Facebook and the epistemic logic of friendship

Extended Abstract

Jeremy Seligman
Department of Philosophy
The University of Auckland
Auckland, New Zealand

Fenrong Liu
Department of Philosophy
Tsinghua University
Beijing, China

Patrick Girard
Department of Philosophy
The University of Auckland
Auckland, New Zealand

This paper presents a two-dimensional modal logic for reasoning about the changing patterns of knowledge and social relationships in networks organised on the basis of a symmetric 'friendship' relation, providing a precise language for exploring 'logic in the community' [7]. Agents are placed in the model, allowing us to express such indexical facts as 'I am your friend' and 'You, my friends, are in danger'.

We investigate a number of conceptual and technical issues that arise when considering communication between agents in such networks, both from one agent to another, and broadcasts to socially-defined groups of agents, such as the group of my friends.

The technical framework for this work is [4]. We have had to develop that to cope with the present two-dimensional setting, but many of the technical details have been omitted from this abstract, to allow us to focus on the interesting conceptual distinctions that we have encountered. A more complete treatment will be given in the full paper.

In particular, we extend the treatment of announcements to questions, in which agents are taken to be sincere and cooperative interlocutors, consider network structure changing operations such as adding and deleting friends (with the permission of other agents) and the effects of all this on the concept of common knowledge, which is more varied and rich in the social network setting.

This is just a starting point. The range of epistemic subtleties that emerge only from consideration of two relations (epistemic indistinguishability and symmetric friendship) suggest many interest explorations into the function of belief and preference, along the lines begun in our work [6].

1. A LANGUAGE OF SOCIAL KNOWING

We start with a language $\mathcal{L}$ of *epistemic friendship logic* EFL based on atoms of two types: propositional variables $\rho \in \mathsf{Prop}$ representing indexical propositions such as 'I am in danger', and agent nominals $n \in \mathsf{ANom}$ which stand for indexical propositions asserting identity, such as 'I am Dr. Livingstone'. The language is then inductively defined as:

$$\begin{aligned}\varphi \quad ::= \quad & \rho \mid n \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K\varphi \mid F\varphi \mid \\ & U_A\varphi \mid U_W\varphi \mid \ \downarrow\! n\ \varphi\end{aligned}$$

We read K as 'I know that' and F as 'all my friends', U_A as 'every agent' and U_W as 'in every state'. The binder $\downarrow n$ provides a way of referring to 'me' inside the scope of other operators. For example, $\downarrow n\ FK\langle F\rangle n$ says that all my friends know they are friends with me. Models for this language consist of tuples $M = \langle W, A, k, f, g, V\rangle$, where W is a set of states, A a set of agents, and

1. k is a family of equivalence relations k_a for each agent $a \in A$, representing the ignorance of a in distinguishing epistemic possibilities. Note that, for each agent a, the structure $M_a = \langle W, k_a, V_a\rangle$ where $V_a(\rho) = \{w \mid \langle w, a\rangle \in V(\rho)\}$ is a model for standard S5-epistemic logic.
2. f is a family of symmetric and irreflexive relations f_w for each $w \in W$, representing the friendship relation in state w. As above, for each w, we have a structure $M_w = \langle A, f_w, g, V_w\rangle$ with $V_w(\rho) = \{a \mid \langle w, a\rangle \in V(\rho)\}$.
3. g is an indexed set of agents g_n for each agent nominal $n \in \mathsf{ANom}$; in context we abbreviate this as $\underline{n}$.
4. V is a family of valuation functions $V_a : \mathsf{Prop} \to \mathcal{P}W$ for each agent $a \in A$, representing the proposition expressed by p when referring to a.

Models are used to interpret $\mathcal{L}$ in a double-indexical way. The salient clauses are:

$M, w, a \models \rho$	iff	$(w, a) \in V(p)$, for $\rho \in \mathsf{Prop} \cup \mathsf{ANom}$
$M, w, a \models K\varphi$	iff	$M, v, a \models \varphi$ for every $v \in W$ such that $\langle w, v\rangle \in k_a(w)$
$M, w, a \models F\varphi$	iff	$M, w, b \models \varphi$ for every $b \in A$ such that $\langle a, b\rangle \in f_w(a)$
$M, w, a \models U_A\varphi$	iff	$M, w, b \models \varphi$ for every $b \in A$.
$M, w, a \models U_W\varphi$	iff	$M, v, a \models \varphi$ for every $v \in W$.
$M, w, a \models\downarrow n\ \varphi$	iff	$M[^n_a], w, a \models \varphi$.

Where $M[^n_a] = \langle W, A, k, f, g[^n_a], V\rangle$ and $g[^n_a]_m = a$ if $m = n$ and g_m, otherwise. As is usual in modal logic, we can define the duals of the operators, which we write inside angle brackets: $\langle K\rangle = \neg K \neg$ 'it is epistemically possible for me that', $\langle F\rangle = \neg F\neg$ 'I have a friend who', $E_A = \langle U_A\rangle = \neg U_A\neg$ 'there is someone who', and $E_W = \langle U_W\rangle = \neg U_W\neg$ 'in some state'. The English glosses are not so exact and require some manipulation to get proper translations, because of the way pronouns work in English. For example, if d represents 'I am in danger' then $\langle F\rangle Kd$ means 'I have a friend who knows that he is in danger' rather than 'I have a friend who I know that I am in danger' which is not even grammatically correct. E_W and U_W are particularly difficult to translate because they quantify over states including those that I know not to be the case: they are states that some other agent may consider epistemically possible.

We also use abbreviations for the universal modality $U = U_AU_W$ (equivalently, U_WU_A) and the hybrid logic operators $@_n\varphi = U_A(n \to \varphi)$ (equivalently, $E_A(n \wedge \varphi)$). If $\underline{n}$ is Charlie then the agent operator $@_n$ simply shifts the indexical subject to Charlie, so that, for example $@_n d$ means 'Charlie is in danger'.

Models in which every agent has a name (i.e., g is surjective) are called *named agent* models.

Relations and change.

We will use the framework of general dynamic dynamic logic ($\mathsf{G}DDL$ [4]) to represent actions within a social network, and the resulting changes in what is known and by whom. We define a class of operators $\mathcal{D}$ and corresponding actions on models such that for each $\Delta \in \mathcal{D}$ and each M model for $\mathcal{L}$, there is a $\mathcal{L}$ model ΔM, and for each state w of M, a state Δw of ΔM. We then extend $\mathcal{L}$ to a language $\mathcal{L}(\mathcal{D})$ of *dynamic epistemic friendship logic* (DEFL) by adding the elements of $\mathcal{D}$ as propositional operators and defining

$$M, w, a \models \Delta\varphi \qquad \text{iff} \quad \Delta M, \Delta w, a \models \varphi$$

The full definition of $\mathcal{D}$ will be given in Section 6; for now, an informal presentation will suffice. The operators are built out of the terms/programs of $\mathsf{P}DL$, with atomic programs k, f and e_A corresponding to the modal operators K, F and U_A of $\mathcal{L}$. These are interpreted as binary relations on the product $W \times A$. Specifically,

1. e_W is interpreted as relating $\langle w, a\rangle$ with $\langle v, b\rangle$ just in case $a = b$,
2. k is interpreted as relating $\langle w, a\rangle$ with $\langle v, b\rangle$ just in case $a = b$ and $k_a(w, v)$ in M,
3. e_A is interpreted as relating $\langle w, a\rangle$ with $\langle v, b\rangle$ just in case $w = v$, and
4. f is interpreted as relating $\langle w, a\rangle$ with $\langle v, b\rangle$ just in case $w = v$ and $f_w(a, b)$ in M.

Notice that each of these basic relations is a sub-relation of either e_W or e_A. This is an important property that plays a key role in the technical development in Section 6. Complex terms, such are built up in the usual way: $(\pi_1; \pi_2)$ for the relational composition of π_1 and π_2, $(\pi_1 \cup \pi_2)$ for their union, $\varphi?$ for the 'test' consisting of a link from $\langle w, a\rangle$ to itself iff $M, w, a \models \varphi$, and π^* for the reflexive, transitive closure of π.

The simplest operators, called $\mathsf{P}DL$*-transformations* are built from $\mathsf{P}DL$-terms using assignment statements, such as $[k := n?; k]$. The effect of this operator on a model M is to produce a new model $[k := n?; k]M$ in which k is reinterpreted as relating $\langle w, a\rangle$ with $\langle v, b\rangle$ iff $\langle w, a\rangle$ passes the test $n?$ (which holds only when $\underline{n} = a$) and the two pairs are also related by k in M.[1] In other words, world w and v are in the k_b relation in the new model (indistinguishable for agent b) just in case $b = a$ and they are indistinguishable for agent a. This operator gives complete knowledge to every agent other than a.

More complicated operators can be constructed from finite relational structures whose elements are each associated with a $\mathsf{P}DL$-transformation, and whose combined effect on the a model is calculated by 'integrating' them according to a further such transformation. These are the full $\mathsf{G}DDL$ operators. A precise definition will be given in Section 6, but for now an example will suffice. Consider the operator Δ given by

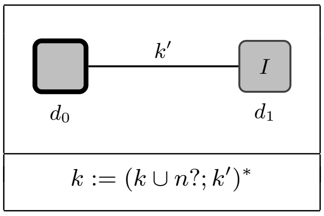

This represent an action d_0 (highlighted as the action that is actual performed) whose effect on the model is given by the $\mathsf{P}DL$-transformation we considered before: $[k := n?; k]$. There is also an action d_1 which does not change the model at all, and this is shown by labelling it with I, the identity transformation. Then, the relationship between d_0 and d_1 is labelled k', which is a new symbol initially without interpretation in the model, but which comes to represent the ignorance of n about whether or not action d_0 has occurred. Thus, this operator represent the result of giving complete information to all agents accept n, and furthermore, allowing n to remain ignorant about the fact that this has happened. The assigned statement in the lower half of the diagram shows how ignorance about which actions has occurred is integrated. In this case, we say simply that ignorance in the final model is either ignorance in one of the resulting two models (which have states $W \times \{d_0\}$ and $W \times \{d_1\}$, respectively) or ignorance about which of these parts of the composite model one is in, and this kind of ignorance is restricted to the agent n. We will see specific examples of such transformations below but the reader is invited to look at Section 6 for precise definitions.

An Example.

For a more useful example, consider the $\mathsf{P}DL$-term $\mathsf{cut}_k(\varphi)$ defined by

$$(\varphi?; k; \varphi?) \cup (\neg\varphi?; k; \neg\varphi?)$$

This relates $\langle w, a\rangle$ to $\langle v, b\rangle$ iff $a = b$, $k_a(w, v)$, and either φ is true of a in both states w and v or false of a in both states. Thus the operator $[k := cutk\varphi]$ produces a new model $[k := cutk\varphi]M$ from M by removing the k_a links between states with conflicting values for φ (about a). Effectively, this 'reveals' to each agent whether or not φ holds (for them). This operator, first introduced in [2], is very interesting epistemologically. We will abbreviate is as $[\mathsf{R}\varphi]$ in this paper.

[1] The application of this operator (and all $\mathsf{P}DL$-transformations) to states is just the identity, so $[k := n?; k]w = w$.)

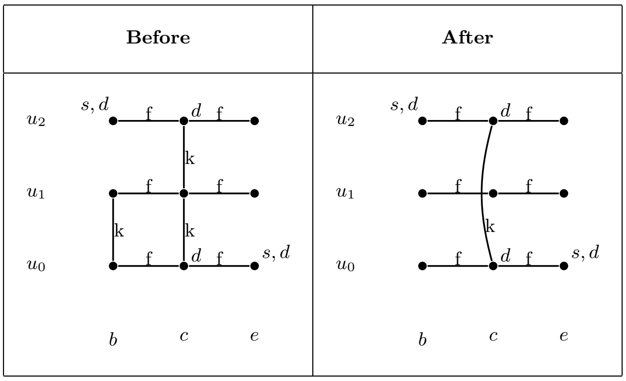

Figure 1: Spy Network

To take a Cold War example, suppose we are reasoning about the effect of a spy network being exposed.

> Charlie (c) is friends with Bella (b) and Erik (e), neither of whom are friends with each other. Unknown to the others is whether Erik is a spy (s). The others are not spies, and Erik knows that. Bella knows that Charlie is not a spy, but Charlie does not know whether or not Bella is. Charlie, however, knows that not all of his friends are spies. After the network is exposed, all the spies and their friends will be interrogated by the police. But just before this happens a message is relayed to all agents revealing whether or not they are in danger (d), that is, whether they are a spy (which they would know in any case) or a friend of a spy.

A model M of the initial situation is depicted in Figure 1, with reflexive and transitive links omitted for readability. Notice the crucial role of the friendship component of our modelling in analysing this examples. An agent is in danger, according to the story, if she is spy or is friends with a spy. It's for the second disjunct that our friendship operators get full credit! A model M of the initial situation is depicted in Figure 1, with reflexive and transitive links omitted for readability. The reader should read the diagram in a two-dimensional way, with propositions such as d and s assigned to world-agent pairs. So for instance, Bella is in danger in state $\langle u_2, b\rangle$ because she is a spy there, and Charlie is in danger in state $\langle u_o, c\rangle$, because his friend Erik is a spy there. In $\mathcal{L}$ we can state pertinent facts such as $@_c(K\neg s \wedge \neg K\langle F\rangle s)$ 'Charlie knows that he is not a spy but doesn't know if a friend of his is a spy'. From the assumed equivalence of d 'I am in danger' with $(s \vee \langle F\rangle s)$ 'either I'm a spy or I have a spy as a friend', we also have that $@_c(d \wedge \neg Kd)$ 'Charlie is in danger but doesn't know it', whereas $@_b K\neg d$ 'Bella knows that she is not in danger'. After the revelation $[\mathsf{R}d]$, the new model $[\mathsf{R}d]M$ is as depicted in the right part of Figure 1. Notice that that links between $\langle u_o, c\rangle$ and $\langle u_1, c\rangle$ and between $\langle u_1, c\rangle$ and $\langle u_2, c\rangle$ are removed as a result of the announcement, but the link between $\langle u_0, c\rangle$ and $\langle u_2, c\rangle$ survives. Charlie now knows that he is in danger ($@_c Kd$) and so that one of his friends is a spy ($@_c K\langle F\rangle s$) but he does not know which ($@_c\neg(K@_b s \vee K@_e s)$).
Moreover, in the language of DEFL we can represent reasoning about these changes, such as the validity of

$$[\mathsf{R}d]U_A(Kd \vee K\neg d)$$

which states the valid principle that after it is revealed who is in danger, everyone knows whether or not s/he is in danger.

2. SOCIAL ANNOUNCEMENTS

We now turn to direct communications, or 'announcements', within a social network. In general, an announcement consists of an agent (the sender) transmitting some information (the message) to one or more other agents (the receivers). Yet there are many subtleties concerning the knowledge of different agents about what happens. We will define a basic act of communication in which a message ψ is sent to a group of agents θ by

$$[k := (\theta?; \mathsf{cut}_k(\psi)) \cup (\neg\theta?; k)]$$

The action reveals the truth or falsity of ψ (which may be different for different agents) to all agents satisfying θ, and leaves the k_a relation unchanged for agents a not satisfying θ. The effect of making announcement depends on whether the message is indexical with respect to the sender, e.g. 'I am in danger', or with respect to the receiver, e.g., 'you are in danger', although, as we will show, the first is a special case of the second.

Announcements about the sender.

We first define $[n \triangleleft \psi! : \theta]\varphi$, the statement that φ holds after agent n announces message ψ (about n) to agents satisfying θ as

$$(@_n K\psi \to [k := (\theta?; \mathsf{cut}_k(@_n\psi)) \cup (\neg\theta?; k)]\varphi)$$

To make sense of this, we will look at a progression of simpler cases. First, with $\theta = \top$, the formula $[n \triangleleft \psi! : \top]\varphi$ means that φ holds after agent n publicly announces that ψ, noting that it simplifies to

$$(@_n K\psi \to [k := \mathsf{cut}_k(@_n\psi)]\varphi)$$

The presupposition that n knows that ψ is captured by the antecedent $@_n K\psi$. This form of public announcement is slightly different from PAL in that the announcement is taken to be made by an (possibly less than omniscient) agent in the community, but otherwise produces very similar results. For example, that everyone knows I am in danger after I announce this publicly is captured by the validity of $\downarrow n\ [n \triangleleft d! : \top]U_A K@_n d$.
In the second case, with $\theta = m$, the formula $[n \triangleleft \psi! : m]\varphi$ means that φ holds after agent n announces to m that ψ. For the moment, we will assume that these announcement are only semi-private, in that other agents know conditionally, that if the sender n knows that ψ then this is announced to m. Fully private announcements will be considered later in this section. Replacing θ by m in the definition gives us

$$(@_n K\psi \to [k := (m?; \mathsf{cut}_k(@_n\psi)) \cup (\neg m?; k)]\varphi)$$

in which we see the restriction of the update in knowledge to the receiver m. For agents other that m, the k_m relation is unchanged.
Finally, we consider the case of $\theta = \langle F\rangle n$, the friends of the sender. This models the announcement of φ by (and about) n to her friends, so that $[n \triangleleft \psi! : \langle F\rangle n]\varphi$ is given by

$$(@_n K\psi \to [k := (\langle F\rangle n?; \mathsf{cut}_k(@_n\psi)) \cup (\neg\langle F\rangle n?; k)]\varphi)$$

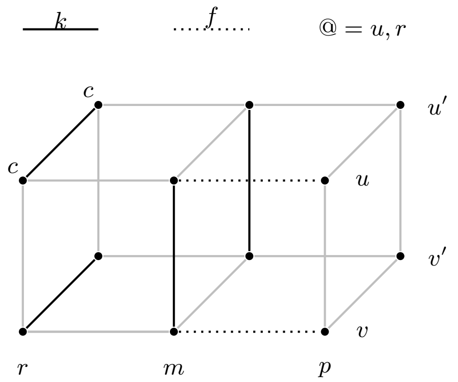

Figure 2: Roger's Quandry

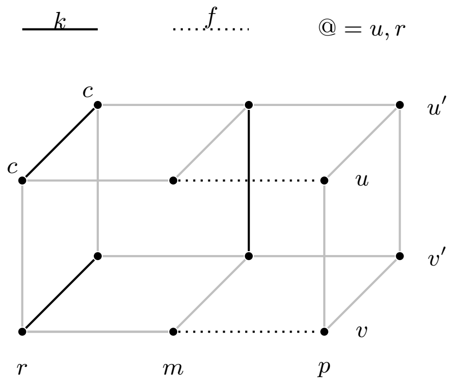

Figure 3: After Peggy's gossip

Let's see how this works with an everyday example of infidelity and gossip.

> Peggy (p) knows that Roger (r) is cheating (c) on his wife, Mona (m). What's more, Roger knows that Peggy knows, because they met accidentally while he was with his mistress. Mona does not know about the affair, and both Peggy and Roger know this. The situation (for Roger) deteriorates when he discovers that Peggy is a terrible gossip. She is bound to have told all her friends about his affair. What Roger does not know is whether Mona is a friend of Peggy (she is).

We can represent the epistemic state of this network before Peggy's announcement with the model depicted in Figure 2. At u, r (Roger in the actual world, represented by @ in the picture), the statements listed in Table 1 are all true. That

p	I'm cheating
$\downarrow n\ K(@_p K @_n p \wedge @_m \neg K @_n p)$	I know that Peggy (but not Mona) knows I am cheating
$\downarrow n\ @_p K @_n K @_p K @_n p$	Peggy knows I know she knows I am cheating
$\neg K @_m \langle F \rangle p$	I don't know whether Peggy and Mona are friends.
$\downarrow n\ @_p K @_n \neg K @_m \langle F \rangle p$	Peggy knows I don't know whether she and Mona are friends.

(For common knowledge, see Section 5.)

Table 1: Facts about Roger

some proposition φ holds after the announcement 'Roger is cheating!' that Peggy makes to her friends is given by $[p \lhd @_r c! : \langle F \rangle p]\varphi$, which expands and simplifies to

$$(@_p K @_r c \rightarrow [k := (\langle F \rangle p?; \mathsf{cut}_k(@_r c)) \cup (\neg \langle F \rangle p?; k)]\varphi)$$

When evaluated at u, p (Peggy in the actual world), the presupposition that Peggy knows that Roger is cheating is satisfied, and so the formula φ is evaluated in the transformed model shown in Figure 3. (Note the missing vertical line in the middle.) As a result, we can compute that at r, w (Roger in the actual world) of the original model, the formula

$$\downarrow n\ [p \lhd @_n c! : \langle F \rangle p] @_m K @_n c$$

is true, i.e., "I don't know that Mona will know about my cheating after Peggy tells her friends about it."

Announcements about the receivers.

Announcements that are indexical about the receiver such as 'you are in danger' (announced to Charlie) or 'you are my friends' (announced by Peggy to her friends) can be expressed with a slight change that captures the different preconditions for announcements. We define $[n : \psi! \rhd \theta]\varphi$, the statement that φ holds after agent n announces message ψ (about θ) to agents satisfying θ as

$$(@_n K U_A(\theta \rightarrow \psi) \rightarrow [k := (\theta?; \mathsf{cut}_k(\psi)) \cup (\neg\theta?; k)]\varphi)$$

Again, we first consider the simple case of public announcement, represented by $[n : \psi! \rhd \top]\varphi$, which can be seen to be equivalent to

$$(@_n K U_A \psi \rightarrow [k := \mathsf{cut}_k(\psi)]\varphi)$$

Consider, for example, my announcing to everyone 'you are in danger'. The precondition for this is that I know that everyone is in danger, captured by the antecedent $KU_A d$, and a consequence is that after the announcement everyone knows that she is in danger, as is represented by the validity of $\downarrow n\ [n : d! \rhd \top] U_A K d$.

The case of agent-to-agent announcement displays a nice symmetry between the two kinds of indexical message. Agent n announcing 'you are in danger' to agent m is equivalent to announcing (again to m) that m is in danger. More generally, the following equivalence is valid

$$[n : \psi! \rhd m]\varphi \ \leftrightarrow\ [n \lhd @_m \psi! : m]\varphi$$

For announcement to friends, an interesting new phenomenon arises. Consider the case of my announcing 'you are my friend' to my friends. That φ holds after such an announcement is represented by $[n : \langle F \rangle n! \rhd \langle F \rangle n]$. The message is the same as the description of the set of receivers, so when this is expanded, we find that the precondition for the announcement is $\downarrow n\ KU_A(\langle F \rangle n \rightarrow \langle F \rangle n)$, which is valid, so the announcement can always be made, by anyone. But nonetheless, it can be informative, as can be seen from the validity of $\downarrow n\ [n : \langle F\rangle n! \triangleright \langle F\rangle n]FK\langle F\rangle n$, which says that after my making this announcement, my friends all know that they are my friends, something they may not have known before.

Finally, we note that any sender-indexical announcement is equivalent to a receiver-indexical announcement in the case that there is at least one receiver ($U_A\neg\theta$ is false). The trick is that the statement ψ about n (the sender) is then equivalent to the statement $@_n\psi$ about any (every) receiver. More formally, the following is valid:[2]

$$(\neg U_A\neg\theta \to [n \triangleleft \psi! : \theta]\varphi \ \leftrightarrow\ [n : @_n\psi! \triangleright \theta]\varphi)$$

Private announcements.

Communications of the form $[n \triangleleft \psi! : \theta]$ and $[n : \psi! \triangleright \theta]$ are only semi-private. Their effect on the model ensures that every agent will know that the announcement has occurred, if the sender satisfies the precondition, so, for example,

$$\downarrow n\ [n \triangleleft d! : m]U_AK(@_nKd \to @_mK@_nd)$$

is valid: after I announce to m that I am in danger, everyone will know that if I know I am in danger then m also knows I am in danger. This is (typically) an unjustified violation of the privacy of the communication between me and m.

To get more private announcements, we need to use a product construction in the spirit initiated in [1], and within our framework, this is done with GDDL operators. First we define the PDL-transformation $\mathsf{send}^n_\theta(\psi)$ to be $[k := (\theta?; \mathsf{cut}_k(@_n\psi)) \cup (\neg\theta?; k)]$, so that $[n \triangleleft \psi! : \theta]\varphi$ is just $(@_nK\psi \to \mathsf{send}^n_\theta(\psi)\varphi)$. To make this action private, the transformation $\mathsf{send}^n_\theta(\psi)$ must be embedded in a GDDL operator, and then given the same precondition. Thus φ holding after the private announcement of ψ (about n) by n to agents θ can be represented as

$$(@_nK\psi \quad \to \quad \boxed{\begin{array}{c} \underset{d}{\boxed{\mathsf{send}^n_\theta(\psi)}} \overset{k'}{\text{———}} \underset{e}{\boxed{I}} \\ \hline k := (k \cup (\neg\theta?; k'))^* \end{array}} \quad \varphi)$$

Call this formula $[\![n \triangleleft \psi! : \theta]\!]\varphi$. Inside the G$DDL$ operator, the internal relation k' represents ignorance about whether the communication $\mathsf{send}^n_\theta(\psi)$ has occurred or not, the latter possibility represented by the identity transformation, I. The integrating transformation $k := (k \cup (\neg\theta?; k'))^*$ restricts ignorance of the k' kind to agents other than θ and factors this in to the new epistemic relation. The $*$ is needed to ensure that the result is an equivalence relation.

We will illustrate the application of this operator by returning to Roger's little problem.

> Before returning home to face Mona, Roger is uneasy. He would really like to know whether or not she knows about his affair. He already knows that she knows if and only if she is friends with Peggy. So if Peggy told him that they are friends, he would be prepared for Mona's fury. But for his planned excuses to be convincing, Mona must not know that he knows she knows (about the affair). It is therefore very important that Peggy tells him in private.

[2]The key observation here is that the precondition for the sender-indexical announcement is $@_nK\psi$, which is equivalent to the precondition $@_nKU_A(\theta \to @_n\psi)$ when $U_A\neg\theta$ is false.

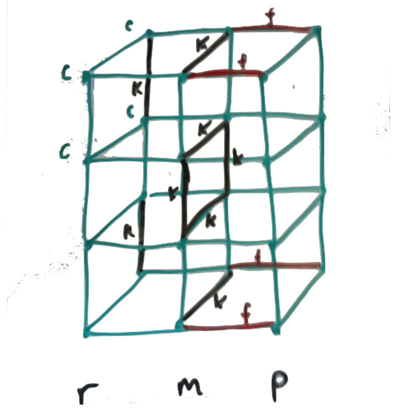

Figure 4: Peggy to Roger, privately.

Now let us suppose that the ever-loquacious Peggy announces to Richard privately that Mona is her friend. This can be represented as $[\![p \triangleleft \langle F\rangle m! : r]\!]$. Since Peggy does know that she and Mona are friends (as we might expect), whether the crucial proposition φ

$$(@_rK@_mK@_rc \wedge \neg @_mK@_rK@_mK@_rc)$$

(that Roger knows Mona knows he has been cheating but Mona doesn't know that he knows) holds must be determined by evaluating it in the model obtained by transforming the one in Figure 3 using the GDDL operator

$$\boxed{\begin{array}{c} \underset{d}{\boxed{\mathsf{send}^p_r(\langle F\rangle m)}} \overset{k'}{\text{———}} \underset{e}{\boxed{I}} \\ \hline k := (k \cup (\neg r?; k'))^* \end{array}}$$

The result is shown in Figure 4. In this diagram, the points on the front face represent states $\langle w, d\rangle$ (for $w \in W$), in which action d occurs, as described by $\mathsf{send}^p_r(\langle F\rangle m)$, whereas those on the back face represent states $\langle w, e\rangle$ (for $w \in W$) in which action e occurs, as described by I, the identity. The back face is just a copy of the model in Figure 3, whereas in the front face is missing two k_r links, which are cut when Roger learns that Mona is Peggy's friend. In this model φ holds of r in state u and Roger can meet Mona prepared.[3]

[3]Even the additional level of privacy offered here is still not perfect, as it involves some change in Mona's knowledge. She goes from knowing that Roger doesn't know that she is friends with Peggy to not knowing this. However, one may just think that privacy is a matter of degree.

3. ASKING QUESTIONS

As well as making announcements, agents in a social network can ask questions. Our approach to modelling questions will assume that agents are cooperative to the extent that they answer those questions to which they know the answer. A more elaborate model would consider the preferences of agents, but that is beyond the scope of the current paper. With this assumption, the effect of asking whether ψ of an agent a who knows that ψ is the same as an announcement by a that ψ. Likewise, the effect of asking whether ψ of an agent a who knows that $\neg\psi$ is the same as an announcement by a that $\neg\psi$. In the case that a does not know whether ψ, we assume that this also is communicated (possibly by the mere absence of an expected reply). With this in mind, we define $[n\colon\psi?\colon m]\varphi$, the proposition that φ holds after agent n asks agent m whether ψ as

$$([m \triangleleft \psi!\colon n]\varphi \wedge [m \triangleleft \neg\psi!\colon n]\varphi \wedge [m \triangleleft \neg(K\psi \vee K\neg\psi)!\colon n]\varphi$$

In other words, φ holds after n asks m whether ψ just in case φ holds after in all three cases: (1) m answers 'yes', so announcing ψ to n (2) m answers 'no', so announcing $\neg\psi$ to n and (3) m answers 'I don't know', so announcing $\neg(K\psi \vee K\neg\psi)$ to n. This ensures that the following are valid:

$$\begin{array}{l} (@_m K @_n p \to [n\colon p?\colon m]@_n Kp) \\ (@_m K @_n \neg p \to [n\colon p?\colon m]@_n K\neg p) \\ (@_m \neg(K @_n p \vee K @_n \neg p) \\ \quad \to [n\colon p?\colon m]@_n K @_m \neg(K @_n p \vee K @_n \neg p)) \end{array}$$

So, for example, after Charlie c asks Erik e whether he (Charlie) is in danger, d, he will either know that he is in danger Kd or know that he is not in danger $K\neg d$, or know that Erik doesn't know whether or not he (Charlie) is in danger, $\downarrow n\ K@_e\neg(K@_n d \vee K@_n\neg d)$.

As with announcements, this model of questions assumes that the answers are only semi-private. For example, after Charlie asks Erik whether he is in danger, a third-party will know that Charlie either knows whether he is in danger or knows that Erik doesn't know the answer. To make questioning more private, we need private announcements too. Here we will give one simple example.

> Roger approaches Peggy in private and asks her directly whether or not she and Mona are friends. Being sincere and cooperative, Peggy answers that they are. Mona, of course, knows nothing of their conversation.

This private question $[\![r\colon\langle F\rangle m?\colon p]\!]$ is defined by direct analogy with the semi-private question $[r\colon\langle F\rangle m?\colon p]$ so that φ holds after the question is asked just in case

$$([\![p \triangleleft \langle F\rangle m!\colon r]\!]\varphi \wedge [\![p \triangleleft \neg\langle F\rangle m!\colon r]\!]\varphi \wedge [\![p \triangleleft \neg(K\langle F\rangle m \vee K\neg\langle F\rangle m)!\colon r]\!]\varphi$$

In this case, only the precondition of $[\![p \triangleleft \langle F\rangle m!\colon r]\!]$ is satisfied, and so the results are just as depicted in Firgure 4.

4. CHANGING THE NETWORK

What makes networking intriguing is the dynamics of network changes. You can be friends with someone one day on Facebook, but you may drop him as a friend the following day or add someone else. Those acts, though simple, have a direct impact on information flow in communities. Let us consider the following:

> Roger, scared of the possibility that Mona will find out about his affair from Peggy, does all that he can to distance them. His smear campaign is designed to break their friendship and so protect his information.

To define the operation of deleting a friendship link, we first define the result of cutting the friendship link between agent n and m in one direction

$$\mathsf{cut}_f(n,m) = (\neg n?; f) \cup (f; \neg m?)$$

Then, to deleting the link between n and m we need to cut in both directions:[4]

$$[-f_{n,m}] = [f := \mathsf{cut}_f(n,m)][f := \mathsf{cut}_f(m,n)]$$

It is then fairly easy to show that $[\![f]\!]^{[-f_{nm}]M} = [\![f]\!]^M \setminus \{\langle n,m\rangle, \langle m,n\rangle\}$, as required.[5]

Now how is this going to help Roger? Well, after the application of $[-f_{mp}]$ to the model of Figure **??**, Peggy's announcement to her friends that Roger is cheating has no effect; in fact, she has no friends to receive the message. So the model is unchanged. In other words, in this original model, it is true for Roger that

$$[-f_{mp}] \downarrow n\ [p \triangleleft @_n c!\colon \langle F\rangle p]@_m \neg K @_n c$$

'after Peggy loses Mona as a friend, even after she tells her friends that I am cheating, Mona won't know.'

Next we consider adding a friend. In the basic case, we can define the operation $[+f_{n,m}]$ by analogy with deletion, but more simply, as

$$[f =\colon f \cup (n?; e_A; m?)]$$

But a more interesting model of adding friends follows the protocol of Facebook and other online social networks, whereby one must first request friendship. To capture this aspect of network change, we need to represent whether or not an agent *wants* to be friends with another agent. In a fuller account, this could be done with a preference order, showing that the agent prefers states in which they are friends to those in which they are not. But for now, suppose that there is some additional indexical relation d_w in our models, with $d_w(a,b)$ interpreted to mean that in state w, agent a wants to become friends with agent b.

The question 'do you want to be my friend?' from n to m is thus represented by $[n\colon\langle d\rangle n?\colon m]$, but as a *request* we interpret this as involving an action: if the answer is 'yes' then we become friends; otherwise, there is no change to the social network, thought there are consequent epistemic changes, such as my learning that you don't want to be my friend. That φ holds after this 'friend request' is therefore represented by

$$\begin{array}{rl} [\mathsf{add}(m)]\varphi = & \downarrow n\ [n\colon\langle d\rangle n?\colon m]((K@_m\langle d\rangle n \wedge [+f_{n,m}]\varphi) \\ & \vee(\neg K@_m\langle d\rangle n \wedge \varphi)) \end{array}$$

A private version of this operation can be obtained by replacing the announcement and network change by a $\mathsf{G}DDL$-based version.

[4]It is also interesting to consider asymmetric relationships such as "following" on Twitter or "subscribing" on Facebook, as studied in [5].

[5]This follows from the fact that $a[\![f]\!]^{[f:=\mathsf{cut}_f(n,m)]M}b$ iff $a[\![f]\!]^M b$ and $\langle a,b\rangle \neq \langle n,m\rangle$.

The following validity shows some of the epistemic consequence of friend requests:

$$\downarrow n\ ((\neg\langle F\rangle m \wedge \neg K@_m\langle d\rangle n) \rightarrow [\mathsf{add}(m)]((K@_m K\langle d\rangle n \wedge \langle F\rangle m) \vee(K@_m\neg K\langle d\rangle n \wedge \neg\langle F\rangle m))$$

5. COMMON KNOWLEDGE

In the context of social networks or communities, common knowledge is clearly an important notion. One can easily imagine the situations in which we want to reason about whether or not something is commonly known in some community or among my friends. There are at least two subtleties involved in making this precise. The first has to do with identifying the group of agents who are said to have common knowledge. This may be by means of a specific list ('Charlie, Bella, and Erik'), or a description ('Charlie's friends') or even an indexical description ('friends of mine'). Secondly, the information that is shared may be rigid ('it is common knowledge that Charlie is not a spy') or indexical (e.g. 'it is common knowledge among Charlie's friends that I am in danger' or 'it is common knowledge among my friends that they are in danger.')

To capture all these cases, first define $\overline{k}_a$ to be $(e_A; a?; k)$. Then $[\overline{k}_a]\varphi$ means that agent a knows that φ, as justified by the following equivalence:

$$M, u, b \models [\overline{k}_a]\varphi \qquad \text{iff} \quad M, v, a \models \varphi \text{ for all } v \in W \text{ such that } k_a(u, v).$$

Here φ could be an indexical proposition, so, for example, 'Charlie knows that he is not a spy' would be represented by $[\overline{k}_c]\neg s$, whereas 'Bella knows that Charlie is not a spy' would have to be represented as $[\overline{k}_b]@_c\neg s$. Now, for common knowledge, define

$$\mathsf{c}_\theta = (e_A; \theta?; k)^*; e_A; \theta?$$

and interpret $[\mathsf{c}_\theta]\varphi$ to mean, roughly, that there is common knowledge among θ-agents that φ. So this enables us to talk, in our formal language, about the common knowledge of some group.This definition seems more general than the standard notion of common knowledge (see e.g. [3]) . It is justified by the following applications, each of which can be suitably generalised.

1. Common knowledge among an enumerated set of agents about a non-indexical proposition. For example, that there is common knowledge between Bella (b) and Charlie (c) that Charlie is not a spy (s) can be represented by $[\mathsf{c}_{(b\vee c)}]@_c\neg s$.[6] To justify this claim, first note the standard way of defining common knowledge for a group of agents G is to introduce a new operator C_G such that

   $$M, w, a \models C_G\varphi \qquad \text{iff} \quad M, v, a \models \varphi \text{ for all } \langle u, v\rangle \in (\textstyle\bigcup_{a'\in G} k_{a'})^*$$

   We can then prove that, for example, $[\mathsf{c}_{(b\vee c)}]@_c\neg s$ is equivalent to $C_{\{b,c\}}@_c\neg s$.[7]

2. Common knowledge among a non-indexically described group of agents about a non-indexical proposition. For example, that it is common knowledge among Peggy's (p) friends that Roger (r) is cheating (c) can be represented as $[\mathsf{c}_{\langle F\rangle p}]@_r c$. This implies that every friend of Peggy knows that Roger is cheating ($@_pFK@_rc$), but also that each of them knows that all of Peggy's friends know this ($@_pFK@_pFK@_rc$), and that each of them knows they all know *that* ($@_pFK@_pFK@_pFK@_rc$), and so on. As such, it is not equivalent to any statement of the form $C_G\varphi$. In particular, if, say, Peggy's only friends are Mona (m) and Nancy (n), it may *not* have the same truth value as $C_{\{m,n\}}@_rc$, which is compatible with Mona's and Nancy's ignorance about what Peggy's friends (in general) know.

3. Common knowledge among a non-indexically described group of agents about a proposition that is indexical with respect to each member of the group. This is the subtlest case. As an example, after the spy network has been exposed, that it is common knowledge among Erik's (e) friends that they are in danger (d) is represented by $[\mathsf{c}_{\langle F\rangle e}]d$. This implies that every friend of Erik (the spy) knows that s/he is in danger ($@_eFKd$), that each of them knows they all know this ($@_eFK@_eFKd$), and so on. Again, this is compatible with their ignorance about the friendship relation, so long as all epistemically indistinguishable state are ones in which the friends of Erik (whoever they may be) are still in danger. The reason to have the final part $e_A; \theta?$ in the above definition of c_θ is this: when φ is indexical, we need to ensure that it is about the members of θ. But when φ is not indexical, this part is basically redundant.

4. Common knowledge among an indexically described group of agents about a non-indexical proposition. For example, that it is common knowledge among my friends that Roger is cheating is represented by $\downarrow n\ [\mathsf{c}_{\langle F\rangle n}]@_rc$. This is a straightforward generalisation of the previous case to an indexically specified description, with the $\langle F\rangle n$ using the nominal n, which is bound to the speaker by $\downarrow n$.

5. Common knowledge among an indexically described group of agents about a proposition that is indexical with respect to the speaker. For example, that there is common knowledge among my friends that I am not a spy is represented by $\downarrow n\ [\mathsf{c}_{\langle F\rangle n}]@_n\neg s$. This is really no more complicated than the last case. Again, the indexical work is all done by $\downarrow n$ in creating a temporary name 'n' for the speaker. Within that context, both the description of group ($\langle F\rangle n$) and the content of the common knowledge $@_n\neg s$ are both non-indexical.

6. Common knowledge among an indexically described group of agents about a proposition that is indexical with respect to each member of the group. For example, that it is common knowledge among my friends that they are in danger represented by $\downarrow n\ [\mathsf{c}_{\langle F\rangle n}]d$.

[6] Another concrete and interesting area of application is our ordinary email exchange, see an interesting analysis in [8].

[7] The argument is simple. First note that $(e_A; (b\vee c)?; k)^*$ is equivalent to $(\overline{k}_b \cup \overline{k}_c)^*$. Also, since $@_c\neg s$ is non-indexical, $[e_A; (b\vee c)?]@_c\neg s$ is equivalent to $@_c\neg s$. Thus $[\mathsf{c}_{(b\vee c)}]@_c\neg s$ is equivalent to $[(\overline{k}_b \cup \overline{k}_c)^*]@_c\neg s$, which is obviously equivalent to $C_{\{b,c\}}@_c\neg s$.

$[\![\rho]\!]^M$	$=$	$V(\rho)$, for $\rho \in P$
$[\![(\varphi \wedge \psi)]\!]^M$	$=$	$[\![\varphi]\!]^M \cap [\![\psi]\!]^M$
$[\![\neg\varphi]\!]^M$	$=$	$W \setminus [\![\varphi]\!]^M$
$[\![\langle\pi\rangle\varphi]\!]^M$	$=$	$\{w \in W \mid u[\![\pi]\!]^M v$ and $v \in [\![\varphi]\!]^M$ for some $v \in W\}$
$[\![\tau]\!]^M$	$=$	$V(\tau)$, for $\tau \in R$
$[\![\varphi?]\!]^M$	$=$	$\{\langle u, u\rangle \mid u \in [\![\varphi]\!]^M\}$
$[\![\pi_1; \pi_2]\!]^M$	$=$	$\{\langle u, v\rangle \mid u[\![\pi_1]\!]^M w$ and $w[\![\pi_2]\!]^M v$ for some $w \in W\}$
$[\![\pi_1 \cup \pi_2]\!]^M$	$=$	$[\![\pi_1]\!]^M \cup [\![\pi_2]\!]^M$
$[\![\pi^*]\!]^M$	$=$	$\{\langle u, v\rangle \mid u = v$ or $u_i[\![\pi]\!]^M u_{i+1}$ for some $n \geq 0$, $u_0, \ldots, u_n \in W$, $u_0 = u$ and $u_n = v\}$

Table 2: Semantics of PDL

Again, this is an obvious generalisation of the previous cases.

Other useful specifications of groups of agents as the subjects of common knowledge include 'common knowledge of φ in my community' ($\downarrow n\ [\mathsf{c}_{\langle f^*\rangle n}]\varphi$), 'common knowledge of φ among those who know they are in danger' ($[\mathsf{c}_{Kd}]\varphi$), 'common knowledge of φ among those who know they are my friends' ($\downarrow n\ [\mathsf{c}_{K\langle F\rangle n}]\varphi$).

6. TECHNICAL DETAILS

Technical details in this extended abstract will be confined to a basic outline, which will be expanded in the final paper.

General Dynamic Dynamic Logic.

Our starting point is the framework GDDL of [4], in which the language of propositional dynamic logic (PDL) is extended with model-changing operators. As indicted in the previous sections, we augment the language $\mathcal{L}$ with operators defined in GDDL. First, we need to define the basics of GDDL in a more general setting.

Given a signature $\langle P, R\rangle$ of propositional symbols P and relation symbols R, we define the set $L(P, R)$ of PDL-formulas and the set $T(P, R)$ of PDL-terms in the usual way and interpret these over Kripke models $M = \langle W, V\rangle$, in which $V(\rho)$ is a subset of W for each $\rho \in P$ and $V(r)$ is a binary relation on W for each $r \in R$. The semantics clauses as as usual. (Shown in Table 2.)

We are particularly interested in the signature $\langle P_{\mathsf{EFL}}, R_{\mathsf{EFL}}\rangle$ given by $P_{\mathsf{EFL}} = \mathsf{Prop} \cup \mathsf{ANom}$ and $R = \{f, k, e_A, e_W\}$.

For now, we refer the reader to [4] for the definition of PDL*-transformation and* GDDL*-operator.*

We extend the set of formula $L(P, R)$ by adding the GDDL operators to get $G(P, R)$

THEOREM 1. *There is an algorithm for computing a formula $\varphi^\dagger$ in $L(P, R)$ from each formula φ in $G(P, R)$ such that*

$$(\varphi \leftrightarrow \varphi)^\dagger$$

is valid.

PROOF. *Refer to [4]* □

Separable expressions and models.

In this part, we introduce the concept of separability, which characterises those models of our PDL*language $L(P_{\mathsf{EFL}}, R_{\mathsf{EFL}})$ which are equivalent to models for $\mathcal{L}$ and over which our dynamic operators must be well-behaved.*

Given a model M of the form $\langle W_0 \times W_1, V\rangle$ whose domain of states is a product $W_0 \times W_1$, we classify program as 'left' or 'right' or 'mixed' as follows:

π is *left* in M	iff $\langle u_0, u_1\rangle[\![\pi]\!]^M\langle v_0, v_1\rangle$ implies $u_1 = v_1$
π is *right* in M	iff $\langle u_0, u_1\rangle[\![\pi]\!]^M\langle v_0, v_1\rangle$ implies $u_0 = v_0$
π is *mixed* in M	iff π is neither left nor right in M

A model M is *separable* iff no basic program r in R is mixed. More specifically, given a subset $S \subseteq R$, M is S*-separable* if the programs in S are all right in M and those in $R \setminus S$ are left. To extend this division to complex programs, we define subsets $T_0(P, R)$ and $T_1(P, R)$ of programs $T(P, R)$ and a subset $F_S(P, R)$ of formulas $F(P, R)$ as follows:

$T_{S_0}(P, R)$	$\pi^0 ::= \tau \mid \varphi? \mid (\pi^0; \pi^0) \mid (\pi^0 \cup \pi^0) \mid \pi^{0*}$ (for $\tau \in R \setminus S$)
$T_{S_1}(P, R)$	$\pi^1 ::= \tau \mid \varphi? \mid (\pi^1; \pi^1) \mid (\pi^1 \cup \pi^1) \mid \pi^{1*}$ (for $\tau \in S$)
$F_S(P, R)$	$\varphi ::= \rho \mid \neg\varphi \mid (\varphi \vee \varphi) \mid \langle\pi\rangle^0\varphi \mid \langle\pi\rangle^1\varphi$ (for $\rho \in P$)

LEMMA 1. *If M is S-separable then every program in $T_{S_0}(P, R)$ is left in M and every program in $T_{S_1}(P, R)$ is right in M*

In our case, we will be especially interested in the case of $\{f, e_A\}$-separability.

THEOREM 2. *The problem of whether a given formula of $L_S(P, R)$ is satisfiable in an S-separable model is decidable.*

PROOF. We give a full proof in the main paper. The idea is that separable formulas can be translated into into an expanded language (with 'witnessing' constants for the elements of W_2) such that any model of the translated formula can be used to construct a separable model for the original formula. □

Dynamic epistemic friendship logic.

Not all GDDL operators are fit for including in the set $\mathcal{D}$ of operators we add to $\mathcal{L}$ to get our language of dynamic epistemic friendship logic, $\mathcal{L}(\mathcal{D})$. Say that an operator Δ is *suitable* if ΔM is $\{f, e_A\}$-separable whenever M is, and moreover, that ΔM does not change the value of g in M.

LEMMA 2. *If Δ is suitable then $M, w, a \models \downarrow n\ \Delta\varphi$ iff $M, w, a \models \Delta \downarrow n\ \varphi$*

This lemma allows us to pull all downarrows outside the scope of dynamic operators, when we are restricting our attention to named-agent models. Otherwise, our results will be restricted to those formula that do not contain downarrows inside the scope of dynamic operators.

LEMMA 3. *For each formula φ of $\mathcal{L}(\mathcal{D})$ that does not contain downarrows in the scope of dynamic operators, there is an $\{f, e_A\}$-separable formula φ' of $L(P_{\mathsf{EFL}}, R_{\mathsf{EFL}})$ and for each model M of $\mathcal{L}$, there is a model M' of signature $\langle P_{\mathsf{EFL}}, R_{\mathsf{EFL}}\rangle$ such that*

$$M, w, a \models \varphi \text{ iff } M', \langle w, a\rangle \models \varphi'$$

The Main Result.
Combining the above lemmas, we show, by a sequence of translations:

THEOREM 3. *The problems of satisfiability for formulas of $\mathcal{L}(\mathcal{D})$ is decidable.*

7. CONCLUDING REMARKS

What has emerged from this study is an appreciation of the diversity of subtle logic distinctions when combining epistemic and social relations, especially when we allow indexical propositions, as are very common in the social setting. Although Facebook was an inspiration for this work, we have only scratched the surface. The real Facebook offers many interesting features that would be good to model, such as the wall, subscriptions, commenting, liking. Real world networks, such as the spy networks and gossips networks of the running examples in this paper, are also worthy of further study, especially when diluting knowledge to belief and adding preference.

8. REFERENCES

[1] A. Baltag, L. S. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicious. Technical Report SEN-R9922, CWI, Amsterdam, 1999.

[2] J. v. Benthem and F. Liu. Dynamic logic of preference upgrade. *Journal of Applied Non-Classical Logic*, 17:157–182, 2007.

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

[4] P. Girard, J. Seligman, and F. Liu. General dynamic dynamic logic. In T. Bolander, T. Braüner, S. Ghilardi, and L. S. Moss, editors, *Advances in Modal Logics Volume 9*, pages 239–260, 2012.

[5] J. Ruan and M. Thielscher. A logic for knowledge flow in social networks. In *Australasian Conference on Artificial Intelligence*, pages 511–520, 2011.

[6] J. Seligman, P. Girard, and F. Liu. Logical dynamics of belief change in the community. Manuscript, 2012.

[7] J. Seligman, F. Liu, and P. Girard. Logic in the community. In M. Banerjee and A. Seth, editors, *ICLA*, volume 6521 of *Lecture Notes in Computer Science*, pages 178–188, 2011.

[8] F. Sietsma and K. Apt. Common knowledge in email exchanges. In J. Eijck and R. Verbrugge, editors, *Proceedings of the Workshop on Reasoning About Other Minds: Logical and Cognitive Perspectives*, pages 5–19, 2011.