar notion in network games. Because of the differences between social choice theory and network formation games, a proper anonymity condition requires some reconsideration. First we define *agent permutations*.

Definition 4.2.1 (Agent permutation)**.** An agent permutation π on a social network (A, F) is a bijection $\pi : A \to A$. For the application of a permutation π on a relation $F : A \times A$ we write $\pi(F) = \{(\pi(i), \pi(j)) \mid (i, j) \in F\}$. Recall that action profiles are tuples of networks. For readability, we extend the notation for permutations to also work on such tuples: for any action profile for $G \subseteq A$: σ_G, we write $\pi(\sigma_A)$ for $(\pi(\sigma_i))_{i \in A}$.

With respect to choice sets, anonymity entails that whenever an agent can submit an action, any other agent should also be able to submit it. In our particular setting, the choices that agents have are made up of these same agents — the networks between which the agents A have to choose are defined over A. Consequently, there is another requirement for network formation games to be truly anonymous: that the particular choice sets are not defined in terms of relations between particular individuals, i.e. that the choice sets are closed under permutations over agents. We say that choice sets, and games defined over these choice sets, satisfy *pure anonymity of choice*, when they satisfy both of these anonymity conditions

Definition 4.2.2 (Pure anonymity of choice)**.** A choice set Σ_A is purely anonymous iff for any agent $a \in A$, and for any permutation π

1. $\sigma \in \Sigma_a \iff \sigma \in \Sigma_{\pi(a)}$ (i.e. for all $b \in A$: $\Sigma_a = \Sigma_b$), and
2. $\sigma \in \Sigma_a \iff \pi(\sigma) \in \Sigma_a$.

A game satisfies *pure anonymity of choice* when its choice set is purely anonymous.

Condition (1.) states that if an agent can submit a certain network, then every agent can. Condition (2.) states that an agent a should not be limited in their network choices solely because of the names of the agents that are (dis)connected in them.

Pure anonymity of choice is a "literal" translation of the concept of anonymity of social choice to the choice sets of network games. It, however, is often too strong for network games. Consider the following example.

Example 4.2.1 (Star shaped network games)**.** A *star-shaped network* is a social network $F \in \mathbf{F}_A$ for which there is an $a \in A$ s.t. for all $(i, j) \in F$, $j = a$. Let the set of all star shaped networks be denoted by $\mathbf{F}_A^\star \subseteq \mathbf{F}_A$. Let the choice set Σ_A be such that for all $a \in A$, $\Sigma_a = \mathbf{F}_A^\star$.

A star shaped network centered at $a \in A$ is a network F such that for all $(i, j) \in F$, $j = a$. Let the set of star shaped networks centered at a be denoted by $\mathbf{F}_A^{\star(a)} \subseteq \mathbf{F}_A^\star$.

Consider the choice set Σ'_A where all agents $a \in A$ can choose the star shaped networks centered around a particular $x \in A$, i.e. for all $a \in A$, $\Sigma'_a = \mathbf{F}_A{}^{\star(x)}$. Next consider the choice set Σ''_A where all agents $a \in A$ can choose the star shaped networks centered around themselves, i.e. for all $a \in A$, $\Sigma''_a = \mathbf{F}_A^{\star(a)}$.

The choice set Σ_A is purely anonymous, but neither Σ'_A nor $\Sigma_{A''}$ are. For Σ'_A, condition (2.) fails — as it would impose on Σ'_A that any agent can choose all star shaped networks, centered around any agent. For Σ''_A, condition (1.) and (2.) fail: (2.) fails for the same reason as for Σ'_A, and (1.) would impose that any agents can choose all star shaped networks with any agent in the center.

Intuitively, Σ'_A is indeed not anonymous. However, Σ''_A *is* anonymous in some sense: all agents are treated equally regardless of their label or name, in the sense that any agent is allowed to choose all star shaped networks with themselves in the center. A more natural anonymity requirement for network games, therefore, is that when an agent a can choose a network shape relative to a, then any other agent b should be able to choose that shape relative to b. We introduce a different notion of anonymity to capture this intuition: *pseudo-anonymity*. Pseudo-anonymity requires that a can choose σ when $\pi(a)$ can choose $\pi(\sigma)$ for all permutations π.

Definition 4.2.3 (Pseudo-anonymity of choice)**.** A choice set Σ_A is pseudo-anonymous iff for all permutations π:

$$\sigma_a \in \Sigma_a \iff \pi(\sigma_a) \in \Sigma_{\pi(a)}.$$

A game satisfies *pseudo-anonymity of choice* when Σ_A is pseudo-anonymous.

For the sake of clarity, we will give an overview of the definitions presented in this section. For a summary of the definitions in this subsection, see Table 4.1.

Name	Description
Pure anonymity of choice	Choice set not determined by agent name.
Pseudo-anonymity of choice	Choice determined by name only in reference to self.

Table 4.1: Summary of game axiom definitions

4.2.2 Relations Between Choice Set and Outcome Function

To force network games to represent games where agents have choice in shaping the network, an intuitive requirement is that no social relation in the network will come about spontaneously: (i, j) should only be in the outcome of a game if at least some agent had a choice in the matter. The simplest such requirement would be that if $(i, j) \in o(\sigma_A)$ then there is an agent $x \in A$ such that $(i, j) \in \sigma_x$. Formally, a game $G = (A, \Sigma_A, o, S)$ must be so that:

$$(i, j) \in o(\sigma_A) \Longrightarrow \exists x \in A \text{ s.t. } (i, j) \in \sigma_x.$$

This does not ensure that x actually had a choice in including (i, j) in their σ_x; the pair (i, j) could be present in all choices of agent x (i.e., in all networks in Σ_x). Therefore, we present a stricter version of such an axiom, that ensures that x could have chosen an action that does not include (i, j). *Weak positive choice* reflects that when a social tie exists in the outcome of a game, at least some agent made an attempt for the tie to be in the network and that attempt was a choice for that player — they could have chosen not to make the attempt.

Definition 4.2.4 (Weak positive choice)**.** A network game $\mathcal{G} = (A, \Sigma_A, o, S)$ satisfies weak positive choice when:

$$(i, j) \in o(\sigma_A) \Longrightarrow \exists x \in A \text{ s.t. } (i, j) \in \sigma_x \text{ and } \exists \sigma'_x \in \Sigma_x \text{ s.t. } (i, j) \notin \sigma'_x$$

Note that in weak positive choice, an agent has a choice in including (i, j) or not if they can submit a network σ_x^+ and σ_x^- such that $(i, j) \in \sigma_x^+$ and $(i, j) \notin \sigma_x^-$. However, this choice is not pure in that the agent does not have a choice between a network σ including (i, j) and that same network without (i, j): $\sigma \setminus \{(i, j)\}$. If they do, we say that the game has *strong positive choice.*

Definition 4.2.5 (Strong positive choice)**.** A network game $\mathcal{G} = (A, \Sigma_A, o, S)$ satisfies strong positive choice when:

$$(i, j) \in o(\sigma_A) \Longrightarrow \exists x \in A \text{ s.t. } (i, j) \in \sigma_x \text{ and } \sigma_x \setminus \{(i, j)\} \in \Sigma_a$$

To illustrate the difference between weak and strong positive choice, regard the following example.

Example 4.2.2 (Bidirectional game)**.** For all $a \in A$: $\Sigma_a = \{F \in \mathbf{F}_A \mid (i, j) \in F \text{ iff } (j, i) \in F\}$. Paired with the outcome function:

$$(i, j) \in o(\sigma) \text{ iff } \exists x \in A \text{ s.t. } (i, j) \in \sigma_x$$

The bidirectional game satisfies weak positive choice, but not strong positive choice: when an agent x proposes a σ_x such that $(i, j) \in \sigma$, then $(i, j) \in o(\sigma)$; however, x does not have a pure choice between σ_x such that $(i, j) \in \sigma_x$ and $\sigma_x \setminus \{(i, j)\}$.

Positive choice has a natural dual: *negative choice.* Again, we present a weak and a strong version. *Weak negative choice* states that if a social tie is excluded from the outcome of a game, then at least some agent made an attempt for the tie not to be in the network *and* that attempt was a choice for that player — they could have chosen not to make it.

Name	Description
Weak positive choice	Edge can exist only when an agent had a choice in the matter.
Strong positive choice	Edge can exist only when an agent had a pure choice in the matter.
Weak negative choice	Edge cannot exist only when an agent had a choice in the matter.
Strong negative choice	Edge cannot exist only when an agent had a pure choice in the matter.

Table 4.2: Summary of game axiom definitions

Definition 4.2.6 (Weak negative choice)**.** A network game $\mathcal{G} = (A, \Sigma_A, o, S)$ satisfies weak negative choice when:

$$(i,j) \notin o(\sigma_A) \implies \exists x \text{ s.t. } (i,j) \notin \sigma_x \text{ and } \exists \sigma'_x \in \Sigma_x \text{ s.t. } (i,j) \in \sigma'_x$$

As with strong positive choice, strong negative choice is the variant of weak negative choice where "choice" is interpreted more drastically.

Definition 4.2.7 (Strong negative choice)**.** A network game $\mathcal{G} = (A, \Sigma_A, o, S)$ satisfies strong negative choice when:

$$(i,j) \notin o(\sigma_A) \implies \exists x \text{ s.t. } (i,j) \notin \sigma_x \text{ and } \sigma_x \cup \{(i,j)\} \in \Sigma_x$$

Negative choice is arguably less applicable to network formation games, as it could well be that the formation of a link is limited by some outside factor. Negative choice characterises settings in which such limitations are not present.

The game in Example 4.2.2 satisfies negative choice, but does not satisfy strong negative choice: when no agent proposes a σ_x such that $(i,j) \in \sigma$, then $(i,j) \notin o(\sigma)$; however, no agent has a pure choice between a σ_x such that $(i,j) \notin \sigma_x$ and $\sigma_x \cup \{(i,j)\}$.

For a summary of the definitions introduced in this subsection, and their purposes, see Table 4.2.

4.2.3 Properties of the Outcome Function

A natural requirement for outcome functions of network games is *monotonicity* with respect to edge existence. A game satisfies *positive monotonicity*[5] if more agents attempting to form a network topology with (i,j) in it will not remove (i,j) from the outcome network. That is, if the social relation (i,j) is in the outcome of a game with action profile σ, and if a player x that submitted a network σ_x without (i,j) in it switches to $\sigma_x \cup \{(i,j)\}$, then (i,j) is still in the outcome of the game.

Definition 4.2.8 (Weak positive monotonicity)**.** A network game $\mathcal{G} = (A, \Sigma_A, o, S)$ satisfies weak positive monotonicity when: for any $x \in A$, if $(i,j) \in o(\sigma)$, $(i,j) \notin \sigma_x$, and $\sigma_x \cup \{(i,j)\} \in \Sigma_x$, then $(i,j) \in o(\sigma_{-x}, \sigma_x \cup \{(i,j)\})$.

To illustrate weak positive monotonicity, regard the following example.

Example 4.2.3 (Interval game)**.** Consider the following game of intervals. Each player can choose between any network: $\Sigma_x = \mathbf{F}_A$. In this game, the agents are shy and need a certain number of agents to encourage them to form a social link. However, if too many agents encourage them, the agent will be over-encouraged, and they will not form the social

[5] This is closely related to *Maskin monotonicity* in social choice theory: if x is a winner of a vote, and a voter a switches to a ballot in favour of x, then x remains the winner [21]. Still, the concept of weak versus strong monotonicity in social choice theory is distinct from ours.

relation. The number of required encouragement must fall into a certain interval. Formally this can be represented as follows:

$$(i,j) \in o(\sigma) \text{ iff } l \leq |\{a \in A \mid (i,j) \in \sigma_a\}| \leq u.$$

This game is not weakly positive monotone: if exactly u agents submit a profile with (i,j) and there is an agent x who did not submit (i,j), then when x does submit i,j, (i,j) will not be included in the outcome.

The above formulation of positive monotonicity is weak in the following sense: given a profile σ that produces an outcome $o(\sigma)$ with (i,j) in it, weak positive monotonicity dictates that (i,j) must be in the outcome $o(\sigma')$ for any σ' that differs from σ in that *only the choice on (i,j) is changed* in some agent's choice set. Instead, we can require that (i,j) should be in the outcome of $o(\sigma')$ for any profile σ' that differs from σ in some agent's choice set, such that the agent now includes (i,j) in their choice, *no matter what that agent submits for the other edges.* We call this requirement *strong positive monotonicity.*

Definition 4.2.9 (Strong Positive monotonicity)**.** A network game $\mathcal{G} = (A, \Sigma_A, o, S)$ satisfies strong positive monotonicity when: For any $x \in A$, if $(i,j) \in o(\sigma)$ and $(i,j) \notin \sigma_x$ then for all $\sigma'_x \in \Sigma_x$ such that $(i,j) \in \sigma'_x$, $(i,j) \in o(\sigma_{-x}, \sigma'_x)$.

Note that strong positive monotonicity implies its weak variant, take $\sigma'_x = \sigma_a \cup \{(i,j)\}$. To illustrate strong positive monotonicity, regard the following example.

Example 4.2.4 (Triadic Closure)**.** Consider a setting of mutual consent, where a social tie (i,j) is formed only when both i and j include it in their submitted network choice. For simplicity's sake, let all agents be able to choose any network: for all $a \in A$ $\Sigma_a = \mathbf{F}_A$. First, regard an outcome function of mutual consent called m, where the outcome contains an edge only if both parties agree.

$$(i,j) \in m(\sigma_A) \text{ iff } (i,j) \in \sigma_i \text{ and } (i,j) \in \sigma_j.$$

An empirical phenomenon observed in the shape of social networks is that of triadic closure [28]. Often, this phenomenon is regarded in a bidirectional setting, but we can translate triadic closure to a directional setting as follows: two agents a, b are more likely to have a social tie with each other when they have a friend in common — the link $a\ F\ b$ is more likely to form when there is an agent c such that $a\ F\ c$ and $b\ F\ c$. Consider the effects of triadic closure, in its most extreme form, on the setting of mutual consent described above. Define an outcome function o of the triadic closure of m:

$$(i,j) \in o(\sigma_A) \text{ iff } (i,j) \in m(\sigma_A) \text{ or } \exists x \in A \text{ s.t. } (x,i) \in m(\sigma_A) \text{ and } (x,j) \in m(\sigma_A)$$

For a submitted σ, $o(\sigma)$ contains an edge (i,j) iff both agents agree to this, or there is an agent x that mutually consents with i and j respectively to be reachable from i and j.

The game of triadic closure satisfies weak positive monotonicity, but not strong positive monotonicity. Consider a set of players of three agents $A = \{a,b,c\}$. Let $\sigma_a = \{(a,c)\}$, $\sigma_b = \{(b,c)\}$, $\sigma_c = \{(c,a),(c,b)\}$. Then $o(\sigma_A) = \{(a,b),(b,a),(a,c),(c,a),(b,c),(c,b)\}$. Furthermore, o is trivially weakly monotone: adding a single edge will never delete an edge in the outcome. Finally, for σ'_A such that $\sigma'_a = \{(a,b)\}$, $\sigma'_b = \sigma_b$, and $\sigma'_c = \sigma_c$ $(a,b) \notin o(\sigma'_A)$. Therefore, o is not strongly positively monotone.

The dual of positive monotonicity is *negative monotonicity*: negative monotonicity reflects that the outcome function acts monotone with respect to the attempts made *against* the existence of a social tie.

Definition 4.2.10 (Weak negative monotonicity)**.** A network game $\mathcal{G} = (A, \Sigma_A, o, S)$ satisfies weak negative monotonicity when: for any $x \in A$, if $(i,j) \notin o(\sigma)$, $(i,j) \in \sigma_x$, and $\sigma_x \setminus \{(i,j)\} \in \Sigma_x$ then, $(i,j) \notin o(\sigma_{-x}, \sigma_x \setminus \{(i,j)\})$.

An example of an outcome function that is not weakly negatively monotone is the interval game in Example 4.2.3. If Σ_A is such that $|\{a \in A \mid (i,j) \in \sigma_a\}| = u+1$, then $(i,j) \notin o(\sigma_A)$. Furthermore, for any agent such that $(i,j) \in \sigma_i$ we have that $(i,j) \in o(\sigma_{-x}, \sigma_x \setminus \{(i,j)\})$.

As with positive monotonicity, we can also define a stronger variant of negative monotonicity.

Definition 4.2.11 (Strong negative monotonicity)**.** A network game $\mathcal{G} = (A, \Sigma_A, o, S)$ satisfies weak positive monotonicity when: for any $x \in A$, if $(i,j) \notin o(\sigma)$ and $(i,j) \in \sigma_x$ then for all $\sigma'_x \in \Sigma_x$ such that $(i,j) \notin \sigma'_x$, $(i,j) \notin o(\sigma_{-x}, \sigma'_x)$.

An example of an outcome function that is not weakly negatively monotone is the example of triadic closure in Example 4.2.4. If we consider the profiles σ_A, σ'_A from the example, i.e. $\sigma_a = \{(a,c)\}$, $\sigma_b = \{(c,a),(c,b)\}$, $\sigma_c = \{(c,a),(c,b)\}$, and $\sigma'_a = \{(a,b)\}, \sigma'_b = \sigma_b$, and $\sigma'_c = \sigma_c$, then $(a,b) \notin o(\sigma'_A)$ and $(a,b) \in \sigma'_a$. However, $(a,b) \notin \sigma_a$ and $(a,b) \in o(\sigma_A)$.

Like anonymous choice sets, we can also regard anonymous outcome functions. Akin to anonymity in social choice theory, we require the outcome function to not use the order of the actions in the action profile to determine the outcome — o is anonymous when it is symmetric w.r.t. the ordering of σ_A. First, we define permutations on the ordering of the action profile.

Definition 4.2.12 (Agent order permutation)**.** For an agent permutation π, let $\pi_R(\sigma_A)$ denote the application of π on the ordering of the strategies: $\pi_R(\sigma_A) = (\sigma_{\pi(a)})_{a \in A}$.

As we noted before, the choices and outcomes themselves are defined over the same agents as those that play the game, and therefore also need to be taken into consideration when defining anonymity. This entails that a permutation of the input networks will bring about the same permutation on the output network. The full definition of pure anonymity of outcome therefore is the following.

Definition 4.2.13 (Pure anonymity of outcome)**.** An outcome function o is anonymous for a purely anonymous Σ_A if for any $\sigma_A \in \Sigma_A$ and any agent permutation π:

1. $o(\pi_R(\sigma_A)) = o(\sigma_A)$
2. $o(\pi(\sigma_A)) = \pi(o(\sigma_A))$

A network game satisfies pure anonymity of outcome when its outcome function does.

Condition (1) states that the outcome function must not care for the name of the agent that submitted each choice in the action profile. Condition (2) states that the outcome function must not care for the name of the agents in each submitted network choice.

Again, this definition of anonymity represents a "literal" translation of the concept of anonymity of social choice to network formation games. But this translation does not always align with the intuitive expectations of anonymity in a network formation games. For example, the game in Example 4.2.4 satisfies neither condition (1) nor (2) of pure anonymity, even though all agents are treated equally in that they can all determine who they are connected to modulo triadic closure (which again, treats all agents equally). Therefore, like pseudo-anonymity of choice, we introduce *pseudo-anonymity of outcome*.

Pseudo-anonymity of outcome dictates that, given an outcome $F = o(\sigma_A)$, if the order of the action profiles of agents Σ_A is changed according to a permutation π, and for each agent the names of the agents in the submitted action profiles is changed respectively, then the outcome must be equal to the outcome network whose names are respectively changed as well: $\pi(F)$.[6]

[6] A similar notion of anonymity is discussed in Jackson [37], here in the context of allocation rules and value functions. Additionally, in Galeotti et al. [31], a slighly diferent adaptation of anonymity to the context of networks is discussed. In Jackson [38] such games are further developed, here called *semi-anonymous* graphical games. Semi-anonymity in the context of graphical games implies that agents only care about their neighbours, but in an anonymous way — they only care about how many of their neighbours take certain actions.

Definition 4.2.14 (Pseudo-anonymity of outcome)**.** An outcome function o, defined over a pseudo-anonymous Σ_A, is *pseudo-anonymous* iff for any σ_A and any agent permutation π:

$$o(\pi_R(\pi(\sigma_A)) = \pi(o(\sigma_A)).$$

A network game is pseudo-anonymous if Σ_A is pseudo-anonymous and o is pseudo-anonymous for Σ_A.

Note that the set of outcomes (the image of the outcome function) of a pseudo-anonymous game is closed under permutations. That is, $F = o(\sigma_A)$ iff for some $\sigma'_A \in \Sigma_A$: $\pi(F) = o(\sigma'_A)$. The definition of pseudo-anonymity specifies a way to construct this σ'_A from σ_A (via $\pi_R \circ \pi$).

Let the example games given above sketch some differences between pseudo-anonymity and pure anonymity. The interval game in Example 4.2.3 is purely anonymous and pseudo-anonymous: $\Sigma_x = \mathbf{F}_A$ for all x, o is symmetric, and any application of a permutation to the action profile is reflected in the outcome. The bidirectional game in Example 4.2.4 is pseudo-anonymous and not purely anonymous: reordering the action profile changes the outcome as the action profiles for agents a and b determine whether $(a, b) \in o(\sigma_A)$. It is pseudo-anonymous as if we change the names of the agents in the action profile in accordance with the reordering, then the outcome changes accordingly: $(i, j) \in \sigma_i \Leftrightarrow (\pi(i), \pi(j)) \in \pi(\sigma_i)$ and hence $(i, j) \in b(\sigma_A)$ iff $(\pi(i), \pi(j)) \in \pi(\sigma_{\pi(i)})$.

The following is an example of a game that is purely anonymous and not pseudo-anonymous.

Example 4.2.5 (Purely anonymous and not pseudo-anonymous game)**.** An example of such a rule is a game where the only two outcomes are the empty network or $\{(i, j)\}$, and where the outcome is $\{(i, j)\}$ only when an agent submitted $\{(i, j)\}$. Let $\Sigma_a = \mathbf{F}_A$ for any $a \in A$, and fix an $i, j \in A$. Then:

$$o(\sigma) = \begin{cases} \{i, j\} & \exists a \in A : \sigma_a = \{(i, j)\} \\ \emptyset & \text{otherwise} \end{cases}$$

This game is anonymous as the outcome of the game only depends on whether *any* agent submitted $\{(i, j)\}$. However, this game is not pseudo-anonymous, as the outcome specifies a *specific* network with a link from i to j instead of a structure containing a single link that depends on the submitted action profile. Let σ be such that an agent submitted $\{(i, j)\}$. If this game were to be pseudo-anonymous, then $\{(i, j)\} = o(\pi_R(\pi(\sigma))) = \pi(o(\sigma)) = \pi(\{(i, j)\})$. This clearly does not hold for every permutation π: it does not for example hold for a permutation associating i with k for $k \neq i$.

Finally, there are of course games that are neither purely nor pseudo anonymous:

Example 4.2.6 (Non-pseudo-anonymous and not purely anonymous game)**.** An example of a non-pseudo-anonymous and not purely anonymous game is a dictatorial game. Let $\Sigma_a = \mathbf{F}_A$ for any $a \in A$. Fix some dictator $k \in A$. For any $\sigma \in \Sigma_A$ let $o(\sigma) = \sigma_k$.

Clearly this rule is not purely anonymous. Take a σ where for some $i \in A$ $\sigma_i \neq \sigma_k$. Then for a π that associates k with i: $\sigma_i = o(\pi_R(\sigma)) \neq o(\sigma) = \sigma_k$. This rule is also not pseudo-anonymous by similar reasoning as above, except taking a σ for which $\pi(\sigma_i) \neq \pi(\sigma_k)$ and $\sigma_i \neq \sigma_k$: $\pi(\sigma_i)) = o(\pi_R(\pi(\sigma))) \neq \pi(o(\sigma)) = \pi(\sigma_k)$.

For a summary of the definitions introduced in this subsection, see Table 4.3.

4.2.4 Variability and Fairness

We conclude with an obvious and unequivocal requirement for any (noteworthy) network game, that its outcome is mutable at all. We call such games *variable*. Formally, a game is variable when there are two action profiles that each have a different outcome.

Name	Description
Pure anonymity of outcome	Outcome not determined by agent name.
Pseudo-anonymity of outcome	Outcome determined by name only in reference to self.
Weak positive monotonicity	A pure positive variation towards edge existence will not bring the edge out of existence.
Strong positive monotonicity	Any positive variation towards edge existence will not bring the edge out of existence.
Weak negative monotonicity	A pure negative variation towards edge non-existence will not bring the edge into existence.
Strong negative monotonicity	Any negative variation towards edge non-existence will not bring the edge into existence.

Table 4.3: Summary of game axiom definitions

Definition 4.2.15 (Variability)**.** A game is variable iff there exists $\sigma, \sigma' \in \Sigma_A$ such that $o(\sigma) \neq o(\sigma')$. We call a game *constant* when it is not variable.

Perhaps a more disputable requirement for a network game is that no edge in the network is constant. In other words, for any pair of agents i, j there is an action profile σ_A such that i follows j in $o(\sigma_A)$, and a σ'_A such that i does not follow j in $o(\sigma'_A)$. We call such games *fair.*

Definition 4.2.16 (Fairness)**.** A network game (A, Σ_A, o) is *fair* when for every $i, j \in A$ there are $\sigma_A, \sigma'_A \in \Sigma_A$ such that $(i, j) \in o(\sigma_A)$ and $(i, j) \notin o(\sigma'_A)$.

While fairness is a stronger requirement, for pseudo-anonymous games it is equivalent to variability.

Proposition 4.2.1. *For pseudo anonymous games, a game G is fair iff it is variable.*

Proof. Let G be a game that is pseudo-anonymous. Let Σ_A be the choice set of this game and o the outcome function. We will show that G is variable iff it is fair. The right-to-left direction is trivial. For the left-to-right direction, assume G is variable. We will lay out a method to construct $\sigma^\dagger$ and $\sigma^\ddagger$ for any two $a, b \in A$ such that $(a, b) \in o(\sigma^\dagger)$ and $(a, b) \notin o(\sigma^\ddagger)$.

Since G is variable there are $\sigma, \sigma' \in \Sigma_A$ such that $o(\sigma) \neq o(\sigma')$. Hence, there must be some pair (i, j) that is connected in one of σ, σ' and not connected in the other. Without loss of generality, assume that $(i, j) \in o(\sigma)$ and $(i, j) \notin o(\sigma')$. Define a permutation π that associated i with a and j with b where $i \neq a$ and $j \neq b$, and maps all other agents to themselves. Then $(a, b) \in \pi(o(\sigma))$ and $(a, b) \notin \pi(o(\sigma'))$. Hence, by pseudo-anonymity: $(a, b) \in o(\pi_R(\pi(\sigma)))$ and $(a, b) \notin o(\pi_R(\pi(\sigma')))$. Now let $\sigma^\dagger = \pi_R(\pi(\sigma))$ and $\sigma^\ddagger = \pi_R(\pi(\sigma'))$. ☐

Moreover, any *any* game with at least two players is constant iff its outcome function is such that all action profiles map to either the empty network or the fully connected network. This holds in particular for pseudo-anonymous games, as their outcomes are closed under permutation.

Proposition 4.2.2. *A game with more than two players where the image of its outcome function is closed under permutations is constant iff for all $\sigma \in \Sigma_A : o(\sigma) = \emptyset$ or for all $\sigma \in \Sigma_A$: $o(\sigma) = \{(i, j) \mid i, j \in A\}$.*

Proof. The right-to-left direction is trivial. For the left-to-right direction, we will show the case where the image of o is $\{\emptyset\}$, the proof goes similarly when the image of o is $\{\{(i, j) \mid i, j \in A\}\}$. Let $\mathcal{G} = (A, \Sigma_A, o,)$ be a constant network game whose image is closed

Name	Description
Variability	The game must have at least two distinct outcomes.
Fairness	It is possible for every edge to be and not to be in the outcome.

Table 4.4: Summary of game axiom definitions

under permutation. Assume towards a contradiction that for all $\sigma \in \Sigma_A$, $o(\sigma) \neq \emptyset$ and $\sigma \in \Sigma_A$, $o(\sigma) \neq \emptyset$. Then for all $\sigma \in \Sigma_A$: $o(\sigma) = F$ for some $F \neq \emptyset$. There is at least some $(i, j) \in F$, and $(k, l) \notin F$. Let π be a permutation such that i becomes k and j becomes l, and visa-versa. Then, since the image of o is closed under permutations, there must be a σ' such that $o(\sigma') = \pi(F)$. $F \neq \pi(F)$, hence G is not constant. Contradiction, therefore or all $\sigma \in \Sigma_A$ $o(\sigma) = \emptyset$. □

To summarise, the definitions and their intuitive description of this subsection can be found in Table 4.4.

4.3 Concrete Single Shot Network Formation Games

In this section we will discuss some concrete single-shot network formation games, driven by the axioms defined in the last section. Again, this consists of two parts: determining an actual choice set and an actual outcome function. In the previous section we mainly discussed closure properties of Σ_A. We have yet to determine the actual contents of Σ_A. A very simple way to define Σ_A such that it satisfies any closure property is to allow agents to choose any network. We have assumed such a maximal Σ_A in all the example in the previous section. Such games with *universal domain*, are games such that for all $a \in A$:

$$\Sigma_a = \mathbf{F}_A.$$

Network games with universal domain have many applications. For example, in an intercity public transport network, the existence of a connection from a certain city to another can have influence on the flow of passengers of the entire network. It is therefore reasonable to allow every agent (a city's public transport) to have influence on the global network structure. However, the main characteristic of *social* network formation is that agents form social ties between each-other (somewhat) autonomously. The most extreme of such a setting is one where agents that make up a social relation have full control over whether that relation exists or not. This perspective is taken by the vast majority of game-theoretic studies on social network formation [51, p. 296]. For an overview of such studies see Jackson and Zenou [39], and Bloch and Dutta [21].[7] Goyal [32] poses two reasons for this: complexity and tractability. They state:

> One of the principal attractions of networks is that they are amenable to subtle and quick transformations via local, and small-scale, linking and de-linking activity. In line with this observation, it has been felt that invoking the incentives, coordination, and agreements among large groups is not the "right" way to approach a positive theory of network formation. Another important reason is tractability; networks are complicated objects and even with single- or two-person moves the analysis of network formation is quite intricate and general characterisation results have been difficult to obtain. ([32, p. 144])

As such, we will mainly develop a local approach to network formation.

[7] Other major studies in this field that are particularly interesting include Jackson [38, ch. 11], Goyal [32], and Ray [51].

4.3.1 Games With a Local Domain

In the simplest and most extreme settings of local influence, the control of adding or removing a relation between two agents a and b in a social network lays entirely with a and b themselves — agents only have a say in whom they follow and who follows them. Formally, this is reflected in both the outcome function and the choice set. With respect to the outcome function, locality entails that the outcome function should only take into account the proposed networks of the agents a and b when deciding on the existence of the edge (a, b) in the outcome network. We call games that abide by this rule *locally decided*. With respect to the choice set, locality requires that agents are allowed to only propose who they follow or who will follow them: $\Sigma_a \subseteq \mathscr{P}\{(i,j) \in A \times A \mid i = a \text{ or } j = a\}$. We call such Σ_a local to a. The most obvious pseudo-anonymous local choice set is the maximal one, where each agent can propose any network that is local to them. We call locally decided games defined over such a maximally local choice: *local games*.

Definition 4.3.1 (Local games)**.** A game is local when:

1. for every $a \in A$, $\Sigma_a = \mathscr{P}\{(i,j) \in A \times A \mid i = a \text{ or } j = a\}$ (local choice), and
2. for $\sigma, \sigma' \in \Sigma_A$ such that $\sigma_i = \sigma'_i$ and $\sigma_j = \sigma'_j$: $(i,j) \in o(\sigma)$ iff $(i,j) \in o(\sigma')$ (local decidedness).

Locally decided games do not fully determine when exactly an edge exists in the outcome of an action profile, only that the outcome function should decide this based only on the action profiles of the agents that the edge in question connects.

In addition to locality, common ground between all studies on social network formation is that the outcome of a game should follow the decision of the players when they both agree on the existence of an edge between them. We call such games *strict majority following*.[8]

Definition 4.3.2 (Strict majority following)**.** A local network game $\mathcal{G}$ satisfies strict majority following iff for any $i, j \in A$ and any $\sigma \in \Sigma_A$ the following holds:

(i) if $(i,j) \in \sigma_i$ and $(i,j) \in \sigma_j$ then $(i,j) \in o(\sigma)$, and

(ii) if $(i,j) \notin \sigma_i$ and $(i,j) \notin \sigma_j$ then $(i,j) \notin o(\sigma)$.

This implies that the agents are able to (not) form any edge between them and another agent when the other agent tries to (not) form this edge as well. We can characterise strict majority following games by the earlier introduced axioms: under a local setting, the restrictions of weak choice and that of weak monotonicity and fairness are both equivalent to the restrictions of strict majority following

Proposition 4.3.1. *For a local network game $\mathcal{G} = (A, \Sigma_A, o)$ the following are equivalent:*

1. *$\mathcal{G}$ satisfies strict majority following;*
2. *$\mathcal{G}$ satisfies weak positive and negative monotonicity, and fairness;*
3. *$\mathcal{G}$ satisfies weak positive and negative choice;*
4. *$\mathcal{G}$ satisfies weak positive and negative monotonicity, variability, and is pseudo-anonymous.*

[8]Strict majority following has a connection to the *Pareto Principle* of social choice. We can understand the local network formation game as an aggregate of elections on the existence of an edge, where the candidates are "existence" and "non-existence". Each agent can vote only on the election of edges they are a part of. The Pareto Principle dictates that if both i and j (strictly) prefer (i,j) (not) in the outcome network, then (i,j) must (not) be in the outcome network.

Proof. We will prove that for local games, the second and third statement are equivalent to the first. Thereby, we also prove that they are all equivalent to the fourth, as by Proposition 4.2.1, variability and fairness are equivalent under pseudo-anonymous games.

(2. $\Rightarrow$ 1.) We will show the contrapositive. Assume that G does not satisfy strict majority following. Distinguish two cases:

(1) For some $i, j \in A$ there is a $\sigma \in \Sigma_A$ such that: $(i,j) \in \sigma_i$, $(i,j) \in \sigma_j$, and $(i,j) \notin o(\sigma)$. We will show that in this case weak negative monotonicity and local decidedness imply non-fairness. Let σ' be the same as σ except that $(i,j) \notin \sigma'_i$, let σ'' be the same as σ except that $(i,j) \notin \sigma''_j$, and let σ''' be the same as σ expect that $(i,j) \notin \sigma'''_i$ and $(i,j) \notin \sigma'''_j$. By local choice $\sigma', \sigma'', \sigma''' \in \Sigma_A$. By weak negative monotonicity: $(i,j) \notin o(\sigma')$, $(i,j) \notin o(\sigma'')$, and $(i,j) \notin o(\sigma''')$. Finally, by local decidedness, any $\sigma^\dagger$ not equal to σ, σ', σ'', and σ''' must have the same outcome as one of σ, σ', σ'', or σ'''. Hence, $(i,j) \notin o(\sigma)$ for any $\sigma \in \Sigma_A$.

(2) For some $i, j \in A$ there is a σ such that: $(i,j) \notin \sigma_i$, $(i,j) \notin \sigma_j$, and $(i,j) \in o(\sigma)$. By a similar argument as (1.), weak positive monotonicity and local decidedness imply non-fairness.

(1. $\Rightarrow$ 2.) We will show the contrapositive: if G does not satisfy one of weak positive monotonicity, weak negative monotonicity, or fairness, then it is not strict majority following.

1. If G does not satisfy weak positive monotonicity, then there are $x, i, j \in A$ and $\sigma \in \Sigma_A$ such that: (a) $(i,j) \in o(\sigma)$, (b) $(i,j) \notin \sigma_x$, (c) $\sigma_x \cup \{(i,j)\} \in \Sigma_x$, and (d) $(i,j) \notin o(\sigma_{-x}, \sigma_x \cup \{(i,j)\})$. (c) together with local choice implies that $x = i$ or $x = j$. Without loss of generality, assume that $x = i$. There are two possibilities: either $(i,j) \notin \sigma_j$, or $(i,j) \in \sigma_j$. As by (b) $(i,j) \notin \sigma_i$ and by (a) $(i,j) \in o(\sigma)$, if $(i,j) \notin \sigma_j$ then G does not satisfy strict majority following. If $(i,j) \in \sigma_j$, then for σ' such that $\sigma'_{-i} = \sigma_{-i}$ and $\sigma'_i = \sigma_i \cup \{(i,j)\}$: $(i,j) \in \sigma'_i$ and $(i,j) \in \sigma'_j$. By (c) $\sigma' \in \Sigma_A$, and by (d) $(i,j) \notin o(\sigma')$. Therefore, G does not satisfy strict majority following.
2. If G does not satisfy weak negative monotonicity, then we can repeat the argument of (1.), this time assuming that $(i,j) \notin o(\sigma)$, $(i,j) \in \sigma_x$, $\sigma_x \setminus \{(i,j)\} \in \Sigma_x$, and $(i,j) \in o(\sigma_{-x}, \sigma_x \setminus \{(i,j)\})$.
3. If G does not satisfy fairness, then there are $i, j \in A$ such that either for all $\sigma \in \Sigma_A$: $(i,j) \in o(\sigma)$, or for all $\sigma \in \Sigma_A$: $(i,j) \notin o(\sigma)$. By local choice there is a $\sigma^+ \in \Sigma_A$ and $\sigma^- \in \Sigma_A$ such that $(i,j) \in \sigma^+_i$, $(i,j) \in \sigma^+_j$, $(i,j) \notin \sigma^-_i$, and $(i,j) \notin \sigma^-_j$. It must either hold that $(i,j) \notin o(\sigma^+)$ or $(i,j) \in o(\sigma^-)$. Therefore, G does not satisfy strict majority following.

(3. $\Rightarrow$ 1.) Assume towards a contradiction that there is local game that satisfies weak positive and negative choice that is not strict majority following. Distinguish two cases:

(i) For some $i, j \in A$, there is a $\sigma \in \Sigma_A$ such that $(i,j) \in \sigma_i$, $(i,j) \in \sigma_j$, and $(i,j) \notin o(\sigma)$. By weak negative choice, there is an x such that $(i,j) \notin \sigma_x$ and a $\sigma'_x \in \Sigma_x$ such that $(i,j) \in \sigma'_x$. By local choice, this x can only be i or j; hence $(i,j) \notin \sigma_i$ or $(i,j) \notin \sigma_j$. Contradiction, therefore the game does satisfy strict majority following.

(ii) For some $i, j \in A$, there is a $\sigma \in \Sigma_A$ such that $(i,j) \notin \sigma_i$, $(i,j) \notin \sigma_j$, and $(i,j) \in o(\sigma)$. By weak positive choice, there is an x such that $(i,j) \in \sigma_x$ and a $\sigma'_x \in \Sigma_x$ such that $(i,j) \notin \sigma'_x$. By local choice, this x can only be i or j; hence $(i,j) \in \sigma_i$ or $(i,j) \in \sigma_j$. Contradiction, therefore the game does satisfy strict majority following.

(1. $\Rightarrow$ 3.) We will show the contrapositive. That is, if G does not satisfy either weak positive or negative choice, then it is not strict majority following.

1. If G does not satisfy weak negative choice, then there are $i, j \in A$ and $\sigma \in \Sigma_A$ such that $(i, j) \notin o(\sigma)$, and for every $x \in A$ either (a) $(i, j) \in \sigma_x$, or (b) for all $\sigma'_x \in \Sigma_x$, $(i, j) \notin \sigma'_x$. By local choice (b), cannot hold for $x = i$ or $x = j$. Hence, in these cases (a) must hold. Hence, in particular: $(i, j) \in \sigma_i$ and $(i, j) \in \sigma_j$. Recall that $(i, j) \notin o(\sigma)$. Therefore, G does not satisfy strict majority following.

2. If G does not satisfy weak positive choice, then a similar argument to (1.) can be made. This time assuming that there are $i, j \in A$ and $\sigma \in \Sigma_A$ such that $(i, j) \in o(\sigma)$, and for every $x \in A$ either $(i, j) \in \sigma_x$, or for all $\sigma'_x \in \Sigma_x$, $(i, j) \notin \sigma'_x$.

□

The conditions of strict majority following still leave open what happens when the sending and receiving party conflict in choice on whether there should be a link between them. It is exactly here that the studies on social network formation diverge. We will discuss some simple interpretations of majority following games and the contexts in which they are reasonable.

In a context where agents can choose to listen to any communication that takes place in the entire network, e.g. public communication or situations where agents are able to maliciously listen in on communications of others, the receiving party decides whether they listen to (or read) the information of the sending party. Examples of such situations are: publications of papers, digital social network feeds, publications on blogs and websites, or eavesdropping, among others. These cases call for a strict majority following game where the receiving agent determines the existence of an edge from the sender to the receiver. Formally, define an outcome function o_r such that:

$$(i, j) \in o_r(\sigma) \text{ iff } (i, j) \in \sigma_i.$$

Such *receiver-decided games* are often the implicit assumption in studies on directed social network formation [21]. For example, in the one-way flow setting in Bala and Goyal [10] agents can choose agents to form edges to, obtaining their benefit (e.g. information, opportunities in trade-networks etc.). Similar unilateral games are discussed in Goyal [32] and Jackson [38, ch. 11.3].

In studies on bidirectional social relations it is the norm to require consent from both parties involved in order to form a link [21]. For local games, this is exactly equal to the game described in Myerson [42]. Similar games are also discussed in the bidirectional/bilateral settings in Jackson [37], Jackson [38], and Goyal [32]. In such games, agents need permission in order to create relations with other agents. Established social relations can be broken by both parties individually. We call such games *consensual following*:[9] let $o_{s\wedge r}$ be the outcome function such that:

$$(i, j) \in o_{s\wedge r}(\sigma) \text{ iff } (i, j) \in \sigma_i \text{ and } (i, j) \in \sigma_j.$$

Digital social network platforms, for example, implement such a rule by a request-and-accept dynamic. Aumann and Myerson [8] formalises such a dynamic explicitly as an extensive game; $o_{s\wedge r}$ is its single-shot counterpart.

The outcome functions o_r and $o_{r\wedge s}$ are the most common assumption in game-theoretic studies on social networks. They have natural duals. o_s, the dual of o_r, is an outcome function where agents can decide who they send information to. This is true for forms of obtrusive communication, for example for direct messages on digital social networks, or sending letters or emails. Formally, define the outcome function o_s, such that:

$$(i, j) \in o_s(\sigma) \text{ iff } (i, j) \in \sigma_j.$$

[9]Related to what is called "obtainable via deviations" in Jackson [37], (p.28): g' is obtainable from g' by C iff for any added link both parties are member of C and for any severed link one of them is in C.

$o_{s\vee r}$, the dual of $o_{s\wedge r}$, is an outcome function where both the sending and receiving agent can decide to form an edge between each other. But only the sender and receiver together can decide to not form this edge. Formally, define $o_{s\vee r}$, such that:

$$(i, j) \in o_{s\vee r}(\sigma) \text{ iff } (i, j) \in \sigma_i \text{ or } (i, j) \in \sigma_j.$$

These duals are not discussed by economic studies in a truly directed setting. The reason for this most likely is that the underlying directed networks in economic studies represent flow of benefit. Therefore, there is no real incentive for the party that "sends" benefit to initiate a directed relation. Only in a setting with so called *two-way flow of benefit*, where benefit is bidirectional (and cost is unilateral), do such games occur — that is, when $(i, j) \in o(\sigma)$ implies that i bears the cost of maintaining the edge, while both i and j obtain benefit from it. For example in the two-way flow setting in Bala and Goyal [10] and Haller and Sarangi [34] (a generalisation of the former).

In an epistemic setting of communication, where the underlying network is a relation of epistemic following, benefit is more ambiguous. It doesn't always flow in the same direction as to the underlying social relation. Yet, a certain aspect of this flow is two-way: being able to *send* and being able to *receive* information are both beneficial, albeit in a qualitatively different way. This is distinct from the two-way flow in economic studies, where it implies *equal* flow in both direction, both qualitatively and quantitatively, Therefore, games with outcome function o_s or $o_{s\vee r}$ are of interest in an epistemic setting.

4.4 Social Influence and Extensive Network Formation

The local games of the last section depict network formation in its simplest and most autonomous form. Because of this, they lack many of the intricacies involved in social network formation. In reality, as many studies have shown , social relations are influenced by the agent's social environment, social norms, and other underlying processes of *social influence*. In this section we will extend strategic games and network games to incorporate social influence in its most general form: we let the outcome of a network game be dependent on the current social network.

Moreover, in a social setting, networks are not formed at a certain moment, but are continually changing. We will see that such continual games immediately follow from the structure of the network games with social influence.

4.4.1 Game State

Formalizing any process of social influence already assumes that the agents find themselves in a social environment or network. We can already implicitly take such a social environment into account in single-shot games: assume a social relation F over the agents, and (implicitly) define the outcome function and choice set over this F. A problem with such an implicit approach is that, then, the axioms cannot take this network into account. This issue is most prominent with the axiom of pseudo-anonymity. Pseudo-anonymity forbid any form of social influence that is not global, i.e. that is restricted by the current social network: an outcome function that treats an agent a different from b because of their placement in F is not pseudo-anonymous. The first step in incorporating social influence, therefore, is to make the social setting explicit. We will partake this explication first from the abstract. Later, we will make this concrete as network games.

Abstractly, taking into account the social network in a network game entails that there is a game state that determines some part of the game's logic. Explicating the social network, then, entails that we explicate the state over which the game is played. Denote this state by s_0. For social influence on network games, the game state s_0 itself is exactly what is modified by the players when playing the game — the outcome of a game is again a game state, s_0 is an element of S. The outcome function o, then, takes the shape of a *state transition function* that dictates how the agents can modify the state s_0.

With the explication of the current game state, we can also explicate how game state determines game logic. Instead of specifying a state and the state transition from this state, we specify a transition function o for any state in S. We associate a transition game with each state of S. We call such a structure a *game frame*.

Definition 4.4.1 (Game frame)**.** Let Γ_S^A denote the set of all possible strategic games played by A with (outcome) states S. A game frame is a tuple $GF := (\gamma, \Sigma, S)$, where S is a non-empty set of possible game states and $\gamma : S \to \Gamma_S^A$ is a function that associates any game state with a strategic game.

To easily reference the logic of a game on state S, we use the following notation: let $\mathbf{\Sigma}_A^\gamma(s)$ be the choice set of the game in state s, and let $\mathbf{o}^\gamma(s)$ be the outcome function of the game in state s.

4.4.2 Extensive Form Games With Simultaneous Play

A game frame becomes a game (form) when combined with a current state s_0. Out of these "pointed" game frames, *extensive games (game forms)* immediately follow: given a starting state s_0, for $n \in \mathbb{N}$ let $\sigma_n \in \mathbf{\Sigma}_A^\gamma(s_n)$ and

$$s_{n+1} = \mathbf{o}^\gamma(s_n)(\sigma_n).$$

When γ is a partial function, this structure is a state automaton, with alphabet $\Sigma = \bigcup_{s \in S} \mathbf{\Sigma_A}^\gamma(s)$, states S, initial state s_0, state transition function $\mathbf{o}^\gamma$, and terminating states consisting of the states not in the domain of γ. Taking into account that Σ_A are choices of agents in a strategic game, such a structure is in essence what in Osborne et al. [44, Section 7.1] is called an *extensive form game with simultaneous play.*[47][10] Note that we also get finite automata and extensive form games with total functions γ by assigning a dummy game to terminating states, as noted in Pauly [47].

Normally, instead of actions, play in extensive games (with perfect information) is analysed through a player's *strategy*. Such strategy is a function that assigns a choice for each *history*: a possible sequence of profiles and their states. We will not develop these strategies here. Instead, we will present the outlines of a logic of ability (effectivity) for these games in the next section.

4.4.3 Extensive Games of Network Modification

To make the game forms concrete as network games, we must associate a social network with each game state. Let (γ, S) be a game frame such that for all $s \in S$, $\mathbf{\Sigma_A}^\gamma(s) \subseteq \mathbf{F}_A$. Let $\nu : S \to \mathbf{F}_A$ be a function that assigns a social network to each game state. A *network game frame* is a tuple:

$$NF := (\gamma, S, \nu).$$

If we care only about the networks associated with state and not the states themselves, then we care only about the *network games* associated with each state by γ: the structure $(\nu \circ \mathbf{o}^\gamma, \mathbf{\Sigma}_A^\gamma, \mathbf{F}_A)$. Let this structure be denoted by $\gamma^\nu(s)$. ν is an isomorphism between

[10] These games are different from repeating games: in repeating games, a player receives a payoff each round of the game, whereas in extensive form games, a payoff is only given at a terminal node. Furthermore, repeating games repeat the same game multiple times, whereas in an extensive form game, the actions and results change each round [44, ch. 14]. Hence, game frames correspond to repeating games when Σ_A and o are constant over the domain of games Γ_S^A.

Extensive form games with simultaneous play are a generalisation of the standard definition of *extensive form games*, as these extensive form games without simultaneous play can be modeled as simultaneous ones by a game with subsequent rounds where, at each round, a particular player has a choice set containing more than one choice and the other players all have a choice set containing only one choice.[47]

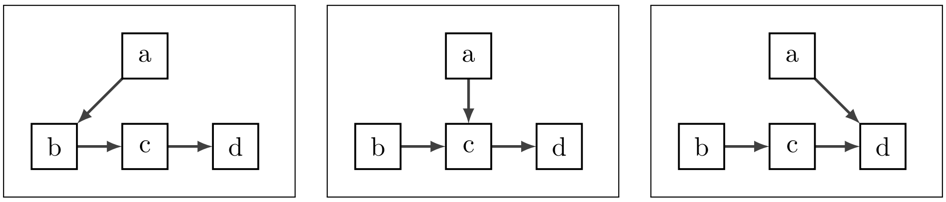

(a) F_0 (b) $o(F_0)(\sigma_A)$ (c) $o(o(F_0)(\sigma_A))(\sigma'_A)$

Figure 4.1

this structure and $(\mathbf{o}^\gamma, \mathbf{\Sigma}^\gamma_A, S)$ iff it is a bijection from S to $\mathbf{F}_A$. We can then simplify the structure of network games as follows:

Definition 4.4.2 (Bijective network game frame)**.** Let N_A be the set of all network games with players A. A network game frame is a game frame where the set of states $S = \mathbf{F}_A$ and $\gamma : S \to N^A$.

Out of pointed network game frames, network games (game forms) follow. Subsequent plays of such network games (game forms) bring about extensive network games. The use of such subsequent plays is illustrated by the following example.

Example 4.4.1. Take a game generated from any of the local rules in the previous section, say o_r. Add the middleman requirement: i can propose a network with (i,j) in it when i follows someone who follows j. Likewise, j can propose a network with (i,j) in it when j follows someone who follows i. Formally, for any $a \in A$, $\mathbf{\Sigma}^\gamma_a(F) = \{(a,i) \mid \exists x \in A \text{ s.t. } a\ F\ x\ F\ i\} \cup \{(i,a) \mid \exists x \in A \text{ s.t. } a\ F\ x\ F\ i\}$.

Let the initial network F_0 be as depicted in Figure 4.1a. Assume that a wants a connection to d. By the middleman requirement this is not directly possible. However, a does have a strategy to make the connection (a,d) come about: first connect with c, for which b is a middleman, and then connect with d through the middleman c. Let σ_A be such that $\sigma_a = \{(a,c)\}$, $\sigma_b = (b,c)$, $\sigma_c = (c,d)$, and $\sigma_d = \emptyset$. Then $F_1 = o(F_0)(\sigma_A)$ is as depicted in Figure 4.1b. Then, let $\sigma'_a = \{(a,d)\}$, and let $\sigma'_x = \sigma_x$ for $x \in A \setminus \{a\}$. The outcome of $o(F_1)(\sigma_A)$ is depicted in Figure 4.1c.

4.4.4 Axioms

Besides uncovering an extensive form of network formation, network game frames allow us to specify properties of the relation between game logic (i.e. outcome function and choice set) game state, and network state. In particular, we can extend the axioms of single-shot network games, introduced in the previous section, to also take into account these relations between game state (the social network) and game logic (the games played on this social network).

The choice axioms dictate that edges are formed by the choice of at least some agent. In the extensive case of network modification, this implies that an edge changes state, i.e. flips from existing to non-existing or visa versa, only when at least some agent made the choice to do so. For positive choice this entails that if (i,j) is in the outcome, then an agent must have chosen (i,j) to be in the outcome, but only when (i,j) was *not* already in the current network. As before, the distinction between weak and strong positive choice come about by different interpretations of whether an agent had a choice in including or excluding an edge.

Definition 4.4.3 (Weak positive choice)**.** $\forall s \in S$, $i,j \in A$, $\sigma_A \in \mathbf{\Sigma}^\gamma_A(s)$: $(i,j) \in \mathbf{o}^\gamma(s)(\sigma_A)$ and $(i,j) \notin \nu(s) \Longrightarrow \exists x \in A$ s.t. $(i,j) \in \sigma_X$ and $\exists \sigma'_x \in \mathbf{\Sigma}^\gamma_x(s)$ s.t. $(i,j) \notin \sigma_x$.

Definition 4.4.4 (Strong positive choice)**.** $\forall s \in S$, $i,j \in A$, $\sigma_A \in \mathbf{\Sigma}^\gamma_A(s)$: $(i,j) \in \mathbf{o}^\gamma(s)(\sigma_A)$ and $(i,j) \notin \nu(s) \Longrightarrow \exists x \in A$ s.t. $(i,j) \in \sigma_X$ and $\sigma_x \setminus \{(i,j)\} \in \mathbf{\Sigma}^\gamma_x(s)$.

Similarly for negative choice, if (i,j) is *not* in the outcome, then an agent must have chosen (i,j) to *not* be in the outcome only when (i,j) was already in the current network.

Definition 4.4.5 (Weak negative choice)**.** $\forall s \in S$, $i,j \in A$, $\sigma_A \in \mathbf{\Sigma}_A^\gamma(s)$: $(i,j) \notin \mathbf{o}^\gamma(s)(\sigma_A)$ and $(i,j) \in \nu(s) \Longrightarrow \exists x \in A$ s.t. $(i,j) \notin \sigma_X$ and $\exists \sigma'_x \in \mathbf{\Sigma}_x^\gamma(s)$ s.t. $(i,j) \in \sigma_x$.

Definition 4.4.6 (Strong negative choice)**.** $\forall s \in S$, $i,j \in A$, $\sigma_A \in \mathbf{\Sigma}_A^\gamma(s)$: $(i,j) \notin \mathbf{o}^\gamma(s)(\sigma_A)$ and $(i,j) \in \nu(s) \Longrightarrow \exists x \in A$ s.t. $(i,j) \notin \sigma_X$ and $\sigma_x \cup \{(i,j)\} \in \mathbf{\Sigma}_x^\gamma(s)$.

For extensive games, monotonicity should still hold for all rounds of the game: all games $\gamma(F)$ for $F \in \mathbf{F}_A$ should be monotone. Monotonicity should also hold relative to the current network F: if agents can bring about (i,j) in a network, then submitting the same profile on a network with (i,j) in it must also bring about a network with (i,j) in it. That is, if an edge is in the outcome of a game on network F, where the agents have submitted a σ_A, then if the agents can submit σ_A in the game on network $F \cup \{(i,j)\}$ as well, (i,j) must also be in the outcome.

Definition 4.4.7 (Weak positive monotonicity)**.** $\forall s \in S$:

1. $\nu \circ \mathbf{o}^\gamma(s)$ is weakly positively monotone for $\mathbf{\Sigma}_A^\gamma(s)$, and
2. $\forall \sigma_A \in \mathbf{\Sigma}_A^\gamma(s)$, if $(i,j) \in \mathbf{o}^\gamma(s)(\sigma_A)$ and $(i,j) \notin \nu(s)$ then: $\forall t \in S$ s.t. $\nu(t) = F \cup \{(i,j)\}$: if $\sigma_A \in \mathbf{\Sigma}_A^\gamma(t)$ then $(i,j) \in \mathbf{o}(t)(\sigma_A)$.

Definition 4.4.8 (Strong positive monotonicity)**.** $\forall s \in S$:

1. $\nu \circ \mathbf{o}^\gamma(s)$ is strongly positively monotone for $\mathbf{\Sigma}_A^\gamma(s)$, and
2. $\forall \sigma_A \in \mathbf{\Sigma}_A^\gamma(s)$, if $(i,j) \in \mathbf{o}^\gamma(s)(\sigma_A)$ and $(i,j) \notin \nu(s)$ then for any $t \in S$ s.t. $(i,j) \in \nu(t)$: if $\sigma_A \in \mathbf{\Sigma}_A^\gamma(t)$ then $(i,j) \in \mathbf{o}^\gamma(t)(\sigma_A)$.

Similarly, for the negative variants of monotonicity.

Definition 4.4.9 (Weak negative monotonicity)**.** $\forall s \in S$:

1. $\nu \circ \mathbf{o}^\gamma(s)$ is weakly negatively monotone for $\mathbf{\Sigma}_A^\gamma(s)$, and
2. $\forall \sigma_A \in \mathbf{\Sigma}_A^\gamma(s)$, if $(i,j) \notin \mathbf{o}^\gamma(s)(\sigma_A)$, and $(i,j) \in \nu(s)$ then: $\forall t \in S$ s.t. $\nu(t) = F \setminus \{(i,j)\}$, if $\sigma_A \in \mathbf{\Sigma}_A^\gamma(t)$ then $(i,j) \notin \mathbf{o}^\gamma(t)(\sigma_A)$.

Definition 4.4.10 (Strong negative monotonicity)**.** $\forall s \in S$:

1. $\nu \circ \mathbf{o}^\gamma(s)$ is strongly negatively monotone for $\mathbf{\Sigma}_A^\gamma(s)$, and
2. $\forall \sigma_A \in \mathbf{\Sigma}_A^\gamma(s)$: if $(i,j) \notin \mathbf{o}^\gamma(s)(\sigma_A)$, and $(i,j) \in \nu(s)$ then for any $t \in S$ s.t. $(i,j) \notin \nu(t)$: if $\sigma_A \in \mathbf{\Sigma}_A^\gamma(t)$ then $(i,j) \notin \mathbf{o}^\gamma(t)(\sigma_A)$.

Recall that pseudo-anonymity dictates that the names of the agents do not matter in the formation of the social network. Hence, for any agent permutation π, if an agent a can choose network σ_a in network F, then $\pi(a)$ must be able to choose $\pi(\sigma_a)$ in network $\pi(F)$.

Definition 4.4.11 (Pseudo-anonymity of choice)**.** $\mathbf{\Sigma}_A^\gamma$ is pseudo-anonymous iff for all permutations π and all $s,t \in S$ s.t. $\nu(s) = \pi(\nu(t))$:

$$\sigma_a \in \mathbf{\Sigma}_A^\gamma(s) \Longleftrightarrow \pi(\sigma_a) \in \mathbf{\Sigma}_{\pi(A)}^\gamma(t)$$

Similarly, pseudo-anonymity of outcome should dictate that if F_{i+1} is the outcome of a network game on a state with network F_i where the players have submitted σ_A, then $\pi(F_{i+1})$ should be the outcome of all games with a network $\pi(F_i)$ where the players have submitted $\pi(\sigma_A)$.

Definition 4.4.12 (Pseudo-anonymity of outcome). $\mathbf{o}^\gamma$ is pseudo-anonymous for a pseudo-anonymous choice set function $\mathbf{\Sigma}^\gamma_A$ iff: for all $s, t \in S$ s.t. $\nu(s) = \pi(\nu(t))$

$$\nu(\mathbf{o}^\gamma(t)(\pi_R(\pi(\sigma_A)))) = \pi(\nu(\mathbf{o}^\gamma(s)(\sigma_A))).$$

Note that if a game frame satisfies pseudo-anonymity of choice and pseudo-anonymity of outcome, then in particular $\nu(s) = \nu(t) \Rightarrow \mathbf{\Sigma}^\gamma_A(s) = \mathbf{\Sigma}^\gamma_A(t)$, and $\nu(s) = \nu(t) \Rightarrow \nu \circ \mathbf{o}^\gamma(s) = \nu \circ \mathbf{o}^\gamma(t)$. I.e. states with equal networks are associated with equal network games.[11]

The variability and fairness axioms can be extended to game frames by requiring the single-shot conditions of all games associated with a state.

Definition 4.4.13 (Variability). For all $s \in S$, $\gamma^\nu(s)$ is variable.

Definition 4.4.14 (Fairness). For all $s \in S$, $\gamma^\nu(s)$ is fair.

4.5 Network Formation Logic

We have yet to touch upon any proper game theory, in that we have not given players any incentive, and have not discussed solution concepts. In game theory, such incentives take the form of orderings over outcome states. Often such an order is made concrete by score. As the goal of the game-theoretic framework presented here is to reason about network formation, and particularly to analyse agent's *ability*, we do not present such incentive or solution concepts.[12] Instead, we will present the outlines of a logic to reason about ability of agents in network formation games. A full study of such a logic is beyond the scope of a thesis.

4.5.1 Coalition Logic

A natural candidate for logics of ability is *Coalition Logic* [47]. Even more so for the network game frames of the last section, as Coalition Logic is the modal logic of coalitional ability, or *effectivity*, in extensive form games with simultaneous play.[13] We will review Coalition Logic here, and adapt some of its syntax and semantics to better match network formation games. Formally, effectivity is represented by a function $E : \mathscr{P}(A) \to \mathscr{P}(\mathscr{P}(S))$ that associates each coalition $C \subseteq A$ with a set of states. For $X \subseteq A$, if $E(C) = X$, then C is effective for X, meaning that C has a strategy to always bring about a state inside X. In such cases, we say that C can force the set of outcomes X.

Pauly [47] presents a connection between a certain class of effectivity functions and strategic games, *playable* effectivity functions. These are effectivity functions E that satisfy *Liveness*, *Safety*, *A-maximality*, *Outcome monotonicity*, and *Superadditivity*. A theorem of effectivity functions is that they are playable iff they are the effectivity function of some strategic game — if there is a strategic game G such that:

$$X \in E_G(C) \text{ iff } \exists \sigma_C \forall \sigma_{\overline{C}}\ o(\sigma_C, \sigma_{\overline{C}}) \in X$$

By this connection between effectivity functions and strategic games, coalition logic can be immediately employed as a logic of effectivity in network game frames.

The language of coalition logic then is generated from the following BNF. To better align with games of network formation, we fix a set of network propositions instead of a set of generic propositions,

$$\varphi ::= F_{a,b} \mid \neg\varphi \mid \varphi \vee \varphi \mid [C]\varphi$$

[11] This does not mean that states associated with unequal networks are always associated with unequal network games.

[12] A way to add incentive to our framework is to assign goals to agents, such as the strategic positions described in the previous chapter. Agents can, for example, rank the set of states on whether they can block information in it between specific sets. Maybe even ranking those worlds higher where they are a more unique such connector or blocking set.

[13] More precisely, α-effectivity [1].

As with Communication Logic, $F_{a,b}$ denotes that a follows b. $[C]\varphi$ indicates that a set of agents C, when operating as a coalition, is effective for φ, meaning that C has a strategy such that φ will be true no matter what other agents do. More precisely, if $\mathfrak{M}, w \Vdash [C]\varphi$ then C is effective for the set $\varphi^{\mathfrak{M}} := \{w \in W \mid \mathfrak{M}, w \Vdash \varphi\}$.

Coalition Logic is evaluated on effective models. If we interchange its evaluation function with the network assignment function ν, we get the structure of the "effective model" of a network game frame.

Definition 4.5.1 (Network Formation Model)**.** A network model $\mathfrak{M}$ is a tuple (S, ν, E), where S is a set of states, ν is the network function that associates each state with a network $\nu : S \to \mathbf{F}_A$, and E is a function that associates each state with a playable effectivity function: $E : S \to \mathscr{P}(A) \to \mathscr{P}(\mathscr{P}(S))$.

The interesting parts of the semantics of coalition logic are as follows, keeping in mind that the "valuation function" of a network formation model is a network function ν, associating states with a social network.

$$\begin{array}{lll} \mathfrak{M}, s \Vdash F_{a,b} & \text{iff} & (a, b) \in \nu(s) \\ \mathfrak{M}, s \Vdash [C]\varphi & \text{iff} & \varphi^{\mathfrak{M}} \in E(s)(C). \end{array}$$

where $\varphi^{\mathfrak{M}} = \{s \in S \mid \mathfrak{M}, s \Vdash \varphi\}$. For an axiomatisation of this logic, see Pauly [47].

4.5.2 Towards Network Formation Logics

Coalition logic is the logic of effectivity of *generic* network formation games. As pointed out before, however, not all network formation games make sense. We attempted to consolidate the class of network formation games to consider, by presenting "axioms" of the process of network formation. A natural next step is to phrase these axioms in Coalition Logic. Such axioms, then, bring about logics of effectivity of, say, pseudo-anonymous network formation games, positively monotone network formation games, and so on. These logics will give more insight into the effects such axioms have on the ability of coalitions in shaping the network in certain ways.

Recall the correspondence between effectivity functions and strategic games presented by Pauly [47]; a playable effectivity function E *of* a game G — denoted by E_G — is such that:

$$X \in E_G(C) \text{ iff } \exists \sigma_C \forall \sigma_{\overline{C}}\ o(\sigma_C, \sigma_{\overline{C}}) \in X$$

Note that the effectivity function E_G is not unique to the particular strategic game G; as the coalitional ability between two different games can be the same, different strategic games can share the same effectivity function. This is not explored in Pauly [47]. But this connection between specific strategic games and their effectivity function is crucial in understanding the relation between the axioms of network game frames and the behaviour of effectivity functions. We will briefly go over this relation here.

Particular effectivity functions identify particular "effective-equivalence" classes of single-shot games. And particular network formation models identify particular "effective-equivalence" classes of extensive games with simultaneous play. Let the function $E^\gamma : S \to \mathscr{P}(A) \to \mathscr{P}(\mathscr{P}(S))$ be such that $E^\gamma(s)$ is the effectivity function of the game $\gamma(s)$: $E^\gamma(s) = E_{\gamma(s)}$.

Definition 4.5.2 (Effective-equivalence)**.** We say that a game frame $G = (S, \nu, \gamma)$ is effective-equivalent to a game frame $G' = (S', \nu', \gamma')$, notation $G \simeq_E G'$, iff $E^\gamma = E^{\gamma'}$, $S = S'$, and $\nu = \nu'$. The effective-equivalence class of a game frame G is the class of game frames that are effective-equivalent to G.

Note that two games $G = (S, \nu, \gamma)$ and $G' = (S', \nu', \gamma')$ can be effective-equivalent, even if $\gamma \neq \gamma'$. The "axioms" presented in the previous chapter are not necessarily closed under such effective-equivalent classes. Thus, it is impossible to identify a class of, say, pseudo-anonymous effectivity functions that correspond to pseudo-anonymous games. Yet, we can still formulate pseudo-anonymous properties of network formation models. If $\pi(\nu(s)) = \nu(t)$, and C is effective "for F"[14] in s, then $\pi(C)$ is effective "for $\pi(F)$" in t.

Definition 4.5.3 (Effectively pseudo-anonymous)**.** A network formation model $\mathfrak{M} = (S, \nu, E)$ satisfies pseudo-anonymity when, for any permutation π, $s, t \in S$ s.t. $\nu(s) = \pi(\nu(s'))$, network $F \in \mathbf{F}_A$, and coalition $C \subseteq A$:

$$\varphi_F^{\mathfrak{M}} \in E(s)(C) \text{ iff } \varphi_{\pi(F)}^{\mathfrak{M}} \in E(t)(\pi(C))$$

The axiom of Definition 4.5.3 identifies the class of pseudo-anonymous Network Formation Models.

Proposition 4.5.1. *If the game frame $G = (S, \gamma, \nu)$ is pseudo-anonymous, then the model $\mathfrak{M} = (S, \nu, E^\gamma)$ is effectively pseudo-anonymity.*

Proof. Assume that G is pseudo-anonymous. Take an arbitrary $F \in \mathbf{F}_A$, permutation π, $s, t \in S$ s.t. $\nu(s) = \pi(\nu(t))$, $C \subseteq A$. To show: $\varphi_F^{\mathfrak{M}} \in E(s)(C)$ iff $\varphi_{\pi(F)}^{\mathfrak{M}} \in E(t)(\pi(C))$. We will show the left-to-right direction, the right-to-left direction follows similarly. $\varphi_F^{\mathfrak{M}} \in E(s)(C)$ implies that $\exists \sigma_C \in \Sigma^\gamma(s) \forall \sigma_{\overline{C}} \in \Sigma^\gamma(s) o^\gamma(s)(\sigma_C, \sigma_{\overline{C}}) \in \varphi_F^{\mathfrak{M}}$, i.e. $\nu(o^\gamma(s)(\sigma_C, \sigma_{\overline{C}})) = F$. Let σ denote the action profile where C chooses σ_C and $\overline{C}$ chooses $\overline{C}$. Let $\sigma' = \pi_R(\pi(\sigma))$. By pseudo-anonymity of G, $\nu(o^\gamma(t)(\sigma')) = \pi(F)$, and $\sigma' \in \Sigma_A^\gamma(t)$. Thus, $\exists \sigma_{\pi(C)} \in \Sigma_{\pi(C)}^\gamma(t) \forall \sigma_{\overline{\pi(C)}} \in \Sigma_{\overline{\pi(C)}}^\gamma o^\gamma(t)(\sigma_{\pi(C)}, \sigma_{\overline{\pi(C)}}) \in \varphi_{\pi(F)}^{\mathfrak{M}}$. Therefore, $\varphi_{\pi(F)}^{\mathfrak{M}} \in E(t)(\pi(C))$. ☐

The right-to-left direction of this proposition does not hold, as the axioms of pseudo-anonymity in the previous chapter are not closed under effective-equivalence classes. This proposition implies that, if the effective-equivalence class of a game $G = (S, \gamma, \nu)$ contains a pseudo-anonymous game, then the model $\mathfrak{M} = (S, \nu, E^\gamma)$ is pseudo-anonymous. Whether such effectivity functions also *identify* such classes, i.e. whether this also holds in the other direction, is unknown.

We will not make any further attempt to phrase the axioms in the previous chapter for network formation models, and we will not give an analysis of effective-equivalent classes, as this is beyond the scope of this thesis. We will conclude by phrasing effective pseudo-anonymity as a property of Network Formation Logic.

Proposition 4.5.2 (`GP`)**.** *Let the* network formula *of a network $F \in \mathbf{F}_A$ be denoted by*

$$\varphi_F := \bigwedge_{a,b \in A} \begin{cases} F_{a,b} & \textit{if } (a,b) \in F \\ \neg F_{a,b} & \textit{otherwise.} \end{cases}$$

A model $\mathfrak{M}$ is effectively pseudo-anonymous iff it satisfies `GP`*: For all $F, F' \in \mathbf{F}_A$, and $C \subseteq A$:*

$$(\texttt{GP}) \ \Vdash \varphi_F \to [C]\varphi_{F'} \ \textit{ iff } \ \Vdash \varphi_{\pi(F)} \to [\pi(C)]\varphi_{\pi(F')}$$

Proof. The left-to-right direction follows from the semantics of $[C]$. We will prove the right-to-left direction by contrapositive. Assume that $\mathfrak{M}$ is not pseudo-anonymous. Then there are $F \in \mathbf{F}_A$, $C \subseteq A$, $s, t \in S$ s.t. $\nu(s) = \pi(\nu(t))$ and either: 1. $\varphi_F^{\mathfrak{M}} \in E(s)(C)$ but $\varphi_{\pi(F)}^{\mathfrak{M}} \notin E(t)(C)$; or 2. $\varphi_F^{\mathfrak{M}} \notin E(s)(C)$ but $\varphi_{\pi(F)}^{\mathfrak{M}} \in E(t)(C)$. (i) implies that $\mathfrak{M}, s \Vdash [C]\varphi_F$ and $\mathfrak{M}, t \nVdash [C]$. (ii) implies that $\mathfrak{M}, s \nVdash [C]\varphi_F$ and $\mathfrak{M}, t \Vdash [C]$. As $\mathfrak{M}, s \Vdash \varphi_{\nu(s)}$ and $\mathfrak{M}, t \Vdash \varphi_{\pi(\nu(s))}$, (i) disproves the left-to-right direction of `GP` and (ii) disproves the right-to-left direction of `GP`. ☐

[14] Effective for states s such that $\nu(s) = F$.

Chapter 5

Conclusion

In this chapter we will summarize the contributions of this study to the field. We have already discussed many such contributions, specifically in Section 3.5. We will go over the broad lines again, and add some new perspectives on the contributions. After, we outline a unification of the two frameworks presented in this thesis: the dynamic epistemic logic of full communication in Chapter 2, Communication Logic, and the game-theoretic framework of network dynamics in Chapter 4. Finally, we discuss other directions of future work.

5.1 Contributions

With Communication Logic, the study on communication as the process of knowledge propagation and its mediating positions, and the formalisation of a game-theoretic framework on network dynamics, we laid the groundwork for a further study on the interplay between epistemics, communication, social networks, crucial positions, and network change.

Logics of information propagation through a network often regard the propagation of bits, packets, or singular pieces of information. This passive and atomic form of exchange can by no means be equated to the intricate process of the propagation of knowledge over a network. Much more suited to such epistemic categories is the framework of distributed knowledge and communication as its concentration, actualisation, or realisation. Our study gives such an epistemic treatment of information propagation. With our directional treatment of distributed knowledge realisation — the realisation of that what is distributively known by the sender and receiver, to knowledge by the receiver — we gave an epistemic account of the process of information propagation over a network.

Social network logics are, of course, not new. However, most such logics take formulas of a language as the objects of communication, alike *PAL*. Communication Logic functions as a full communication logic that treats communication, instead, as the transmission of similarity of worlds; abstracting away from the specifics regarding the language in which the knowledge is phrased, and the order in which formulas are announced.

The resolution of distributed knowledge, its actualisation, is the subject of full communication logics and resolution logics [3; 16]. Our contribution to this field is the formalisation of a logic that restricts such resolution operators by a network, that is explicitly included in the model. As such it further develops the reading map modalities of Baltag and Smets [16], modalities that model restricted resolution by a reading map stated in the syntax of the logic itself. Communication Logic, instead, treats the social network as an object of its models: the logic is evaluated *on* a social network, and full communication modalities model the resolution of distributed knowledge relative to this social network.

With social networks and their science come mediating positions. These positions are crucial to the propagation of knowledge over the network. Belardinelli [19] gave a logical study of these categories of the social sciences, specifically of Gatekeeping Theory. In

Chapter 3, our work extended this direction of study. We showed that connectors enable the resolution of knowledge among senders and receivers to knowledge among the receivers. Their negation, blocking sets, blocks this resolution. Firstly, this chapter functions as a generalisation of the work of Belardinelli [19] to a setting of *directional* network. Secondly, it gives an *epistemic* account of the crucial positions and their implications for knowledge propagation. Our work expresses these concepts in a logic that is arguably more standard than that of Belardinelli [19], because of its possible world semantics.

In Chapter 4, we constructed a game-theoretic framework to reason about network formation and change. The perspective often taken in logical studies of network formation is to present a descriptive logic: "φ will be true after the network is changed in some way". Our study provides a different perspective to this field of qualitative network formation, one where agents themselves form and change the network structure, one that focuses on *ability* rather than possibility.

5.2 A Logic of Communication and Network Formation

A most evident further work is to bring the two parts of this thesis together into a single logic. This requires two steps: presenting a logic of network games, and combining this logic with Communication Logic. We outlined how to construct the former in Section 4.5. Here, we will outline the latter.

5.2.1 Unifying Communication Logic and Network Formation Logic

A straightforward combination of the structure of coalition logic and communication logic provides us with models containing a set of worlds, a valuation function, a similarity relation between worlds, the network function $\nu : W \to \mathbf{F}_A$, and a function associating a playable effectivity function with each world. By the nature of the effectivity function, *any* world can result from a strategic game, also worlds that differ in, say, their evaluation of propositions. It must therefore be ensured that playing a strategic game only changes the state of the network. This can be done by, in each world, forcing the empty coalition to be effective for the *EL*-equivalence class of the world.[1] Particular care must also be given to the agent's knowledge about the network. In Communication Logic, we assumed that the network was common knowledge. This can also be assumed here.[2]

In the unified logic, the communication model update of Communication Logic must update the model relative to the social network of the current world. In order to give a complete axiomatisation using the reduction method, something must be said about the relation between this communication modality and the coalition modality. This is not straightforward, as such a relation is not expressible for a communication modality that is *implicitly* relative to the current network. One solution is to introduce a communication modality that is explicitly relative to a network.[3] Such a modality is related to the *reading map* modalities in Baltag and Smets [16].[4]

Finally, the semantics of knowledge about effectivity must be worked out. Relevant here is the Epistemic Coalition Logic in Ågotnes and Alechina [2], and its discussion on different

[1]If the *EL*-equivalence class of w is $|w|_{EL}$, then the models must satisfy the following condition: **NG** $\forall w \in W : |w|_{EL} \in E(w)(\emptyset)$. In the language of Network Formation Logic, this is expressible as follows; for any $\varphi^{EL} \in \mathcal{L}_{EL}$: $\Vdash_{\mathcal{L}_{NFL}} \varphi^{EL} \to [\emptyset]\varphi^{EL}$.

[2]The models must satisfy the following; **KF** $\forall w, v \in W$ and $\forall a \in A : w \sim_a v \Rightarrow \nu(w) = \nu(v)$, i.e. $\Vdash_{\mathcal{L}_{NFL}} F_{a,b} \to D_G F_{a,b}$.

[3]Given a communication modality relative to a network F, $[!G(F)]$, the communication modality relative to the current network $[!G]$ are related as follows: $\Vdash [!G]\varphi \leftrightarrow \bigvee_{F \in \mathbf{F}_A} (\varphi_F \wedge [!G(F)]\varphi)$. The truth value of the right-hand-side of this equivalence is independent of the state of the current social network. As such, this equivalence can properly function as a reduction axiom.

[4]Defined in terms of reading maps, a communication update relative to a network F, $!G(F)$, is defined by $(a : \mathcal{F}|_G^+(a, F))_{a \in A}$; where $\mathcal{F}|_G^+(a, F)$ is set that a follows in G, in network F.

forms of knowledge about effectivity.

5.2.2 Communication and Network Dynamics

A unified logic of coalitional ability and communication can help analyse the interplay between network formation and communication. A question that must be answered in such an analusis is how this interplay takes shape: is the network formed first, after which it stays static and communication happens, or do network formation and communication alternate each other in some way? For the former, one must identify, not connectors and blocking sets themselves, but coalitions that are effective for these network-structural concepts. Such an analysis has a temporal element, as players can be ineffective for something in the current state, but effective for it in a later state. For the latter, notions of connectors and blocking sets must be worked out that take into account the network dynamics. A simple sufficient requirement for connectors, then, is that a set C not is a connector now, but C stays a connector for some time — long enough for the information from the sender to reach the receiver. The required length of such "stability" is related to the connector's latency. Their actual necessary requirements are more complex: when the network changes as communication happens, connectors don't have to form an entire path from the receiver to the sender at any given time. Instead, they must be able to make sure that the right parts of this path exist at the right time. In this way, their structural requirements can be relaxed, and "streched over time".

The coalitional modality for network change only represents one-step actions; single plays of a network formation game. Any finite formula of in such a modal language can only represent effectivity for finite plays of the network formation game. More complex temporal statements of effectivity require more complex temporal modalities, such as the "since" and "until" modalities of temporal logic. Fixed-point operators are needed to talk about what can be achieved "in the long run". Effectivity "in the long run" is discussed in Pauly [46], where they formalise notions of effectivity for being able to make something come about in all terminating states (partial effectivity), possibly combined with being able to terminate the game (totally effective), and forcing something to be true over the whole game frame (globally effective). Relevant also is Alternating Time Logic (ATL) [4], with temporal modalities for "until", "always", and "next". And because of the similarity between automata and extensive games with simultaneous play: fixed-point logics, co-algebras, and μ-calculi [66].

5.3 Other Future Work

As the fields of studies on social networks, communication, and epistemic are large, there are a lot of other directions to go from here. In this section, we will hint towards topics that we think are interesting to explore in the future. By no means are these the only directions one could go from here.

Common Knowledge

Distributed knowledge was a sufficient tool to analyse the effects of full communication on the knowledge of agents. It served as a sufficient preconditions for agents to know certain things after communication. However, one might also be interested in the effects of full communication on another important result of communication, its effects on common knowledge within a group. Adding common knowledge to Communication Logic allows for an analysis of a process of distributed knowledge realisation resulting in common knowledge within a set of agents, In particular, one can formulate actions and preconditions required for such knowledge realisation.

Adding common knowledge to Communication Logic is interesting in itself, as axiomatisations of logics with common knowledge are more involved: a reduction of common knowledge is not possible without extending the basic language. Relevant to such a project is the axiomatisation of the logic of full communication (without an explicit social network) that includes common knowledge presented in Baltag and Smets [16], because of the similarity between the reading map modalities and the full communication modality of Communication Logic.

Network Uncertainty

Throughout this thesis, we assumed that the social network is common knowledge. This is a strong assumption. Dropping it lead to many interesting questions. For one, it brings to light the question of the epistemic preconditions of communication: can an agent communicate though a channel only if they know about its existence? This also has implications for connectors and blocking sets. Particularly interesting are settings where agents have certainty about a specific fragment of their network. The most obvious of which is the neighbourhood of the agents, i.e. who they follow, and possibly, who follows them. Of relevance is the concept of the "sphere of sight" in Baltag et al. [17].

More Social Relations

In this study we considered a social relation that is directional. Already, this is an added complexity in comparison to many studies on the same subject. Still, we have assumed that there is a single social relation, something that is not always a given in studies of social networks. Agents interpret the information they receive from relations of different types in different ways. For example, agents might not take for granted all that is said by relations which are less reliable with respects to the truth, say non-academic relations. Social networks typically represent "positive" relation. A concept of "negative" relations is missing. Related to doxastics, such negative relations can influence your beliefs negatively: after a negative relations says that p, you might start believing that not p.

Typed networks also influence network formation. Signed networks drive tendencies towards balance: an enemy of an ones enemy tends to be ones friend etc. Networks of weak and strong relationships influence the scope of tendencies of triadic closure, where strong relations are triadically closed, whereas weak relations might not. A logical analysis of tendencies such as balance, polarisation, and strong triadic closure is given in Pedersen [48].

Communication Sequences

We have made a simplification in our treatment of iterated communication. We only considered iterated communication by a set G. A more general iteration is that of a sequence $(G_i)_{i=0}^n$ of possibly distinct sets. As mentioned in Chapter 3, the communicational formulas of connectors are too chatty. For communication over an n-connector C from S to R n communication by all of $S \cup C \cup R$ are not required. This sequential iterated communication would help formulate formulas that are more precise with respects to their required iterated communication.

Compositionality of Connectors

Something not touched upon in this thesis, but something to note is the compositionality of connectors. If C is a $\forall^s\exists^r$-n-connector from S to R, and C' is a $\forall^s\exists^r$-m-connector from R to R', then $C \cup R \cup C'$ is a $\forall^s\exists^r$-$n+m$-connector from S to R'. This does not hold for all types. For example, if C and C' are $\exists^s\exists^r$-connectors then $C \cup R \cup C'$ is not necessarily a $\exists^s\exists^r$-$(m+n)$-connector. These relations between the concepts sketched out in Chapter 3 deserve a more thorough analysis.

Bibliography

[1] Joseph Abdou and Hans Keiding. *Effectivity functions in social choice*. Vol. 8. Springer Science & Business Media, 2012. ISBN: 978-94-011-3448-4. DOI: https://doi.org/10.1007/978-94-011-3448-4.

[2] Thomas Ågotnes and Natasha Alechina. "Coalition logic with individual, distributed and common knowledge". In: *J. Log. Comput.* 29.7 (2019), pp. 1041–1069. DOI: 10.1093/logcom/exv085.

[3] Thomas Ågotnes and Yì N. Wáng. "Resolving distributed knowledge". In: *Artif. Intell.* 252 (2017), pp. 1–21. DOI: 10.1016/j.artint.2017.07.002.

[4] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. "Alternating-time temporal logic". In: *J. ACM* 49.5 (2002), pp. 672–713. DOI: 10.1145/585265.585270.

[5] Carlos Areces and Balder ten Cate. "Hybrid logics". In: *Handbook of Modal Logic*. Ed. by Patrick Blackburn, J. F. A. K. van Benthem, and Frank Wolter. Vol. 3. Studies in logic and practical reasoning. Elsevier, 2007, pp. 821–868. ISBN: 978-0-444-51690-9. DOI: 10.1016/s1570-2464(07)80017-6.

[6] Carlos Areces, Raul Fervari, and Guillaume Hoffmann. "Relation-changing modal operators". In: *Logic Journal of the IGPL* 23.4 (Apr. 2015), pp. 601–627. ISSN: 1367-0751. eprint: https://academic.oup.com/jigpal/article-pdf/23/4/601/5069526/jzv020.pdf. DOI: 10.1093/jigpal/jzv020.

[7] Guillaume Aucher, Johan van Benthem, and Davide Grossi. "Modal logics of sabotage revisited". In: *J. Log. Comput.* 28.2 (2018), pp. 269–303. DOI: 10.1093/logcom/exx034.

[8] Robert J Aumann and Roger B Myerson. "Endogenous formation of links between players and of coalitions: an application of the Shapley value". eng. In: *The Shapley Value*. Cambridge University Press, 1988, pp. 175–192. ISBN: 9780521021333.

[9] Sanaz Azimipour and Pavel Naumov. "Lighthouse Principle for Diffusion in Social Networks". In: *FLAP* 5.1 (2018), pp. 97–120. URL: http://collegepublications.co.uk/ifcolog/?00021.

[10] Venkatesh Bala and Sanjeev Goyal. "A Noncooperative Model of Network Formation". In: *Econometrica* 68.5 (2000), pp. 1181–1229. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/2999447.

[11] Alexandru Baltag. "What is DEL good for?" Slides presented at the ESSLLI 2010 workshop on Logic, Rationality and Intelligent Interaction. 2010. URL: http://ai.stanford.edu/~epacuit/lograt/esslli2010-slides/copenhagenesslli.pdf.

[12] Alexandru Baltag, Rachel Boddy, and Sonja Smets. "Group knowledge in interrogative epistemology". In: *Jaakko Hintikka on Knowledge and Game-Theoretical Semantics*. Springer, 2018, pp. 131–164. DOI: 10.1007/978-3-319-62864-6_5.

[13] Alexandru Baltag and Lawrence S. Moss. "Logics for Epistemic Programs". In: *Synth.* 139.2 (2004), pp. 165–224. DOI: 10.1023/B:SYNT.0000024912.56773.5e.

[14] Alexandru Baltag, Lawrence S. Moss, and Slawomir Solecki. "The Logic of Public Announcements and Common Knowledge and Private Suspicions". In: *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98), Evanston, IL, USA, July 22-24, 1998*. Ed. by Itzhak Gilboa. Morgan Kaufmann, 1998, pp. 43–56. ISBN: 1-55860-563-0.

[15] Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, eds. *Logic, Rationality, and Interaction - 6th International Workshop, LORI 2017, Sapporo, Japan, September 11-14, 2017, Proceedings*. Vol. 10455. Lecture Notes in Computer Science. Springer, 2017. ISBN: 978-3-662-55664-1. DOI: 10.1007/978-3-662-55665-8.

[16] Alexandru Baltag and Sonja Smets. "Learning What Others Know". In: *LPAR 2020: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Alicante, Spain, May 22-27, 2020*. Ed. by Elvira Albert and Laura Kovács. Vol. 73. EPiC Series in Computing. EasyChair, 2020, pp. 90–119. URL: https://easychair.org/publications/paper/V8Jp.

[17] Alexandru Baltag et al. "Dynamic Epistemic Logics of Diffusion and Prediction in Social Networks". In: *Studia Logica* 107.3 (2019), pp. 489–531. DOI: 10.1007/s11225-018-9804-x.

[18] Karine Barzilai-Nahon. "Toward a theory of network gatekeeping: A framework for exploring information control". In: *J. Assoc. Inf. Sci. Technol.* 59.9 (2008), pp. 1493–1512. DOI: 10.1002/asi.20857.

[19] Gaia Belardinelli. "Gatekeepers in Social Networks: Logics for Communicative Actions". MSc Thesis. ILLC-UvA, 2019. URL: https://eprints.illc.uva.nl/id/eprint/1717/.

[20] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Vol. 53. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001. ISBN: 978-1-10705088-4. DOI: 10.1017/CBO9781107050884.

[21] Francis Bloch and Bhaskar Dutta. "Formation of networks and coalitions". In: *Handbook of social economics*. Ed. by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson. Vol. 1. Handbook of Social Economics. Elsevier, 2011, pp. 729–779. DOI: 10.1016/B978-0-444-53187-2.00016-4.

[22] Felix Brandt et al., eds. *Handbook of Computational Social Choice*. Cambridge University Press, 2016. DOI: 10.1017/CBO9781107446984.

[23] Ronald S Burt. *Structural holes: The social structure of competition*. Harvard university press, 1995. ISBN: 0-674-84371-1.

[24] Robert M. Carrington. "Learning and Knowledge in Social Networks". MSc Thesis. ILLC-UvA, 2013. URL: https://eprints.illc.uva.nl/id/eprint/907/.

[25] Zoé Christoff. "A Logic for Social Influence through Communication". In: *Proceedings of the Eleventh European Workshop on Multi-Agent Systems (EUMAS 2013), Toulouse, France, December 12-13, 2013*. Ed. by Emiliano Lorini. Vol. 1113. CEUR Workshop Proceedings. CEUR-WS.org, 2013, pp. 31–39. URL: http://ceur-ws.org/Vol-1113/paper3.pdf.

[26] Zoé Christoff and Pavel Naumov. "Diffusion in social networks with recalcitrant agents". In: *J. Log. Comput.* 29.1 (2019), pp. 53–70. DOI: 10.1093/logcom/exy037.

[27] Zoé Christoff and Rasmus K Rendsvig. "Dynamic logics for threshold models and their epistemic extension". In: *Epistemic logic for individual, social, and interactive epistemology workshop*. 2014.

[28] David A. Easley and Jon M. Kleinberg. *Networks, Crowds, and Markets - Reasoning About a Highly Connected World*. Cambridge University Press, 2010. ISBN: 978-0-521-19533-1. DOI: 10.1017/CBO9780511761942.

[29] Ronald Fagin et al. *Reasoning about knowledge*. MIT Press, 1995. ISBN: 0262061627. URL: http://www.worldcat.org/oclc/31413117.

[30] Peter C. Fishburn. *The Theory of Social Choice*. Princeton University Press, 2015. ISBN: 9781400868339. DOI: 10.1515/9781400868339.

[31] Andrea Galeotti et al. "Network Games". In: *The Review of Economic Studies* 77.1 (2010), pp. 218–244. ISSN: 00346527, 1467937X. URL: http://www.jstor.org/stable/40587626.

[32] Sanjeev Goyal. *Connections. An Introduction to the Economics of Networks*. Princeton University Press, 2007. ISBN: 9781400829163. DOI: 10.1515/9781400829163.

[33] Peter Gärdenfors. "Rights, Games and Social Choice". In: *Noûs* 15.3 (1981), pp. 341–356. ISSN: 00294624, 14680068. URL: http://www.jstor.org/stable/2215437.

[34] Hans Haller and Sudipta Sarangi. "Nash networks with heterogeneous links". In: *Math. Soc. Sci.* 50.2 (2005), pp. 181–201. DOI: 10.1016/j.mathsocsci.2005.02.003.

[35] Joseph Y. Halpern and Yoram Moses. "A Guide to Completeness and Complexity for Modal Logics of Knowledge and Belief". In: *Artif. Intell.* 54.2 (1992), pp. 319–379. DOI: 10.1016/0004-3702(92)90049-4.

[36] Jaakko Hintikka. "Individuals, Possible Worlds, and Epistemic Logic". In: *Noûs* 1.1 (1967), pp. 33–62. ISSN: 0029-4624. DOI: 10.2307/2214711.

[37] Matthew O. Jackson. "A Survey of Network Formation Models: Stability and Efficiency". In: *Group Formation in Economics: Networks, Clubs, and Coalitions*. Ed. by Gabrielle Demange and MyrnaEditors Wooders. Cambridge University Press, 2005, 11–57. DOI: 10.1017/CBO9780511614385.002.

[38] Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, 2010. ISBN: 9780691134406. DOI: 10.2307/j.ctvcm4gh1.

[39] Matthew O. Jackson and Yves Zenou. "Games on Networks". In: ed. by H. Peyton Young and Shmuel Zamir. Vol. 4. Handbook of Game Theory with Economic Applications. Elsevier, 2015, pp. 95–163. DOI: 10.1016/B978-0-444-53766-9.00003-3.

[40] Andrés Occhipinti Liberman and Rasmus K. Rendsvig. "Dynamic Term-Modal Logic for Epistemic Social Network Dynamics". In: Springer, 2019, 168–182. DOI: 10.1007/978-3-662-60292-8_13.

[41] Robert J Mokken. "Cliques, clubs and clans". In: *Quality & Quantity* 13.2 (1979), pp. 161–173. DOI: 10.1007/BF00139635.

[42] Roger B. Myerson. "Graphs and Cooperation in Games". In: *Math. Oper. Res.* 2.3 (1977), pp. 225–229. DOI: 10.1287/moor.2.3.225.

[43] Mark E. J. Newman. *Networks: An Introduction*. Oxford University Press, 2010. ISBN: 978-0-19920665-0. DOI: 10.1093/ACPROF:OSO/9780199206650.001.0001.

[44] Martin J Osborne et al. *An introduction to game theory*. Vol. 3. 3. Oxford university press New York, 2004.

[45] Eric Pacuit and Rohit Parikh. "Reasoning about Communication Graphs". In: *Interactive Logic: Selected Papers from the 7th Augustus de Morgan Workshop, London*. Amsterdam University Press, 2007, pp. 135–158. ISBN: 9789053563564. URL: http://www.jstor.org/stable/j.ctt45kdbf.9.

[46] Marc Pauly. "A Logical Framework for Coalitional Effectivity in Dynamic Procedures". In: *Bulletin of Economic Research* 53.4 (2001), 305–324. ISSN: 1467-8586. DOI: 10.1111/1467-8586.00136.

[47] Marc Pauly. "A Modal Logic for Coalitional Power in Games". In: *J. Log. Comput.* 12.1 (2002), pp. 149–166. DOI: 10.1093/logcom/12.1.149.

[48] Mina Young Pedersen. “Polarization and Echo Chambers: A Logical Analyses of Balance and Triadic Closure in Social Networks”. MSc Thesis. ILLC-UvA, 2019. URL: https://eprints.illc.uva.nl/id/eprint/1700.

[49] Mina Young Pedersen, Sonja Smets, and Thomas Ågotnes. “Analyzing Echo Chambers: A Logic of Strong and Weak Ties”. In: *Logic, Rationality, and Interaction - 7th International Workshop, LORI 2019, Chongqing, China, October 18-21, 2019, Proceedings.* Ed. by Patrick Blackburn, Emiliano Lorini, and Meiyun Guo. Vol. 11813. Lecture Notes in Computer Science. Springer, 2019, pp. 183–198. ISBN: 978-3-662-60291-1. DOI: 10.1007/978-3-662-60292-8_14.

[50] Hans Peters. *Game theory: a multi-leveled approach.* Springer, 2015. ISBN: 978-3-662-46949-1. DOI: 10.1007/978-3-662-46950-7.

[51] Debraj Ray. *A Game-Theoretic Perspective on Coalition Formation.* Oxford University Press, 2007. ISBN: 9780199207954. DOI: 10.1093/acprof:oso/9780199207954.001.0001.

[52] Floris Roelofsen. “Exploring logical perspectives on distributed information and its dynamics”. MSc Thesis. ILLC-UvA, 2005. URL: https://eprints.illc.uva.nl/id/eprint/759/.

[53] Ji Ruan and Michael Thielscher. “A Logic for Knowledge Flow in Social Networks”. In: *AI 2011: Advances in Artificial Intelligence - 24th Australasian Joint Conference, Perth, Australia, December 5-8, 2011. Proceedings.* Ed. by Dianhui Wang and Mark Reynolds. Vol. 7106. Lecture Notes in Computer Science. Springer, 2011, pp. 511–520. ISBN: 978-3-642-25831-2. DOI: 10.1007/978-3-642-25832-9_52.

[54] Katsuhiko Sano and Satoshi Tojo. “Dynamic Epistemic Logic for Channel-Based Agent Communication”. In: *Logic and Its Applications, 5th Indian Conference, ICLA 2013, Chennai, India, January 10-12, 2013. Proceedings.* Ed. by Kamal Lodaya. Vol. 7750. Lecture Notes in Computer Science. Springer, 2013, pp. 109–120. ISBN: 978-3-642-36038-1. DOI: 10.1007/978-3-642-36039-8_10.

[55] Jeremy Seligman, Fenrong Liu, and Patrick Girard. “Facebook and the Epistemic Logic of Friendship”. In: *CoRR* abs/1310.6440 (2013). URL: http://arxiv.org/abs/1310.6440. arXiv: 1310.6440.

[56] Jeremy Seligman, Fenrong Liu, and Patrick Girard. “Logic in the Community”. In: *Logic and Its Applications - 4th Indian Conference, ICLA 2011, Delhi, India, January 5-11, 2011. Proceedings.* Ed. by Mohua Banerjee and Anil Seth. Vol. 6521. Lecture Notes in Computer Science. Springer, 2011, pp. 178–188. ISBN: 978-3-642-18025-5. DOI: 10.1007/978-3-642-18026-2_15.

[57] Sonja Smets and Fernando R. Velázquez-Quesada. “A Logical Perspective On Social Group Creation”. In: *The Logica Yearbook 2017.* Ed. by Pavel Arazim and Tomáš Lávička. College Publications, 2017, pp. 271–288. ISBN: 978-1-84890-218-7.

[58] Sonja Smets and Fernando R. Velázquez-Quesada. “How to Make Friends: A Logical Approach to Social Group Creation”. In: *Logic, Rationality, and Interaction - 6th International Workshop, LORI 2017, Sapporo, Japan, September 11-14, 2017, Proceedings.* Ed. by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada. Vol. 10455. Lecture Notes in Computer Science. Springer, 2017, pp. 377–390. ISBN: 978-3-662-55664-1. DOI: 10.1007/978-3-662-55665-8_26.

[59] Johan van Benthem. “An Essay on Sabotage and Obstruction”. In: *Mechanizing Mathematical Reasoning, Essays in Honor of Jörg H. Siekmann on the Occasion of His 60th Birthday.* Ed. by Dieter Hutter and Werner Stephan. Vol. 2605. Lecture Notes in Computer Science. Springer, 2005, pp. 268–276. ISBN: 3-540-25051-4. DOI: 10.1007/978-3-540-32254-2_16.

[60] Johan van Benthem. *Logic in Games*. The MIT Press, 2014. ISBN: 9780262320290. DOI: 10.7551/mitpress/9674.001.0001.

[61] Johan van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2011. DOI: 10.1017/CBO9780511974533.

[62] Johan van Benthem. "One is a lonely number: on the logic of communication". In: *Logic colloquium*. Vol. 2, pp. 96–129. URL: https://staff.fnwi.uva.nl/j.vanbenthem/Muenster.pdf.

[63] Wiebe van der Hoek, Bernd van Linder, and John-Jules Meyer. "Group knowledge is not always distributed (neither is it always implicit)". In: *Mathematical Social Sciences* 38.2 (1999), pp. 215–240. ISSN: 0165-4896. DOI: https://doi.org/10.1016/S0165-4896(99)00013-X.

[64] Hans van Ditmarsch and Barteld P. Kooi. "The Secret of My Success". In: *Synthese* 151.2 (2006), pp. 201–232. DOI: 10.1007/s11229-005-3384-9.

[65] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld P. Kooi. "Knowing More - From Global to Local Correspondence". In: *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*. Ed. by Craig Boutilier. 2009, pp. 955–960. URL: http://ijcai.org/Proceedings/09/Papers/162.pdf.

[66] Yde Venema. "Automata and Fixed Point Logics for Coalgebras". In: *Electronic Notes in Theoretical Computer Science* 106 (2004), 355–375. ISSN: 1571-0661. DOI: 10.1016/j.entcs.2004.02.038.

[67] Yanjing Wang and Qinxiang Cao. "On axiomatizations of public announcement logic". In: *Synth.* 190.Supplement-1 (2013), pp. 103–134. DOI: 10.1007/s11229-012-0233-5.

[68] Yì N. Wáng and Thomas Ågotnes. "Public announcement logic with distributed knowledge: expressivity, completeness and complexity". In: *Synth.* 190.Supplement-1 (2013), pp. 135–162. DOI: 10.1007/s11229-012-0243-3.

[69] Yì N. Wáng and Thomas Ågotnes. "Simpler completeness proofs for modal logics with intersection". In: *CoRR* abs/2004.02120 (2020). URL: https://arxiv.org/abs/2004.02120. arXiv: 2004.02120.

[70] Heinrich Wansing. "Agency and Deontic Logic, J.F. Horty". In: *J. Log. Lang. Inf.* 13.3 (2004), pp. 379–381. DOI: 10.1023/B:JLLI.0000028421.66183.6b.

[71] Zuojun Xiong et al. "Towards a Logic of Tweeting". In: *Logic, Rationality, and Interaction - 6th International Workshop, LORI 2017, Sapporo, Japan, September 11-14, 2017, Proceedings*. Ed. by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada. Vol. 10455. Lecture Notes in Computer Science. Springer, 2017, pp. 49–64. ISBN: 978-3-662-55664-1. DOI: 10.1007/978-3-662-55665-8_4.