Towards a Logic of Tweeting

Zuojun Xiong[1,2(✉)], Thomas Ågotnes[2,3], Jeremy Seligman[4], and Rui Zhu[4]

[1] Institute of Logic and Intelligence, Southwest University, Chongqing, China
xiongzuojun@gmail.com
[2] Department of Information Science and Media Studies, University of Bergen, Bergen, Norway
thomas.agotnes@infomedia.uib.no
[3] Center for the Study of Language and Cognition, Zhejiang University, Hangzhou, China
[4] Department of Philosophy, University of Auckland, Auckland, New Zealand
j.seligman@auckland.ac.nz, zrui956@aucklanduni.ac.nz

Abstract. In this paper we study the logical principles of a common type of network communication events that haven't been studied from a logical perspective before, namely *network announcements*, or *tweeting*, i.e., simultaneously sending a message to all your friends in a social network. In particular, we develop and study a minimal modal logic for reasoning about propositional network announcements. The logical formalisation helps elucidate core logical principles of network announcements, as well as a number of assumptions that must be made in such reasoning. The main results are sound and complete axiomatisations.

1 Introduction

Formalising reasoning about different types of interaction between multiple agents is an active research topic in logic, artificial intelligence, knowledge representation and reasoning, multi-agent systems, formal specification and verification, and other fields. Most modern approaches are based on modal logic. Despite their ubiquitousness, relatively little attention has been given to formalising reasoning about interaction in *social networks* – with some notable exceptions [2–5,8–14].

In this paper we deal with an issue that has not yet been studied in the sparse but growing literature on logics for social networks. Existing works broadly speaking fall in two main categories; those using formal logic to characterise "global" network phenomena such as cascades (e.g., [2]), and those using formal logic to capture the often subtle details of "local" social network events such as message passing (e.g., [11]). Works in the latter category, in which the current paper falls, have mostly been motivated by capturing events typical in *Facebook*-like applications, such as privately sending a message to a friend (one-to-one messaging). In this paper we formalise reasoning about (what we call) *network announcements* in social networks, the primary communication event on, e.g., *Twitter*: the sending by one agent of a message which received simultaneously

A. Baltag et al. (Eds.): LORI 2017, LNCS 10455, pp. 49–64, 2017.
DOI: 10.1007/978-3-662-55665-8_4

by a number of other agents (the sender's *followers*), determined by the network structure. We will also refer to the act of making network announcements as *tweeting*.

We introduce a minimal modal logic for reasoning about network announcements, having expressions of the form $\langle a : \theta \rangle \varphi$ with the intuitive meaning that if agent a tweets θ, φ will become true. This tweeting operator is not completely original; a very similar operator $[F!\varphi]$ with the meaning "after I announce φ to my friends" was defined semantically in [10] (and also mentioned in [11]), but not systematically studied. In particular, the *logic of network announcements*, their logical principles, axiomatic basis, and so on, has not been studied. Certain aspects of what we call network announcements have been studied in computer science under the term (one-to-many) *multi-cast messaging* [6,7], but not using formal logic.

In this paper we restrict the beliefs of the agents, and thus also the messages they can tweet, to be about basic propositional facts only, as opposed to higher-order beliefs, beliefs about who follows whom, beliefs about who said what, etc. We also make a number of additional idealising assumptions:

Sincerity	Agents only tweet what they believe.
Credulity	Agents believe the messages they receive.
Conservatism	Agents never stop believing what they believed before.
Network stability	Who follows whom after a tweet is the same as before.
Rationality	Agents only believe what follows logically from their previous beliefs and the messages they receive.
Doxastic Omniscience	Agents believe all the logical consequences of what they believe.

These assumptions limit the applicability of the logic, but also allow us to focus on the core concepts of network announcement epistemology. These need to be understood before other more complex issues are addressed. We will see that already several interesting phenomena emerge. In Sect. 6 we discuss the prospect for extensions.

One remaining, natural assumption is the consistency of each agent's beliefs. Instead of building this into our models from the beginning, we develop the logic without any assumption of consistency and then characterise classes of models in which various consistency assumptions hold. In this paper, we will consider two:

Weak Coherence	Each agent has consistent beliefs.
Global Coherence	Agents have mutually consistent beliefs.

The paper is structured as follows. In the next section we introduce the syntax and semantics of the logic, and illustrate what it can express and discuss some of its properties. In Sect. 3 we study logical properties in the form of valid formulas, and the relationship between them, enabling us to form a Hilbert-style axiomatic system that is shown to be complete in Sect. 4. A red thread, and indeed the crux of the completeness proof, is how formulas expressing some agents' beliefs or ignorance after some tweet implicitly contains information about the network structure. In Sect. 5 we look at completeness results for some variants of the logic, and we conclude in Sect. 6.

2 Propositional Network Announcements

Let Agnt and Prop be non-empty sets of agent names and atomic propositional letters, respectively.

Definition 1 (Language). The language of *propositional network announcement logic* is defined by the following grammar, where $p \in \mathsf{Prop}$ and $a \in \mathsf{Agnt}$:

$$\theta ::= p \mid \neg\theta \mid \theta \wedge \theta \qquad \varphi ::= B_a\theta \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle a : \theta\rangle\varphi.$$

Expressions of type θ are called *message formulas* or just *messages*; those of type φ are called *formulas*. The usual derived propositional connectives are used, as well as $[a : \theta]$ for $\neg\langle a : \theta\rangle\neg$. The intended meaning of $B_a\theta$ is that agent a believes θ, while $\langle a : \theta\rangle\varphi$ means that a can tweet θ, after which φ is the case. Formulas of the form $B_a\theta$ are called *belief formulas*.

A model for our language has two parts: an assignment of belief states to each agent and a "following" relation between agents. Recall that a propositional logic *valuation* is a function from Prop to truth values. We denote the set of all valuations Val. We model an agent's belief state by a subset of Val, with no further restrictions. Each message θ determines the set $[\![\theta]\!]$ of those valuations that make it true, according to the usual semantics of propositional logic.

Definition 2 (Models). A *propositional network announcement model* over Agnt and Prop is a pair (F, ω), where the *following relation* F is a binary relation on Agnt and the *belief state function* $\omega\colon \mathsf{Agnt} \to \mathrm{pow}(\mathsf{Val})$ assigns each agent a (possibly empty) set of valuations. We write Fa for the set $\{b \mid bFa\}$ of followers of a.

Note that the subset ordering of belief states is inverse to the strength of the state, so we define $\omega_1 \leq \omega_2$ iff $\omega_2(a) \subseteq \omega_1(a)$ for every a. Any belief of a's in state ω_1 is also a belief in state ω_2.

Definition 3 (Updates). When (only) agent a's belief state is updated with θ, the result is the belief state function $[a \uparrow \theta]\omega$. More generally, the result of simultaneously updating all the agents in a set C of agents with θ is $[C \uparrow \theta]\omega$, where

$$[C \uparrow \theta]\omega(b) = \begin{cases} \omega(b) \cap [\![\theta]\!] & \text{if } b \in C \\ \omega(b) & \text{otherwise} \end{cases}$$

Note that updating is monotonic: $\omega \leq [C \uparrow \theta]\omega$. The language is interpreted in these models as follows.

Definition 4 (Satisfaction). A formula φ of the language of propositional network announcement logic is *satisfied* by a model (F, ω), written $F, \omega \models \varphi$, as follows:

$$\begin{array}{lll} F,\omega \models B_a\theta & \text{iff} & \omega(a) \subseteq [\![\theta]\!] \\ F,\omega \models \neg\varphi & \text{iff} & F,\omega \not\models \varphi \\ F,\omega \models \varphi \wedge \psi & \text{iff} & F,\omega \models \varphi \text{ and } F,\omega \models \psi \\ F,\omega \models \langle a : \theta\rangle\varphi & \text{iff} & F,\omega \models B_a\theta \text{ and } F,[Fa{\uparrow}\theta]\omega \models \varphi \end{array}$$

As usual we say that a formula is *valid* iff it is satisfied by every model. We will also be interested in the class of models in which agents' beliefs are mutually consistent.

Definition 5. A model (F, ω) is *weakly coherent* iff $\omega(a) \neq \emptyset$ for every $a \in \mathsf{Agnt}$. It is *globally coherent* iff $\bigcap_{a\in\mathsf{Agnt}} \omega(a) \neq \emptyset$.

Clearly global coherence implies weak coherence but not vice versa.

2.1 A Simple Example

Figure 1 shows a model of three agents with the following beliefs (this is *all* they believe, modulo logical consequence): Claire believes that the party will be at Anna's place (q); Bill believes that Anna's mother is in town (p); and Anna believes that if her mother is in town the party will not be at her (Anna's) place ($p \to \neg q$). Note that the beliefs are mutually inconsistent (the model is not globally coherent).

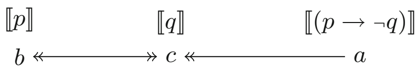

Fig. 1. A simple model. An arrow from a to b means that a follows b.

At this point, each of the three friends have a consistent belief state (the model is *weakly* coherent). Claire can tweet q, after which Anna and Bill will also believe the party is at Anna's place. That's described by the formula $\langle c : q\rangle(B_a q \wedge B_b q)$. Moreover, Bill can tweet p, but the formula $\langle b : p\rangle\neg B_a p$ tells us that still Anna would not believe p. That's because only Claire is following Bill, and so only she will get the message, and before she does she cannot tweet it to Anna: $\neg\langle c : p\rangle B_a p$.

After Claire receives Bill's tweet she can retweet it, and since Anna is following her, she will then believe her mother is in town: $\langle b : p\rangle\langle c : p\rangle B_a p$. Anna also believes $(p \to \neg q)$, so when she receives this message she will also believe $\neg q$, that the party is not going to be at her place. Problems arise for Anna when she receives both of Claire's tweets: one indicating that the party will be at her place, and the other that it will not. Since she is completely credulous, this will leave her in an inconsistent state: $\langle c : q\rangle\langle b : p\rangle\langle c : p\rangle B_a\bot$. Neither Bill nor Claire suffer the same fate. In fact, both will still believe that the party is at Anna's place: $\langle c : q \rangle \langle b : p \rangle \langle c : p \rangle (B_b q \wedge B_c q)$.

The possibility of inconsistent belief states can be regarded as a limitation of our model due to the sometimes unrealistic assumption of credulity: our agents have no way of revising their beliefs. But it can also be regarded as a feature. Even if the agents' beliefs are globally inconsistent, the network structure will allow that inconsistency to emerge in some places but not in others, and this can be described by formulas of our language.

2.2 Conditional Tweeting and Relational Semantics

The dual tweeting operator $[a : \theta]$, can be seen to be a conditional:

Proposition 1. *$F, \omega \models [a : \theta]\varphi$ iff if $F, \omega \models B_a\theta$ then $F, [Fa \uparrow \theta]\omega \models \varphi$.*

Our notation for $\langle a : \theta \rangle$ and $[a : \theta]$ is no accident. Just like the diamond and box of ordinary modal logic, they can be given a relational ("Kripke") semantics. (We omit the straightforward proof of the following.)

Proposition 2 (Relational Semantics). *Define the relation F^a_θ between belief state functions by: $F^a_\theta(\omega_1, \omega_2)$ iff $\omega_1(a) \subseteq [\![\theta]\!]$ and $\omega_2 = [Fa \uparrow \theta]\omega_1$. Then:*

$$F, \omega \models \langle a : \theta \rangle \varphi \quad \textit{iff} \quad F, \nu \models \varphi \textit{ for some } \nu \textit{ such that } F^a_\theta(\omega, \nu)$$
$$F, \omega \models [a : \theta] \varphi \quad \textit{iff} \quad F, \nu \models \varphi \textit{ for every } \nu \textit{ such that } F^a_\theta(\omega, \nu)$$

Define a function V from belief formulas to sets of belief state functions, by $V(B_a\theta) = \{\omega \mid \omega(a) \subseteq [\![\theta]\!]\}$. Let W be the set of belief state functions and let $M(F)$ be the (multi-)modal model (W, F, V) and take our language to be a language of propositional modal logic, with each belief formula considered as a propositional variable and each announcement operator as a modal operator. Then $F, \omega \models \varphi$ iff $M(F), \omega \models \varphi$.

2.3 Potential Belief, and Tracking Ghosts

An agent's *potential beliefs* are those she may acquire as a result of communications from other agents. To describe these clearly we need some notation. Given agents $c_0, \ldots, c_n$ and messages $\theta_0, \ldots, \theta_n$, let $\langle c_0 : \theta_0, \ldots, c_n : \theta_n \rangle$ be an abbreviation for the sequence of tweets $\langle c_0 : \theta_0 \rangle \ldots \langle c_n : \theta_n \rangle$. Let $\vec{c}$ be a variable over expressions of the form $c_0 : \theta_0, \ldots, c_n : \theta_n$, so we can also write the (possibly empty) sequence of tweets as $\langle \vec{c} \rangle$. As for the basic language, we define $[\vec{c}]$ as $\neg \langle \vec{c} \rangle \neg$. The *reversal* of $\vec{c}$, denoted $\overleftarrow{c}$, is the reverse sequence of tweets $c_n : \theta_n, \ldots, c_0 : \theta_0$.

Definition 6. A formula of the form $\langle \vec{c} \rangle B_a\theta$ is a *potential belief formula*. Agent a has a *potential belief* that θ iff $F, \omega \models \langle \vec{c} \rangle B_a\theta$ for some $\vec{c}$. (F, ω) is *potentially equivalent* to (F', ω') iff every agent has the same potential beliefs in (F, ω) as in (F', ω').

In the example, Bill's and Claire's beliefs are always consistent, whereas Anna's are not. She potentially believes a contradiction. The formula $\langle c : q\rangle\langle b : p\rangle\langle c : p\rangle B_a\bot$ that expresses a way in which Anna's beliefs can become inconsistent is an example of a potential belief formula. The concept of potential inconsistency gives a finer-grained picture of the initial distribution of beliefs than merely saying they are mutually inconsistent.

Lemma 1. *If* (F, ω) *is potentially equivalent to* (F', ω') *and* $\omega(a) \subseteq [\![\theta]\!]$*, then* $(F, [F_a\!\uparrow\!\theta]\omega)$ *is potentially equivalent to* $(F', [F'a\!\uparrow\!\theta]\omega')$.

Proof. Assume that (F, ω) is potentially equivalent to (F', ω') and $\omega(a) \subseteq [\![\theta]\!]$. By Definition 6, $\omega'(a) \subseteq [\![\theta]\!]$. Then for any $\vec{c}$, χ, and agent b, the following are equivalent:

$F, [Fa\!\uparrow\!\theta]\omega \models \langle\vec{c}\rangle B_b\chi$	
$F, \omega \models \langle a : \theta\rangle\langle\vec{c}\rangle B_b\chi$	by Definition 4 and $\omega(a) \subseteq [\![\theta]\!]$
$F', \omega' \models \langle a : \theta\rangle\langle\vec{c}\rangle B_b\chi$	by Definition 6 and assumption
$F', [F'a\!\uparrow\!\theta]\omega' \models \langle\vec{c}\rangle B_b\chi$	by Definition 4 and $\omega'(a) \subseteq [\![\theta]\!]$.

Theorem 1. *Propositional network announcement models satisfying the same potential belief formulas are indistinguishable: they satisfy all the same formulas.*

Proof. We prove that for any pair of potentially equivalent models (F, ω) and (F', ω'), $F, \omega \models \varphi$ iff $F', \omega' \models \varphi$ for any φ by induction on φ. (i). The atomic case, $F, \omega \models B_a\theta$ iff $F', \omega' \models B_a\theta$, follows directly from Definition 6. (ii). The boolean cases are straightforward. (iii). For $\varphi = \langle a : \theta\rangle\psi$, the induction hypothesis is that for any pair of potentially equivalent models (F, ω) and (F', ω'), $F, \omega \models \chi$ iff $F', \omega' \models \chi$ for any subformula χ of φ – in particular for $\chi = \psi$. The following are equivalent:

$F, \omega \models \langle a : \theta\rangle\psi$	
$\omega(a) \subseteq [\![\theta]\!]$ and $F, [Fa\!\uparrow\!\theta]\omega \models \psi$	by Definition 4
$\omega'(a) \subseteq [\![\theta]\!]$ and $F', [F'a\!\uparrow\!\theta]\omega' \models \psi$	by (i), i.h. and Lemma 1
$F', \omega' \models \langle a : \theta\rangle\psi$	by Definition 4.

A closely related idea is that of "tracking".

Definition 7. Agent b *tracks* agent a in a model (F, ω) iff $F, \omega \models \langle\vec{c}\rangle B_a\theta \to \langle\vec{c}\rangle B_b\theta$ for every $\vec{c}$ and θ.

If b tracks a then any potential belief of a is also a potential belief of b, but more than that, their potential beliefs are synchronised, in the sense that whenever a acquires a belief, b either acquires it at the same time or already has it.

The interesting point about tracking is that it obscures the following relation. It is possible for an agent to track without following, perhaps by coincidence or just because she already believes every one of some other agent's potential beliefs. If agent b tracks agent a, it is impossible to detect, using the logical language, whether b follows a. We say that b is a *ghost follower* of a if b tracks a without following a. Ghost followers are indistinguishable from real followers.

3 The Logic

We consider various valid principles and properties of propositional network announcement logic, working towards an axiomatisation (shown to be complete in the next section): the system pNAL, shown in Fig. 3. We use $\vdash$ to represent derivability in pNAL. We now introduce the axioms step by step, and along the way present some additional properties that are derivable by the axioms we have introduced "so far".

3.1 The Modal Base

We have observed that the announcement operators have a relational semantics (Proposition 2). It follows that their logic must be an extension of the modal logic K. The belief operator B also has most of the properties of a normal modal operator, except that substitution of propositional variables is restricted to message formulas. By the *modal base* of our logic, we mean the system in Fig. 2. The following is straightforward.

Taut	$\vdash \varphi$	φ subst. inst. of prop. tautology	MP	if $\vdash \varphi \to \psi$ and $\vdash \varphi$ then $\vdash \psi$
K_B	$\vdash B_a(\theta \to \chi) \to (B_a\theta \to B_a\chi)$		Nec_B	if $\vdash_0 \theta$ then $\vdash B_a\theta$
$\mathsf{K}_:$	$\vdash [a:\theta](\varphi \to \psi) \to ([a:\theta]\varphi \to [a:\theta]\psi)$		$\mathsf{Nec}_:$	if $\vdash \varphi$ then $\vdash [a:\theta]\varphi$

Fig. 2. The modal base of pNAL. $\vdash_0$ represents derivability in classical propositional logic.

Proposition 3. *The modal base is sound: every derivable formula is valid.*

Because of this modal base, we can do standard normal modal logic reasoning in our logic (with the syntactic restriction that θ must be propositional in $B_a\theta$), which we will make frequent use of in the following. The base is also enough to show that equivalent formulas can be swapped. The proof (omitted here) is a standard induction on formulas.

Proposition 4 (Replacement of Logical Equivalents).

RLE *if* $\vdash_0 \theta \leftrightarrow \chi$ *then* $\vdash \Theta(\theta) \leftrightarrow \Theta(\chi)$ *if* $\vdash \varphi \leftrightarrow \psi$ *then* $\vdash \Phi(\varphi) \leftrightarrow \Phi(\psi)$

where $\Theta(\chi)$ *is the formula obtained from* $\Theta(\theta)$ *by replacing some instances of* θ *by* χ *or vice versa, and similarly for formulas* $\Phi(\varphi)$ *and* $\Phi(\psi)$.

3.2 Duality and Sincerity

The following dualities are derivable using propositional logic (and replacement of equivalents) alone:

Proposition 5. Dual $\vdash [a:\theta]\neg\varphi \leftrightarrow \neg\langle a:\theta\rangle\varphi \qquad \vdash \neg[a:\theta]\varphi \leftrightarrow \langle a:\theta\rangle\neg\varphi$

But the operators are also linked by our assumption of sincerity, which says that $B_a\theta$ is presupposed by $\langle a:\theta\rangle$ and serves as the antecedent of $[a:\theta]$:

$$\mathsf{Sinc} \quad [a:\theta]\varphi \leftrightarrow (B_a\theta \to \langle a:\theta\rangle\varphi) \qquad \langle a:\theta\rangle\varphi \leftrightarrow (B_a\theta \wedge [a:\theta]\varphi)$$

These are inter-derivable using the modal base; the first is included as an axiom in pNAL. Validity can be easily checked directly from the semantical definitions.

Proposition 6. Sinc *is valid.*

Sinc implies that when the precondition is satisfied, the two operators are equivalent:

Proposition 7. Swap $\vdash B_a\theta \to ([a:\theta]\varphi \leftrightarrow \langle a:\theta\rangle\varphi)$

Using only the precondition principle Sinc and the modal base, we can show the existence of normal forms.

Theorem 2 (Normal form). *Every formula is provably equivalent to one in conjunctive normal form or disjunctive normal form, with atoms generated by*

$$\varphi ::= B_a\theta \mid [a:\theta]\varphi \mid \langle a:\theta\rangle\varphi$$

which is to say: sequences of diamonds and boxes ending in a belief formula. Moreover, the modal depth of the normal form is no greater than the depth of the original formula.

Proof. Given an arbitrary formula, first rewrite (by expanding or introducing abbreviations) so that the only operators are $\neg$, $\wedge$, B_a and $[a:\theta]$. Then, given RLE, we only need note that:

$$\mathsf{Red}_\neg \quad \vdash [a:\theta]\neg\varphi \leftrightarrow \neg(B_a\theta \wedge [a:\theta]\varphi)$$
$$\mathsf{Red}_\wedge \quad \vdash [a:\theta](\varphi\wedge\psi) \leftrightarrow ([a:\theta]\varphi \wedge [a:\theta]\psi)$$

($\mathsf{Red}_\neg$ is tautologically equivalent to an instance of Sinc and $\mathsf{Red}_\wedge$ is just the modally derivable distribution of box over conjunction.)

To see that there is no increase of modal depth (nesting of tweets), it is enough to note that the formulas on either side of the equivalences $\mathsf{Red}_\neg$ and $\mathsf{Red}_\wedge$ are of the same depth.

3.3 Rational Conservative Updating

The direct effect of a tweet is captured by two axioms:

$$\mathsf{Cnsv} \quad B_b\chi \to [a:\theta]B_b\chi \qquad \mathsf{Rat} \quad \langle a:\theta\rangle B_b\chi \to B_b(\theta\to\chi)$$

Conservatism (Cnsv) is our assumption that old beliefs are retained when receiving new information. Rationality (Rat) is our assumption that agents only believe what follows logically from their old beliefs and the content of the tweets they receive. Note that the soundness of these axioms rely on the fact that there are no higher-order beliefs.

Proposition 8. *Both* Cnsv *and* Rat *are valid.*

Proof. For Cnsv recall the previously mentioned monotonicity of updating: $\omega \leq [C \uparrow \theta]\omega$. Since tweeting is a special case of updating, it is also monotonic. For Rat consider two cases. If b is a follower of a, then b's updated state $[F_a \uparrow \theta]\omega(b) = \omega(b) \cap \theta$. But $\omega(b) \cap \theta \subseteq [\![\chi]\!]$ iff $\omega(b) \subseteq [\![\theta \to \chi]\!]$. If b is not a follower of a then $[F_a \uparrow \theta]\omega(b) = \omega(b)$ and $\omega(b) \subseteq [\![\chi]\!]$ implies $\omega(b) \subseteq [\![\theta \to \chi]\!]$.

The implication in Rat can be turned into an equivalence under the assumption that b believes what a tweets. The following can be proved using Rat and Cnsv.

Proposition 9. Up $\vdash \langle a : \theta \rangle B_b \theta \to (\langle a : \theta \rangle B_b \chi \leftrightarrow B_b(\theta \to \chi))$

Proof. One half of the equivalence is a weakening of Rat. For the other, note that $[a : \theta]B_b(\theta \to \chi) \to (\langle a : \theta \rangle B_b \theta \to \langle a : \theta \rangle B_b \chi)$ is derivable by purely modal reasoning. And $B_b(\theta \to \chi) \to [a : \theta]B_b(\theta \to \chi)$ is an instance of Cnsv.

3.4 Following

Given the problem of ghosts described in Sect. 2.3, there can be no formula that exactly defines the relation of following. Nonetheless, some sufficient and necessary conditions are expressible.

Proposition 10. *Given a and b, for any $\vec{c}$ and any θ,*

(Sufficient) If $F, \omega \models \langle \vec{c} \rangle (\neg B_b \chi \wedge \langle a : \theta \rangle B_b \chi)$ then bFa.
(Necessary) If bFa then $F, \omega \models [\vec{c}][a : \theta]B_b \theta$

Proof. For the sufficient condition, suppose $F, \omega \models \langle \vec{c} \rangle (\neg B_b \chi \wedge \langle a : \theta \rangle B_b \chi)$. Let ω' be the result of updating ω according to $\langle \vec{c} \rangle$. Then $F, \omega' \not\models B_b \chi$ but $F, [Fa \uparrow \theta]\omega' \models B_b \chi$. So $\omega'(b) \neq [Fa \uparrow \theta]\omega'(b)$ and so $b \in Fa$.

For the necessary condition, suppose bFa. Then either one of the preconditions in evaluating $[\vec{c}]$ fails, and in which case the formula is satisfied, or they all succeed. In that case, let ω' be as before. Since $b \in Fa$, $[Fa \uparrow \theta]\omega'(b) = \omega'(b) \cap [\![\theta]\!]$ and so $F, [Fa \uparrow \theta]\omega' \models B_b \theta$. Thus $F, \omega \models [\vec{c}][a : \theta]B_b \theta$.

Our approach to the logic, then, is to include an axiom saying that the sufficient condition implies the necessary condition:

$$\text{Foll} \quad \langle \vec{c} \rangle (\neg B_b \chi \wedge \langle a : \chi' \rangle B_b \chi) \to [\vec{e}][a : \theta]B_b \theta$$

3.5 Network Stability

The following axiom captures the assumption that the only thing that changes anything is tweeting events, and that an "empty" tweet (of a tautology) changes nothing.

$$\text{Null} \quad \text{if} \vdash_0 \theta \text{ then} \vdash \varphi \leftrightarrow \langle a : \theta \rangle \varphi$$

Null is in fact the last axiom we need in order to get a complete axiomatic system, as we shall see in the next section. We end this section with mentioning two additional properties related to network stability. They are both already derivable (follows from the completeness result in the next section).

First, tweeting does not affect the network structure; in fact, the following relation is kept fixed. That's our assumption of network stability. Consider the following property:

$$\mathsf{Stab}\quad \langle b:\chi\rangle B_c\delta \to [a:\theta]\langle b:\chi\rangle B_c\delta$$

Stab is a necessary but insufficient condition of network stability; and that's the best we can do. There is no sufficient condition available. Fortunately, we don't need one to get a complete axiomatisation, as we shall see in the next section.

Network stability is needed for many principles involving the iteration of tweets; in particular, for moving a conditional tweet to the beginning of a sequence of tweets:

$$\mathsf{Perm}\quad \langle \vec{a}\,\rangle[b:\chi]\varphi \to [b:\chi]\langle \vec{a}\,\rangle\varphi$$

Proposition 11. Null*,* Stab *and* Perm *are valid.*

Proof. The case for Null is trivial. We show the case for Stab, Perm is (also) straightforward. Suppose $F,\omega \models \langle b:\chi\rangle B_c\delta$ then (1) $F,\omega \models B_b\chi$ and (2) $F,[Fb{\uparrow}\chi]\omega \models B_c\delta$. Now suppose $F,\omega \models B_a\theta$. Then (3) $F,[Fa\uparrow\theta]\omega \models B_b\chi$ since (1) and $[Fa\uparrow\theta]\omega(b) \subseteq \omega(b)$. We also have (4) $F,[Fb\uparrow\chi]([Fa\uparrow\theta]\omega) \models B_c\delta$ since $[Fb{\uparrow}\chi]([Fa{\uparrow}\theta]\omega) = [Fa{\uparrow}\theta]([Fb{\uparrow}\chi]\omega) \subseteq [Fb{\uparrow}\chi]\omega$ and (2). Thus by (3) and (4), $F,\omega \models [a:\theta]\langle b:\chi\rangle B_c\delta$.

Taut	if $\vdash_0 \varphi$ then $\vdash \varphi$	MP	if $\vdash \varphi \to \psi$ and $\vdash \varphi$ then $\vdash \psi$
K_B	$\vdash B_a(\theta \to \chi) \to (B_a\theta \to B_a\chi)$	$\mathsf{K}_:$	$\vdash [a:\theta](\varphi \to \psi) \to ([a:\theta]\varphi \to [a:\theta]\psi)$
Nec_B	if $\vdash_0 \theta$ then $\vdash B_a\theta$	$\mathsf{Nec}_:$	if $\vdash \varphi$ then $\vdash [a:\theta]\varphi$
Sinc	$\vdash [a:\theta]\varphi \leftrightarrow (B_a\theta \to \langle a:\theta\rangle\varphi)$	Cnsv	$\vdash B_b\chi \to [a:\theta]B_b\chi$
Rat	$\vdash \langle a:\theta\rangle B_b\chi \to B_b(\theta \to \chi)$	Foll	$\vdash \langle\vec{c}\rangle(\neg B_b\chi \wedge \langle a:\chi'\rangle B_b\chi) \to [\vec{e}][a:\theta]B_b\theta$
Null	if $\vdash_0 \theta$ then $\vdash \varphi \leftrightarrow \langle a:\theta\rangle\varphi$		

Fig. 3. Axioms and rules of pNAL.

4 Completeness

We show that the system pNAL, displayed in Fig. 3 is (strongly) complete. We assume the usual concepts of consistency, maximal consistency, and logical closure. The Lindenbaum result that any consistent set of formulas can be extended to a maximal consistent set holds for standard reasons. We have shown the derivability of various additional principles (Dual, Swap, and Up) that will be used below (Propositions 5, 7 and 9). All that remains is to use a maximal consistent set of formulas Γ to construct a following relation F_Γ and a belief state function ω_Γ for which we can prove a Truth Lemma (that F_Γ, ω_Γ satisfies every formula in Γ). So we define:

$$bF_\Gamma a \text{ iff } [\vec{c}][a:\theta]B_b\theta \in \Gamma \text{ for all } \vec{c} \text{ and } \theta \qquad \omega_\Gamma(a) = \bigcap\{[\![\theta]\!] \mid B_a\theta \in \Gamma\}.$$

Note that the definition of F_Γ uses the set of all formulas providing *necessary* conditions for following, as identified in Proposition 10. In a given model, any ghost follower of a will also satisfy all these conditions. Our approach is to build a model in which all trackers of a are taken to be followers.

Regarding belief formulas, the definition of ω_Γ does it's job.

Lemma 2. *$\omega_\Gamma(a) \subseteq [\![\theta]\!]$ iff $B_a\theta \in \Gamma$ for any a and θ,*

Proof. Right-to-left is immediate since $\omega_\Gamma(a) = \bigcap\{[\![\chi]\!] \mid B_a\chi \in \Gamma\}$. For the other direction, by completeness of propositional logic, $\vdash_0 (\chi_1 \wedge \ldots \wedge \chi_n) \to \theta$ for some $B_a\chi_1, \ldots B_a\chi_n \in \Gamma$. Then $\vdash B_a((\chi_1 \wedge \ldots \wedge \chi_n) \to \theta)$ by Nec_B and so $\vdash (B_a\chi_1 \wedge \ldots \wedge B_a\chi_n) \to B_a\theta$ by more modal reasoning using K_B and Nec_B. Thus $B_a\theta \in \Gamma$.

Lemma 2 is the obvious base case of an attempt to prove the Truth Lemma by induction on the structure of formulas. But such a direct approach doesn't work because the clause for $\langle a:\theta\rangle$ requires us to update the model. There are several options here. A first thought is to try to construct the set of formulas satisfied by the updated model, i.e., to find a maximal consistent set Γ' such that $F_{\Gamma'} = F_\Gamma$ and $\omega_{\Gamma'} = [F_\Gamma a \uparrow \theta]\omega_\Gamma$. But the search for something satisfying the first of these conditions is plagued by ghosts: each time the model is updated, new ghost followers may appear. Instead, we'll construct a Γ' to meet only the second condition, using a syntactic update operation:

$$\langle a:\theta\rangle\Gamma = \{\varphi \mid \langle a:\theta\rangle\varphi \in \Gamma\}.$$

Our proof of the Truth Lemma (Lemma 5) will involve a strengthening of it that quantifies over sets obtained by repeated applications of syntactic update. This requires new notation and a some technical lemmas.

Definition 8. Define the relation $\trianglelefteq$ between sets of formulas as follows: $\Gamma \trianglelefteq \Gamma'$ iff $B_a\theta \in \Gamma$ and $\Gamma' = \langle a:\theta\rangle\Gamma$ for some a and θ. Let $\leq$ be the transitive closure of $\trianglelefteq$.

Lemma 3. *If Γ is a maximal consistent set and $\Gamma \leq \Gamma'$ then*

1. *Γ' is also a maximal consistent set and*
2. *there is a $\vec{c}$ such that: (a) $\Gamma' = \langle\vec{c}\rangle\Gamma$, and (b) $[\overleftarrow{c}]\varphi \in \Gamma$ iff $\varphi \in \Gamma'$ for all φ, where $\overleftarrow{c}$ is the reversal of $\vec{c}$.*

Proof. By induction on the length of the shortest chain $\Gamma \trianglelefteq \ldots \trianglelefteq \Gamma'$. In the base case, $\Gamma' = \Gamma$ and we can take $\vec{c}$ to be the empty sequence. So now suppose the length of the shortest chain is strictly positive. Then there is a Γ'' such that $\Gamma \trianglelefteq \Gamma'' \leq \Gamma'$. Note that the chain $\Gamma'' \trianglelefteq \ldots \trianglelefteq \Gamma'$ is shorter. By definition of $\trianglelefteq$, there are a and θ such that $B_a\theta \in \Gamma$ and $\Gamma'' = \langle a : \theta\rangle\Gamma$.

1. We first show that Γ'' is a maximal consistent set: $\neg\varphi \in \Gamma''$ iff $\neg\varphi \in \langle a : \theta\rangle\Gamma$ (by $\Gamma'' = \langle a : \theta\rangle\Gamma$) iff $\langle a : \theta\rangle\neg\varphi \in \Gamma$ (by def. of $\langle a : \theta\rangle\Gamma$) iff $\neg[a : \theta]\varphi \in \Gamma$ (by Dual) iff $\neg\langle a : \theta\rangle\varphi \in \Gamma$ (by Swap and $B_a\theta \in \Gamma$) iff $\langle a : \theta\rangle\varphi \notin \Gamma$ iff $\varphi \notin \langle a : \theta\rangle\Gamma$ (by def. of $\langle a : \theta\rangle\Gamma$) iff $\varphi \notin \Gamma''$ ($\Gamma'' = \langle a : \theta\rangle\Gamma$). Now, since there is a shorter chain from Γ'' to Γ' and Γ'' is a maximal consistent set we can apply the induction hypothesis, to get that Γ' is a maximal consistent set.
2. Also from the induction hypothesis: there is $\vec{c}$ such that (1) $\Gamma' = \langle\vec{c}\rangle\Gamma''$ and (2) $[\overleftarrow{c}]\varphi \in \Gamma''$ iff $\varphi \in \Gamma'$ for all φ. So let $\vec{e} = \vec{c}, a : \theta$. Then from (1): $\Gamma' = \langle\vec{c}\rangle\Gamma'' = \langle\vec{c}\rangle\langle a : \theta\rangle\Gamma = \langle\vec{c}, a : \theta\rangle\Gamma = \langle\vec{e}\rangle\Gamma$. We get the following equivalences: $[\overleftarrow{e}]\varphi \in \Gamma$ iff $[a : \theta, \overleftarrow{c}]\varphi \in \Gamma$ (def. of $\overleftarrow{e}$) iff $[a : \theta][\overleftarrow{c}]\varphi \in \Gamma$ (definition of $[a : \theta, \overleftarrow{c}]$) iff $\langle a : \theta\rangle[\overleftarrow{c}]\varphi \in \Gamma$ (Swap, $B_a\theta \in \Gamma$) iff $[\overleftarrow{c}]\varphi \in \langle a : \theta\rangle\Gamma$ (def. of $\langle a : \theta\rangle\Gamma$) iff $[\overleftarrow{c}]\varphi \in \Gamma''$ ($\Gamma'' = \langle a : \theta\rangle\Gamma$) iff $\varphi \in \Gamma'$ (2).

Lemma 3 enables us to show that belief state functions behave properly under updates.

Lemma 4. *For any m.c.s. Γ, if $\Gamma \leq \Gamma'$ and $B_a\theta \in \Gamma'$ then $[F_\Gamma a \uparrow \theta]\omega_{\Gamma'} = \omega_{\langle a:\theta\rangle\Gamma'}$.*

Proof. Suppose $\Gamma \leq \Gamma'$ and $B_a\theta \in \Gamma'$. We first need a fact about propositional logic:

Claim: $\bigcap\{[\![\chi]\!] \mid \theta \to \chi \in \Delta\} = \bigcap\{[\![\chi]\!] \mid \chi \in \Delta\} \cap [\![\theta]\!]$
for any formula θ and any logically closed set Δ of formulas (of propositional logic). It can be proved easily from the deduction theorem.

By Lemma 3, $\Gamma' = \langle\vec{c}\rangle\Gamma$ for some $\vec{c}$ and this is a maximal consistent set. Let b be in Agnt. We will show that $\omega_{\langle a:\theta\rangle\Gamma'}(b) = [F_\Gamma a \uparrow \theta]\omega_{\Gamma'}(b)$. We have two cases depending on whether or not b is tracking a:

$b \in F_\Gamma a$ Then $[\overleftarrow{c}][a : \theta]B_b\theta$ is in Γ. By Lemma 3.2, $[a : \theta]B_b\theta$ is in $\langle\vec{c}\rangle\Gamma$. But $B_a\theta$ is in Γ' and so by Swap so is $\langle a : \theta\rangle B_b\theta$. From this, Up tells us that $\langle a : \theta\rangle B_b\chi \leftrightarrow B_b(\theta \to \chi)$ is in Γ' for any χ, but also by definition of $\langle a : \theta\rangle\Gamma'$, we know that $\langle a : \theta\rangle B_b\chi \in \Gamma'$ iff $B_b\chi \in \langle a : \theta\rangle\Gamma'$. Putting these together: $B_b\chi \in \langle a : \theta\rangle\Gamma'$ iff $B_b(\theta \to \chi) \in \Gamma'$.

So
$$\begin{aligned}\bigcap\{[\![\chi]\!] \mid B_b\chi \in \langle a : \theta\rangle\Gamma'\} &= \bigcap\{[\![\chi]\!] \mid B_b(\theta \to \chi) \in \Gamma'\}\\ &= \bigcap\{[\![\chi]\!] \mid B_b\chi \in \Gamma'\} \cap [\![\theta]\!] \qquad \text{by Claim, with } \Delta = \{\chi \mid B_b\chi \in \Gamma'\}\end{aligned}$$

Thus
$$\omega_{\langle a:\theta\rangle\Gamma'}(b) = [F_\Gamma a \uparrow \theta]\omega_{\Gamma'}(b)$$

$b \notin F_\Gamma a$ Then there is some $\vec{e}$ and some θ' for which $[\vec{e}][a : \theta']B_b\theta$ is not in Γ. Foll then tells us that $\langle\overleftarrow{c}\rangle(\neg B_b\chi \wedge \langle a : \theta\rangle B_b\chi)$ is not in Γ for any χ. But Γ is a maximal consistent set so it does contain $[\overleftarrow{c}](\langle a : \theta\rangle B_b\chi \to B_b\chi)$. So by Lemma 2, $\langle a : \theta\rangle B_b\chi \to B_b\chi$ is in $\langle\vec{c}\rangle\Gamma = \Gamma'$. We also have that $B_b\chi \to \langle a : \theta\rangle B_b\chi$ is in Γ'. (This is by Cnsv and Swap since $B_a\theta \in \Gamma'$.) Thus:

$$\begin{array}{ll} & B_b\chi \in \Gamma' \text{ iff } \langle a:\theta\rangle B_b\chi \in \Gamma' \\ & \qquad\qquad \text{iff } B_b\chi \in \langle a:\theta\rangle\Gamma' \quad \text{Defn. } \langle a:\theta\rangle\Gamma' \\ \text{Hence} & \bigcap\{[\![\chi]\!] \mid B_b\chi \in \langle a:\theta\rangle\Gamma'\} = \bigcap\{[\![\chi]\!] \mid B_b\chi \in \Gamma'\} \\ \text{Thus} & \omega_{\langle a:\theta\rangle\Gamma'}(b) = [F_\Gamma a\!\uparrow\!\theta]\omega_{\Gamma'}(b) \end{array}$$

Lemma 5 (Truth Lemma). *$F_\Gamma, \omega_\Gamma \models \varphi$ iff $\varphi \in \Gamma$, for any formula φ and m.c.s. Γ.*

Proof. Let Γ be a maximal consistent set. We prove that for any Γ', if $\Gamma \leq \Gamma'$ then $F_\Gamma, \omega_{\Gamma'} \models \psi$ iff $\psi \in \Gamma'$, by induction on ψ. The base case $\psi = B_a\theta$ follows from Lemma 2, and the cases for negation and conjunction are straightforward.

Consider the case that $\psi = \langle a : \theta\rangle\varphi$. Note that Γ' is a maximal consistent set by Lemma 3.1. The following are equivalent:

$F_\Gamma, \omega_{\Gamma'} \models \langle a : \theta\rangle\varphi$	
$F_\Gamma, \omega_{\Gamma'} \models B_a\theta$ and $F_\Gamma, [F_\Gamma a\!\uparrow\!\theta]\omega_{\Gamma'} \models \varphi$	by semantics (Definition 4)
$B_a\theta \in \Gamma'$ and $F_\Gamma, [F_\Gamma a\!\uparrow\!\theta]\omega_{\Gamma'} \models \varphi$	by the $B_a\theta$ case, above
$B_a\theta \in \Gamma'$ and $F_\Gamma, \omega_{\langle a:\theta\rangle\Gamma'} \models \varphi$	by Lemma 4, since $\Gamma \leq \Gamma'$
$B_a\theta \in \Gamma'$ and $\varphi \in \langle a : \theta\rangle\Gamma'$	by I.H., since $\Gamma \leq \langle a : \theta\rangle\Gamma'$
$B_a\theta \in \Gamma'$ and $\langle a : \theta\rangle\varphi \in \Gamma'$	by definition of $\langle a : \theta\rangle\Gamma'$
$\langle a : \theta\rangle\varphi \in \Gamma'$	by Sinc and closure of Γ'

This completes the induction. Finally, let $\top$ be any tautology. Then $B_a\top$ by Nec_B and $\Gamma = \langle a : \top\rangle\Gamma$ by Null. Thus $\Gamma \trianglelefteq \Gamma$ and so $\Gamma \leq \Gamma$, and the result follows.

When Γ is a set of formulas and φ a formula, we write $\Gamma \models \varphi$ to mean that any model satisfying Γ also satisfies φ, and $\Gamma \vdash \varphi$ to mean that there is a theorem $(\varphi_1 \wedge \ldots \wedge \varphi_n) \to \varphi$ of pNAL for some finite sequence of formulas $\varphi_1, \ldots, \varphi_n$ in Γ.

Theorem 3 (Soundness and Completeness). *For any Γ and φ, $\Gamma \models \varphi$ iff $\Gamma \vdash \varphi$.*

Proof. Soundness follows from validity of the axioms and rules (Propositions 3, 6, 8, 10 and 11). (Strong) completeness follows from the Truth Lemma.

5 Variants

We have a brief look at some natural variants of the logic.

5.1 Irreflexivity

Do agents follow themselves? We have not assumed that they *don't*; we allow models were agents do follow themselves. Indeed, the canonical model in the previous section is reflexive – all agents always follow themselves. However, it is easy to see that self-following cannot be expressed in our logical language. I.e., we have the following property.

Proposition 12. *For any formula* φ *and any model* (F,ω), $F,\omega \models \varphi$ *iff* $F^-,\omega \models \varphi$*, where* F^- *is the largest irreflexive submodel of* F *(i.e.,* $F^- = F \setminus \{(a,a) : a \in \mathsf{Agnt}\}$*).*

We thus immediately get the following corollary of the results in the previous section (by taking F to be the (reflexive) canonical following relation).

Corollary 1. pNAL *is sound and strongly complete with respect to the class of models with an irreflexive following relation.*

5.2 Coherence

Our logic can easily be extended to axiomatise the class of weakly or globally coherent models. Consider the following schemata (n ranges over positive natural numbers):

WCoh if $\vdash_0 \neg\theta$ then $\vdash \neg B_a\theta$
GCoh if $\vdash_0 \neg(\theta_1 \wedge \ldots \wedge \theta_n)$ then $\vdash \neg(B_{a_1}\theta_1 \wedge \cdots \wedge B_{a_n}\theta_n)$

Let the wCpNAL be the axiom system pNAL extended with WCoh and let gCpNAL be the axiom system pNAL extended with GCoh. The following can be easily checked.

Lemma 6. *A model is weakly (globally) coherent iff it satisfies all inst. of* WCoh *(*GCoh*).*

From this, and the fact that the rules of the logic preserve validity on the classes of weakly and globally coherent models, respectively, we immediately get the following.

Theorem 4. wCpNAL *and* gCpNAL *are sound and strongly complete with respect to the classes of weakly and globally coherent models, respectively.*

6 Discussion

In this paper we laid the groundwork for formal reasoning about network announcements in social networks ("tweeting"). We defined a minimal modal logic based on (not necessarily consistent) propositional beliefs and a "tweeting" modality $\langle a : \theta \rangle$, and studied the logic in detail. We believe that this detailed study lays a solid foundation for richer frameworks to be studied in the future. For example, the technique we used for encoding network structure using logical formulas in the completeness proof is general. We made several assumptions clear in the beginning of the paper, some of which showed up again as axioms of the logic. It could be interesting to investigate a weakening of some of these assumptions starting from a syntactic angle, by weakening the axioms. Regarding coherence, there is a third, natural, form that we haven't considered in this paper: no agent can enter an inconsistent belief state as a result of network announcements ("local coherence").

There are two main and orthogonal directions for future work. The first is extending the *semantics* to model agents with higher-order beliefs and possibly even beliefs about the network structure. Higher-order beliefs would introduce a number of subtleties and complications and would require a number of assumptions. For example, with higher order beliefs assumptions would have to be made about different agents' beliefs about the possibility of tweeting events taking place, belief states are would no longer be monotonic under tweeting, the belief state of the tweeter would not be static under tweeting, and so on. With incomplete information about the network structure, new beliefs about that structure could actually be formed as a result of receiving tweets in certain situations. The second direction is extending the *syntax*. One natural and interesting possibility for enriching the language is to add modalities of the form $\langle a \rangle$ (where a is an agent) *quantifying* over tweets, known from group announcement logic [1], where a formula of the form $\langle a \rangle \varphi$ would intuitively mean that a can make φ true by tweeting some message. Such operators can potentially be used to capture many interesting phenomena related to the information flow in social networks.

Acknowledgments. The first author is supported by Projects of the National Social Science Foundation of China under research no. 15AZX020.

References

1. Ågotnes, T., Balbiani, P., van Ditmarsch, H., Seban, P.: Group announcement logic. J. Appl. Logic **8**(1), 62–81 (2010)
2. Baltag, A., Christoff, Z., Ulrik Hansen, J., Smets, S.: Logical Models of Informational Cascades. Studies in Logic. College Publications, London (2013)
3. Christoff, Z.: Dynamic Logics of Networks. Ph.D. thesis, University of Amsterdam (2016)
4. Christoff, Z., Ulrik Hansen, J.: A logic for diffusion in social networks. J. Appl. Logic **13**(1), 48–77 (2015)
5. Ulrik Hansen, J.: Reasoning about opinion dynamics in social networks. J. Logic Comput. (2015). https://doi.org/10.1093/logcom/exv083
6. Harte, L.: Introduction to Data Multicasting. Althos Publishing, Raleigh (2008)
7. Hosszú, G.: Mediacommunication based on application-layer multi-cast. In: Encyclopedia of Virtual Communities and Technologies, pp. 302–307. IGI Global (2006)
8. Pacuit, E., Parikh, R.: The logic of communication graphs. In: Leite, J., Omicini, A., Torroni, P., Yolum, I. (eds.) DALT 2004. LNCS, vol. 3476, pp. 256–269. Springer, Heidelberg (2005). doi:10.1007/11493402_15
9. Ruan, J., Thielscher, M.: A logic for knowledge flow in social networks. In: Wang, D., Reynolds, M. (eds.) AI 2011. LNCS, vol. 7106, pp. 511–520. Springer, Heidelberg (2011). doi:10.1007/978-3-642-25832-9_52
10. Seligman, J., Liu, F., Girard, P.: Logic in the community. In: Banerjee, M., Seth, A. (eds.) ICLA 2011. LNCS, vol. 6521, pp. 178–188. Springer, Heidelberg (2011). doi:10.1007/978-3-642-18026-2_15
11. Seligman, J., Liu, F., Girard, P.: Facebook and the epistemic logic of friendship. In: Schipper, B.C. (ed.) Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge, Chennai, India, pp. 229–238 (2013)

12. van Ditmarsch, H., van Eijck, J., Pardo, P., Ramezanian, R., Schwarzentruber, F.: Gossip in dynamic networks. Liber Amicorum Alberti. A Tribute to Albert Visser, pp. 91–98. College Publications, London (2016)
13. van Eijck, J., Sietsma, F.: Message-generated kripke semantics. In: The 10th International Conference on Autonomous Agents and Multiagent Systems, vol. 3, pp. 1183–1184 (2011)
14. Wang, Y., Sietsma, F., Eijck, J.: Logic of information flow on communication channels. In: Omicini, A., Sardina, S., Vasconcelos, W. (eds.) DALT 2010. LNCS, vol. 6619, pp. 130–147. Springer, Heidelberg (2011). doi:10.1007/978-3-642-20715-0_8