Visibility and Exploitation in Social Networks*

Rustam Galimullin† Mina Young Pedersen‡

Abstract

Social media is not a neutral channel. How visible information posted online is, depends on many factors such as the network structure, the emotional volatility of the content and the design of the social media platform. In this paper, we use formal methods to study the visibility of agents and information in a social network, as well as how vulnerable the network is to exploitation. We introduce a modal logic to reason about a social network of agents that can follow each other, post and share information. We show that by imposing some simple rules on the system, a potentially malicious agent can take advantage of the network construction to post an unpopular opinion that may reach many agents. The network is presented both in static and dynamic forms. We prove completeness, expressivity and model checking problem complexity results for the corresponding logical systems.

1 Introduction

Social media is not a neutral channel for information distribution. How visible information posted online is, and how many users in a social network it can reach, depends on many factors. These include the network structure [27], the emotional volatility of the content [10], past exposure to similar information [34] and the design and recommendation algorithms of the particular social media platform [33]. The design of the social media platform might also determine how vulnerable the platform is to exploitation: Is it possible to act tactically to increase the visibility of a post?

This paper contributes to the study of social networks and the measurable impact social media platforms have on their users. Social networks have been studied using numerous methods in numerous disciplines. Formal logic methods for representing and reasoning about social networks have been used to analyse opinion diffusion and social influence [13, 14, 15, 32, 3], social bots [37, 38], group polarisation [40, 41], gatekeepers [7], echo chambers [39] and informational cascades [5], among other phenomena. Our work is positioned within this literature.

*This paper is an extended and revised version of a paper by [22], first published in the proceedings of the 28th Workshop on Logic, Language, Information and Computation (WoLLIC).

†University of Bergen, Bergen, Norway; rustam.galimullin@uib.no

‡University of Bergen, Bergen, Norway; mina.pedersen@uib.no

We are here concerned with the problem of applying **formal methods to study the visibility of agents and information in a social network**. In addition to having structural properties, a number of agents and how they are connected, a social network also has other properties connected to the visibility of an agent, such as: which interests and opinions the agents have, what they are communicating and how the network changes through time. Furthermore, given a particular social network with rules inspired by real-life behavior, we aim to **analyse the safety of a network: whether it is vulnerable to exploitation by a potential malicious agent**. It is our position that logic-based methods are needed to complement empirical methods to reach a full understanding of these properties of social networks.

The notion of visibility is rooted in the idea of being seen. One of our main motivations is to present an analysis of visibility that captures a complex view of what it means to be visible in a social network, one that extends merely counting the followers of an agent. To do this, we introduce a modal logic for representing agents, their opinions and interactions in a social network.

Our social network consists of a set of agents and two sets of relations between them: one represents followers, and the other represents posts that pass through the network. We turn to Figure 1 for an intuitive explanation of the network.

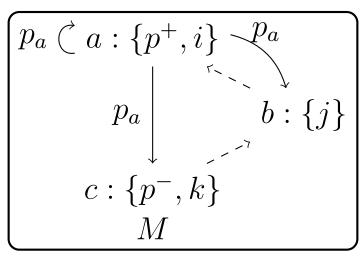

Figure 1: Model M with the followership relation depicted by dashed arrows.

The network M consists of three agents a, b and c. Dashed arrows represent a followership relation: c follows b and b follows a. The situation concerns a post on a particular topic, called p. Agent a is in favour of, *pro*, p, denoted p^+, whereas c is *contra* p, denoted p^-. Agent b has no opinion about p. Furthermore, agent a has posted on p, represented by a reflexive loop denoted p_a and agents b and c have seen the post, denoted by p_a-arrows from a.

The intuition behind our models is to observe a situation of posting and sharing a post after it has happened. Posting, sharing, following and unfollowing adhere to some simple rules of the system:

1. When an agent posts, all her followers can see the post.

2. If an agent sees a post on a topic she likes, she will reshare the post and follow the original poster.

3. If an agent sees a post on a topic she dislikes, she does not reshare it and unfollows the agent from whom she has seen the post.

4. If an agent sees a post on a topic she is indifferent to, she does not do anything.

Knowing the rules of the system, we can return to M in Figure 1 and observe that c likely has unfollowed a after a posted on p.

These rules are an oversimplification of a real-life network, but we believe they capture some key notions of a social network that we can use to analyze situations that may occur in an actual network setting. Although it might be unrealistic that an agent would for instance always unfollow when seeing a post she disagrees with, these simplified rules of the system capture some basic notions that can be found in existing networks: namely, hostility towards agents that we disagree with, and friendliness towards agents whom we agree with. It is also a point to be made that even with such simple rules, we can model interesting situations in which similar mechanisms actually happen.

We first present a logic that specifies a static network, as seen in Figure 1. The purpose of this logic is not to define what visibility is, but to allow us to discuss different qualitative and quantitative measures of visibility and formalise some of them in the logic. Next, we extend the framework into a dynamic setting where we step-wise observe what happens when information is posted in the network. We show that according to the rules of the system, the interests of the agents' followers matter a lot to what information is shared and seen. We also show that a malicious agent could take advantage of the network construction to post an unpopular opinion that will reach many agents. Then, we extend the logic further to include operators to analyse tactical actions from an agent's perspective. This lets us formally reason about whether an agent can act in a certain way to increase the visibility of their posts. We believe these observations can be useful in understanding how agents in a network contribute to spreading controversial information such as misinformation.

We are also interested in the mathematical properties of the three logics we present in this paper: static visibility logic (SVL), visibility logic (VL) and arbitrary visibility logic (AVL). We give formulas corresponding to the rules of the system and show that SVL is complete with respect to the models with these rules. The model checking problem for SVL is in P. With the first dynamic extension, we show that the language of VL is strictly more expressive than SVL. We also prove that the model checking problem for VL is PSPACE-complete. When extending VL to AVL, we show that adding quantification over actions to the dynamic language results in a new strictly more expressive language. The model checking problem for AVL is, however, also PSPACE-complete.

The contribution of the paper is the following:

- We introduce three novel logics to analyse posting and sharing information in a social network and prove mathematical results about these formal systems.

- We propose quantitative and qualitative measures of visibility and reachability, and formalise some of the properties as logical formulas.

- We use our formal system to reason about mechanisms that might occur in real-life online social networks, specifically we formalise how a potentially malicious agent could take advantage of the network construction to post a controversial opinion that will reach many agents.

- Motivated by analysing safety and exploitation in our system, we introduce quantification over actions to formally study tactical actions from an agent's perspective.

The paper is structured as follows. In Section 2, we give an overview of work in social network analysis on reachability and visibility. In Section 3, we present static visibility logic (SVL). We specify mathematical properties of the logic, give some logical formulas corresponding to measures of visibility, and prove soundness and completeness of SVL. In Section 4, we extend SVL with a dynamic operator and name it visibility logic (VL). We give a motivating example where we show that one can exploit the network structure to expose more agents to a controversial opinion. We also prove an expressivity result and give the complexity for the model checking problem. Then, in Section 5, we extend VL with an action operator and name it arbitrary visibility logic (AVL). We show another expressivity result as well as the complexity of the model checking problem for AVL. In Section 6 we give an account of other dynamic hybrid logics for social networks and position our logics within this literature. In Section 7, we summarise our paper and outline directions for future work.

2 Visibility and Reachability

Visibility in social networks is yet to be explicitly explored from a formal logical perspective. The concept has however been researched in the social network analysis literature. We present a selected collection of this work to learn how this related field has attempted to measure visibility. There seems to be no consensus in the literature on what it means to be visible in a social network, which motivates the usefulness of further study on this topic. This is confirmed in the literature review by [49], which focuses on communication visibility in computer-mediated communication.

Closely related to visibility in social networks is the notion of reachability. What exactly reachability is, or how closely related it is to visibility, is not agreed upon, which is illustrated by the different measures seen in this section. Visibility and reachability are presented as properties of both networks, agents and posts, when relevant we specify which in the following.

In a book known to be part of the canonical literature in social network analysis, [18] describe the reachability properties of a network in terms of identifying which agents are reachable from which others through connected paths of edges.

[44] distinguish reachability and visibility in an online social network, where the first measure is dependent on the second. The network is represented as an undirected graph where nodes represent agents and the relation between them represents one of three non-overlapping relations: trusted friends, acquaintances or distrusted agents. Agents can

post information with four different visibility settings: trusted friends, trusted friends and acquaintances, all friends and public. The visibility of an agent is therefore measured with respect to what relation the viewers of the post have to the agent that posts. The reachability factor of a post is defined in terms of a function: $d(v_1, v_2) = \frac{|e(v_1,v_2)|}{\sqrt{|v_1|} \times \sqrt{|v_2|}}$. In this function, v_1 is the set of agents in the network that have seen the post and v_2 is the set of agents that have not seen the post. $e(v_1, v_2)$ is the set of relations between agents across v_1 and v_2 specified with respect to the relations in the network graph. The reachability factor is dependent on the visibility settings of the agent who posts; the set v_1 increases and v_2 decreases when the visibility settings include a higher number of agents.

[48] present a temporal characterisation of reachability. In this work, the reachability is measured between two given nodes in a time interval in the network. The network is presented as a series of undirected graphs that represent how a network changes through time. The nodes in the network can be regarded as agents and the relation between them as information channels. Node j is reachable in the time interval $[t_{min}, t_{max}]$ from node i if a message can be delivered through the information channel in that time interval.

[43] define two types of visibility of an agent in an online social network: topological and behavioral visibility. Although it is mentioned that this could be a generic social network, the examples refer to the microblogging network Twitter, at the time of writing now called X, which is represented as a directed graph of agents who can post and follow each other. In the model presented by [43], a tweet embodies at least one topic from a set of interests S. Each agent in the network also has some specified interests from S. Topological visibility of an agent is calculated based on the number of followers of the agent and the clustering coefficient of the network. The clustering coefficient is usually defined in the literature in terms of directed graphs and is meant to give a view of the network structure. The higher the number, the more highly connected the network is. It is not specified which definition of clustering coefficient is used by [43]. The behavioral visibility of an agent is defined as the average of the visibility of all the tweets that are shared by the user in a time interval Δt. The visibility of a tweet represents the number of users influenced by the tweet and is proportional to the number of followers whose interests match the topics of the tweet.

[31] propose a framework to compute the privacy score of users in online social networks. In this framework, the more visible the information is in the network, the higher the privacy risk. As part of computing privacy scores, an estimation of the visibility of information is also made. The visibility $V(i, j)$ denotes the visibility of an item of information i for a user j, and is calculated as the probability that j has made the information associated with i publicly available.

3 Reasoning About Visibility in a Static Setting

We are ready to present some of the main concepts underlying our intuitions about visibility in a formal setting. We begin by introducing the language and semantics of *static visibility logic* (SVL), which serve as a basis for the logics presented in later sections.

3.1 Language and Semantics of SVL

Let $\mathsf{Nom} = \{i, j, k, ...\}$ be a countable set of nominals, and $\mathsf{Top} = \{p, q, r, ...\}$ be a countable set of topics, such that $\mathsf{Nom} \cap \mathsf{Top} = \emptyset$.

Definition 1. We define the well-formed formulas of *the language of the static fragment of visibility logic* $\mathbb{SVL}$ to be generated by the following grammar:

$$\varphi ::= p^+ \mid p^- \mid i \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Diamond_{i:p}\varphi \mid \Diamond_{i:p}^{-1}\varphi \mid \blacklozenge\varphi \mid \blacklozenge^{-1}\varphi \mid @_i\varphi$$

where $p \in \mathsf{Top}$ and $i \in \mathsf{Nom}$. We define propositional connectives like $\vee, \rightarrow$ and the formulas $\top, \bot$ as usual and the duals as standard $\Box := \neg\Diamond\neg$, $\Box^{-1} := \neg\Diamond^{-1}\neg$, $\blacksquare := \neg\blacklozenge\neg$, and $\blacksquare^{-1} := \neg\blacklozenge^{-1}\neg$.

Given a formula $\varphi \in \mathbb{SVL}$, we can recursively define the *modal depth* of the formula $md(\varphi)$ in the following way: $md(p^+) = md(p^-) = md(i) = 0$, $md(\neg\varphi) = md(@_i\varphi) = md(\varphi)$, $md(\varphi \wedge \psi) = \max\{md(\varphi), md(\psi)\}$, and $md(\Diamond_{i:p}\varphi) = md(\Diamond_{i:p}^{-1}\varphi) = md(\blacklozenge\varphi) = md(\blacklozenge^{-1}\varphi) = md(\varphi)+1$. The *size* of φ, denoted $|\varphi|$, is defined as follows: $|p^+| = |p^-| = |i| = 1$, $|\neg\varphi| = |\Diamond_{i:p}\varphi| = |\Diamond_{i:p}^{-1}\varphi| = |\blacklozenge\varphi| = |\blacklozenge^{-1}\varphi| = |@_i\varphi| = |\varphi| + 1$, and $|\varphi \wedge \psi| = |\varphi| + |\psi| + 1$.

In our language, similar to other approaches to logic-based analysis of social networks (see, e.g., [14]), we distinguish three possible dispositions of an agent to a topic $p \in \mathsf{Top}$. The agent may be *pro* p, which we express with p^+, *contra* p, expressed by p^-, or *indifferent to* p, if the agent is neither *pro* nor *contra* p. The agent cannot be both *pro* and *contra* p.

Constructs $\blacklozenge\varphi$ and $\blacklozenge^{-1}\varphi$ express that 'the current agent follows an agent satisfying φ' and 'the current agent is followed by an agent that satisfies φ' respectively. Formulas $\Diamond_{i:p}\varphi$ and $\Diamond_{i:p}^{-1}\varphi$ mean that 'there is an agent satisfying φ who sees the (re)post by the current agent on topic p (originally posted by an agent named i)' and 'there is an agent satisfying φ whose (re)post on topic p (originally posted by an agent named i) is seen by the current one'.

Formulas of SVL are defined on relational visibility models.

Definition 2. A *visibility model* (or a *model*) M is a tuple $(A, F, +, -, V, R)$, where

- A is a non-empty set of agents;
- $F : A \rightarrow 2^A$ is an irreflexive followership relation,
- $+ : A \rightarrow 2^{\mathsf{Top}}$ assigns to each agent a set of topics she is *pro*,
- $- : A \rightarrow 2^{\mathsf{Top}}$ assigns to each agent a set of topics she is *contra* such that for all agents $a \in A$, it holds that $+(a) \cap -(a) = \emptyset$,
- $V : \mathsf{Nom} \rightarrow 2^A$ is a valuation such that for all $i \in \mathsf{Nom}$: $|V(i)| = 1$,
- $R : \mathsf{Top} \times A \rightarrow 2^{A \times A}$ is a visibility relation for each topic and each agent satisfying the following conditions, where $p \in \mathsf{Top}$ and $a, b, c \in A$:

1. If $(a, b) \in R(p, c)$, then $(a, a) \in R(p, c)$.
2. If $(a, a) \in R(p, c)$, then $(a, b) \in R(p, c)$ for all b such that $b \in F(a)$.
3. If $(a, b) \in R(p, c)$, $p \in +(b)$, and $b \neq c$, then $(b, b) \in R(p, c)$ and $b \in F(c)$.
4. If $(a, b) \in R(p, c)$, $p \in -(b)$, and $a \neq b$, then $(b, b) \notin R(p, c)$ and $b \notin F(a)$.
5. If $(a, b) \in R(p, c)$, $p \notin +(b)$, $p \notin -(b)$, and $a \neq b$, then $(b, b) \notin R(p, c)$.

A *pointed visibility model* M_a is a model M with a distinguished point $a \in A$ where evaluation takes place. If necessary, we refer to the elements of the tuple as A_M,F_M,$+_M$,$-_M$,V_M, and R_M. A visibility model such that for all $a \in A$ there is some $i \in \mathsf{Nom}$ such that $V(i) = \{a\}$ is called *named*. All models we will be dealing with in the paper are named. Let $Nom(a) := \{i \in \mathsf{Nom} \mid a \in V(i)\}$ be a set of all nomimals assigned to an agent, and $Top(a) := \{p \in \mathsf{Top} \mid R(p, a)\}$ be a set of all topics that an agent posted. A visibility model M is *finite* if all of A, $\bigcup\{+(a) \mid a \in A\}$, $\bigcup\{-(a) \mid a \in A\}$, $\bigcup\{Nom(a) \mid a \in A\}$, and $\bigcup\{Top(a) \mid a \in A\}$ are finite.

Let $M = (A, F, +, -, V, R)$ be a finite visibility model. The *size* of M equals to

$$\text{card}(A) + \text{card}(F) + \sum_{a \in A} \begin{pmatrix} \text{card}(+(a)) \\ \text{card}(-(a)) \\ \text{card}(Nom(a)) \\ \sum_{p \in Top(a)} \text{card}(R(p, a)) \end{pmatrix}.$$

In Definition 2 above, the first condition on R states that if agent b sees a post, which was originally posted by agent c on topic p, from agent a, then a herself can see the post. The second condition ensures that if an agent posts a post, all her followers can see the post. Condition number three specifies that if an agent sees a post on a topic she likes, she will reshare the post and follow the original poster. The fourth condition says that if an agent sees a post on a topic she dislikes, she does not reshare it and unfollows the agent from whom she has seen the post. Finally, the last condition stipulates that if an agent sees a post on a topic she is indifferent to, she does not reshare the post.

Note that our definition of R does not preclude situations where agents may have seen a post on a topic they dislike from an agent they do not follow. How such situations may come about will be the focus of the next section.

Definition 3. Let $M = (A, F, +, -, V, R)$ be a model, $a, b, c \in A$, $p \in \mathsf{Top}$, $i \in \mathsf{Nom}$, and

$\varphi, \psi \in \mathbb{SVL}$. *The semantics of SVL* is recursively defined as follows:

$$
\begin{array}{lll}
M_a \models p^+ & \text{iff} & p \in +(a) \\
M_a \models p^- & \text{iff} & p \in -(a) \\
M_a \models i & \text{iff} & a \in V(i) \\
M_a \models \neg\varphi & \text{iff} & M_a \not\models \varphi \\
M_a \models \varphi \wedge \psi & \text{iff} & M_a \models \varphi \text{ and } M_a \models \psi \\
M_a \models \Diamond_{i:p}\varphi & \text{iff} & \exists b, c \in A : (a, b) \in R(p, c) \text{ and } V(i) = \{c\} \text{ and } M_b \models \varphi \\
M_a \models \Diamond_{i:p}^{-1}\varphi & \text{iff} & \exists b, c \in A : (b, a) \in R(p, c) \text{ and } V(i) = \{c\} \text{ and } M_b \models \varphi \\
M_a \models \blacklozenge\varphi & \text{iff} & \exists b \in A : a \in F(b) \text{ and } M_b \models \varphi \\
M_a \models \blacklozenge^{-1}\varphi & \text{iff} & \exists b \in A : b \in F(a) \text{ and } M_b \models \varphi \\
M_a \models @_i\varphi & \text{iff} & M_b \models \varphi \text{ and } \{b\} = V(i)
\end{array}
$$

Observe that if $M_a \not\models p^+$ then we have that either $p \in -(a)$ or not. This corresponds to the intuition that agent a is not *pro* p if she actively dislikes the topic (she is *contra* p), or if she is indifferent to it. Similarly, for $M_a \not\models p^-$.

The valuation function V is such that for all $i \in \mathsf{Nom} : |V(i)| = 1$. In other words, a name can only be true for one agent. However, note that one agent can have several names, i.e. it can be the case that two nominals i and j are forced at the same agent a. Furthermore, we allow for two agents to post on the same topic, that is, regardless of whether two nominals i and j refer to the same agent, both $\Diamond_{i:p}\top$ and $\Diamond_{j:p}\top$ can be forced simultaneously.

Recall the example from Figure 1. In the figure, we have that $M_a \models p^+$, $M_c \models p^-$, and $M_b \models \neg p^+ \wedge \neg p^-$, meaning that agent a is *pro* topic p, agent b is indifferent towards the topic, and c is *contra* p. Moreover, we have, for example, that $M_c \models \Diamond_{i:p}^{-1}\top \wedge \blacksquare\neg p^+$, meaning that agent c has seen a post by the agent with name i on topic p, and that all agents that c follows are not *pro* p.

We can formalise some of the notions of visibility and reachability as formulas in SVL. Let $M = (A, F, +, -, V, R)$ be a model. Some quantitative amounts related to visibility that we can count in finite models are:

- How many followers the agent called i has: $|\{a \in A \mid M_a \models \blacklozenge i\}|$
- How many agents have seen the agent called i's post on p:

  $|\{a \in A \mid M_a \models \Diamond_{i:p}^{-1}\top\}|$
- How many agents that are pro p have seen the agent called i's post on p:

  $|\{a \in A \mid M_a \models p^+ \wedge \Diamond_{i:p}^{-1}\top\}|$

We also present some formulas corresponding to qualitative properties of agents in the network. The following formulas are forced at an agent iff the property holds of that agent:

- The current agent i is the original poster of a post on p: $\quad i \wedge \Diamond_{i:p}\top$
- The current agent has seen i's post on p: $\quad \Diamond_{i:p}^{-1}\top$

- All the followers of the current agent i have shared i's post on p:
  $i \wedge \blacksquare^{-1} \Diamond_{i:p} \top$
- The current agent i shared a post to a follower j, but j also saw the post from another source: $i \wedge \blacklozenge^{-1}(j \wedge \Diamond_{i:p}^{-1} i \wedge \Diamond_{i:p}^{-1}(\neg i \wedge \neg j))$
- The current agent i has gained a follower who is *pro* p, after i posted on p:
  $i \wedge \Diamond_{i:p} \top \wedge \blacklozenge^{-1}(p^+ \wedge \Diamond_{i:p}^{-1} i)$
- The current agent i has reached the agent j with i's post on p in no more than 3 steps: $i \wedge \Diamond_{i:p} \Diamond_{i:p} \Diamond_{i:p} j$

Definition 4. A formula φ is called *valid* if for all models M_a, we have that $M_a \models \varphi$. Formulas φ and ψ are *equivalent*, if for all models M_a, it holds that $M_a \models \varphi$ if and only if $M_a \models \psi$.

Definition 5. Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two languages. We say that $\mathcal{L}_2$ *is more expressive than* $\mathcal{L}_1$ if for each $\varphi \in \mathcal{L}_1$ there is an equivalent $\psi \in \mathcal{L}_2$, and there is a $\chi \in \mathcal{L}_2$ for which there is no equivalent $\tau \in \mathcal{L}_1$.

The following notion of bisimulation is based on hybrid bisimulation [1] and on bisimulation for logics with 'backward-looking' modalities (see, e.g., [29]).

Definition 6. Let $M = (A_M, F_M, +_M, -_M, V_M, R_M)$ and $N = (A_N, F_N, +_N, -_N, V_N, R_N)$ be visibility models, and $Q \subseteq \mathsf{Nom}$. We say that M and N are *Q-bisimilar* (denoted $M \leftrightarrows_Q N$) if there is a non-empty relation $B \subseteq A_M \times A_N$, called *Q-bisimulation* such that the following conditions are satisfied:

Atoms$^+$ If $B(a, b)$, then for all $p \in \mathsf{Top}$: $p \in +_M(a)$ iff $p \in +_N(b)$,

Atoms$^-$ If $B(a, b)$, then for all $p \in \mathsf{Top}$: $p \in -_M(a)$ iff $p \in -_N(b)$,

Nominals 1 If $B(a, b)$, then for all $i \in Q$: $a \in V_M(i)$ iff $b \in V_N(i)$,

Nominals 2 For all $i \in Q$, if $V_M(i) = \{a\}$ and $V_N(i) = \{b\}$, then $B(a, b)$,

Forth $\Diamond$ If $B(a, b)$ and $(a, a') \in R_M(p, c)$, then there is a $b' \in A_N$ such that $(b, b') \in R_N(p, c)$ and $B(a', b')$,

Back $\Diamond$ If $B(a, b)$ and $(b, b') \in R_N(p, c)$, then there is an $a' \in A_M$ such that $(a, a') \in R_M(p, c)$ and $B(a', b')$,

Forth $\Diamond^{-1}$ If $B(a, b)$ and $(a', a) \in R_M(p, c)$, then there is a $b' \in A_N$ such that $(b, b') \in R_N(p, c)$ and $B(a', b')$,

Back $\Diamond^{-1}$ If $B(a, b)$ and $(b', b) \in R_N(p, c)$, then there is an $a' \in A_M$ such that $(a', a) \in R_M(p, c)$ and $B(a', b')$,

Forth $\blacklozenge$ If $B(a,b)$ and $a \in F_M(a')$, then there is a $b' \in A_N$ such that $b \in F_N(b')$ and $B(a',b')$,

Back $\blacklozenge$ If $B(a,b)$ and $b' \in F_N(b)$, then there is an $a' \in A_M$ such that $a' \in F_M(a)$ and $B(a',b')$,

Forth $\blacklozenge^{-1}$ If $B(a,b)$ and $a' \in F_M(a)$, then there is a $b' \in A_N$ such that $b' \in F_N(b)$ and $B(a',b')$,

Back $\blacklozenge^{-1}$ If $B(a,b)$ and $b' \in F_N(b)$, then there is an $a' \in A_M$ such that $a' \in F_M(a)$ and $B(a',b')$.

We say that M_a and N_b are Q-bisimilar, and denote this by $M_a \leftrightarrows_Q N_b$, if there is a bisimulation linking agents a and b. If $Q = \mathsf{Nom}$, we say that M_a and N_b are bisimilar, and write $M_a \leftrightarrows N_b$.

The following theorem is a standard result in hybrid modal logic [1].

Theorem 7. Let M_a and N_b be two models. If $M_a \leftrightarrows_Q N_b$, then for all $\varphi \in \mathbb{SVL}$ such that φ includes only nominals from Q, $M_a \models \varphi$ if and only if $N_b \models \varphi$.

Definition 8. Let $M = (A_M, F_M, +_M, -_M, V_M, R_M)$ and $N = (A_N, F_N, +_N, -_N, V_N, R_N)$ be visibility models. We say that M and N are *n-bisimilar* (denoted $M \leftrightarrows_Q^n N$) if there is a sequence of relations $B_n \subseteq \ldots B_0$, called *n-bisimulation* that satisfies the following. First, a non-empty relation $B_0 \subseteq A_M \times A_N$ is a *0-bisimulation* if for all (a,b) such that $B(a,b)$, clauses **Atoms**$^+$, **Atoms**$^-$, **Nominals 1**, and **Nominals 2** hold. Then, for $n > 0$, relation B_n is an *n-bisimulation*, if there is an $(n-1)$-bisimulation $B_{n-1} \subseteq B_n$ such that

Forth $\lozenge$ If $B_n(a,b)$ and $(a,a') \in R_M(p,c)$, then there is a $b' \in A_N$ such that $(b,b') \in R_N(p,c)$ and $B_{n-1}(a',b')$,

Back $\lozenge$**, Forth** $\lozenge^{-1}$**, Forth** $\blacklozenge$**, Back** $\blacklozenge$**, Back** $\lozenge^{-1}$**, Forth** $\blacklozenge^{-1}$**, Back** $\blacklozenge^{-1}$ Similar to the cases in Definition 6 with subscripts n and $n-1$ for B as for the case **Forth** $\lozenge$.

We say that M_a and N_b are n-bisimilar and denote this by $M_a \leftrightarrows^n N_b$ if there is an n-bisimulation linking agents a and b. If, additionally, n-bisimulation between M_a and N_b is restricted to nominals from $Q \subseteq \mathsf{Nom}$, we say that M_a and N_b are Q-n-bisimilar and write $M_a \leftrightarrows_Q^n N_b$.

One of the classic results in standard modal logic is a restricted version of Theorem 7 stating that n-bisimulation implies satisfaction of the same modal formulas up to modal depth n [24]. This, however, is not the case for hybrid modal logic, since operator $@_i$ allows reaching an agent named i no matter 'how far' the agent is from the current agent. On the other hand, the result holds for hybrid modal logic if we only concentrate on formulas that do not include nominals that are used in models at hand. Granted this is an even more restrictive case, the fact will be useful later in the paper.

Theorem 9. Let M_a and N_b be two models, and let $Q \subseteq \mathsf{Nom} \setminus \{i \in \mathsf{Nom} \mid V_M(i) \neq \emptyset \text{ or } V_N(i) \neq \emptyset\}$. If $M_a \leftrightarrows^n_Q N_b$, then for all $\varphi \in \mathbb{SVL}$ such that $md(\varphi) \leqslant n$, and φ includes only nominals from Q, $M_a \models \varphi$ if and only if $N_b \models \varphi$.

The proof of the theorem is standard noting that all nominals appearing in φ have the empty denotation.

3.2 A Sound and Complete Axiomatisation of SVL

In this section, we present a sound and complete axiomatisation of SVL. The main direction of our proofs follows the strategy of establishing completeness of hybrid logic that can be found in works by [8] and [9]. We omit proofs of lemmas if they can be adopted from the cited literature with minimal changes.

Definition 10. The proof system of *SVL*, **SVL**, comprises axioms and rules of inference from Tables 1 and 2.

Table 1: Hybrid and bidirectional axioms

Prop	Propositional tautologies	T$_{\Box^{-1}\Diamond}$	$\varphi \to \Box^{-1}_{i:p}\Diamond_{i:p}\varphi$
K$_{\Box}$	$\Box_{i:p}(\varphi \to \psi) \to (\Box_{i:p}\varphi \to \Box_{i:p}\psi)$	T$_{\blacksquare\blacklozenge^{-1}}$	$\varphi \to \blacksquare\blacklozenge^{-1}\varphi$
K$_{\Box^{-1}}$	$\Box^{-1}_{i:p}(\varphi \to \psi) \to (\Box^{-1}_{i:p}\varphi \to \Box^{-1}_{i:p}\psi)$	T$_{\blacksquare^{-1}\blacklozenge}$	$\varphi \to \blacksquare^{-1}\blacklozenge\varphi$
K$_{\blacksquare}$	$\blacksquare(\varphi \to \psi) \to (\blacksquare\varphi \to \blacksquare\psi)$	HT$_{@\Box\Diamond^{-1}}$	$@_i\Box_{i:p}\Diamond^{-1}_{i:p}i$
K$_{\blacksquare^{-1}}$	$\blacksquare^{-1}(\varphi \to \psi) \to (\blacksquare^{-1}\varphi \to \blacksquare^{-1}\psi)$	HT$_{@\Box^{-1}\Diamond}$	$@_i\Box^{-1}_{i:p}\Diamond_{i:p}i$
K$_{@}$	$@_i(\varphi \to \psi) \to (@_i\varphi \to @_i\psi)$	HT$_{@\blacksquare\blacklozenge^{-1}}$	$@_i\blacksquare\blacklozenge^{-1}i$
Self-Dual	$@_i\varphi \leftrightarrow \neg @_i \neg\varphi$	HT$_{@\blacksquare^{-1}\blacklozenge}$	$@_i\blacksquare^{-1}\blacklozenge i$
Ref	$@_i i$	MP	From $\varphi \to \psi, \varphi,$ infer ψ
Agree	$@_i @_j \varphi \leftrightarrow @_j\varphi$	Nec$_{\Box}$	From φ, infer $\Box_{i:p}\varphi$
Intro	$i \wedge \varphi \to @_i\varphi$	Nec$_{\Box^{-1}}$	From φ, infer $\Box^{-1}_{i:p}\varphi$
Back$_{\Diamond}$	$\Diamond_{i:p}@_i\varphi \to @_i\varphi$	Nec$_{\blacksquare}$	From φ, infer $\blacksquare\varphi$
Back$_{\Diamond^{-1}}$	$\Diamond^{-1}_{i:p}@_i\varphi \to @_i\varphi$	Nec$_{\blacksquare^{-1}}$	From φ, infer $\blacksquare^{-1}\varphi$
Back$_{\blacklozenge}$	$\blacklozenge @_i\varphi \to @_i\varphi$	Nec$_{@}$	From φ, infer $@_i\varphi$
Back$_{\blacklozenge^{-1}}$	$\blacklozenge^{-1}@_i\varphi \to @_i\varphi$	Sub	From φ, infer φ^{σ}
T$_{\Box\Diamond^{-1}}$	$\varphi \to \Box_{i:p}\Diamond^{-1}_{i:p}\varphi$	Name	From $@_i\varphi, i \notin \varphi$, infer φ

In rule Sub, σ is a *substitution* that uniformly replaces nominals by nominals, and topics with formulas.

Table 2: Followership and visibility axioms

Irref	$@_i\neg\blacklozenge i$	Foll	$@_i((p^+ \wedge \Diamond^{-1}_{j:p}\top \wedge \neg j) \to (\Diamond_{j:p}i \wedge \blacklozenge j))$
Cons	$p^+ \wedge p^- \to \bot$	Unfoll	$@_i((p^- \wedge \Diamond^{-1}_{j:p}(\neg i \wedge k)) \to (\neg\Diamond_{j:p}\top \wedge \neg\blacklozenge k))$
Ag-See	$@_i(\Diamond_{j:p}\top \to \Diamond_{j:p}i)$	Indiff	$@_i((\neg p^- \wedge \neg p^+ \wedge \Diamond^{-1}_{j:p}\neg i) \to \neg\Diamond_{j:p}\top)$
Fol-See	$@_i(\Diamond_{j:p}i \to \blacksquare^{-1}\Diamond^{-1}_{j:p}i)$		

The axiomatisation of SVL combines the axiomatisations of hybrid logic [9, Section 7.3], hybrid tense logic [8], and additional novel axioms and rules of inference for followership and visibility.

Reamark 11. Note that we do not claim that **SVL** is the minimal set of axioms and rules of inferences that is sound and complete for the class of visibility models. For the purposes at hand, it is enough that **SVL** is finitary and complete.

Let us consider Table 1 first. Axioms K are standard modal distributivity axioms. Axiom schema SELF-DUAL states that hybrid operator $@_i$ is its own dual. REF states that an agent named i actually satisfies nominal i. According to AGREE, accessing an agent in two @-steps is the same as accessing the same agent in a single @-step. Axiom INTRO allows us to put an arbitrary true formula under the scope of @. Interactions between @ and all diamonds in our language are captured by the set of BACK axiom schemata. Axioms T are standard axioms of tense logic ensuring that our models are bidirectional. The interaction between @ and the fact that our followership and visibility relations are bidirectional is captured by axioms HT. Finally, we have standard rules of inference like modus ponens MP, necessitations NEC, and substitution SUB. Rule NAME states that if we can prove that φ holds at an arbitrary agent i, then we can prove φ.

Now let us turn to the axiom schemata in Table 2. IRREF captures the irreflexivity of the followership relation, and CONS states that an agent's views on a topic are consistent. Axiom FOLL states that if the current agent sees a post from some other agent on a topic she likes, she starts following the original poster and reposts the post. Alternatively, UNFOLL specifies that if the current agent sees a post on the topic she does not like, she unfollows the agent she the post from. Finally, according to INDIFF, if the current agent is indifferent to the topic of the post she sees, then she does nothing.

That **SVL** is sound can be shown by a direct application of the definition of semantics.

Theorem 12. **SVL** is sound.

As an example of an **SVL** derivation, consider a set of theorems, where $\heartsuit \in \{\lozenge_{i:p}, \lozenge_{i:p}^{-1}, \blacklozenge, \blacklozenge^{-1}\}$:

$$\text{BRIDGE}_{\heartsuit}\ \heartsuit i \land @_i\varphi \to \heartsuit\varphi$$

We show how to derive $\text{BRIDGE}_{\blacklozenge}$:

1.	$\blacksquare(\neg\varphi \to \neg i) \to (\blacksquare\neg\varphi \to \blacksquare\neg i)$	$K_{\blacksquare}$
2.	$\blacksquare\neg\varphi \land \blacklozenge i \to \blacklozenge(\neg\varphi \land i)$	Prop. reasoning and dual of $\blacksquare$: 1
3.	$i \land \neg\varphi \to @_i\neg\varphi$	Intro
4.	$\blacksquare(\neg @_i\neg\varphi \to (\neg i \lor \varphi))$	Prop. reasoning and $\text{Nec}_{\blacksquare}$: 3
5.	$\blacksquare\neg @_i\neg\varphi \to \blacksquare(\neg i \lor \varphi)$	$K_{\blacksquare}$ and MP: 4
6.	$\blacklozenge(i \land \neg\varphi) \to \blacklozenge @_i\neg\varphi$	Prop. reasoning and dual of $\blacksquare$: 5
7.	$\blacksquare\neg\varphi \land \blacklozenge i \to \blacklozenge @_i\neg\varphi$	Prop. reasoning: 2 and 6
8.	$\blacklozenge @_i\neg\varphi \to @_i\neg\varphi$	$\text{Back}_{\blacklozenge}$
9.	$\blacksquare\neg\varphi \land \blacklozenge i \to @_i\neg\varphi$	Prop. reasoning and dual of $\blacksquare$: 7 and 8
10.	$@_i\varphi \land \blacklozenge i \to \blacklozenge\varphi$	Self-Dual, prop. reasoning and dual of $\blacksquare$: 9

Now we turn to the proof of the completeness of **SVL**. In our proof sketch we follow [9, Section 7.3] and [8], and omit proofs of lemmas that can be taken 'as is' from the cited literature with straightforward changes.

As usual in the completeness proofs that employ the canonical model construction, we commence with the notion of a maximal consistent set (MCS).

Definition 13. Let Γ be a set of formulas. We call Γ *consistent* if $\Gamma \nvdash \bot$, and Γ is *maximal* if for any φ, either $\varphi \in \Gamma$ or $\neg\varphi \in \Gamma$. If Γ is both maximal and consistent, then we call Γ a *maximal consistent set* (MCS). We say that MCS Γ is *named* if and only if it contains a nominal. For $i \in \mathsf{Nom}$ and MCS Γ, set $\{\varphi \mid @_i\varphi \in \Gamma\}$ is called a *named set yielded by* Γ.

It is straightforward to show that MCSs contain all the instances of the axioms of **SVL** and are closed under MP. From the hybrid perspective, each MCS itself contains a collection of named MCSs with the properties defined in the following lemma.

Lemma 14 ([9, Lemma 7.24])**.** Let Γ be a MCS, and for all $i \in \mathsf{Nom}$, let $\Delta_i = \{\varphi \mid @_i\varphi \in \Gamma\}$. Then the following holds:

1. For all $i \in \mathsf{Nom}$, Δ_i is an MCS and $i \in \Delta_i$.
2. For all $i, j \in \mathsf{Nom}$, if $i \in \Delta_j$, then $\Delta_j = \Delta_i$.
3. For all $i, j \in \mathsf{Nom}$, $@_i\varphi \in \Delta_j$ if and only if $@_i\varphi \in \Gamma$.
4. For all $i \in \mathsf{Nom}$, if $i \in \Gamma$, then $\Gamma = \Delta_i$.

Before we continue, we need a set of additional rules of inference that are called $\text{Paste}_{\heartsuit}$, where $\heartsuit \in \{\Diamond_{i:p}, \Diamond_{i:p}^{-1}, \blacklozenge, \blacklozenge^{-1}\}$.

$$\text{Paste}_{\heartsuit} \quad \text{From } @_i\heartsuit j \land @_j\varphi \to \psi \text{ with } j \neq i \notin \varphi, \psi, \text{ infer } @_i\heartsuit\varphi \to \psi$$

These rules are derivable from **SVL** in the presence of axioms HT [8]. Rules $\text{Paste}_{\heartsuit}$ may not look entirely straightforward. However, for the purposes at hand, we need these rules just for the next definition, and more on the intuition behind Paste rules can be found in [9, Section 7.3].

Definition 15. An MCS Γ is called *pasted* if and only if $@_i\heartsuit\varphi \in \Gamma$ implies that for some $j \in \mathsf{Nom}$, $@_i\heartsuit j \wedge @_j\varphi \in \Gamma$, where $\heartsuit \in \{\lozenge_{i:p}, \lozenge_{i:p}^{-1}, \blacklozenge, \blacklozenge^{-1}\}$.

The next lemma is a hybrid version of Lindenbaum Lemma, and its proof relies on rules $\textsc{Paste}_\heartsuit$.

Lemma 16 ([9, Lemma 7.25])**.** Let Nom^* be a new countable set of nominals such that $\mathsf{Nom} \cap \mathsf{Nom}^* = \emptyset$. Moreover, let $\mathbb{SVL}^*$ be the language obtained from $\mathbb{SVL}$ by adding Nom^*. Then every consistent set of formulas of $\mathbb{SVL}$ can be extended to a named and pasted MCS of formulas of $\mathbb{SVL}^*$.

Now we have all the ingredients for the construction of the canonical model.

Definition 17. Let Γ be a named and pasted MCS. The named and pasted model yielded by Γ, $\mathfrak{M}^\Gamma$, is a tuple $(\mathfrak{A}^\Gamma, \mathfrak{F}^\Gamma, +^\Gamma, -^\Gamma, \mathfrak{V}^\Gamma, \mathfrak{R}^\Gamma)$, where

- $\mathfrak{A}^\Gamma = \{\{\varphi \mid @_i\varphi \in \Gamma\} \mid i \in \mathsf{Nom}\}$ is the set of all named sets yielded by Γ with typical elements denoted $\mathfrak{a}$, $\mathfrak{b}$, and $\mathfrak{c}$,
- $\mathfrak{a} \in \mathfrak{F}^\Gamma(\mathfrak{b})$ if and only if $\forall\varphi : \blacksquare\varphi \in \mathfrak{a}$ implies $\varphi \in \mathfrak{b}$,
- $+^\Gamma(\mathfrak{a}) = \{p \in \mathsf{Top} \mid p^+ \in \mathfrak{a}\}$,
- $-^\Gamma(\mathfrak{a}) = \{p \in \mathsf{Top} \mid p^- \in \mathfrak{a}\}$,
- $\mathfrak{V}^\Gamma(i) = \{\mathfrak{a} \in \mathfrak{A}^\Gamma \mid i \in \mathfrak{a}\}$,
- $(\mathfrak{a}, \mathfrak{b}) \in \mathfrak{R}^\Gamma(p, \mathfrak{c})$ if and only if $\forall\varphi : \Box_{i:p}\varphi \in \mathfrak{a}$ implies $\varphi \in \mathfrak{b}$, where $\mathfrak{V}^\Gamma(i) = \{\mathfrak{c}\}$.

There are several properties that we need to check in order to ensure that $\mathfrak{M}^\Gamma$ is indeed a visibility model. First, observe that by items (1) and (2) of Lemma 14, each nominal names a unique agent in $\mathfrak{A}^\Gamma$. Next, no agent is *pro* and *contra* the same topic. This follows from the fact that each agent $\mathfrak{a} \in \mathfrak{A}^\Gamma$ satisfies all the instances of axiom $\textsc{Cons}$.

The followership and visibility relations of visibility models are bidirectional, which is manifested by the presence of converse modalities in $\mathbb{SVL}$. In the next two lemmas we argue that the corresponding relations of the canonical model $\mathfrak{M}^\Gamma$ are also bidirectional.

Lemma 18. The following definitions of $\mathfrak{a} \in \mathfrak{F}^\Gamma(\mathfrak{b})$ are equivalent for all φ:

1. $\blacksquare\varphi \in \mathfrak{a}$ implies $\varphi \in \mathfrak{b}$
2. $\blacksquare^{-1}\varphi \in \mathfrak{b}$ implies $\varphi \in \mathfrak{a}$
3. $\varphi \in \mathfrak{b}$ implies $\blacklozenge\varphi \in \mathfrak{a}$
4. $\varphi \in \mathfrak{a}$ implies $\blacklozenge^{-1}\varphi \in \mathfrak{b}$

Proof. From (1) to (2). Assume that (1) holds. We show (2) by contraposition, i.e. we demonstrate that $\varphi \notin \mathfrak{a}$ implies $\blacksquare^{-1}\varphi \notin \mathfrak{b}$. Let $\varphi \notin \mathfrak{a}$. Since $\mathfrak{a}$ is an MCS, $\varphi \notin \mathfrak{a}$ if and only if $\neg\varphi \in \mathfrak{a}$. By $\mathrm{T}_{\blacksquare\blacklozenge^{-1}}$ and MP, we also have that $\blacksquare\blacklozenge^{-1}\neg\varphi \in \mathfrak{a}$. By (1), the latter implies that $\blacklozenge^{-1}\neg\varphi \in \mathfrak{b}$, which is equivalent to $\neg\blacksquare^{-1}\varphi \in \mathfrak{b}$. Since $\mathfrak{b}$ is an MCS, it holds that $\blacksquare^{-1}\varphi \notin \mathfrak{b}$. *From (2) to (1).* Similar to above using $\mathrm{T}_{\blacksquare^{-1}\blacklozenge}$.

From (3) to (4). Assume that (3) holds. We show that $\blacklozenge^{-1}\varphi \notin \mathfrak{b}$ implies $\varphi \notin \mathfrak{a}$. To this end, let $\blacklozenge^{-1}\varphi \notin \mathfrak{b}$, which is equivalent to $\blacksquare^{-1}\neg\varphi \in \mathfrak{b}$. By (3), the latter implies that $\blacklozenge\blacksquare^{-1}\neg\varphi \in \mathfrak{a}$. From $\blacklozenge\blacksquare^{-1}\neg\varphi \in \mathfrak{a}$ and contraposition of $\mathrm{T}_{\blacksquare\blacklozenge^{-1}}$ we get $\neg\varphi \in \mathfrak{a}$ by MP. Since $\mathfrak{a}$ is an MCS, $\neg\varphi \in \mathfrak{a}$ is equivalent to the fact that $\varphi \notin \mathfrak{a}$. *From (4) to (3).* Similar to above using $\mathrm{T}_{\blacksquare^{-1}\blacklozenge}$.

(1) is equivalent to (3). By taking a contraposition of (1). □

Lemma 19. The following definitions of $(\mathfrak{a}, \mathfrak{b}) \in \mathfrak{R}^\Gamma(p, \mathfrak{c})$ with $\mathfrak{V}^\Gamma(i) = \{\mathfrak{c}\}$ are equivalent for all φ:

1. $\Box_{i:p}\varphi \in \mathfrak{a}$ implies $\varphi \in \mathfrak{b}$
2. $\Box^{-1}_{i:p}\varphi \in \mathfrak{b}$ implies $\varphi \in \mathfrak{a}$
3. $\varphi \in \mathfrak{b}$ implies $\Diamond_{i:p}\varphi \in \mathfrak{a}$
4. $\varphi \in \mathfrak{a}$ implies $\Diamond^{-1}_{i:p}\varphi \in \mathfrak{b}$

Proof. The proof is similar to the proof of Lemma 18 using $\mathrm{T}_{\Box\Diamond^{-1}}$ and $\mathrm{T}_{\Box^{-1}\Diamond}$. □

Finally, we show that the followership and visibility relations of $\mathfrak{M}^\Gamma$ satisfy the properties from Definition 2, i.e. that $\mathfrak{M}^\Gamma$ is indeed a visibility model. But before we delve into the proof *per se*, we also mention that the following schema is a theorem of **SVL**:

$$\textsc{Elim} \quad i \wedge @_i\varphi \to \varphi$$

Elim can be derived from the contraposition of Intro using Self-Dual.

Lemma 20. Model $\mathfrak{M}^\Gamma$ has an irreflexive followership relation, and its visibility relation satisfies (1)-(5) from Definition 2.

Proof. Irreflexivity. Assume towards a contradiction that there is $\mathfrak{a} \in \mathfrak{A}^\Gamma$ such that $\mathfrak{a} \in \mathfrak{F}^\Gamma(\mathfrak{a})$. Since $\mathfrak{a}$ is a named MCS, there is an $i \in \mathsf{Nom}$ such that $\mathfrak{V}^\Gamma(i) = \{\mathfrak{a}\}$. Moreover, $\mathfrak{a}$ contains all the instances of Irref, and in particular $@_i\neg\blacklozenge i \in \mathfrak{a}$, which is equivalent to $@_i\blacksquare\neg i \in \mathfrak{a}$. From $i \in \mathfrak{a}$ and $@_i\blacksquare\neg i \in \mathfrak{a}$, we can conclude $\blacksquare\neg i \in \mathfrak{a}$ by Elim and MP. By the definition of $\mathfrak{F}^\Gamma$, $\mathfrak{a} \in \mathfrak{F}^\Gamma(\mathfrak{a})$ if and only if $\forall\varphi : \blacksquare\varphi \in \mathfrak{a}$ implies $\varphi \in \mathfrak{a}$. Thus, $\blacksquare\neg i \in \mathfrak{a}$ implies $\neg i \in \mathfrak{a}$, which contradicts $i \in \mathfrak{a}$ and $\mathfrak{a}$ being consistent.

(1). We need to show that if $(\mathfrak{a}, \mathfrak{b}) \in \mathfrak{R}^\Gamma(p, \mathfrak{c})$, then $(\mathfrak{a}, \mathfrak{a}) \in \mathfrak{R}^\Gamma(p, \mathfrak{c})$, where $\mathfrak{a}, \mathfrak{b}, \mathfrak{c} \in \mathfrak{A}^\Gamma$. Thus, assume that $(\mathfrak{a}, \mathfrak{b}) \in \mathfrak{R}^\Gamma(p, \mathfrak{c})$, and let $\mathfrak{V}^\Gamma(i) = \{\mathfrak{a}\}$ and $\mathfrak{V}^\Gamma(j) = \{\mathfrak{c}\}$. Since a is an MCS, it contains all instances of Ag-See, and in particular $@_i(\Diamond_{j:p}\top \to \Diamond_{j:p}i) \in \mathfrak{a}$. By Elim and MP we further have that $\Diamond_{j:p}\top \to \Diamond_{j:p}i \in \mathfrak{a}$. From the fact that $\top \in \mathfrak{b}$ and

$(\mathfrak{a},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$, we have $\lozenge_{j:p}\top \in \mathfrak{a}$. The latter implies that $\lozenge_{j:p}i \in \mathfrak{a}$ by MP. From $i \in \mathfrak{a}$ and $\lozenge_{j:p}i \in \mathfrak{a}$ we conclude, by Lemma 19, that $(\mathfrak{a},\mathfrak{a}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$.

(2). We show that if $(\mathfrak{a},\mathfrak{a}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$, then $(\mathfrak{a},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$ for all $\mathfrak{b} \in \mathfrak{F}^\Gamma(\mathfrak{a})$. Let $\mathfrak{V}^\Gamma(i) = \{\mathfrak{a}\}$, $\mathfrak{V}^\Gamma(j) = \{\mathfrak{c}\}$, and $\mathfrak{b} \in \mathfrak{A}^\Gamma$ be an arbitrary agent such that $\mathfrak{b} \in \mathfrak{F}^\Gamma(\mathfrak{a})$. Moreover, let us assume $(\mathfrak{a},\mathfrak{a}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$. Since $\mathfrak{a}$ is an MCS, it contains the following instance of Fol-See: $@_i(\lozenge_{j:p}i \to \blacksquare^{-1}\lozenge^{-1}_{j:p}i)$. We can use Elim to get $\lozenge_{j:p}i \to \blacksquare^{-1}\lozenge^{-1}_{j:p}i \in \mathfrak{a}$. By the third item of Lemma 19, $(\mathfrak{a},\mathfrak{a}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$ is equivalent to the fact that if $i \in \mathfrak{a}$, then $\lozenge_{j:p}i \in \mathfrak{a}$. We obtain the latter from $i \in \mathfrak{a}$ and MP. Furthermore, from $\lozenge_{j:p}i \to \blacksquare^{-1}\lozenge^{-1}_{j:p}i \in \mathfrak{a}$ and $\lozenge_{j:p}i \in \mathfrak{a}$ we derive $\blacksquare^{-1}\lozenge^{-1}_{j:p}i \in \mathfrak{a}$. From $\mathfrak{b} \in \mathfrak{F}^\Gamma(\mathfrak{a})$ and $\blacksquare^{-1}\lozenge^{-1}_{j:p}i \in \mathfrak{a}$, we have, by item (2) of Lemma 18, $\lozenge^{-1}_{j:p}i \in \mathfrak{b}$. Finally, the fact that $i \in \mathfrak{a}$ and $\lozenge^{-1}_{j:p}i \in \mathfrak{b}$ is equivalent to $(\mathfrak{a},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$ by item (4) of Lemma 19.

(3). We need to show that if $(\mathfrak{a},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$, $p \in +^\Gamma(\mathfrak{b})$, and $\mathfrak{b} \neq \mathfrak{c}$, then $(\mathfrak{b},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$ and $\mathfrak{b} \in \mathfrak{F}^\Gamma(\mathfrak{c})$. Assume that it holds that $(\mathfrak{a},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$, $p \in +^\Gamma(\mathfrak{b})$, and $\mathfrak{b} \neq \mathfrak{c}$, and let $\mathfrak{V}^\Gamma(i) = \{\mathfrak{b}\}$ and $\mathfrak{V}^\Gamma(j) = \{\mathfrak{c}\}$. As $\mathfrak{b} \in \mathfrak{A}^\Gamma$ and thus is an MCS, it contains the following instance of Foll:$@_i((p^+ \wedge \lozenge^{-1}_{j:p}\top \wedge \neg j) \to (\lozenge_{j:p}i \wedge \blacklozenge j))$. From $i \in b$ by Elim and MP we obtain $(p^+ \wedge \lozenge^{-1}_{j:p}\top \wedge \neg j) \to (\lozenge_{j:p}i \wedge \blacklozenge j) \in \mathfrak{b}$. The truth of the antecedent follows from the assumptions that $(\mathfrak{a},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$, $p \in +^\Gamma(\mathfrak{b})$, $\mathfrak{b} \neq \mathfrak{c}$, and $\mathfrak{V}^\Gamma(j) = \{\mathfrak{c}\}$. Hence, $\lozenge_{j:p}i \wedge \blacklozenge j \in \mathfrak{b}$, or, equivalently, $\lozenge_{j:p}i \in \mathfrak{b}$ and $\blacklozenge j \in \mathfrak{b}$. From $\lozenge_{j:p}i \in \mathfrak{b}$ and $i \in \mathfrak{b}$ by item (4) of Lemma 19, we have that $(\mathfrak{b},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$. Similarly, from $\blacklozenge j \in \mathfrak{b}$ and $j \in \mathfrak{c}$ by item (3) of Lemma 18, we obtain $\mathfrak{b} \in \mathfrak{F}^\Gamma(\mathfrak{c})$.

(4). Assume that for some $\mathfrak{a},\mathfrak{b},\mathfrak{c} \in \mathfrak{A}^\Gamma$, we have that $(\mathfrak{a},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$, $p \in -^\Gamma(\mathfrak{b})$, $\mathfrak{a} \neq \mathfrak{b}$, and $\mathfrak{V}^\Gamma(i) = \{\mathfrak{b}\}$, $\mathfrak{V}^\Gamma(j) = \{\mathfrak{c}\}$, and $\mathfrak{V}^\Gamma(k) = \{\mathfrak{a}\}$. We need to show that all of these imply $(\mathfrak{b},\mathfrak{b}) \notin \mathfrak{R}^\Gamma(p,\mathfrak{c})$ and $\mathfrak{b} \notin \mathfrak{F}^\Gamma(\mathfrak{a})$. Assume towards a contradiction that $(\mathfrak{b},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$ or $\mathfrak{b} \in \mathfrak{F}^\Gamma(\mathfrak{a})$. First, let $(\mathfrak{b},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$. Since $\mathfrak{b}$ is an MCS it contains the following instance of Unfoll: $@_i((p^- \wedge \lozenge^{-1}_{j:p}(\neg i \wedge k)) \to (\neg\lozenge_{j:p}\top \wedge \neg\blacklozenge k))$. From $i \in \mathfrak{b}$ and Elim we have $(p^- \wedge \lozenge^{-1}_{j:p}(\neg i \wedge k)) \to (\neg\lozenge_{j:p}\top \wedge \neg\blacklozenge k) \in \mathfrak{b}$. Our assumptions imply the truth of the antecedent. Hence, $\neg\lozenge_{j:p}\top \wedge \neg\blacklozenge k \in \mathfrak{b}$. At the same time, $(\mathfrak{b},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$ is equivalent to $\top \in \mathfrak{b}$ implies $\lozenge_{j:p}\top \in \mathfrak{b}$, by item (3) of Lemma 19. Since $\mathfrak{b}$ is an MCS, we have that $\top \in \mathfrak{b}$, and hence $\lozenge_{j:p}\top \in \mathfrak{b}$, which contradicts $\neg\lozenge_{j:p}\top \in \mathfrak{b}$. Now, let $\mathfrak{b} \in \mathfrak{F}^\Gamma(\mathfrak{a})$. From $k \in \mathfrak{a}$ by item (3) of Lemma 18 we conclude that $\blacklozenge k \in \mathfrak{b}$, which contradicts $\neg\blacklozenge k \in \mathfrak{b}$.

(5). Assume that for some $\mathfrak{a},\mathfrak{b},\mathfrak{c} \in \mathfrak{A}^\Gamma$, we have that $(\mathfrak{a},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$, $p \in +^\Gamma(\mathfrak{b})$, $p \in -^\Gamma(\mathfrak{b})$, $\mathfrak{a} \neq \mathfrak{b}$, and $\mathfrak{V}^\Gamma(i) = \{\mathfrak{b}\}$, $\mathfrak{V}^\Gamma(j) = \{\mathfrak{c}\}$, and $\mathfrak{V}^\Gamma(k) = \{\mathfrak{a}\}$. We demonstrate that $(\mathfrak{b},\mathfrak{b}) \notin \mathfrak{R}^\Gamma(p,\mathfrak{c})$. Assume towards a contradiction that $(\mathfrak{b},\mathfrak{b}) \in \mathfrak{R}^\Gamma(p,\mathfrak{c})$. Since $\mathfrak{b}$ is an MCS, $\top \in \mathfrak{b}$, and by item (3) of Lemma 19, we thus have $\lozenge_{j:p}\top \in \mathfrak{b}$. At the same time, b contains all the instances of Indiff, and in particular, $@_i((\neg p^- \wedge \neg p^+ \wedge \lozenge^{-1}_{j:p}\neg i) \to \neg\lozenge_{j:p}\top) \in \mathfrak{b}$. By $i \in \mathfrak{b}$ and Elim we have $(\neg p^- \wedge \neg p^+ \wedge \lozenge^{-1}_{j:p}\neg i) \to \neg\lozenge_{j:p}\top \in \mathfrak{b}$. The antecedent follows from our assumption. Hence, $\neg\lozenge_{j:p}\top \in \mathfrak{b}$, which contradicts the aforeshown $\lozenge_{j:p}\top \in \mathfrak{b}$. □

Lemma 20 thus establishes that our canonical model $\mathfrak{M}^\Gamma$ is indeed a visibility model.

Lemma 21. Let Γ be a named and pasted MCS, and let $\mathfrak{M}^\Gamma = (\mathfrak{A}^\Gamma, \mathfrak{F}^\Gamma, +^\Gamma, -^\Gamma, \mathfrak{V}^\Gamma, \mathfrak{R}^\Gamma)$ be the named and pasted model yielded by Γ. Then the following holds.

1. Let $\Diamond_{j:p}\varphi \in \mathfrak{a}$ and $\mathfrak{a} \in \mathfrak{A}^\Gamma$. Then there is a $\mathfrak{b} \in \mathfrak{A}^\Gamma$, such that $(\mathfrak{a}, \mathfrak{b}) \in \mathfrak{R}(p, \mathfrak{c})$, $\mathfrak{V}^\Gamma(j) = \{\mathfrak{c}\}$, and $\varphi \in \mathfrak{b}$.
2. Let $\Diamond_{j:p}^{-1}\varphi \in \mathfrak{b}$ and $\mathfrak{b} \in \mathfrak{A}^\Gamma$. Then there is an $\mathfrak{a} \in \mathfrak{A}^\Gamma$, such that $(\mathfrak{a}, \mathfrak{b}) \in \mathfrak{R}(p, \mathfrak{c})$, $\mathfrak{V}^\Gamma(j) = \{\mathfrak{c}\}$, and $\varphi \in \mathfrak{a}$.
3. Let $\blacklozenge\varphi \in \mathfrak{a}$ and $\mathfrak{a} \in \mathfrak{A}^\Gamma$. Then there is a $\mathfrak{b} \in \mathfrak{A}^\Gamma$, such that $\mathfrak{a} \in \mathfrak{F}(\mathfrak{b})$ and $\varphi \in \mathfrak{b}$.
4. Let $\blacklozenge^{-1}\varphi \in \mathfrak{b}$ and $\mathfrak{b} \in \mathfrak{A}^\Gamma$. Then there is an $\mathfrak{a} \in \mathfrak{A}^\Gamma$, such that $\mathfrak{a} \in \mathfrak{F}(\mathfrak{b})$ and $\varphi \in \mathfrak{a}$.

Proof. We prove only item (3) and all other items can be shown analogously. Assume that $\blacklozenge\varphi \in \mathfrak{a}$ and $\mathfrak{V}^\Gamma(i) = \{\mathfrak{a}\}$. By Intro we thus have that $@_i\blacklozenge\varphi \in \mathfrak{a}$. Since Γ is pasted, $@_i\blacklozenge\varphi \in \mathfrak{a}$ implies that there is a $j \in \mathsf{Nom}$ such that $@_i\blacklozenge j \wedge @_j\varphi \in \Gamma$. Let $\mathfrak{b}$ be the MCS with $j \in \mathfrak{b}$. By Elim we thus have that $\blacklozenge j \in \mathfrak{a}$ and $\varphi \in \mathfrak{b}$. From $j \in \mathfrak{b}$ and $\blacklozenge j \in \mathfrak{a}$ we conclude, by item (3) of Lemma 18, that $\mathfrak{a} \in \mathfrak{F}(\mathfrak{b})$. □

With the next lemma, usually called in the literature *The Truth Lemma*, we establish the equivalence between membership of a formula in some $\mathfrak{a} \in \mathfrak{A}^\Gamma$ and truth of the formula at agent $\mathfrak{a}$ of the canonical model.

Lemma 22. Let $\mathfrak{M}^\Gamma = (\mathfrak{A}^\Gamma, \mathfrak{F}^\Gamma, +^\Gamma, -^\Gamma, \mathfrak{V}^\Gamma, \mathfrak{R}^\Gamma)$ be the named and pasted model yielded by a named and posted MCS Γ, and let $\mathfrak{a} \in \mathfrak{A}^\Gamma$. Then for all φ, $\varphi \in \mathfrak{a}$ if and only if $\mathfrak{M}^\Gamma_\mathfrak{a} \models \varphi$.

Proof. The proof is by induction of φ. Cases of topics and nominals follow directly from Definition 17, and Boolean cases are straightforward. Among the modal cases, we show just $\blacklozenge\varphi$ and other modal cases can be proved similarly by using the appropriate items of Lemmas 18, 19, and 21.

Case $\varphi = \blacklozenge\psi$. From left to right. Let $\blacklozenge\psi \in \mathfrak{a}$. By item (3) of Lemma 21, this implies that there is a $\mathfrak{b} \in \mathfrak{A}^\Gamma$ such that $\mathfrak{a} \in \mathfrak{F}(\mathfrak{b})$ and $\psi \in \mathfrak{b}$. By the Induction Hypothesis, the latter is equivalent to the fact that there is a $\mathfrak{b} \in \mathfrak{A}^\Gamma$ such that $\mathfrak{a} \in \mathfrak{F}(\mathfrak{b})$ and $\mathfrak{M}^\Gamma_\mathfrak{b} \models \psi$, which, in turn, is equivalent to $\mathfrak{M}^\Gamma_\mathfrak{a} \models \blacklozenge\psi$ by the definition of semantics.

From right to left. Assume that $\mathfrak{M}^\Gamma_\mathfrak{a} \models \blacklozenge\psi$. By the definition of semantics, this is equivalent to the fact that there is a $\mathfrak{b} \in \mathfrak{A}^\Gamma$ such that $\mathfrak{a} \in \mathfrak{F}(\mathfrak{b})$ and $\mathfrak{M}^\Gamma_\mathfrak{b} \models \psi$. By the Induction Hypothesis, the latter is equivalent to $\mathfrak{a} \in \mathfrak{F}(\mathfrak{b})$ and $\psi \in \mathfrak{b}$, which implies $\blacklozenge\psi \in \mathfrak{a}$ by item (3) of Lemma 18.

Case $\varphi = @_i\psi$. Assume that $\mathfrak{M}^\Gamma_\mathfrak{a} \models @_i\psi$. By the definition of semantics, this is equivalent to $\mathfrak{M}^\Gamma_\mathfrak{b} \models \psi$ for a $\mathfrak{b} \in \mathfrak{A}^\Gamma$ such that $\mathfrak{V}^\Gamma(i) = \{\mathfrak{b}\}$. By the Induction Hypothesis, $\mathfrak{M}^\Gamma_\mathfrak{b} \models \psi$ if and only if $\psi \in \mathfrak{b}$, which implies $@_i\psi \in \mathfrak{b}$ by Intro and MP. By item (3) of Lemma 14, we have that $@_i\psi \in \mathfrak{a}$. The other direction is similar by using Elim instead of Intro. □

We finally have all the ingredients in order to demonstrate that the axiom system **SVL** is complete for the class of visibility models.

Theorem 23. Every consistent set of formulas is satisfiable on a visibility model.

Proof. Let Σ be a consistent set of formulas. By Lemma 16, Σ can be extended to a named and pasted MCS Γ in $\mathbb{SVL}^*$. Let $\mathfrak{M}^\Gamma = (\mathfrak{A}^\Gamma, \mathfrak{F}^\Gamma, +^\Gamma, -^\Gamma, \mathfrak{V}^\Gamma, \mathfrak{R}^\Gamma)$ be the named and pasted model yielded by Γ. Since Γ is named, by item (4) of Lemma 14 and the definition of $\mathfrak{M}^\Gamma$, $\Gamma \in \mathfrak{A}^\Gamma$. By Lemma 22 and the fact that $\Sigma \subseteq \Gamma$, the latter is equivalent to $\mathfrak{M}^\Gamma_\Gamma \models \varphi$ for all $\varphi \in \Sigma$. □

4 Visibility Logic

To reason about the effects of agents posting on various topics, we introduce a dynamic extension of SVL that we call *visibility logic* (VL). Compared to SVL, VL is enriched with dynamic operators $[\pi]\varphi$, where π is an action of the current agent making a post. While defining VL, we follow *dynamic epistemic logics* (DELs) [17], and in particular *action model logic* [6, 17]. We begin with a motivating example.

4.1 Example: Taking the Advantage to be Seen by Many

In some networks, the best tactic for exposing more agents to a controversial opinion is to first post on a popular topic. Consider the follower-network M in Figure 2. For simplicity, we do not include nominals in the Figure. This network consists of 6 agents named in alphabetical order from a to f. Agent a has two followers b and c. Agent b has three followers d, e and f. We assume that agents d, e and f might have some followers that we do not have information about, noted in the figure with dots. Furthermore, agent a is positive in favor of vaccination (abbreviated v in the figure) which is a controversial topic amongst the agents: all agents have an opinion about vaccines and three of the agents a, c and f are *pro* vaccination, whereas b, d and e are *contra* vaccination. The topic of dogs (abbreviated d) on the other hand, is widely liked. All agents like dogs, except a who is indifferent: $d \notin +(a)$ and $d \notin -(a)$.

Imagine that agent a wants to post on vaccines, and wants as many as possible of the other agents in the network to see the post. We show that the best tactic for agent a is to first post on dogs, even though a is indifferent about dogs, and then later post on vaccines. Consider first the scenario in Figure 3 where agent a posts v from the outset. An update happens in two steps. First, we add visibility arrows corresponding to posting and resharing. In the second step, we update the followership relation based on whether the agents who have seen a post are *pro* or *contra* the post. The resulting update is $M^{a:v}$ in Figure 3, where agents b and c have seen a's vaccine post, and only c remains as a's follower.

Then, consider instead the situation where agent a posts on dogs in the update $M^{a:d}$ in Figure 4 before posting on vaccines in the update $M^{a:d,a:v}$ in Figure 5. Note that to make the situation easier to read, we omit the followership arrows in the visibility update and the visibility arrows in the followership update in both figures. After $M^{a:d}$, all agents have seen agent a's post on dogs. Since they like dogs, all agents also follow a after the update.

In $M^{a:d,a:v}$, we see the results after agent a first posted on dogs, and then vaccines. All

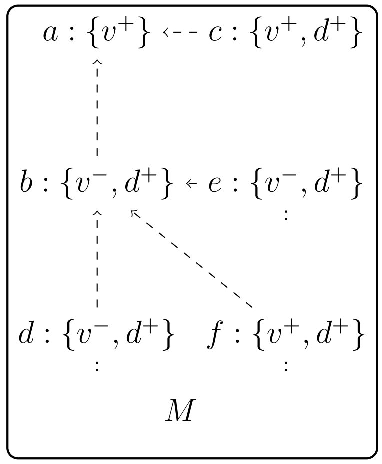

Figure 2: A follower-network M where vaccination is a controversial topic.

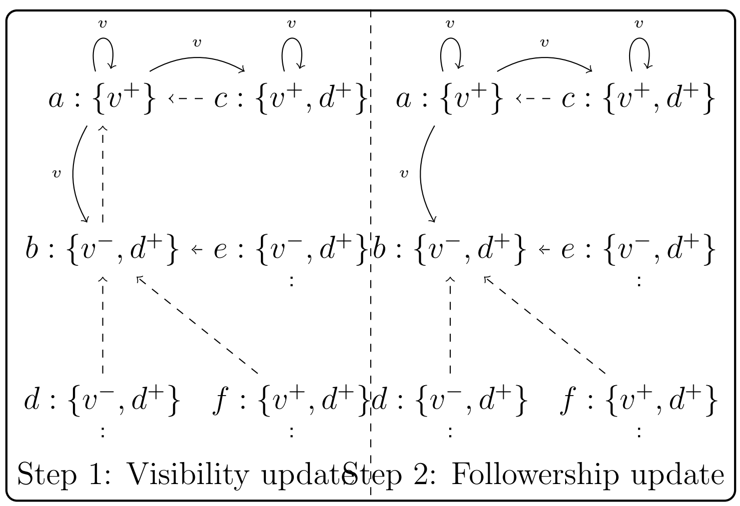

Figure 3: Update $M^{a:v}$ after agent a posts in favor of vaccines.

the agents have now seen a's vaccine-post. Most of them did not like it and unfollowed a, but only after they were exposed to the post. Interestingly, we also notice that agent f, who was not originally a follower of a in the initial network outset, now follows a and has shared the vaccine post to their followers.

There are two tactical reasons for agent a to post on dogs before their more controversial post on vaccines. Firstly, a larger portion of the agents now saw a's post on vaccines since they followed agent a after the dog post. Secondly, a has been able to reach out and expand their network: agent f who is also pro vaccines, has shared the vaccination post to their, for us unknown, followers.

The reason behind a phenomenon such as this is directly connected to an underlying

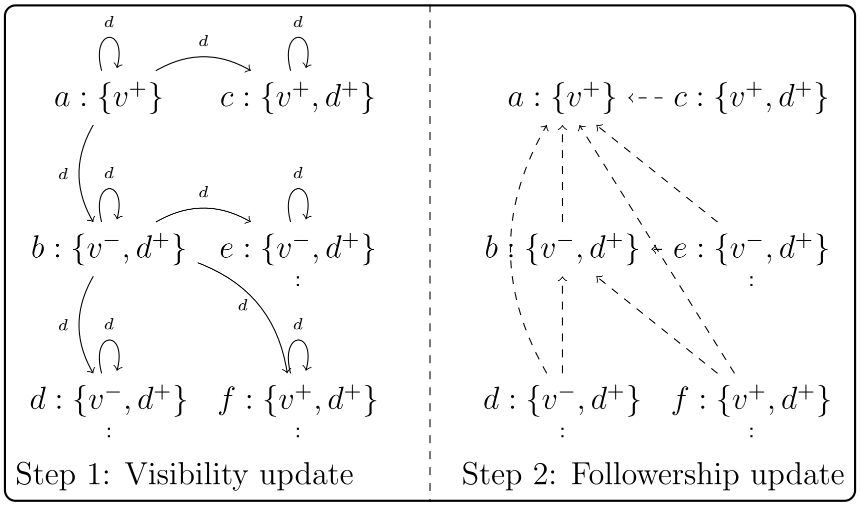

Figure 4: Update $M^{a:d}$ after agent a posts on dogs.

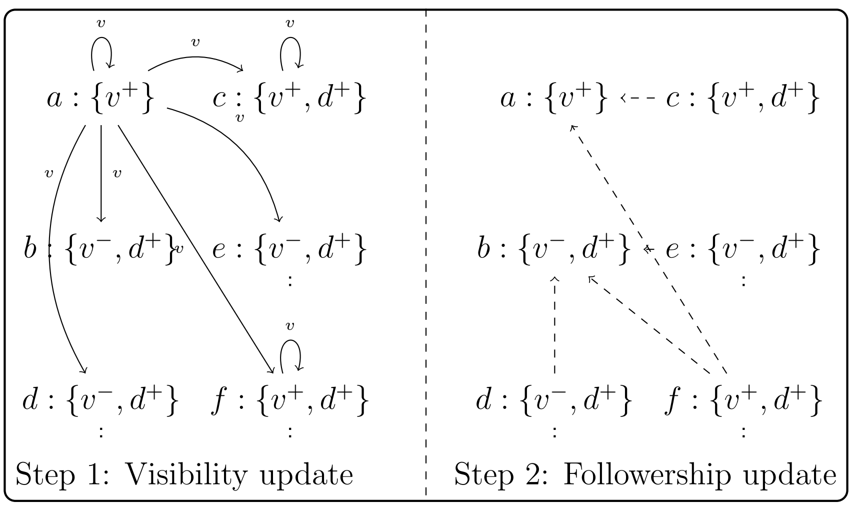

Figure 5: Update $M^{a:d,a:v}$ after agent a posts on vaccination after first posting on dogs.

notion of trust between agents in the network. In our setting, agents follow other agents when the former is exposed to content that they like by the latter. In the example, we imagine agents likely followed a because they wanted to see more dog-friendly content. Agent a misuses the trust of their followers by pretending to be interested in dogs before posting on vaccination.

What becomes clear in this example, is that in our simplified setting of posting and sharing in a social network, the interests of an agent's followers matter a lot to what information is shared and seen. Secondly, the system is vulnerable to exploitation by a potentially malicious agent: there are opportunities to tactically post on popular topics to later expose more agents to a controversial opinion. To reason about dynamic situations such as these, we introduce VL.

4.2 Language, Semantics, and Logical Properties of VL

The language of VL is an extension of the language of SVL.

Definition 24 (Syntax)**.** *The language of visibility logic* $\mathbb{VL}$ is defined recursively by the following BNF:

$$\begin{array}{lcl} \varphi & ::= & p^+ \mid p^- \mid i \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \lozenge_{i:p}\varphi \mid \lozenge_{i:p}^{-1}\varphi \mid \blacklozenge\varphi \mid \blacklozenge^{-1}\varphi \mid @_i\varphi \mid [\pi]\varphi \\ \pi & ::= & p \mid (\pi \cup \pi) \end{array}$$

where $[\pi]\varphi$ is read 'after the current agent executes action π, φ holds'.

Union of actions $(\pi \cup \tau)$ was inherited by DELs from propositional dynamic logic [20], and in the context of visibility formulas, $[p \cup q]\varphi$ mean 'whichever topic the current agent posts on, p or q, φ will be true (in both cases)'.

Given a formula $\varphi \in \mathbb{VL}$, we define *modal depth* and the *size* of the formula similarly to the corresponding definitions for $\mathbb{SVL}$ with the following additional cases: $md([\pi]\varphi) = md(\varphi)$, and $|[\pi]\varphi| = |\pi| + |\varphi| + 1$, where $|\pi \cup \tau| = |\pi| + |\tau| + 1$.

Definition 25 (Semantics)**.** Let $M = (A, F, +, -, V, R)$ be a visibility model, $a \in A$, and $p, q \in \mathsf{Top}$. *The semantics of VL* is the same as in Definition 3 with the following additions:

$$\begin{array}{lll} M_a \models [p]\varphi & \text{iff} & M_a^{a:p} \models \varphi \\ M_a \models [\pi \cup \tau]\varphi & \text{iff} & M_a \models [\pi]\varphi \text{ and } M_a \models [\tau]\varphi \\ M_a \models \langle \pi \cup \tau \rangle\varphi & \text{iff} & M_a \models \langle \pi \rangle\varphi \text{ or } M_a \models \langle \tau \rangle\varphi \end{array}$$

where $M_a^{a:p}$ is defined in the following two steps. First, let $M^* = (A,F,+,-,\ V,\ R^*)$, where $R^*(p, a)$ is the least fixed point of function $f : 2^{A\times A} \to 2^{A\times A}$ defined as

$$\begin{aligned} f(X) =& X \cup \{(a, a)\} \cup \{(b, c) \mid (b, b) \in X \text{ and } c \in F(b)\} \cup \\ & \cup \{(c, c) \mid p \in +(c) \text{ and } \exists b : (b, c) \in X\}. \end{aligned}$$

Informally, intermediate model M^* differs from M only in R in such a way that R^* now contains the fact that a has posted on p, that her post has reached all her followers, and that all followers who are *pro* p reshare the post further to their followers. Secondly, we construct $M^{a:p}$ out of M^* by updating F:

1. $F^{a:p}(a) = F(a) \cup \{b\}$, if $a \neq b$, $p \in +(b)$, and $\exists c : (c, b) \in R^*(p, a)$,

2. $F^{a:p}(b) = F(b) \setminus \{c\}$, if $p \in -(b)$ and $(c, b) \in R^*(p, a)$.

Intuitively, agent b will follow the original poster a if she has seen the post, maybe not even from a, and if she is *pro* the topic. Agent c will stop following anyone from whom she has seen a post on a topic she dislikes.

Returning to the example in the beginning of this section, we can use the new operators to reason about the situation with formulas that hold in the models given in the example. For simplicity, we name the agents with nominals corresponding to their labels in the model. That is, $M_a \vDash a$, $M_b \vDash b$ and so on, where $a, b \in \mathsf{Nom}$. To avoid confusion with the agent named d, we denote the topics "vaccines" and "dogs" with *vacc* and *dogs* instead of v and d as used in the example. The following formulas, with their respective intuitive readings, hold in the initial model M in Figure 2:

- $M_a \vDash [vacc]\blacksquare^{-1}c$

  "After posting on vaccines, a will have only one follower, c."

- $M_a \vDash [vacc](\Diamond_{a:vacc}b \wedge \Diamond_{a:vacc}c \wedge \Box_{a:vacc}(a \vee b \vee c))$

  "After posting on vaccines, only b and c, except for a, will have seen a's post."

- $M_a \vDash [dogs](\blacklozenge^{-1}b \wedge \blacklozenge^{-1}c \wedge \blacklozenge^{-1}d \wedge \blacklozenge^{-1}e \wedge \blacklozenge^{-1}f)$

  "After posting on dogs, all agents from b to f will follow a."

- $M_a \vDash [dogs][vaccines](\Diamond_{a:vacc}b \wedge \Diamond_{a:vacc}c \wedge \Diamond_{a:vacc}d \wedge \Diamond_{a:vacc}e \wedge \Diamond_{a:vacc}f)$

  "After a first posts on dogs, and then on vaccines, all agents from a to f will have seen a's post on vaccines."

To give a further taste of VL, let us provide some properties that are valid or not valid on visibility models. All the validities can be shown by an application of the definition of the semantics.

Proposition 26. Let $p, q \in \mathsf{Top}$ and $\varphi \in \mathbb{VL}$.

1. $\neg[p]\varphi \leftrightarrow [p]\neg\varphi$ is valid.
2. $[\pi \cup \tau]\varphi \leftrightarrow [\pi]\varphi \wedge [\tau]\varphi$ is valid.
3. $\Diamond^{-1}_{i:p}\varphi \leftrightarrow [q]\Diamond^{-1}_{i:p}\varphi$ is valid.
4. $[p]\varphi \rightarrow [p][p]\varphi$ is not valid.
5. $[p][p]\varphi \rightarrow [p]\varphi$ is not valid.

The first formula states that the operator of posting on a topic is its own dual. The second property shows how to eliminate non-deterministic choice. The third item claims that once an agent has seen a post of an agent with the name i on topic p, no further post can revoke this. The fact that formulas four and five are not valid indicates that consecutive posting on the same topic yields different results. A counterexample showing this would include the current agent with the name i posting on p and gaining new followers. Additional posts on the same topic by the same agent will add new p-arrows to those new followers thus resulting in a different updated model that is not guaranteed to satisfy φ.

4.3 Expressivity and Model Checking

Now, we state that VL is more expressive than its static fragment SVL. This result is quite interesting, since many of DELs, for example public announcement logic [42], arrow update logic [28], and action model logic [17, Chapter 6], are equally expressive as the static logic they are built upon. Those expressivity results are usually obtained with the use of so-called reduction axioms that allow one to equivalently rewrite formulas of dynamic extensions to formulas of the static fragment. Thus, the fact that VL is more expressive than SVL also entails that no reduction axioms for VL are possible.

Theorem 27. $\mathbb{SVL} < \mathbb{VL}$.

Proof. Consider a $\mathbb{VL}$ formula $[p]\blacklozenge^{-1}\blacksquare^{-1}\bot$, and assume towards a contradiction that there is an equivalent formula ψ of $\mathbb{SVL}$ with $md(\psi) = n$. Since ψ has a finite size, there is a set of nominals $Q = \{j_1, \dots, j_{n+1}\}$ that are not present in ψ.

Consider models M and N in Figure 6. The models are chains of length $n + 2$ that start with agent a and with each next agent following the previous one. The only difference between the models is that the last agent in the chain in model M is *pro* topic p, and the last agent in the chain in model N is neither *pro* nor *contra* topic p.

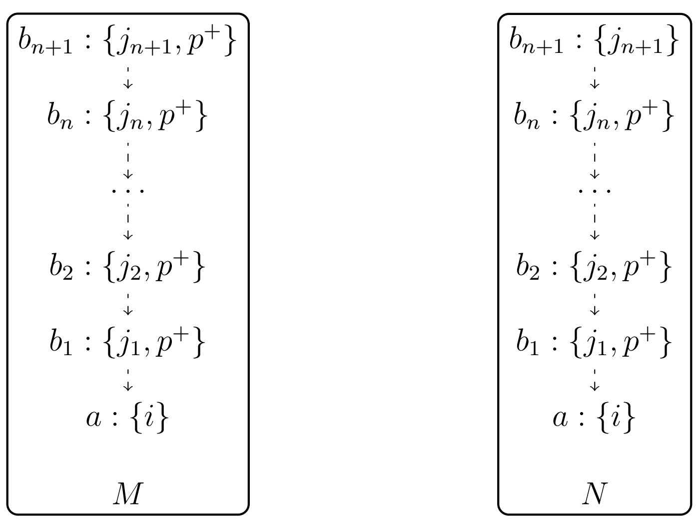

Figure 6: Models M and N.

Now we will argue that $[p]\blacklozenge^{-1}\blacksquare^{-1}\bot$ distinguishes M_a and N_a. In particular, $M_a \models [p]\blacklozenge^{-1}\blacksquare^{-1}\bot$ and $N_a \not\models [p]\blacklozenge^{-1}\blacksquare^{-1}\bot$. Indeed, agent a posting on topic p results in the updated visibility model $M_a^{a:p}$ presented in Figure 7. In the updated model, it holds that $M_{b_{n+1}}^{a:p} \models \blacklozenge i$, i.e. that agent b_{n+1} follows agent a. Moreover, agent b_{n+1} does not have any followers, so $M_{b_{n+1}}^{a:p} \models \blacksquare^{-1}\bot$ is vacuously true. Hence, $M_a \models [p]\blacklozenge^{-1}\blacksquare^{-1}\bot$. To see that $N_a \not\models [p]\blacklozenge^{-1}\blacksquare^{-1}\bot$, it is enough to notice that agent b_{n+1} is not *pro* topic p, and thus they do not follow agent a in the updated model. Updated model $N_a^{a:p}$ is depicted in Figure 7 on the right.

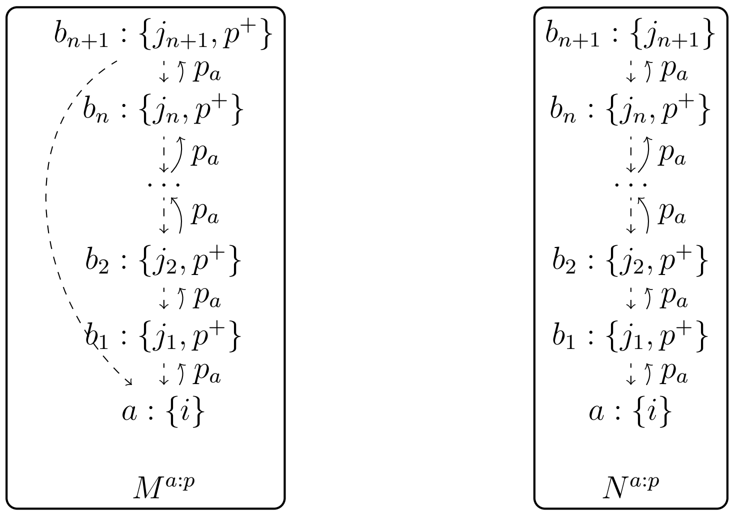

Figure 7: Models $M^{a:p}$ and $N^{a:p}$. Reflexive p_a-arrows and followership arrows from b_k to a for $k \in \{1, ..., n\}$ are omitted for readability.

In order to show that $M_a \models \psi$ if and only if $N_a \models \psi$, observe that M_a and N_a are n-bisimilar (but not $(n+1)$-bisimilar), which implies that the pointed models are $(\mathsf{Nom} \backslash Q)$-$n$-bisimilar. From the fact that $md(\psi) = n$ and Theorem 9, it follows that $M_a \models \psi$ if and only if $N_a \models \psi$. Hence, we have a contradiction with an earlier assumption that ψ is equivalent to $[p]\blacklozenge^{-1}\blacksquare^{-1}\bot$. $\square$

Before we turn to the model checking problem for VL, we mention that the complexity of the model checking problem of SVL is in P. This result follows trivially from the fact that model checking hybrid tense logic with universal modality is in P [21].

Theorem 28. Model checking SVL is in P.

Not only is VL more expressive than SVL, but its model checking problem is also more computationally demanding. We show this by providing a model checking algorithm for VL that runs in polynomial space. For hardness, we use the classic reduction from quantified Boolean formulas.

Theorem 29. The model checking problem for VL is PSPACE-complete.

Proof. To show that the model checking problem for VL is in PSPACE, we present Algorithm 1.

Algorithm 1 An algorithm for model checking VL

procedure MC(M, a, φ)
    case $\varphi = p^+$
        return $p \in +(a)$
    case $\varphi = p^-$
        return $p \in -(a)$
    case $\varphi = i$
        return $a \in V(i)$
    case $\varphi = \neg\psi$
        return not MC(M, a, ψ)
    case $\varphi = \psi \wedge \chi$
        return MC(M, a, ψ) and MC(M, a, χ)
    case $\varphi = \Diamond_{i:p}\psi$
        if $\exists b, c \in A$ such that $(a, b) \in R(p, c)$ with $V(i) = \{c\}$ and MC(M, b, ψ) **then**
            return *true*
        return *false*
    case $\varphi = \Diamond_{i:p}^{-1}\psi$
        if $\exists b, c \in A$ such that $(b, a) \in R(p, c)$ with $V(i) = \{c\}$ and MC(M, b, ψ) **then**
            return *true*
        return *false*
    case $\varphi = \blacklozenge\psi$
        if $\exists b \in A$ such that $a \in F(b)$ with MC(M, b, ψ) **then**
            return *true*
        return *false*
    case $\varphi = \blacklozenge^{-1}\psi$
        if $\exists b \in A$ such that $b \in F(a)$ with MC(M, b, ψ) **then**
            return *true*
        return *false*
    case $\varphi = @_i\psi$
        return MC($M, V(i), \psi$)
    case $\varphi = [p]\psi$
        return MC($M^{a:p}, a, \psi$)
    case $\varphi = [\pi \cup \tau]\psi$
        return MC($M, a, [\pi]\psi$) and MC($M, a, [\tau]\psi$)

The algorithm follows the semantics and its correctness can be shown via induction on φ. Now we argue that the algorithm requires at most polynomial space. The interesting case here is $\varphi = [p]\psi$. Without giving an explicit algorithm for constructing $M^{a:p}$, we note that the size of $M^{a:p}$ is bounded by $\mathcal{O}(|M|^2)$ (the worst-case scenario of $R(p, a)$ and F being universal). Since there are at most $|\varphi|$ symbols in φ, the total space required by the algorithm is bounded by $\mathcal{O}(|\varphi| \cdot |M|^2)$.

To show hardness of the model checking problem we use the classic reduction from the satisfiability of quantified Boolean formulas: given a QBF $\Psi := Q_1 p_1 \ldots Q_n p_n \psi(p_1, \ldots, p_n)$, where $Q_i \in \{\forall, \exists\}$, determine whether Ψ is true. To reduce the satisfiability of QBF Ψ to the model checking of VL, we construct a model M_a and a formula Ψ' of $\mathbb{VL}$ such that Ψ is true if and only if $M_a \models \Psi'$.

More specifically, given a QBF $Q_1 p_1 \ldots Q_n p_n \psi(p_1, \ldots, p_n)$, we construct a visibility model $M = (A, F, +, -, V, R)$, where $A = \{a_0, \ldots, a_n\}$, $F(a_i) = \{a_0\}$ for all $i \neq 0$, $+(a_i) = p_i$ for all $i \neq 0$, $-(a_i) = \emptyset$ for all i, $V(i_j) = \{a_j\}$, and $R(p, a) = \emptyset$ for all $p \in \mathsf{Top}$ and $a \in A$. Additionally, we assume that there is a topic q that no agent is either *pro* or *contra*. Intuitively, M is a model consisting of $n+1$ agents, where everyone follows agent a_0, who follows no one. Each agent, apart from a_0, is *pro* exactly one topic, and no one is *contra* anything. Finally, the translation of the QBF is done recursively as follows:

$$\begin{aligned}
\psi'_0 &:= \psi(\Diamond_{i_0:p_1}(i_1 \wedge \Diamond_{i_0:p_1}\top), \ldots, \Diamond_{i_0:p_n}(i_n \wedge \Diamond_{i_0:p_n}\top)) \\
\psi'_k &:= \begin{cases} [p_k \cup q]\psi'_{k-1} & \text{if } Q_k = \forall \\ \neg[p_k \cup q]\neg\psi'_{k-1} & \text{if } Q_k = \exists \end{cases} \\
\psi' &:= \psi'_n.
\end{aligned}$$

We need to show that

$$Q_1 x_1 \ldots Q_n x_n \psi(p_1, \ldots, p_n) \text{ is satisfiable iff } M_a \models \psi'.$$

Agent a_0 posting on topic p_i means that the truth value of p_i has been set to 1. If agent a_0 posts on topic q, this means that the truth value of the corresponding p_i has been set to 0. Since there are no two agents that are *pro* the same topic, the choice of truth values is unambiguous.

We use non-deterministic choice to model quantifiers. The universal quantifier $\forall p_k$ is emulated with $[p_k \cup q]\psi'_{k-1}$ meaning that no matter what agent a_0 chooses to post on, p_k or q, formula ψ'_{k-1} will be true. Similarly, the existential quantifier $\exists p_k$ is emulated with $\neg[p_k \cup q]\neg\psi'_{k-1}$ meaning that agent a_0 can post on a topic, either p_k or q, to make ψ'_{k-1} true. Finally, propositional variable p_j is translated into the formula $\Diamond_{i_0:p_j}(i_j \wedge \Diamond_{i_0:p_j}\top)$ that is true if and only if there has been a post on p_j, and the corresponding agent a_j, who is *pro* p_j, has reposted it. For all other agents a_k, the formula will not hold. Posting on q instead of p_j results in the fact that $\Diamond_{i_0:p_j}(i_j \wedge \Diamond_{i_0:p_j}\top)$ is not satisfied anywhere in the model, thus corresponding to setting p_j to 0. □

As an example, consider a QBF $\forall p_1 \exists p_2 (p_1 \to p_2)$. The formula is first translated into a formula of $\mathbb{VL}$: $[p_1 \cup q]\neg[p_2 \cup q]\neg(\Diamond_{i_a:p_1}(i_b \wedge \Diamond_{i_a:p_1}\top) \to \Diamond_{i_a:p_2}(i_c \wedge \Diamond_{i_a:p_2}\top))$. The corresponding model M is depicted in Figure 8.

Now, $M_a \models [p_1 \cup q]\neg[p_2 \cup q]\neg(\Diamond_{i_a:p_1}(i_b \wedge \Diamond_{i_a:p_1}\top) \to \Diamond_{i_a:p_2}(i_c \wedge \Diamond_{i_a:p_2}\top))$ if and only if

$$M_a^{a:p_1} \models \neg[p_2 \cup q]\neg(\Diamond_{i_a:p_1}(i_b \wedge \Diamond_{i_a:p_1}\top) \to \Diamond_{i_a:p_2}(i_c \wedge \Diamond_{i_a:p_2}\top))$$

and

$$M_a^{a:q} \models \neg[p_2 \cup q]\neg(\Diamond_{i_a:p_1}(i_b \wedge \Diamond_{i_a:p_1}\top) \to \Diamond_{i_a:p_2}(i_c \wedge \Diamond_{i_a:p_2}\top)).$$

Both updated models $M_a^{i_a:p_1}$ and $M_a^{i_a:q}$ are depicted in Figure 8. Notice that since agent b is *pro* topic p_1, they repost it thus creating a reflexive loop labelled with p_1. Agent c, who

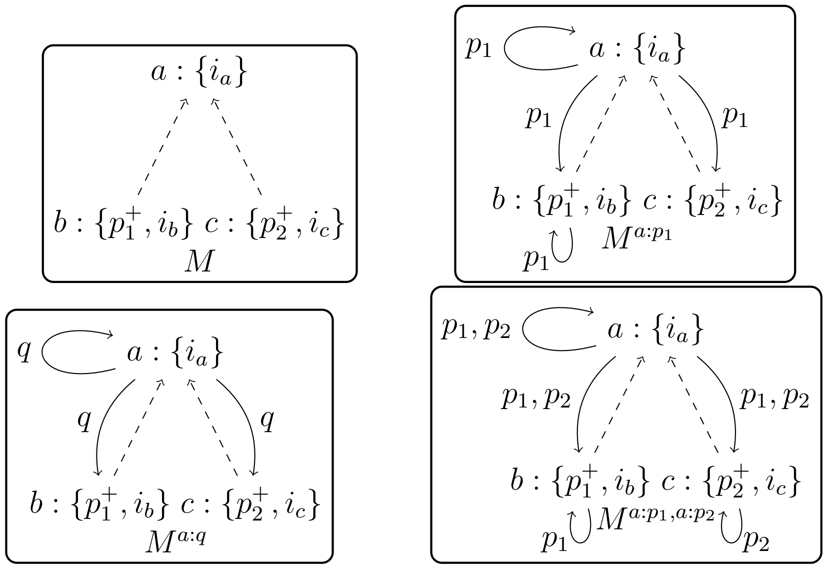

Figure 8: Models M, $M^{a:p_1}$, $M^{a:q}$, and $M^{a:p_1,a:p_2}$.

is not *pro* topic p_1, does not have such a loop. In the case of q, none of the agents have a reflexive loop.

In model $M_a^{a:p_1}$, propositional variable p_1 has been set to 1. Thus, we can choose to update the model further with a's post on p_2 to satisfy the initial formula. Indeed, consider model $M^{a:p_1,a:p_2}$. In the model, we have that $M_a^{a:p_1,a:p_2} \models \Diamond_{i_a:p_1}(i_b \wedge \Diamond_{i_a:p_1}\top) \to \Diamond_{i_a:p_2}(i_c \wedge \Diamond_{i_a:p_2}\top)$ since there is an $a : p_1$-arrow to state satisfying i_b from which there another $a : p_1$-arrow. Similarly for $a : p_2$-arrows. This corresponds to satisfying $p_1 \to p_2$ by setting both variables to 1.

In model $M_a^{a:q}$, propositional variable p_1 has been set to 0. Thus we can choose any further update of the model. Let agent a post on q once again. Since we do not discriminate between different posts on the same topic, model $M_a^{a:q,a:q}$ looks exactly like $M_a^{a:q}$. Moreover, it is clear that $M_a^{a:q,a:q} \models \Diamond_{i_a:p_1}(i_b \wedge \Diamond_{i_a:p_1}\top) \to \Diamond_{i_a:p_2}(i_c \wedge \Diamond_{i_a:p_2}\top)$ since $M_a^{a:q,a:q} \not\models \Diamond_{i_a:p_1}(i_b \wedge \Diamond_{i_a:p_1}\top)$. This corresponds to satisfying $p_1 \to p_2$ by setting both variables to 0.

Reamark 30. Our PSPACE-hardness proof relied on the union operator $\pi \cup \tau$. At the same time, by item (2) of Proposition 26, all formulas with action modalities with unions can be equivalently translated into formulas, where action modalities have only single topics in them. It is not immediately obvious whether our PSPACE-hardness argument can be rewritten without unions. We conjecture that it is indeed possible, taking into account that DEL with unions [2] and DEL without unions are both PSPACE-hard [25]. We note, however, that the PSPACE-hardness argument for DEL without unions is much more complex than the one for DEL with unions.

5 Arbitrary Visibility Logic

The example in Subsection 4.1 demonstrated that what topic an agent posts on matters. Depending on an agent's goals, posting on some topics rather than others can be more productive. For instance, if the goal of the posting agent in the example is to gain as many followers as possible, then posting on dogs is a better strategy than posting on vaccines.

As agents' goals can be reached by posting on one topic, and not reached by posting on another, we can ask a natural general question: given a visibility model M, a goal φ and a current agent a, *is there* a topic such that after posting on it, agent a makes φ true in model M? Observe the existential quantification in the question. A dual question with the universal quantification may be asked about some safety property φ of a visibility model M: is true that *whatever* the current agent posts on, φ will still hold in the updated visibility model?

To capture such reasoning, we extend the language of visibility logic with constructs $[*]\varphi$ that means 'whatever action the current agent executes, φ will be true'. Shifting our perspective from the effects of a particular agent action to the (non-)existence of an action achieving a certain goal is inspired by quantification in dynamic epistemic logic [16] and in particular by arbitrary public announcement logic [4]. It is following the latter that we call VL extended with quantifiers *arbitrary visibility logic* (AVL).

Definition 31. *The language of arbitrary visibility logic* $\mathbb{AVL}$ is defined recursively by the following grammar:

$$\begin{array}{lll} \varphi & ::= & p^+ \mid p^- \mid i \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Diamond_{i:p}\varphi \mid \Diamond^{-1}_{i:p}\varphi \mid \blacklozenge\varphi \mid \blacklozenge^{-1}\varphi \mid @_i\varphi \mid [\pi]\varphi \mid [*]\varphi \\ \pi & ::= & p \mid (\pi \cup \pi) \end{array}$$

where $[*]\varphi$ is read 'whatever action the current agent executes, φ holds'. Its dual, which is defined as $\langle * \rangle\varphi := \neg[*]\neg\varphi$, is read as 'there is an action which the current agent can execute such that φ will hold'.

Given a formula $\varphi \in \mathbb{AVL}$, we define *modal depth* and the *size* of the formula similarly to the corresponding definitions for $\mathbb{VL}$ with the following additional cases: $md([*]\varphi) = md(\varphi)$, and $|[*]\varphi| = |\varphi| + 1$.

Definition 32. Let $M = (A, F, +, -, V, R)$ be a visibility model, $a \in A$, and $p, q \in \mathsf{Top}$. *The semantics of AVL* extends the semantics of VL (Definition 25) with the following dual clauses:

$$\begin{array}{lll} M_a \models [*]\varphi & \text{iff} & \forall \pi : M_a \models [\pi]\varphi \\ M_a \models \langle * \rangle\varphi & \text{iff} & \exists \pi : M_a \models \langle \pi \rangle\varphi \end{array}$$

Note that we quantify not over single topics but over actions that can be more complex.

Returning to the example in Subsection 4.1, we can say that there is a topic on which agent a can post such that agents from b to f will follow her: $M_a \models \langle * \rangle(\blacklozenge^{-1}b \wedge \blacklozenge^{-1}c \wedge \blacklozenge^{-1}d \wedge \blacklozenge^{-1}e \wedge \blacklozenge^{-1}f)$, where $b, ..., f$ in the formula are nominal corresponding to agents $b, ..., f$ in the model. And indeed such a topic is *dogs*. At the same time, agent a, no matter what she posts on, will never lose her follower c, or, formally, $M_a \models [*]\blacklozenge^{-1}c$, where c in the formula is a nominal corresponding to agent c in the model.

Below are some of the valid and not valid formulas of AVL.

Proposition 33. Let $\varphi \in \mathbb{AVL}$.

1. $[*]\varphi \to [\pi]\varphi$ is valid
2. $[*](\varphi \wedge \psi) \to [*]\varphi \wedge [*]\psi$ is valid
3. $[*][*]\varphi \to [*]\varphi$ is not valid
4. $[*]\varphi \to [*][*]\varphi$ is not valid

The first formula states that if after executing any action φ is true, then φ will be true after executing some particular action. The second property is an expected distributivity of a modal box over conjunction. That quantifiers cannot be, in general, collapsed is expressed in formulas three and four.

Next, we show that in general formulas with quantifiers cannot be reduced to equivalent formulas without quantification. This in particular means that AVL is strictly more expressive than VL.

Theorem 34. $\mathbb{VL} < \mathbb{AVL}$

Proof. Consider $\langle * \rangle \blacklozenge^{-1}\blacksquare^{-1}\bot \in \mathbb{AVL}$, and assume towards a contradiction that there is an equivalent $\psi \in \mathbb{VL}$ with $md(\psi) = n$, and such that nominals $Q = \{j_1, \ldots, j_{n+1}\}$ do not appear in ψ, and also none of p, p^+, and p^- appears anywhere in ψ. Now consider models M_a and N_a from Figure 6, for which we argue that $M_a \models \langle * \rangle \blacklozenge^{-1}\blacksquare^{-1}\bot$ and $N_a \not\models \langle * \rangle \blacklozenge^{-1}\blacksquare^{-1}\bot$. For the former, we have that $\exists \pi : M_a \models \langle \pi \rangle \blacklozenge^{-1}\blacksquare^{-1}\bot$ by the semantics, and, letting $\pi := p$, the rest follows as in the proof of Theorem 27. For the latter, we consider $N_a \models [*]\blacksquare^{-1}\blacklozenge^{-1}\top$, which is equivalent to the fact that $\forall \pi : N_a \models [\pi]\blacksquare^{-1}\blacklozenge^{-1}\top$. Now, it is clear from the construction of N_a that for all π that do not contain p, the only agent that will follow a would be agent b_1 that satisfies $\blacklozenge^{-1}\top$. If π contains p, then all agents $b_1, \ldots, b_n$ will follow agent a, and it is easy to check that for all of them $\blacklozenge^{-1}\top$ is satisfied.

To see that $M_a \models \psi$ if and only if $N_a \models \psi$, recall that not only are the models $(\mathsf{Nom} \setminus Q)$-$n$-bisimilar, but ψ does not contain p as well. Thus, the only way for ψ to distinguish the models is to witness states b_{n+1} by a stack of $\blacklozenge^{-1}$, which is impossible due to $md(\psi) = n$ and the models being $(\mathsf{Nom} \setminus Q)$-$n$-bisimilar. □

In Section 4 we have shown that additional expressivity of VL, compared to SVL, comes at a price: the complexity of the model checking problem for VL jumps to PSPACE, compared to P for the case of SVL. Taking into account that AVL is strictly more expressive than VL, a natural question is whether we have to pay with yet another jump in complexity. Interestingly, the answer to the question is ‘no’, and in the rest of the section we argue that the complexity of the model checking problem for AVL is in PSPACE.

First of all, model checking AVL is not entirely straightforward. We cannot directly implement the semantics of the logic in an algorithm since quantifiers $[*]$ and $\langle * \rangle$ quantify over a countably infinite number of topics and their unions. However, we will show that it is enough just to consider the topic used in a finite model and appearing in a given formula.

Recall that $Nom(a) := \{i \in \mathsf{Nom} \mid a \in V(i)\}$ is the set of all nominals assigned to agent a, and $Top(a) := \{p \in \mathsf{Top} \mid R(p, a)\}$ is the set of all topics that agent a has posted.

Definition 35. Let M be a finite visibility model, and φ be a formula. Then *the set of topics appearing in* M, denoted $Var(M)$, is $\bigcup_{a \in A} Top(a)$. Similarly, with $Var(\varphi)$ we denote *the set of topics appearing in formula* φ. Finally, we will write $Var(M, \varphi)$ for $Var(M) \cup Var(\varphi)$.

Definition 36. Let π be an action. We call π *unique* if no topic appears in π twice.

Lemma 37. For each action π there is an equivalent unique action τ.

Proof. Let M_a be a visibility model and let $p \in \mathsf{Top}$. Assume that $M_a \models [\pi \cup p \cup p]\varphi$. By item (2) of Proposition 26, this is equivalent to $M_a \models [\pi]\varphi \wedge [p]\varphi \wedge [p]\varphi$, which in turn is equivalent to $M_a \models [\pi]\varphi \wedge [p]\varphi$. Using again item (2) of Proposition 26, we get $M_a \models [\pi \cup p]\varphi$. □

Definition 38. Let $Var(M, \varphi)$ be given. Then $\Pi(M, \varphi)$ is *the set of all unique actions* that can be built from $p \in Var(M, \varphi)$. Also, $\Pi^*(M, \varphi)$ is the set of all unique actions built from $p \in Var(M, \varphi) \cup \{p^*\}$.

Note that the size of $\Pi(M, \varphi)$ is exponential in the size of $Var(M, \varphi)$. We will have to address this issue later in the analysis of complexity.

Using the set of unique actions, we can redefine the semantics of our quantifiers so that the quantification ranges over a finite set of actions.

Lemma 39. Let M_a be a finite model, and let p^* be a topic such that $p^* \notin Var(M, \varphi)$.

$$\begin{array}{lll} M_a \models [*]\varphi & \text{iff} & \forall \pi \in \Pi^*(M, \varphi) : M_a \models [\pi]\varphi \\ M_a \models \langle * \rangle\varphi & \text{iff} & \exists \pi \in \Pi^*(M, \varphi) : M_a \models \langle \pi \rangle\varphi \end{array}$$

Proof. We show only the second item, and the first item can be proved analogously. Assume that $M_a \models \langle * \rangle\varphi$. By the definition of semantics this is equivalent to $\exists \pi : M_a \models \langle \pi \rangle\varphi$. Let π be an arbitrary action, and, by Lemma 37 we can safely assume that π is unique. There are two cases to consider. First, for all $p \in \pi$ we have that $p \in Var(M, \varphi)$, i.e. π consists only of topics that are explicitly present in the model or formula φ. In this case, the result trivially follows. Now, let us assume that there are some $p_1, \ldots, p_n \in \pi$ such that $p_1, \ldots, p_n \notin Var(M, \varphi)$. Since those topics do not appear in the model, they cannot influence the visibility or posting relations of agents. Hence, posting on each of these topics is equivalent to posting on p^*, which also appears neither in the model nor in the formula φ. Formally, let $\pi = \tau \cup p_1 \cup \ldots \cup p_n$, and thus $M_a \models \langle \tau \cup p_1 \cup \ldots \cup p_n \rangle\varphi$. Using item (2) of Proposition 26 repetitively, we have that $M_a \models \langle \tau \rangle\varphi \vee \langle p_1 \rangle\varphi \wedge \ldots \vee \langle p_n \rangle\varphi$. Since none of $p_1, \dots, p_n$ appear in $Var(M, \varphi)$, $M_a^{a:p_i} \models \varphi$ is different from M_a in that there is a reflexive p_i-arrow at agent a, and p_i-arrows to all of the a's followers. Posting on p^* has exactly the same effect. Since, p_i is not in φ, then $M_a^{a:p_i} \models \varphi$ if and only if $M_a^{a:p^*} \models \varphi$. Thus, we can substitute each p_i with p^* to get $M_a \models \langle\tau\rangle\varphi \vee \langle p^*\rangle\varphi$, which is equivalent to $M_a \models \langle\tau \cup p^*\rangle\varphi$, where $\tau \cup p^* \in \Pi^*(M, \varphi)$. □

Before we provide the algorithm for model checking AVL, we need to take care of $\Pi^*(M, \varphi)$. We mentioned above that the size of $\Pi^*(M, \varphi)$ is exponential in the size of $Var(M, \varphi)$, and thus we cannot keep the whole set in memory if we want to claim that the algorithm is in PSPACE. To deal with this, we introduce an auxiliary function $next(M, a, \pi)$ that, given a model M, an agent a, and an action $\pi \in \Pi^*(M, \varphi)$, returns the next action π' in $\Pi^*(M, \varphi)$.

Without loss of generality, we can assume that $\Pi^*(M, \varphi)$ is ordered based on the order of topics in Top. Hence, the ordering we have in mind looks as follows for $\{p_1, \dots, p_n\} = Var(M, \varphi) \cup \{p^*\}$: p_1, $p_1 \cup p_2$, …, $p_1 \cup p_n$, $p_1 \cup p_2 \cup p_3$, …, $p_1 \cup p_2 \cup p_n$, ..., p_2, $p_2 \cup p_3$, ..., p_n.

Now assume that we are given an arbitrary π in the ordering. To compute $next(M, a, \pi)$, we first check whether the last element in union π of length l can be incremented, i.e. whether it is p_k with $k < n$ or not. If yes, then increment this topic. If not, then check whether the topic at position $l - 1$ can be incremented. If it can be incremented, then we increment it and change the topic at position l to the one which follows the topic at position $l - 1$ in the ordering of Top. If it cannot be incremented, then we recursively proceed until we either produce a new union of length l, or produce a new union of length $l + 1$, or, else, start with a singleton topic that is next in the ordering of Top. If there is no such a topic, i.e. we have reached p_n, then return *end*. The described procedure can be computed in the time (and hence space) polynomial in the sizes of M and φ.

Theorem 40. Model checking AVL is PSPACE-complete.

Proof. The hardness follows trivially from the PSPACE-completeness of the model checking problem for VL. To argue that AVL model checking is in PSPACE, we present Algorithm 2, where all the cases apart from $[*]\psi$ are exactly as in Algorithm 1.

Algorithm 2 An algorithm for model checking AVL

procedure MC(M, a, φ)
  case $\varphi = [*]\psi$
    $\pi \leftarrow p_1$
    while $\pi \neq end$ **do**
      if not MC($M, a, [\pi]\varphi$) **then**
        return *false*
      $\pi \leftarrow next(M, a, \pi)$
    return *true*

The correctness of the algorithm follows from the semantics of AVL and Lemma 39. Regarding the complexity, recall that given M and φ there is a set of all unique actions

$\Pi^*(M, \varphi)$ that are built from $p \in Var(M, \varphi) \cup \{p^*\}$. The size of $\Pi^*(M, \varphi)$ is exponential in the sizes of M, and φ and thus we do not keep it in memory. Instead, we use function $next(M, a, \pi)$, starting from $\pi = p_1$ (line 3), to obtain the next action in $\Pi^*(M, \varphi)$. Computing $next(M, a, \pi)$ can be done in polynomial time (and hence space). Similarly to the case of Algorithm 1, the space required to store the current updated model is bounded by $\mathcal{O}(|M|^2)$. While checking all possible updates takes exponential time, we can reuse space for storing updated models. Hence, the total space required by the algorithm is bounded by $\mathcal{O}(|\varphi| \cdot |M|^2)$. □

6 Dynamic Hybrid Logics for Social Networks

Logics for reasoning about social networks is still a relatively new field. Yet, the community has brought forward many interesting analyses of social phenomena in the past decade. The logical frameworks used widely vary. Our framework is a modal logic with dynamic operators and nominals in the hybrid tradition. In this section, we give an assessment of other social network logics that use versions of dynamic hybrid logics as an underlying framework and compare them to the logics presented in this paper. By social network logic, we mean logics where a relational graph between agents is explicitly modelled.

Many social network logics utilise the modal logic Kripke frame to model a network of agents with a binary relation representing friendship, followership and/or communication channels. In SVL, VL and AVL, the underlying Kripke frame has two relations on the set of agents, the followership relation F and the visibility relation R. As we will see in this section, taking a relation on a Kripke frame to represent followership is not novel in this paper. However, modelling a relation of sharing and re-sharing a post as a binary relation on the set of agents is a new way to model communication using a social network logic. This approach gives us an organised and precise view of complex network situations which we hope can be further implemented in the field. The dynamic operator $[\pi]\phi$ changes both the followership and visibility relation in the updated model, but the valuation function is not updated. That is, in our system agents cannot change their current preferences. The operator $[*]\phi$ included in the language of AVL is particularly interesting from a conceptual view. It lets us reason about whether an agent has a possibility to act such that ϕ holds after the action. Such reasoning about the tactical choices of agents in social networks is a powerful tool that we believe can be useful to analyse the safety of networks, also in future work in the field. We go on to discuss other known dynamic hybrid logics for social networks.

[51] present a logic to model peer pressure within a community of social relationships. The logical model $M = \langle W, A, \sim, \leq, V \rangle$ includes a set of possible states of the world W and a set of agents A. The relation $\sim$ is a symmetric and irreflexive friendship relation on A, where $a \sim_w b$ is read as "agent a is friends with agent b in state w". $\leq$ is a preference relation on W, and $u \leq_a v$ is read as "for agent a, state v is at least as good as state u". The language includes two types of nominals, one for possible states and one for agents. Formulas are evaluated at pairs (w, a) of possible states and agents in the model.

The dynamic components of the logic are the unary operators $[\phi \leq \psi]$ and $[\phi < \psi]$ which update the preferences of an agent in a possible state. Define $[\phi] := \{w \in W \mid M, w, a \vDash \phi\}$. A statement $M, w, a \vDash [\phi \leq \psi]\theta$ holds, roughly, whenever θ holds in the updated model in which agent a's preferences now includes links from all $[\phi]$-states to all $[\psi]$-states. Similarly in the case of $M, w, a \vDash [\phi < \psi]\theta$, but here also preference links from all $[\psi]$-states to all $[\phi]$-states are deleted.

[11, 12, 13, 26, 14] introduce a series of influential social network logic papers using different dynamic hybrid logic frameworks to study diffusion, social influence and opinion dynamics. We give a short account of the frameworks by [13] and [14]. The models in the work by [13], named network models, are tuples $\mathcal{M} = (A, \asymp, g, v)$ where A is a set of agents and $\asymp$ is a binary relation on A representing the social network. g and v are valuation functions with respect to nominals and the features of agents, respectively. Dynamic transformations are modeled in the dynamic modality $[\mathcal{D}]$. The formula $[\mathcal{D}]\phi$ holds at an agent in a network model if and only if ϕ is forced at the agent in the updated model $\mathcal{M}^{\mathcal{D}}$ in which only the valuation v is changed. v' in the model $\mathcal{M}^{\mathcal{D}}$ updates the original valuation given a set of preconditions and post-conditions in $\mathcal{D}$. Simply stated, the post-conditions determine what should hold true in the updated model when the preconditions hold in the original model.

The framework by [14] builds upon and extends the one by [13]. The Logic of Knowledge, Diffusion and Learning (KDL) adds an epistemic dimension to the models and includes a knowledge operator K to the language. The language also includes a dynamic operator $[\mathcal{L}]$ in addition to $[\mathcal{D}]$. The operator $[\mathcal{L}]$ updates what the agents know about their friends. For a finite set of formulas $\mathcal{L}$, an agent a and an epistemic state w, $\mathcal{M}, w, a \vDash [\mathcal{L}]\phi$ is true if and only if ϕ holds in the updated model $\mathcal{M}^{\mathcal{L}}$. The intuition is that after the update, the current agent knows the features of their friends if the features are formulas in $\mathcal{L}$. A running example is given where agents have both an external and an internal hidden opinion. Either type of opinion is pro, contra or neutral. The paper defines some rules such as "if all the friends of an agent express a pro (contra) opinion, the agent will fall in line and express a pro (contra) opinion in the next round", and studies how the network evolves.

The logic known as "Facebook logic" was initially introduced by [46] and further expanded on by [47] and [30]. These papers have been highly influential in the field of social network logics and have been cited as inspiration for papers by authors such as the previously mentioned [51, 13, 14]. [47] present a dynamic version of the logic, called dynamic epistemic friendship logic (DEFL). The models of DEFL are tuples $M = \langle W, A, k, f, V \rangle$ where W is a set of epistemic states and A is a set of agents. k is a family of equivalence relations on W and f is a family of symmetric and irreflexive relations on A. The language includes operators for the two types of relations, nominals and the hybrid binder operator $\downarrow$. With the language, one can state sentences such as "Bella knows that she is not a spy but doesn't know if a friend of hers is a spy", denoted $@_b(K\neg s \wedge \neg K\langle F \rangle s)$. The dynamic operators in DEFL are based on the theory of General Dynamic Dynamic Logic [23] and use details from Propositional Dynamic Logic (PDL). The models can be updated after announcements from agents, which can also be private or public questions from one agent to another. Due to the complex nature of the operators, we will not go into more technical detail. [19] give some further alternatives for dynamic extensions of the framework by [46]. One extension lets agents send and receive asynchronous announcements. Asynchronous announcements are not assumed to be immediately received as they are sent, rather, the message is sent to a queue and can be received at a later stage. Two operators are added to the language: $[n!\phi]$ is read as "agent n sends a message ϕ to the queue" and $[n : r]$ is read as "agent n receives all queued messages sent by her friends".

[50] introduce Dynamic Hybrid Logic for Followership which language is the basic hybrid language added a dynamic operator $[a \uparrow \theta]$. The models are standard hybrid models, with a set of agents and a binary relation representing followership. $[a \uparrow \theta]\phi$ is read as "after a chooses to only follow agents satisfying θ, ϕ holds". Other dynamic social network logics using hybrid elements are included in works by authors [45, 35, 39, 41, 36]. Details of these frameworks are left out for now.

7 Conclusion and Future Work

This work was devoted to the analysis of the concepts of visibility and exploitation in social networks using modal logic. After discussing related work from the perspective of social network analysis, we introduced a logic we named static visibility logic (SVL) and its dynamic extensions, visibility logic (VL) and arbitrary visibility logic (AVL). We did not give a definite answer as to how one should measure visibility, but proposed several quantitative and qualitative measures relevant to our social network models. To motivate VL, we presented an example where we showed how, given some simple rules of the system, a potential malicious agent can take advantage of the network to expose more agents to a controversial opinion. In AVL we introduced operators to reason about whether an agent can act such that a certain outcome holds.

On the mathematical side, we showed soundness and completeness of SVL with respect to social networks that follow our given rules. We also proved that the language of VL is strictly more expressive than the language of SVL, and that the language of AVL is strictly more expressive than the language of VL. The first increase in expressivity, from SVL to VL, also resulted in a significant increase in the complexity of model checking, from P to PSPACE. Interestingly, the second expressivity increase, from VL to AVL, has not resulted in a jump in the complexity of model checking, i.e. the complexity of the model checking problem for AVL is still PSPACE-complete.

As we mention in the paper, an implication of the result $\mathbb{SVL} < \mathbb{VL}$ is that a proof of the completeness of VL using reduction axioms is not possible. Thus one of the open problems is to find a sound and complete axiomatisation of VL. As we also prove that $\mathbb{VL} < \mathbb{AVL}$, completeness of AVL can neither be proved using reduction axioms into SVL nor VL. A sound and complete axiomatisation of AVL is therefore also an open problem.

Another direction for future work is to formalise triggering in social network communication. The idea is that seeing a post on a controversial topic might trigger an agent to post a reaction. To do this, we could expand our framework such that agents can not only post on a topic, but also *pro* or *contra* a topic. This entails letting $\pi ::= p \mid p^+ \mid p^- \mid \pi \cup \pi$ in the dynamic formula $[\pi]\phi$. Then, we could specify particular controversial topics and add a rule stating that if an agent sees a post that is pro the controversial topic and they are themselves contra, then the agent will post contra the topic, or vice versa.

Related to the former point, the social network presented in this paper comes with a set of rules that is an oversimplification of a real-life network. In future work, we would like to add more detailed, and more realistic, rules, which would give us a more complicated system to study other interesting social phenomena with. Furthermore, in our work, we focus on the effects of agents posting on *different topics* rather than *posting different posts on the same topic*. One avenue for further research is to extend the framework to also allow agents to post more than one post on the same topic.

Acknowledgements

Thank you Marija Slavkovik for collaboration on the previous version of this paper, and for helpful suggestions on this extended work. Thank you also to the reviewers for their insightful comments and constructive criticism. This includes the reviewers of the conference proceedings paper which this paper is an extension of.

References

[1] Carlos Areces and Balder ten Cate. Hybrid logics. In Patrick Blackburn, Johan van Benthem, and Frank Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 821–868. Elsevier, 2007.

[2] Guillaume Aucher and François Schwarzentruber. On the complexity of dynamic epistemic logic. In Burkhard C. Schipper, editor, *Proceedings of the 14th TARK*, 2013.

[3] Edoardo Baccini and Zoé Christoff. Comparing social network dynamic operators. In Rineke Verbrugge, editor, *Proceedings of the 19th TARK*, volume 379 of *EPTCS*, pages 66–81, 2023.

[4] Philippe Balbiani, Alexandru Baltag, Hans van Ditmarsch, Andreas Herzig, Tomohiro Hoshi, and Tiago de Lima. 'Knowable' as 'known after an announcement'. *Review of Symbolic Logic*, 1(3):305–334, 2008.

[5] Alexandru Baltag, Zoé Christoff, Rasmus K. Rendsvig, and Sonja Smets. Dynamic epistemic logics of diffusion and prediction in social networks. *Studia Logica*, 107(3):489–531, 2019.

[6] Alexandru Baltag and Lawrence S. Moss. Logics for epistemic programs. *Synthese*, 139(2):165–224, 2004.

[7] Gaia Belardinelli. Gatekeepers in social networks: Logics for communicative actions. Master's thesis, ILLC, University of Amsterdam, 2019.

[8] Patrick Blackburn and Balder ten Cate. Pure extensions, proof rules, and hybrid axiomatics. *Studia Logica*, 84(2):277–322, 2006.

[9] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. CUP, 2001.

[10] William J. Brady, Julian A. Wills, John T. Jost, Joshua A. Tucker, and Jay J. Van Bavel. Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28):7313–7318, 2017.

[11] Zoé Christoff. A logic for social influence through communication. In Emiliano Lorini, editor, *Proceedings of the 11th EUMAS*, volume 1113 of *CEUR Workshop Proceedings*, pages 31–39. CEUR-WS.org, 2013.

[12] Zoé Christoff and Jens Ulrik Hansen. A two-tiered formalization of social influence. In Davide Grossi, Olivier Roy, and Huaxin Huang, editors, *Proceedings of the 4th LORI*, volume 8196 of *LNCS*, pages 68–81. Springer, 2013.

[13] Zoé Christoff and Jens Ulrik Hansen. A logic for diffusion in social networks. *Journal of Applied Logic*, 13(1):48–77, 2015.

[14] Zoé Christoff, Jens Ulrik Hansen, and Carlo Proietti. Reflecting on social influence in networks. *Journal of Logic, Language and Information*, 25(3-4):299–333, 2016.

[15] Louise A Dennis, Yu Fu, and Marija Slavkovik. Markov chain model representation of information diffusion in social networks. *Journal of Logic and Computation*, 32(6):1195–1211, 2022.

[16] Hans van Ditmarsch. To be announced. *Information and Computation*, 292:105026, 2023.

[17] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2008.

[18] David Easley and Jon Kleinberg. *Networks, Crowds and Markets*. CUP, 2010.

[19] Saúl Fernández González. Change in social networks: Some dynamic extensions of social epistemic logic. *Journal of Logic and Computation*, 32(6):1212–1233, 2022.

[20] Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, 1979.

[21] Massimo Franceschet and Maarten de Rijke. Model checking hybrid logics (with an application to semistructured data). *Journal of Appled Logic*, 4(3):279–304, 2006.

[22] Rustam Galimullin, Mina Young Pedersen, and Marija Slavkovik. Logic of visibility in social networks. In Agata Ciabattoni, Elaine Pimentel, and Ruy J. G. B. de Queiroz, editors, *Proceedings of the 28th WoLLIC*, volume 13468 of *LNCS*, pages 190–206. Springer, 2022.

[23] Patrick Girard, Jeremy Seligman, and Fenrong Liu. General dynamic dynamic logic. In Thomas Bolander, Torben Braüner, Silvio Ghilardi, and Lawrence S. Moss, editors, *Proceedings of the 9th AiML*, pages 239–260. College Publications, 2012.

[24] Valentin Goranko and Martin Otto. Model theory of modal logic. In Patrick Blackburn, Johan van Benthem, and Frank Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 249–329. Elsevier, 2007.

[25] Ronald de Haan and Iris van de Pol. On the computational complexity of model checking for dynamic epistemic logic with S5 models. *FLAP*, 8(3):621–658, 2021.

[26] Jens Ulrik Hansen. Reasoning about opinion dynamics in social networks. *Journal of Logic and Computation*, 29(7):1121–1137, 2015.

[27] M. O. Jackson, B. W. Rogers, and Y. Zenou. The economic consequences of social-network structure. *Journal Of Economic Literature*, 55(1):49–95, 2017.

[28] Barteld Kooi and Bryan Renne. Arrow update logic. *Review of Symbolic Logic*, 4(4):536–559, 2011.

[29] Natasha Kurtonina and Maarten de Rijke. Bisimulations for temporal logic. *Journal of Logic, Language and Information*, 6(4):403–425, 1997.

[30] Fenrong Liu, Jeremy Seligman, and Patrick Girard. Logical dynamics of belief change in the community. *Synthese*, 191(11):2403–2431, 2014.

[31] Kun Liu and Evimaria Terzi. A framework for computing the privacy scores of users in online social networks. *ACM Trans. Knowl. Discov. Data*, 5(1), 2010.

[32] Emiliano Lorini and Giovanni Sartor. A STIT logic for reasoning about social influence. *Studia Logica*, 104(4):773–812, 2016.

[33] Philip M. Napoli. Social media and the public interest: Governance of news platforms in the realm of individual and algorithmic gatekeepers. *Telecommunications Policy*, 39(9):751–760, 2015.

[34] C. Thi Nguyen. Echo chambers and epistemic bubbles. *Episteme*, 17(2):141–161, 2020.

[35] Andrés Occhipinti Liberman and Rasmus K. Rendsvig. Dynamic term-modal logic for epistemic social network dynamics. In Patrick Blackburn, Emiliano Lorini, and Meiyun Guo, editors, *Proceedings of the 7th LORI*, volume 11813 of *LNCS*, pages 168–182. Springer, 2019.

[36] Andrés Occhipinti Liberman and Rasmus K Rendsvig. Reasoning about epistemic social network dynamics using dynamic term-modal logic. *Journal of Logic and Computation*, 32(6):1067–1087, 2022.

[37] Mina Young Pedersen, Marija Slavkovik, and Sonja Smets. Social bot detection as a temporal logic model checking problem. In Sujata Ghosh and Thomas Icard, editors, *Proceedings of the 8th LORI*, volume 13039 of *LNCS*, pages 158–173. Springer, 2021.

[38] Mina Young Pedersen, Marija Slavkovik, and Sonja Smets. Detecting bots with temporal logic. *Synthese*, 202(3):79, 2023.

[39] Mina Young Pedersen, Sonja Smets, and Thomas Ågotnes. Analyzing echo chambers: A logic of strong and weak ties. In Patrick Blackburn, Emiliano Lorini, and Meiyun Guo, editors, *Proceedings of the 7th LORI*, volume 11813 of *LNCS*, pages 183–198. Springer, 2019.

[40] Mina Young Pedersen, Sonja Smets, and Thomas Ågotnes. Further steps towards a logic of polarization in social networks. In Mehdi Dastani, Huimin Dong, and Leon van der Torre, editors, *Proceedings of the 3rd CLAR*, volume 12061 of *LNCS*, pages 324–345. Springer, 2020.

[41] Mina Young Pedersen, Sonja Smets, and Thomas Ågotnes. Modal logics and group polarization. *Journal of Logic and Computation*, 31(8):2240–2269, 2021.

[42] Jan Plaza. Logics of public communications. *Synthese*, 158(2):165–179, 2007.

[43] Nemi Chandra Rathore and Somanath Tripathy. Effective visibility prediction on online social network. *IEEE Transactions on Computational Social Systems*, 2021.

[44] Moumita Samanta, Pinakpani Pal, and Abhik Mukherjee. A novel scheme for abatement of privacy concern by controlling the reachability in online social network. In Ajith Abraham, Aswani Kumar Cherukuri, Ana Maria Madureira, and Azah Kamilah Muda, editors, *Proceedings of the 8th SoCPaR*, volume 614 of *Advances in Intelligent Systems and Computing*, pages 686–697. Springer, 2016.

[45] Katsuhiko Sano and Satoshi Tojo. Dynamic epistemic logic for channel-based agent communication. In Kamal Lodaya, editor, *Proceedings of the 5th ICLA*, volume 7750 of *LNCS*, pages 109–120. Springer, 2013.

[46] Jeremy Seligman, Fenrong Liu, and Patrick Girard. Logic in the community. In Mohua Banerjee and Anil Seth, editors, *Proceedings of the 4th ICLA*, volume 6521 of *LNCS*, pages 178–188. Springer, 2011.

[47] Jeremy Seligman, Fenrong Liu, and Patrick Girard. Facebook and the epistemic logic of friendship. In Burkhard C. Schipper, editor, *Proceedings of the 14th TARK*, pages 230–238, 2013.

[48] John Tang, Mirco Musolesi, Cecilia Mascolo, and Vito Latora. Characterising temporal distance and reachability in mobile and online social networks. *Computer Communication Review*, 40(1):118–124, 2010.

[49] Jeffrey W Treem, Paul M Leonardi, and Bart van den Hooff. Computer-Mediated Communication in the Age of Communication Visibility. *Journal of Computer-Mediated Communication*, 25(1):44–59, 01 2020.

[50] Zuojun Xiong and Meiyun Guo. A dynamic hybrid logic for followership. In Patrick Blackburn, Emiliano Lorini, and Meiyun Guo, editors, *Proceedings of the 7th LORI*, volume 11813 of *LNCS*, pages 425–439. Springer, 2019.

[51] Liang Zhen and Jeremy Seligman. A logical model of the dynamics of peer pressure. In Hans van Ditmarsch, David Fernández-Duque, Valentin Goranko, Wojciech Jamroga, and Manuel Ojeda-Aciego, editors, *Proceedings of the 7th M4M and the 4th LAMAS*, volume 278 of *ENTCS*, pages 275–288. Elsevier, 2011.