Diffusion in Social Networks with Recalcitrant Agents

Zoé Christoff
Department of Philosophy
University of Bayreuth
Germany
zoe.christoff@uni-bayreuth.de

Pavel Naumov
Department of Mathematics
Claremont McKenna College
United States
pnaumov@cmc.edu

Abstract

The article generalizes the standard threshold models of diffusion in social networks by introducing the notion of recalcitrant agents, i.e., agents that are fully resistant to the diffusion process. The focus of the article is on capturing a ternary influence relation between groups of agents: agents in one group can indirectly influence agents in another group in spite of the agents in the third group being recalcitrant. The main technical result is a sound and complete axiomatization of this relation.

1 Introduction

1.1 Context

Outside of logic, various disciplines have long recognized the importance of social networks. Social psychologists, economists, sociologists, and political scientists have all used networks models to study social influence [20, 14]. Usually, such models abstract away from individual motivations to regard influence processes as simple *epidemics*, where agents might get infected by the agents they are in contact with, their network-neighbors. Two paradigmatic examples are the models of innovations diffusion [19] and of opinions change under social influence by iterated linear averaging [12].

This article focuses specifically on the "threshold-based" account of social influence [14, 32], according to which agents adopt a new product (or an opinion, trend, or behavior) whenever a *large enough* number of their network neighbors have adopted it already. The *threshold models*, well known since [19, 29], represent the corresponding diffusion dynamics. Typically, such models present an extremely simplified picture, where all agents are taken to be equally influential, and those who are influenced into adopting a product never change their mind to "unadopt" it. Threshold models have received increasing attention in recent literature, mostly outside of the field of logic [1, 14, 17, 21, 22, 24, 25].

In the last few years, logic has also seen a rising number of studies of network phenomena, most of which are based on some variations of threshold models (see [7, Chapter I] for an introduction). The seminal work of Seligman et al. [30, 23, 33] combines tools from dynamic epistemic logic [13] with hybrid logic [6] to create a "Logic of the Community". The resulting two-dimensional settings allow for explicitly capturing both groups structures and individual change under social influence. Since then, such hybrid systems have been extended in various ways [31, 16, 27, 28, 9, 10, 11].

Beyond hybrid logic, two main lines of research appeared. First, logicians have searched for systems *expressive enough* to capture the laws of diffusion dynamics *in the long run*, such as stabilization *vs* oscillation behavior, considering various diffusion dynamics beyond threshold models [5, 8]. Second, they have aimed towards using *less complex* systems to capture diffusion in threshold models. This is the case for instance of [4, 3], each of which proposes a simplified logical modelling of threshold models diffusion. The former, [4], uses a minimal dynamic propositional system (and its epistemic extension) to capture influence between individuals, while the latter, [3] offers a "one binary operator, four axioms" logic, by reducing diffusion to the influencing ability of groups. More precisely, the system of [3] relies on a unique relation between sets of agents: "if all agents in group A are already influenced (infected), all agents in B will eventually be too". The same relation but with an addition of marketing budget was investigated in [26].

At first sight, these two research directions might seem rather orthogonal. Indeed, typically, languages which are expressive enough to capture what happens in the long run also contain fixed point operators (as in the μ-calculus setting of [8]). And logics which are simpler (such as the non-epistemic part of [4]) do not allow for this long run perspective without explicitly listing all possible diffusion sequences. Instead, the setting proposed in this paper (as well as the one of [3]) combines these two directions: by focusing on the influenceability relation between groups, it allows to talk directly about what happens in the long run, while at the same time keeping the logical system extremely simple and its expressions compact.

The idea of modeling influence power between groups also relates to a third line of research, on the judgment aggregation theory side, characterizing network diffusion in terms of iterated parallel synchronous local aggregation processes, as proposed in [18]. As shown by [8], influence powers between sets of agents can be considered in terms of winning and veto coalitions, where winning (respectively, veto) coalitions are sets of influencers who have enough *power* to force (prevent) an agent to adopt the new behavior. This paper relates to this third line of research too, since it proposes a logical system to reason directly about the ability of groups of agents to (perhaps indirectly) influence each other, without having to describe the actual state of any of the individuals or their thresholds. Therefore, our setting can be considered as a logic of coalitions influence power.

On the technical side, this article relates closely to [3]. However, it differs from [3] in two significant ways. First, we extend the language by adding a set of *recalcitrant* agents, i.e., agents who can never be influenced, to the syntax.

The relation used by [3] corresponds therefore to the special case where the set of recalcitrant agents is taken to always be empty. The addition of the set of recalcitrant agents to the influence relation results in a more expressive language in which significantly more complex properties of the influence relation can be captured. The main example of such a property is the Serializability axiom given in Section 3. Our main result is a complete axiomatization of all properties of the influence relation expressible in this richer language. Second, [3] studies properties of influence specific to a given structure of the social network and the logical system obtained there contains a network-specific axiom. The logic proposed in the current article is more general in the sense that it axiomatizes properties of influence that are common to all social networks no matter what their structure is.

1.2 Preliminaries

Consider a situation where a new trend, product, technology, opinion, or behavior is adopted by a few initial people within a community. This often results in others in their "circles" (the agents they are closely connected to in their social network) being influenced into adopting it as well. As the set of adopters sequentially grows, the conformity pressure weighing on the rest of the population increases accordingly, following a "snowballing" dynamics. On an abstract level, it is typically this type of *cascading* or *epidemics* effects, sometimes simply described as "network effects", which drive influence depending processes in networks [14, Chapter 17].

We call an agent *recalcitrant* if she is completely resistant to social influence. She just can never be converted. In the context of the diffusion of a new product, such an agent, for instance, might not be able to use the product at all, or she might have a very strong moral objection against it. In the context of an infection epidemics, the recalcitrant agents are the ones who are immune to the disease, for instance agents who have been previously vaccinated. Recalcitrant agents could also be thought of as a *faulty nodes* in the network, failing to fulfill their role in a diffusion process.

We describe properties of influence in terms of two relations between three groups of agents A, B, and C: $A \rhd_B C$ and $A \lhd_B C$. Statement $A \rhd_B C$ means that agents in the group A have the ability, together, of influencing – directly or indirectly (through intermediate agents) – *every* agent in the group C, and this even if all agents in the group B are recalcitrant. Statement $A \lhd_B C$ means that when all agents in the set B are recalcitrant, then agents in the set A do not have the ability to influence *anybody* in the set C. The relations $A \rhd_B C$ and $A \lhd_B C$ are *dual* in the sense that they are interdefinable:

$$A \rhd_B C \equiv \bigwedge_{c \in C} \neg(A \lhd_B \{c\}),$$

$$A \lhd_B C \equiv \bigwedge_{c \in C} \neg(A \rhd_B \{c\}).$$

It will turn out that some properties of influence are easier to express using the relation $A \rhd_B C$ while others are captured more naturally in terms of the relation $A \lhd_B C$. This article therefore offers an axiomatic system in terms of *both* of these relations, despite their interdefinability.

We use a simple version of threshold models, where agents may have different influence thresholds but all influence links in the social network have equal weight. In this context, $A \rhd_B C$ means that the threshold of all agents in C will be met at some point, whenever all agents in A are influenced (infected), despite the fact that all agents in B are recalcitrant (immune). However, the results presented in this article also hold for threshold models with unequal weights.

Before turning to an example, let us introduce a few notation shortcuts. We write $a_1, \ldots, a_n$ instead of $\{a_1, \ldots, a_n\}$. For instance, we write $a \rhd_b c$ instead of $\{a\} \rhd_{\{b\}} \{c\}$. We also write $A \rhd C$ for $A \rhd_\varnothing C$. Moreover, we denote the union of sets with a comma. For instance, we write $A \rhd_B C, D$ instead of $A \rhd_B (C \cup D)$.

1.3 Example

In this section we illustrate the two relations $A \rhd_B C$ and $A \lhd_B C$ using a simple five-node social network. The formal definitions of these relations are given in Section 2.

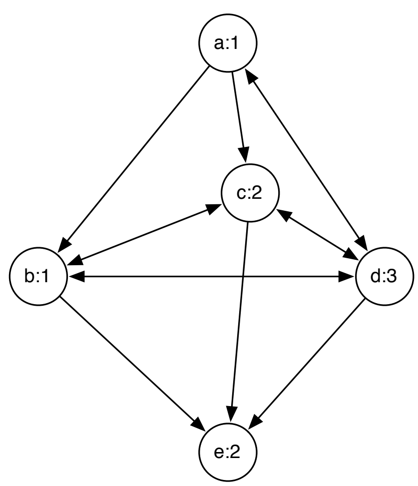

Figure 1: Social Network N

A network containing 5 agents (nodes), named a, b, c, d, and e. Their respective thresholds are indicated next to their name, within the node. The arrows represent the influence relation among agents. For instance, agent a influences agents b, c, and d, but is influenced by agent d only.

Consider the network N depicted in Figure 1. Let us first assume that this network has no recalcitrant agent. Consider the case where agent a is an initial adopter, for instance if she has been given a free sample of a new product. Then, agent b, whose threshold value is only 1, will experience enough peer-pressure (from a) to adopt the product next. At that point, the total pressure on agent c meets her threshold of 2, so she also adopts the product. Finally, agents d and e, whose thresholds are 3 and 2 respectively, will adopt the product too. This means that agent a has the ability to influence *at some point* all other agents, she can trigger a full cascade. Using our notation, this fact is denoted by $a \rhd b, c, d, e$.

Consider now the same network N, but with the assumption that c is a recalcitrant agent. Then, agent a is able to influence agent b, but neither agent d nor e. We denote these facts, respectively, by $a \rhd_c b$ and $a \lhd_c d, e$.

Finally, note that in the network N agent d can influence agent e if b is the only recalcitrant agent: $d \rhd_b e$. Similarly, agent d can influence agent e if c is the only recalcitrant agent: $d \rhd_c e$. At the same time, agent d cannot influence e when b and c are *both* recalcitrant: $d \lhd_{b,c} e$.

This example illustrates how the relations of conditional influenceability ($A \rhd_B C$) and conditional "anti-influenceability" ($A \lhd_B C$) allow to talk about the behavior of diffusion dynamics in the limit, for a given static social network structure.

In the reminder of the article we will introduce a sound and complete axiomatization of the properties of these two relations which hold for all social network structures.

1.4 Outline

We first introduce our syntax and semantics in Section 2. We then give an axiomatization in Section 3. Section 4 contains examples of formal derivations in our logical system. We then prove soundness of the logic in Section 5, and its completeness and decidability in Section 6. Section 7 summarizes and concludes the article.

2 Syntax and Semantics

The language of our logical system is defined for an arbitrary fixed finite set $\mathcal{A}_0$ of names of agents. The role of these names is similar to the one of constants in first order logic. For the sake of simplicity, we will refer to agent names in set $\mathcal{A}_0$ as just "agents".

Although the axioms of our logical system are stated in terms of both relations $A \rhd_B C$ and $A \lhd_B C$, we have chosen to restrict our set of atomic statements to formulae of the form $A \rhd_B C$. We therefore consider $A \lhd_B C$ as an abbreviation for $\bigwedge_{c\in C} \neg(A \rhd_B c)$.

Finally, when we write $A \rhd_B C$, we consider the hypothetical situation where agents in set A are infected (or are initial adopters of the new product) and agents in set B are vaccinated (or are recalcitrant). Thus, any agent in set $A \cap B$ would have to be infected and vaccinated at the same time, which makes the situation impossible. To prevent this, we assume that sets A and B are disjoint and make this into a syntactical restriction on our formulae. Our language is therefore defined as follows:

Definition 1 *Let set of formulae* Φ *be the smallest set such that:*

1. $A \rhd_B C \in \Phi$*, for all* $A, B, C \subseteq \mathcal{A}_0$ *with* $A \cap B = \emptyset$*,*
2. $\neg\varphi \in \Phi$*, for any* $\varphi \in \Phi$*,*
3. $\varphi \to \psi \in \Phi$*, for all* $\varphi, \psi \in \Phi$*.*

In other words, the language Φ is defined by the following BNF grammar: $\varphi := A \rhd_B C \mid \neg\varphi \mid \varphi \to \varphi$. We assume that disjunction and conjunction are defined through implication and negation in the usual way.

Although $\mathcal{A}_0$ is the set of all names of agents that are used in our language Φ, the actual social network might have additional "unnamed" agents. Our approach here again is similar to first order logic where not all elements in the domain might correspond to constants in the language. The role of such additional agents will be exemplified by our canonical model construction in Section 6.

Definition 2 *A social network is a triple* $(\mathcal{A}, \theta, R)$*, where*

1. $\mathcal{A}$ *is a finite superset of* $\mathcal{A}_0$*, called set of* agents*,*
2. $\theta : \mathcal{A} \to \mathbb{N}$ *is a function from set* $\mathcal{A}$ *to non-negative integers, whose values are called* thresholds*, and*
3. $R \subseteq \mathcal{A} \times \mathcal{A}$ *is an arbitrary relation on* $\mathcal{A}$*, called* influence relation.

For example, for the social network N depicted in Figure 1: the set of agents $\mathcal{A}$ is $\{a, b, c, d, e\}$; the function θ is such that $\theta(a) = 1$, $\theta(b) = 1$, $\theta(c) = 2$, $\theta(d) = 3$, and $\theta(e) = 2$; and the relation $\mathcal{R}$ is

$$\{(a, b), (a, c), (a, d), (b, c), (b, d), (b, e), (c, b), (c, d), (c, e), (d, a), (d, b), (d, c), (d, e)\}.$$

As usual with threshold models, two agents are taken to be related in a social network just in case one is susceptible in principle to influence the other by "direct contagion": $(a, b) \in R$ means that a might *directly influence* b (conditional of course on b's threshold being met). Note that we have imposed no constraint on our influence relation, contrary to most other work in logic, typically requiring symmetry and irreflexivity, to capture the usual structure of a "friendship" relationship. See for instance [30, 23, 4].

We write $R^{-1}(a)$ for the set of direct influencers (direct network-neighbors) of an agent a:

$$R^{-1}(a) = \{b \in \mathcal{A} \mid (b, a) \in R\}$$

For example, $R^{-1}(e) = \{b, c, d\}$ in the social network N depicted in Figure 1.

The following two definitions capture, respectively, the step by step diffusion dynamics (who is influenced after n time steps), and its fixed point (who will be influenced in the limit), given that the agents in set B are recalcitrant:

Definition 3 *For any disjoint sets $A, B \subseteq \mathcal{A}$ and any integer $n \geq 0$, let set A_B^n be A if $n = 0$ and otherwise be set*

$$A_B^{n-1} \cup \left\{ a \in \mathcal{A} \setminus B \;\middle|\; |A_B^{n-1} \cap R^{-1}(a)| \geq \theta(a) \right\}.$$

In words, the set of agents who will have adopted the new product *at the next time step* is given by: the agents who have already adopted it together with those who are not recalcitrant and have enough influencers who have adopted it already to meet their respective thresholds. Then the set of agents who will have adopted it *eventually* is defined in a natural way:

Definition 4 *For any disjoint sets $A, B \subseteq \mathcal{A}$, let $A_B^* = \bigcup_{n \geq 0} A_B^n$.*

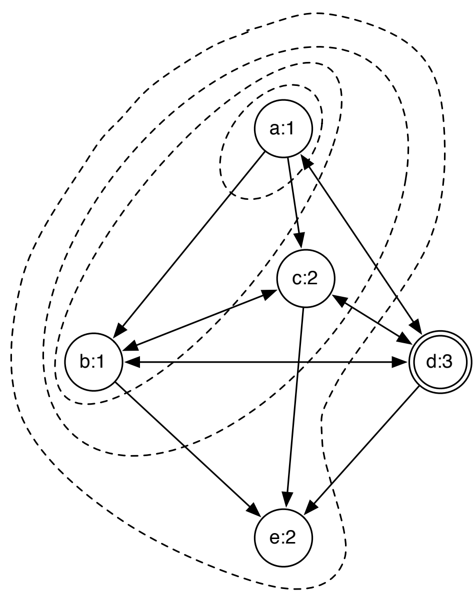

Figure 2: Diffusion Chain in the Social Network N.

The diffusion process in the network depicted in Figure 1, in the case when agent a adopts first, and agent d is recalcitrant. Recalcitrance is denoted by a double circled node, and the successive stages of diffusion are represented by the growing spheres in dotted lines. Influence spreads from agent a to reach all other agents but d.

For example, Figure 2 illustrates with dashed lines the following possible diffusion chain in the social network N:

$$\{a\} = \{a\}_d^0 \subseteq \{a\}_d^1 \subseteq \{a\}_d^2 \subseteq \{a\}_d^3 = \{a\}_d^*,$$

where $\{a\}_d^1 = \{a, b\}$, $\{a\}_d^2 = \{a, b, c\}$, $\{a\}_d^3 = \{a, b, c, e\}$.

In the reminder of the article, whenever we use expressions of the form A^n_B or A^*_B, we assume that those are well-defined, that is $n \geq 0$ and sets A and B are disjoint.

The following six lemmas capture some simple technical properties of our model. Those lemmas will be used later in Sections 4, 5, and 6:

Lemma 1 *There is an integer $m \geq 0$ such that $A^*_B = A^m_B$.*

Proof. The statement of the lemma follows from the monotonicity of the chain $A^0_B \subseteq A^1_B \subseteq A^2_B \subseteq \cdots \subseteq \mathcal{A}$ and the assumption of finiteness of set $\mathcal{A}$ in Definition 2. $\square$

The following properties of sets A^n_B and A^*_B follow from Definitions 3 and 4 in a straightforward way:

Lemma 2 *Sets A^*_B and B are disjoint.*

Lemma 3 *If $A \subseteq A'$, then $A^*_B \subseteq (A')^*_B$.*

Lemma 4 *If $B \subseteq B'$, then $A^*_{B'} \subseteq (A)^*_B$.*

Lemma 5 $(A^*_B)^*_B \subseteq A^*_B$.

Lemma 6 *If sets A^{m-1}_B and D are disjoint, then $A^m_{B \cup D} = A^m_B \setminus D$.*

The following definition captures the key concept of the article. Item (1) formally specifies the meaning of the relation $A \rhd_B C$:

Definition 5 *For any social network $N = (\mathcal{A}, \theta, R)$ and any formula $\varphi \in \Phi$, the satisfaction relation $N \vDash \varphi$ is defined as follows:*

1. *$N \models A \rhd_B C$ if $C \subseteq A^*_B$,*
2. *$N \models \neg\varphi$ if $N \not\models \varphi$,*
3. *$N \models \varphi \to \psi$ if $N \not\models \varphi$ or $N \models \psi$.*

Since the relation $A \lhd_B C$ is defined in our system, its semantics needs not be given in the definition above. Instead, it is captured by the following additional clause:

Lemma 7 *$N \models A \lhd_B C$ iff sets A^*_B and C are disjoint.*

Proof. By the definition of notation $\lhd$, the formula $A \lhd_B C$ is equivalent to the formula $\bigwedge_{c \in C} \neg(A \rhd_B c)$. Thus, by Definition 5, statement $N \models A \lhd_B C$ is equivalent to the statement that $N \not\models A \rhd_B c$ for each $c \in C$. The latter, again by Definition 5, is equivalent to $c \notin A^*_B$ for each $c \in C$. Therefore, $N \models A \lhd_B C$ iff sets A^*_B and C are disjoint. $\square$

3 Axioms

In addition to propositional tautologies in language Φ, our logical system contains the following axioms.

1. Reflexivity: $A \rhd_B C$, where $C \subseteq A$,
2. Augmentation: $A \rhd_B C \to A, D \rhd_B C, D$,
3. Transitivity: $A \rhd_B C \to (C \rhd_B D \to A \rhd_B D)$,
4. Monotonicity: $A \rhd_B C \to A \rhd_D C$, where $D \subseteq B$,
5. Recalcitrance: $\neg(A \rhd_B C)$, where $B \cap C \neq \varnothing$,
6. Serializability: $A \lhd_{B,C} D \to (A \lhd_{B,D} C \to A \lhd_B C, D)$, where $C \cap D = \varnothing$.

The Reflexivity axiom captures the fact that any set can influence any of its subsets. To understand the meaning of the Augmentation axiom, recall that A, D denotes the union of sets A and D. Then, the Augmentation axiom states that if the set A can influence the set C, then the set A, D can influence the set C, D. The Transitivity axiom says that if the set A can influence the set C and the set C can influence the set D, then the set A can influence the set D. This axiom captures the "indirectness" of the influence relation that we consider.

Note that, together, the Reflexivity, the Augmentation, and the Transitivity axioms (written without subscripts) are known in database theory as Armstrong's axioms [15, p. 81], where they give a sound and complete axiomatization of functional dependency [2]. Some versions of these three axioms have also been used in [3] and [26].

The Monotonicity axiom captures monotonicity on the subscript: if the set A can influence the set C when the set B is recalcitrant, then the set A can also influence the set C when only a subset of B is recalcitrant. The Recalcitrance axiom states that if all agents in the set B are recalcitrant, and the set C overlaps with B, then it is not possible for a set A to influence all agents in C.

The Serializability axiom states that if a set A cannot influence anyone in D when B and C are recalcitrant, and A cannot influence anyone in C when B and D are recalcitrant either, then A cannot influence anyone in either C or D when only B is recalcitrant (if sets C and D are disjoint). Intuitively, this axiom is true because if agents in A could eventually influence someone in C or D when B was recalcitrant, then this influence would have to reach either someone in C "first" (i.e, even when D is recalcitrant), or someone in D "first" (i.e, even when C is recalcitrant). We call this axiom "serializability" because of this "one of two must go first" intuition. This is the key axiom of our system.

We write $X \vdash \varphi$ if formula φ is provable from the formulae in a set X, propositional tautologies, and the additional axioms above, using the Modus Ponens inference rule.

4 Examples of Derivations

We prove the soundness of our logical system in Section 5. Below, we give two examples of formal proofs in this system. These examples are used later in the proof of the completeness. The first of these examples generalizes the Serializability axiom from two to an arbitrary number of assumptions.

Lemma 8 *For any $n \geq 1$ and any disjoint sets $C_1, \dots, C_n \subseteq \mathcal{A}_0$,*

$$\vdash \bigwedge_{1 \leq i \leq n} A \lhd_{B, C_1, \dots, C_{i-1}, C_{i+1}, \dots, C_n} C_i \to A \lhd_B C_1, \dots, C_n.$$

Proof. By induction on integer n. If $n = 1$, then the required statement is a propositional tautology. If $n > 1$, then, by the induction hypothesis, where B is taken to be set $B \cup C_n$,

$$\vdash \bigwedge_{1 \leq i \leq n-1} A \lhd_{B, C_1, \dots, C_{i-1}, C_{i+1}, \dots, C_n} C_i \to A \lhd_{B, C_n} C_1, \dots, C_{n-1}.$$

At the same time, by the Serializability axiom,

$$\vdash A \lhd_{B, C_n} C_1, \dots, C_{n-1} \to (A \lhd_{B, C_1, \dots, C_{n-1}} C_n \to A \lhd_B C_1, \dots, C_n).$$

The required statement then follows from the two formulae above by propositional reasoning. $\Box$

The next lemma states a version of the same serializability principle using relation $\rhd$ instead of relation $\lhd$. This is the form in which this principle will be used in the proof of the completeness. Informally, it says that if set A influences an element c_k of set $\{c_1, \dots, c_n\}$, then there must be an element of this set which is influenced "first".

Lemma 9 *For any $n \geq 1$, any distinct $c_1, \dots, c_n$, and any $k \leq n$,*

$$\vdash A \rhd_B c_k \to \bigvee_{1 \leq i \leq n} A \rhd_{B, c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_n} c_i.$$

Proof. By Lemma 8, assuming $C_1 = \{c_1\}, \dots, C_n = \{c_n\}$,

$$\vdash \bigwedge_{1 \leq i \leq n} A \lhd_{B, c_1, \dots, c_{i-1}, i_{i+1}, \dots, c_n} c_i \to A \lhd_B c_1, \dots, c_n.$$

Thus, by the definition of relation $\lhd$,

$$\vdash \bigwedge_{1 \leq i \leq n} \neg(A \rhd_{B, c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_n} c_i) \to \bigwedge_{1 \leq i \leq n} \neg(A \rhd_B c_i).$$

Hence, by the laws of propositional reasoning,

$$\vdash \bigvee_{1 \leq i \leq n} A \rhd_B c_i \to \bigvee_{1 \leq i \leq n} A \rhd_{B, c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_n} c_i.$$

The latter implies the statement of the lemma by the laws of propositional reasoning. $\Box$

5 Soundness

In this section we prove the soundness of our logical system. It suffices to prove the soundness of each of our axioms, which we prove as separate lemmas. To improve readability, in those lemmas we explicitly recall the syntactical restriction from Definition 1 that sets A and B are disjoint for all well-formed formulae $A \rhd_B C$ and $A \lhd_B C$.

Lemma 10 (Reflexivity) *$N \models A \rhd_B C$ for all disjoint sets A and B such that $C \subseteq A$.*

Proof. Assume that $C \subseteq A$. Then $C^*_B \subseteq A^*_B$, by Lemma 3. By Definitions 3 and 4, $C \subseteq C^*_B$. Therefore, $C \subseteq A^*_B$. Hence, $N \models A \rhd_B C$ by Definition 5. $\square$

Lemma 11 (Augmentation) *If sets A and D are disjoint with set B and $N \models A \rhd_B C$, then $N \models A, D \rhd_B C, D$.*

Proof. Assume that $N \models A \rhd_B C$. Thus, $C \subseteq A^*_B$ by Definition 5. Hence, $C \subseteq A^*_B \subseteq (A \cup D)^*_B$ by Lemma 3.

In addition, $D \subseteq A \cup D = (A \cup D)^0_B \subseteq (A \cup D)^*_B$ by Definitions 3 and 4. Thus, $C \cup D \subseteq (A \cup D)^*_B \cup (A \cup D)^*_B = (A \cup D)^*_B$. Therefore, $N \models A, D \rhd_B C, D$ by Definition 5. $\square$

Lemma 12 (Transitivity) *If $N \models A \rhd_B C$ and $N \models C \rhd_B D$, then $N \models A \rhd_B D$, where sets A and C are disjoint with set B.*

Proof. Suppose that $N \models A \rhd_B C$. Thus, $C \subseteq A^*_B$ by Definition 5. Hence, $C^*_B \subseteq (A^*_B)^*_B$ by Lemma 3. Then, $C^*_B \subseteq A^*_B$ by Lemma 5.

Next, assume $N \models C \rhd_B D$. Thus, $D \subseteq C^*_B$, by Definition 5. Hence, $D \subseteq C^*_B \subseteq A^*_B$, because $C^*_B \subseteq A^*_B$. Therefore, $N \models A \rhd_B D$ by Definition 5. $\square$

Lemma 13 (Monotonicity) *If $N \models A \rhd_B C$ and $D \subseteq B$, then $N \models A \rhd_D C$, where sets A and B are disjoint.*

Proof. Suppose $N \models A \rhd_B C$. Thus, $C \subseteq A^*_B$ by Definition 5. Hence, $C \subseteq A^*_B \subseteq A^*_D$ by Lemma 4. Therefore, $N \models A \rhd_D C$ by Definition 5. $\square$

Lemma 14 (Recalcitrance) *If $A \cap B = \varnothing$ and $B \cap C \neq \varnothing$, then $N \not\models A \rhd_B C$.*

Proof. Consider any $a \in B \cap C$. Then, $a \in B$ and $a \in C$. The former implies that $a \notin A^*_B$ by Lemma 2. Thus, $a \in C \setminus A^*_B$. Therefore, $N \not\models A \rhd_B C$ by Definition 5. $\square$

Lemma 15 (Serializability) *If $N \models A \lhd_{B,C} D$ and $N \models A \lhd_{B,D} C$, then it follows that $N \models A \lhd_B C, D$, where (i) set A is disjoint with sets B, C, and D and (ii) sets C and D are disjoint.*

Proof. By contraposition. Suppose that $N \not\models A \lhd_B C, D$. Thus, $A_B^* \cap (C \cup D) \neq \varnothing$ by Lemma 7. Hence, by Definition 4, there is $n \geq 0$ such that $A_B^n \cap (C \cup D) \neq \varnothing$. Note that $A_B^0 \cap (C \cup D) = \varnothing$ because $A_B^0 = A$ by Definition 3 and because set A is disjoint with sets C and D by assumption (ii). Let m be the smallest $k \geq 1$ such that $A_B^k \cap (C \cup D) \neq \varnothing$. Thus,

$$A_B^{m-1} \cap (C \cup D) = \varnothing \tag{1}$$

and $A_B^m \cap (C \cup D) \neq \varnothing$. Without loss of generality, we assume $A_B^m \cap C \neq \varnothing$. At the same time, $A_B^{m-1} \cap D = \varnothing$ due to statement (1). Hence, $A_B^m = A_{B\cup D}^m \setminus D$ by Lemma 6. Then, $(A_{B\cup D}^m \setminus D) \cap C \neq \varnothing$. Hence, $A_{B\cup D}^m \cap C \neq \varnothing$. Therefore, $N \not\models A \lhd_{B,D} C$ by Lemma 7. □

Together, the six lemmas above prove the soundness of our logical system.

6 Completeness

In this section we prove the completeness of our logical system by constructing a canonical social network $N(X)$ for any maximal consistent set of formulae X, in such a way that $N(X) \models \varphi$ if and only if $\varphi \in X$ for each formula φ. As one would expect, the main challenge is in being able to define the structure (the influence relation) and the threshold function θ of the canonical network $N(X)$ based on the maximal consistent set X.

Let us start by discussing the informal idea behind the construction of the canonical network $N(X)$. Consider any set of agents $D \subseteq \mathcal{A}_0$ and any agent $c \in \mathcal{A}_0 \setminus D$. Note that formula $D \rhd_{\mathcal{A}_0\setminus(D\cup\{c\})} c$ means that set D can influence agent c while all other agents in the network are recalcitrant. Thus, one might suppose that if set X contains formula $D \rhd_{\mathcal{A}_0\setminus(D\cup\{c\})} c$, then in the canonocal social network $N(X)$ set D must be able to influence agent c *directly*. To achieve this, one might add all agents from set D to the set $R^{-1}(c)$ of the influencers of agent c and set threshold value of agent c to be the size of set D as shown in Figure 3.

Unfortunately, this simple idea fails. Indeed, consider the counterexample in Figure 4. Here we assume that $X \vdash D \rhd_{\mathcal{A}_0\setminus(D\cup\{c\})} c$ and $X \vdash E \rhd_{\mathcal{A}_0\setminus(E\cup\{c\})} c$. Since $\theta(c) = 2$ and $|D| = |E|$, in this diagram agents $\{d_1, d_2\}$ can directly influence agent c and, also, agents $\{e_1, e_2\}$ can directly influence agent c. These are desired properties because we started with the assumptions $X \vdash D \rhd_{\mathcal{A}_0\setminus(D\cup\{c\})} c$ and $X \vdash E \rhd_{\mathcal{A}_0\setminus(E\cup\{c\})} c$. At the same time, the way the network in Figure 4 is constructed, we also have that agents d_1 and e_1 together can directly influence agent c because $|\{d_1, e_1\}| = 2 = \theta(c)$. This is a side effect of our construction and it might be inconsistent with the maximal consistent set X, potentially containing formula $\neg(D \rhd_{\mathcal{A}_0\setminus(\{d_1,e_1\}\cup\{c\})} c)$.

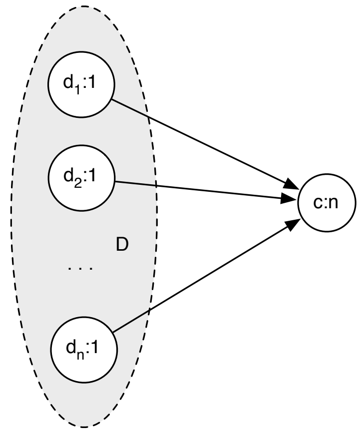

Figure 3: Naïve Construction

For every set D and every agent c such that $X \vdash D \rhd_{\mathcal{A}_0\setminus(D\cup\{c\})} c$, threshold function θ and influence relation R could be potentially defined as shown on the diagram.

To overcome this technical difficulty, one solution is to introduce a new agent x between set D and agent c, and a new agent y between set E and agent c, to set the threshold of x and y to be 2, and the threshold of c to be 1, see Figure 5. This way, only sets D and E can influence c.

We use auxiliary agents to construct our canonical network $N(X)$: for every set $D \subseteq \mathcal{A}_0$ and every agent $c \in \mathcal{A}_0 \setminus D$ such that $X \vdash D \rhd_{\mathcal{A}_0\setminus(D\cup\{c\})} c$, we introduce a distinctive new agent $a(D, c)$ as illustrated in Figure 6.

Definition 6 *For any maximal consistent subset X of Φ we define the canonical social network $N(X) = (\mathcal{A}, \theta, R)$ as follows*

$$\begin{aligned}
\mathcal{A} &= \mathcal{A}_0 \cup \{a(D,c) \mid D \subseteq \mathcal{A}_0, c \notin D, \text{ and } X \vdash D \rhd_{\mathcal{A}_0\setminus(D\cup\{c\})} c\},\\
\theta(a) &= \begin{cases} |D|, & \text{if } a = a(D,c),\\ 1, & \text{if } a \in \mathcal{A}_0,\end{cases}\\
R &= \{(d, a(D,c)) \mid a(D,c) \in \mathcal{A}\setminus\mathcal{A}_0, d \in D\}\\
&\quad \cup \{(a(D,c), c) \mid a(D,c) \in \mathcal{A}\setminus\mathcal{A}_0\}.
\end{aligned}$$

Lemma 16 *Set $\mathcal{A}$ is finite.*

Proof. The statement of the lemma follows from the assumption of finiteness of set $\mathcal{A}_0$ in Section 2. □

The next two lemma follow from Definition 6 in a straightforward way.

Lemma 17 *If $a(A^m_B, c) \in \mathcal{A}$, then $a(A^m_B, c) \in A^{m+1}_B$.*

Lemma 18 *If $a(D, c) \in A^m_B$, then $c \in A^{m+1}_B$.*

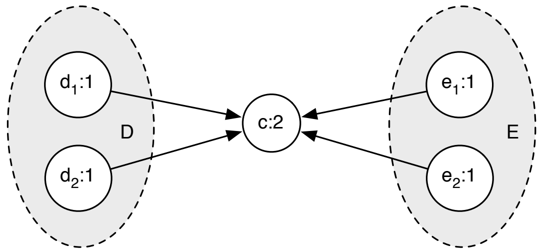

Figure 4: Failure of the Naïve Construction

According to the naïve construction, given that the threshold of agent c is 2, not only can sets D and E influence c, as intended, but sets $\{d_1, e_1\}$, $\{d_2, e_2\}$, $\{d_1, e_2\}$, and $\{d_2, e_1\}$ can also influence c.

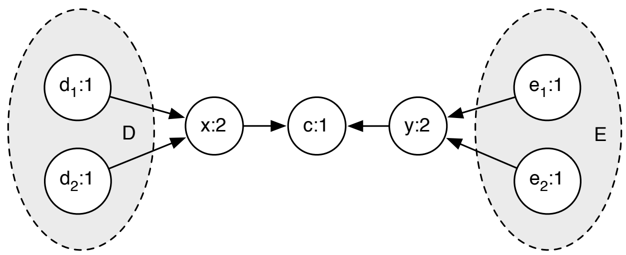

Figure 5: A Solution

New agents x and y are added between sets D and agent c and between set E and agent c, respectively. The thresholds of x and y are set to 2 and the threshold of c is changed to 1. As a result, only sets D and E can influence c, as intended.

Lemma 19 *If $X \vdash A \rhd_B C$, then sets B and C are disjoint.*

Proof. Suppose that $B \cap C \neq \varnothing$. Thus, by the Recalcitrance axiom, $\vdash \neg(A \rhd_B C)$, which contradicts the assumption $X \vdash A \rhd_B C$ and the consistency of set X. $\square$

The key lemma in the proof of completeness is the truth lemma, Lemma 26, that connects derivability with validity in the canonical social network. The following lemma establishes one direction of the base case of Lemma 26.

Lemma 20 *If $X \vdash A \rhd_B C$, then $N(X) \models A \rhd_B C$.*

Proof. Suppose $N(X) \not\models A \rhd_B C$, thus $C \nsubseteq A^*_B$ by Definition 5. Hence, there is $c \in C \subseteq \mathcal{A}_0$ such that $c \notin A^*_B$. Then, $c \notin B$ by Lemma 19 and due to the assumption $X \vdash A \rhd_B C$ of the lemma. Thus,

$$c \in \mathcal{A}_0 \setminus (A^*_B \cup B). \tag{2}$$

Note that $A = A^0_B \subseteq A^*_B$ by Definition 3 and Definition 4. Also, sets A^*_B and B are disjoint by Lemma 2. Thus, $A^*_B \rhd_B A$ is an instance of the Reflexivity

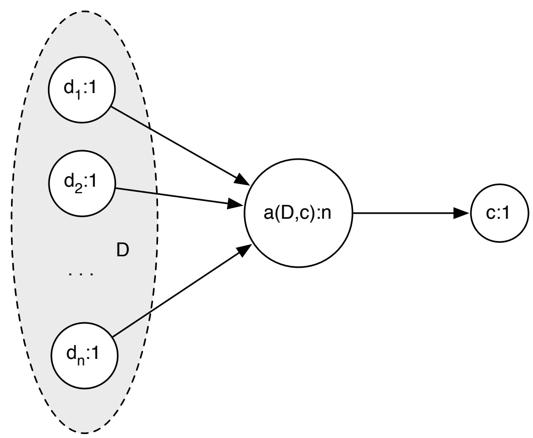

Figure 6: Fragment of the Canonical Network

For every set D and every agent c such that $X \vdash D \rhd_{\mathcal{A}_0 \setminus (D \cup \{c\})} c$, we introduce a distinctive new agent $a(D,c)$. Threshold function θ and influence relation R are defined as shown on the diagram.

axiom. Hence, by the Transitivity axiom and the assumption $X \vdash A \rhd_B C$ of the lemma,

$$X \vdash A^*_B \rhd_B C. \tag{3}$$

Sets B and C are disjoint by Lemma 19. Thus, $C \rhd_B c$ is an instance of the Reflexivity axiom because $c \in C$. Hence, by the Transitivity axiom and due to statement (3),

$$X \vdash A^*_B \rhd_B c \tag{4}$$

Let $c_1, \dots, c_n$ be an enumeration of all agents in set $\mathcal{A}_0 \setminus (A^*_B \cup B)$. Then, by Lemma 9, Lemma 19, and due to statement (2),

$$X \vdash A^*_B \rhd_B c \to \bigvee_{1 \le i \le n} A^*_B \rhd_{B, c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_n} c_i.$$

Thus, statement (4) implies that

$$X \vdash \bigvee_{1 \le i \le n} A^*_B \rhd_{B, c_1, \dots, c_{i-1}, c_{i+1}, \dots, c_n} c_i.$$

Hence, by the maximality of set X, there exists $j \le n$ such that

$$X \vdash A^*_B \rhd_{B, c_1, \dots, c_{j-1}, c_{j+1}, \dots, c_n} c_j.$$

Recall that $c_1, \dots, c_n$ is an enumeration of set $\mathcal{A}_0 \setminus (A^*_B \cup B)$ and that sets A^*_B and B are disjoint by Lemma 2. Hence, set $\mathcal{A}$ contains agent $a(A^*_B, c_j)$ by the definition of set $\mathcal{A}$. Thus, by Lemma 1, there exists $m \ge 0$ such that $a(A^m_B, c_j) = a(A^*_B, c_j) \in \mathcal{A}$. Hence, by Lemma 17,

$$a(A^m_B, c_j) \in A^{m+1}_B. \tag{5}$$

Recall that $c_j \in \mathcal{A}_0 \setminus (A_B^* \cup B)$ by the choice of agents $c_1, \dots, c_n$. Thus, $c_j \notin B$. Hence, statement (5) by Lemma 18 implies that $c_j \in A_B^{m+2}$. Therefore, $c_j \in A_B^*$ by Definition 4, which contradicts the choice of agents $c_1, \dots, c_n$. $\square$

The previous lemma established one direction of the base case of Lemma 26. The other direction is proven later in Lemma 25. Lemmas 21 through 24 are technical lemmas that will be used in the proof of Lemma 25.

Lemma 21 *$X \vdash A_B^{n-1} \rhd_B c$ for each $n \geq 1$ and each $c \in A_B^n \cap \mathcal{A}_0$.*

Proof. If $c \in A$, then $\{c\} \subseteq A \subseteq A_B^{n-1}$ by Definition 3. Therefore, $\vdash A_B^{n-1} \rhd_B c$ by the Reflexivity axiom and due to Lemma 2.

Suppose now that $c \notin A = A_B^0$. Let $m > 0$ be the smallest integer m such that $c \in A_B^m$. Thus, $|A_B^{m-1} \cap R^{-1}(c)| \geq \theta(c)$ by Definition 3. Hence, $|A_B^{m-1} \cap R^{-1}(c)| \geq 1$ by Definition 6. Thus, by Definition 6, there must exist a set $D \subseteq \mathcal{A}_0$ such that

$$X \vdash D \rhd_{\mathcal{A}_0 \setminus (D \cup \{c\})} c \tag{6}$$

and an agent $a(D, c) \in A_B^{m-1}$. The latter implies $D \subseteq A_B^{m-2}$ by Definitions 6 and 3. Thus, by the Reflexivity axiom and Lemma 2,

$$\vdash A_B^{m-2} \rhd_B D. \tag{7}$$

Note that $D \subseteq A_B^{m-2}$ also implies, by Lemma 2, that sets D and B are disjoint. Similarly, assumption $c \in A_B^n$ implies that $c \notin B$. Thus, $B \subseteq \mathcal{A}_0 \setminus (D \cup \{c\})$. Hence, by the Monotonicity axiom, statement (6) implies that $X \vdash D \rhd_B c$. Thus,

$$X \vdash A_B^{m-2} \rhd_B c \tag{8}$$

by the Transitivity axiom due to statement (7).

Finally, note that $m \leq n$ by the choice of integer m and due to the assumption $c \in A_B^n$ of the lemma. Hence, $A_B^{m-2} \subseteq A_B^{n-1}$ by Definition 3. Thus, $\vdash A_B^{n-1} \rhd_B A_B^{m-2}$ by the Reflexivity axiom due to Lemma 2. Therefore, $X \vdash A_B^{n-1} \rhd_B c$ by the Transitivity axiom and due to statement (8). $\square$

Lemma 22 *For any $n \geq 0$ and any agents $c_1, \dots, c_n \in \mathcal{A}_0$, if $X \vdash A \rhd_B c_i$ for each $i \leq n$, then $X \vdash A \rhd_B c_1, \dots, c_n$.*

Proof. Induction on n. If $n = 0$, then we need to show that $X \vdash A \rhd_B \varnothing$, which is an instance of the Reflexivity axiom.

Suppose now that $X \vdash A \rhd_B c_1, \dots, c_{n-1}$. Thus, by the Augmentation axiom, $X \vdash A, c_n \rhd_B c_1, \dots, c_{n-1}, c_n$. At the same time, $X \vdash A \rhd_B c_n$ by the assumption of the lemma. Hence, also by the Augmentation axiom, $X \vdash A \rhd_B A, c_n$. Therefore, by the Transitivity axiom, $X \vdash A \rhd_B c_1, \dots, c_n$. $\square$

Lemma 23 *$X \vdash A_B^{n-1} \rhd_B A_B^n$ for each $n \geq 1$.*

Proof. Note that $A^n_B \subseteq \mathcal{A}$ by Definition 3. Thus, set A^n_B is finite by Lemma 16. Therefore, the required follows from the combination of Lemma 21 and Lemma 22. $\square$

Lemma 24 *$X \vdash A \rhd_B A^n_B$ for each $n \geq 0$.*

Proof. Induction on n. if $n = 0$, then $A^n_B = A$ by Definition3. Thus, $A \rhd_B A^n_B$ is an instance of the Reflexivity axiom.

Suppose now that $X \vdash A \rhd_B A^{n-1}_B$. Therefore, $X \vdash A \rhd_B A^n_B$ by Lemma 23 and the Transitivity axiom. $\square$

Lemma 25 *If $N(X) \models A \rhd_B C$, then $X \vdash A \rhd_B C$.*

Proof. Suppose that $N(X) \models A \rhd_B C$. Thus, $C \subseteq A^*_B$ by Definition 5. Hence, by Lemma 1, there is $n \geq 0$ such that $C \subseteq A^n_B$. Thus, $\vdash A^n_B \rhd_B C$ by the Reflexivity axiom. Therefore, $X \vdash A \rhd_B C$ by the Transitivity axiom and Lemma 24. $\square$

Lemma 26 *$\varphi \in X$ iff $N(X) \models \varphi$, for each formula $\varphi \in \Phi$.*

Proof. We prove this lemma by induction on structural complexity of formula φ. The base case follows from Lemma 20 and Lemma 25. The induction step follows from Definition 5 and the assumption of the maximality and consistency of the set X in the standard way. $\square$

Theorem 1 *For any $\varphi \in \Phi$, if $N \models \varphi$ for each social network N, then $\vdash \varphi$.*

Proof. Suppose that $\nvdash \varphi$. Let X be a maximal consistent set containing formula $\neg\varphi$. Thus, $N(X) \models \neg\varphi$ by Lemma 26. Therefore, $N(X) \not\models \varphi$ by Definition 5. $\square$

By Definition 2, each social network has finitely many agents. Additionally, since we only consider natural numbers as threshold values (as opposed to, say, real numbers), descriptions of social networks could also be assumed to be finite. Hence, the decidability of our logic follows from our completeness theorem in the standard way:

Theorem 2 *The set of all $\varphi \in \Phi$ such that $\vdash \varphi$ is decidable.* $\square$

7 Conclusion

We proposed a sound and complete logical system to reason about influence in social networks, using simple threshold models with recalcitrant agents. The properties of influence were defined in terms of two relations: $\rhd$ and $\lhd$. The most interesting axiom is the Serializability axiom which captures the step-by-step consecutive nature of network diffusion processes. The most original aspect of the article resides in the construction of the canonical model: by invoking a new “in-between” agent each time a group is able to influence an individual outside this group, we are able to represent indirect influence between groups as direct influence between individuals.

Let us briefly look back at our setting's design choices. Of course, recalcitrant agents can be viewed as having very high threshold values and diffusion can be modeled as a step-by-step update process, as was done in [4]. In comparison, the strength of this new logic lies in its ability to express, in a compact form, the indirect and limit behavior of threshold based network diffusion, with (or without) faulty nodes.

Finally, our results are more general than it might appear at first sight: while we have restricted our class of networks to equal weight among influencers, it is easy to see that our axioms are sound for the more general class of weighted networks. Since equally-weighted networks is a subclass of weighted networks, the completeness theorem proved in this article implies completeness with respect to the larger class of weighted networks.

Acknowledgements. Zoé Christoff acknowledges support for this research from the Deutsche Forschungsgemeinschaft (DFG) and Grantová agentura České republiky (GAČR) joint project RO 4548/6–1 “From Shared Evidence to Group Attitudes”.

References

[1] Krzysztof R. Apt and Evangelos Markakis. Diffusion in Social Networks with Competing Products. In G. Persiano, editor, *SAGT 2011*, LNCS 6982, pages 212–223. Springer, 2011.

[2] W. W. Armstrong. Dependency structures of data base relationships. In *Information Processing 74 (Proc. IFIP Congress, Stockholm, 1974)*, pages 580–583. North-Holland, Amsterdam, 1974.

[3] Sanaz Azimipour and Pavel Naumov. Lighthouse principle for diffusion in social networks. *Journal of Logics and their Applications*, 5(1), 2018.

[4] Alexandru Baltag, Zoé Christoff, Rasmus K. Rendsvig, and Sonja Smets. Dynamic epistemic logics of diffusion and prediction in social networks. *Studia Logica*, (online first) July 2018.

[5] Johan van Benthem. Oscillations, logic, and dynamical systems. In Sujata Ghosh and Jakub Szymanik, editors, *The Facts Matter. Essays on Logic and Cognition in Honour of Rineke Verbrugge*, pages 9–22. College Publications, 2015.

[6] Torben Braüner. Hybrid logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2017 edition, 2017.

[7] Zoé Christoff. *Dynamic Logics of Networks: Information Flow and the Spread of Opinion*. PhD thesis, Institute for logic, Language and Computation, University of Amsterdam, Amsterdam, The Netherlands, 2016. ILLC Dissertation Series DS-2016-02.

[8] Zoé Christoff and Davide Grossi. Stability in binary opinion diffusion. In Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, editors, *Logic, Rationality, and Interaction: 6th International Workshop, LORI 2017, Sapporo, Japan, September 11-14, 2017, Proceedings*, pages 166–180, Berlin, Heidelberg, 2017. Springer Berlin Heidelberg.

[9] Zoé Christoff and Jens Ulrik Hansen. A two-tiered formalization of social influence. In Huaxin Huang, Davide Grossi, and Olivier Roy, editors, *Logic, Rationality and Interaction, Proceedings of the Fourth International Workshop (LORI 2013)*, volume 8196 of *Lecture Notes in Computer Science*, pages 68–81. Springer, 2013.

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

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

[12] Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[13] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Syntese Library volume 337. Springer, The Netherlands, 2008.

[14] David Easley and Jon Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, New York, USA, 2010.

[15] Hector Garcia-Molina, Jeffrey Ullman, and Jennifer Widom. *Database Systems: The Complete Book*. Prentice-Hall, second edition, 2009.

[16] Patrick Girard, Jeremy Seligman, and Fenrong Liu. General dynamic dynamic logic. In Thomas Bolander, Torben Brauner, Silvio Ghilardi, and Lawrence Moss, editors, *Advances in Modal Logic, Volume 9*, pages 239–260. College Publication, 2012.

[17] Sanjeev Goyal. Interaction Structure and Social Change. *Journal of Institutional and Theoretical Economics*, 152(3):472–494, 1996.

[18] Umberto Grandi, Emiliano Lorini, and Laurent Perrussel. Propositional opinion diffusion. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 989–997, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.

[19] Mark Granovetter. Threshold Models of Collective Behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978.

[20] Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, 2010.

[21] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *KDD 03: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146. ACM, New York, NY, USA, 2003.

[22] Fa-Hsien Li, Cheng-Te Li, and Man-Kwan Shan. Labeled Influence Maximization in Social Networks for Target Marketing. *2011 IEEE Third Int'l Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third Int'l Conference on Social Computing*, pages 560–563, October 2011.

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

[24] Cheng Long and Raymond Chi-Wing Wong. Minimizing Seed Set for Viral Marketing Paper. *2011 IEEE 11th International Conference on Data Mining*, pages 427–436, December 2011.

[25] Stephen Morris. Contagion. *Review of Economic Studies*, 67:57–78, 2000.

[26] Pavel Naumov and Jia Tao. Marketing impact on diffusion in social networks. *Journal of Applied Logic*, 20:49–74, 2017.

[27] Ji Ruan and Michael Thielscher. A logic for knowledge flow in social networks. In Dianhui Wang and Mark Reynolds, editors, *AI 2011: Advances in Artificial Intelligence*, volume 7106 of *Lecture Notes in Computer Science*, pages 511–520. Springer Berlin Heidelberg, 2011.

[28] Katsuhiko Sano and Satoshi Tojo. Dynamic epistemic logic for channel-based agent communication. In Kamal Lodaya, editor, *Logic and Its Applications*, volume 7750 of *Lecture Notes in Computer Science*, pages 109–120. Springer Berlin Heidelberg, 2013.

[29] Thomas C. Schelling. Models of segregation. *The American Economic Review*, 59(2):488–493, 1969.

[30] Jeremy Seligman, Fenrong Liu, and Patrick Girard. Logic in the community. In Mohua Banerjee and Anil Seth, editors, *Logic and Its Applications*, volume 6521 of *Lecture Notes in Computer Science*, pages 178–188. Springer, 2011.

[31] Jeremy Seligman, Fenrong Liu, and Patrick Girard. Facebook and the epistemic logic of friendship. In Burkhard C. Schipper, editor, *Proceedings of the 14th Conference on Theoretical Aspects of Reasoning about Knowledge*, TARK 2013, pages 207–222, 2013.

[32] Pramesh Singh, Sameet Sreenivason, Boleslaw K. Szymanski, and Gyorgy Korniss. Threshold-limited spreading in social networks with multiple initiators. *Scientific Reports*, 3(2330), 2013.

[33] Liang Zhen and Jeremy Seligman. A logical model of the dynamics of peer pressure. *Electronic Notes in Theoretical Computer Science*, 278:275–288, 2011. Proceedings of the 7th Workshop on Methods for Modalities (M4M 2011) and the 4th Workshop on Logical Aspects of Multi-Agent Systems (LAMAS 2011).