Social Influence without Transparency

Zoé Christoff
Joint work with Jens Ulrik Hansen

Institute for Logic, Language and Computation, University of Amsterdam

Logic Tea Seminar
Amsterdam, October 22, 2013

Outline

Social networks, influence, and belief revision à la Girard, Liu & Seligman

The basics

Social networks

- Agents situated in a social network (a symmetric graph)
- The edges represent a "friendship" relation

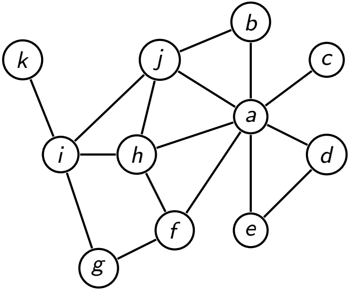

The basics

Social networks

- Agents situated in a social network (a symmetric graph)
- The edges represent a "friendship" relation

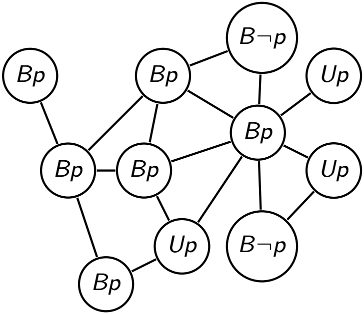

Simple beliefs

The agents can be in 3 possible states:

- Bp
- $B\neg p$
- $Up := \neg Bp \wedge \neg B\neg p$

Social influence and belief revision

Main assumptions

- Agents are influenced by their friends and only by their friends.
- Simple "peer pressure principle": I tend to align my beliefs with the ones of my friends.
- Agents can "see" what their friends believe.
  - I.e. an agent is "transparent" to all of his friends.

Belief revision induced by social influence

Strong influence

When all of my friends believe that p, I (successfully) *revise* with p.
When all of my friends believe that $\neg p$, I (successfully) *revise* with $\neg p$.

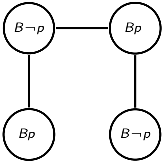

Belief revision induced by social influence

Strong influence

When all of my friends believe that p, I (successfully) *revise* with p.
When all of my friends believe that $\neg p$, I (successfully) *revise* with $\neg p$.

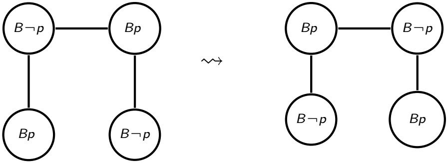

Belief revision induced by social influence

Strong influence

When all of my friends believe that p, I (successfully) *revise* with p.
When all of my friends believe that $\neg p$, I (successfully) *revise* with $\neg p$.

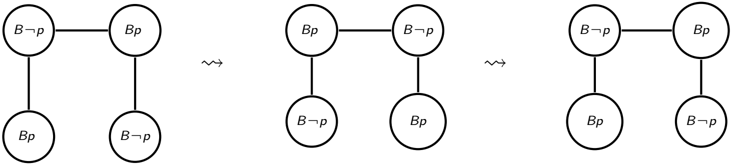

Belief contraction induced by social influence

Weak influence

When none of my friends supports my belief in $\neg p$ and some believe that p, I (successfully) *contract* it.
When none of my friends supports my belief in p and some believe that $\neg p$, I (successfully) *contract* it.

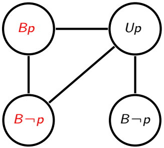

Belief contraction induced by social influence

Weak influence

When none of my friends supports my belief in $\neg p$ and some believe that p, I (successfully) *contract* it.
When none of my friends supports my belief in p and some believe that $\neg p$, I (successfully) *contract* it.

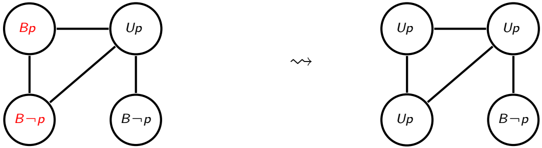

Influence dynamics

The influence operator $\mathcal{I}$

For each agent *a*:

- If *a* is strongly influenced to believe p ($\neg p$), then *a* will believe p ($\neg p$).
- If *a* is only weakly influenced to believe p ($\neg p$) and believes $\neg p$ (p), then *a* will become undecided.
- Otherwise, *a* keeps her current belief state.

Pluralistic ignorance

A phenomenon from social psychology

Different takes on pluralistic ignorance

- Everyone believes the same thing, but mistakenly believes that everyone else believes something else.
- An error of social comparison: Individuals mistakenly believes that others are different from themselves (even though they might act similar).
- Everyone publicly supports a norm they privately reject.
- A discrepancy between private beliefs and public beliefs/behavior.

A phenomenon from social psychology

Different takes on pluralistic ignorance

- Everyone believes the same thing, but mistakenly believes that everyone else believes something else.
- An error of social comparison: Individuals mistakenly believes that others are different from themselves (even though they might act similar).
- Everyone publicly supports a norm they privately reject.
- A discrepancy between private beliefs and public beliefs/behavior.

Examples

- The Emperor's New Clothes
- A silent classroom
- Campus drinking

The dynamics of pluralistic ignorance

Robustness

- Pluralistic ignorance is robust.
- If the environment stays the same the phenomenon often persists.
- For example, the college students might keep obeying an unwanted drinking norm for generations.

The dynamics of pluralistic ignorance

Robustness

- Pluralistic ignorance is robust.
- If the environment stays the same the phenomenon often persists.
- For example, the college students might keep obeying an unwanted drinking norm for generations.

Fragility

- Pluralistic ignorance is also often reported as fragile.
- A small change in the environment might be enough to dissolve the phenomenon (for instance, the announcement of one of the agents' true belief).
- If just one student of the classroom example starts to ask questions about the difficult lecture the rest of the students might soon follow.

Limitations of the framework of Girard, Liu & Seligman

- Possible difference between what one privately believes and what attitude one displays publicly
  - This is incompatible the transparency assumption in the framework of Girard, Liu & Seligman
- Peer pressure may only operate on the public level.
  - I might feel pressured into displaying the same belief as my friends, but I might keep might private belief.

Adding a layer

Two layers

We distinguish what an agent privately believes from what he seems to believe:

- *Private belief*, which we name "inner belief" (I_B) and
- *Public (or observable) behavior*, which we name "expressed belief" (E_B).

Adding a layer

Two layers

We distinguish what an agent privately believes from what he seems to believe:

- *Private belief*, which we name "inner belief" (I_B) and
- *Public (or observable) behavior*, which we name "expressed belief" (E_B).

New assumptions

- Agents can only observe the expressed beliefs of other agents
- Peer pressure only affects the agents' public beliefs, i.e. their visible behavior.

Adding a layer

Two layers

We distinguish what an agent privately believes from what he seems to believe:

- *Private belief*, which we name "inner belief" (I_B) and
- *Public (or observable) behavior*, which we name "expressed belief" (E_B).

New assumptions

- Agents can only observe the expressed beliefs of other agents
- Peer pressure only affects the agents' public beliefs, i.e. their visible behavior.

Changes needed

- Now 9 belief states (2 layers, 3 values each): $I_B p$, $I_B \neg p$ $I_U p$, $E_B p$, $E_B \neg p$, $E_U p$
- A new notion of social influence

A multi-layer version of social influence

What influences an agent

The behavior of an influenced agent in our 2-layer case now depends on:

- What he himself privately believes (3 possible situations)
- What his friends express (8 possible situations corresponding to whether at least one $E_B p$, at least one $E_B \neg p$, or at least one $E_U p$)

A multi-layer version of social influence

What influences an agent

The behavior of an influenced agent in our 2-layer case now depends on:

- What he himself privately believes (3 possible situations)
- What his friends express (8 possible situations corresponding to whether at least one $E_B p$, at least one $E_B \neg p$, or at least one $E_U p$)

Consequences

- A fundamental asymmetry between first and third person perspectives
- Each agent is in one of 24 possible situations.
- Our influence operator $\mathcal{I}$ assigns a postcondition to each of these situations (dictating what the state of the agent will be next).

A new notion of social influence

A first approximation of social influence

- When all of my friends express belief in p ($\neg p$), I will align and express a belief in p ($\neg p$)
  - "strong influence" works as before on the expressed level
- When I am weakly influenced in believing p ($\neg p$) , what I express now depends on my private belief concerning p.

A new notion of social influence

A first approximation of social influence

- When all of my friends express belief in p $(\neg p)$, I will align and express a belief in p $(\neg p)$
  - "strong influence" works as before on the expressed level
- When I am weakly influenced in believing p $(\neg p)$, what I express now depends on my private belief concerning p.

Subtle issues – several types of agents

- What to do when some friends express support and others express conflict, for instance?
- Defining strong and weak influence is no longer enough
- There are several options leading to several different types of agents
- We consider several of these in the paper

A new notion of social influence

A first approximation of social influence

- When all of my friends express belief in p ($\neg p$), I will align and express a belief in p ($\neg p$)
  - "strong influence" works as before on the expressed level
- When I am weakly influenced in believing p ($\neg p$) , what I express now depends on my private belief concerning p.

Subtle issues – several types of agents

- What to do when some friends express support and others express conflict, for instance?
- Defining strong and weak influence is no longer enough
- There are several options leading to several different types of agents
- We consider several of these in the paper
- For now we simply assume that an agent a expresses her private belief, iff
  - some of my friends expresses the same belief ("support") or
  - none of my friends expresses a belief in the negation of what I privately believes ("no conflict")

Our results on pluralistic ignorance

A formal language to talk about 2-layer social influence

Network structures and models

- A social network is a pair $(A, \asymp)$, where $\asymp$ is symmetric and irreflexive binary relation on the set of agents A.
- A model is a social network structure $(A, \asymp)$ together with a function g assigning one agent to each nominal and a function ν assigning exactly one of the 3 possible "values" Bp, $B\neg p$, or Up to each of the 2 layers for each agent in A.

A formal language to talk about 2-layer social influence

Network structures and models

- A social network is a pair $(A, \asymp)$, where $\asymp$ is symmetric and irreflexive binary relation on the set of agents A.
- A model is a social network structure $(A, \asymp)$ together with a function g assigning one agent to each nominal and a function ν assigning exactly one of the 3 possible "values" Bp, $B\neg p$, or Up to each of the 2 layers for each agent in A.

A Formal language

- Six propositional variables $I_B p$, $I_B \neg p$, $I_U p$, $E_B p$, $E_B \neg p$, and $E_U p$.
- A set of nominals $i, j, k, \ldots$.
- Shift operators $@_i, @_j, @_k, \ldots$.
- A "friendship" modality F.
- A global modality G.
- A modal "influence" operator $[\mathcal{I}]$.

A formal language to talk about 2-layer social influence

Network structures and models

- A social network is a pair $(A, \asymp)$, where $\asymp$ is symmetric and irreflexive binary relation on the set of agents A.
- A model is a social network structure $(A, \asymp)$ together with a function g assigning one agent to each nominal and a function ν assigning exactly one of the 3 possible "values" Bp, $B\neg p$, or Up to each of the 2 layers for each agent in A.

A Formal language

- Six propositional variables $I_B p$, $I_B \neg p$, $I_U p$, $E_B p$, $E_B \neg p$, and $E_U p$.
- A set of nominals $i, j, k, \ldots$.
- Shift operators $@_i, @_j, @_k, \ldots$.
- A "friendship" modality F.
- A global modality G.
- A modal "influence" operator $[\mathcal{I}]$.

$$\varphi \quad ::= \quad P \in PROP \mid i \mid \neg\varphi \mid \varphi \wedge \varphi \mid F\varphi \mid G\varphi \mid @_i\varphi \mid [\mathcal{I}]\varphi$$

Formal semantics

Given $\mathcal{M} = (A, \asymp, g, \nu)$ (g assigns an agent to each nominal) and $a \in A$:

$$
\begin{array}{lll}
\mathcal{M}, a \models P & \text{iff} & \nu(a) = P \\
\mathcal{M}, a \models i & \text{iff} & g(i) = a \\
\mathcal{M}, a \models \neg\varphi & \text{iff} & \text{it is not the case that } \mathcal{M}, a \models \varphi \\
\mathcal{M}, a \models \varphi \wedge \psi & \text{iff} & \mathcal{M}, a \models \varphi \text{ and } \mathcal{M}, a \models \psi \\
\mathcal{M}, a \models G\varphi & \text{iff} & \text{for all } b \in A; \mathcal{M}, b \models \varphi \\
\mathcal{M}, a \models F\varphi & \text{iff} & \text{for all } b \in A; a \asymp b \text{ implies } \mathcal{M}, b \models \varphi \\
\mathcal{M}, a \models @_i\varphi & \text{iff} & \mathcal{M}, g(i) \models \varphi \\
\mathcal{M}, a \models [\mathcal{I}]\varphi & \text{iff} & \mathcal{M}^{\mathcal{I}}, a \models \varphi,
\end{array}
$$

where $\mathcal{M}^{\mathcal{I}}$ is the new model obtained by updating ν in accordance to our previous definition of the influence operator.

Formal semantics

Given $\mathcal{M} = (A, \asymp, g, \nu)$ (g assigns an agent to each nominal) and $a \in A$:

$$
\begin{array}{lll}
\mathcal{M}, a \models P & \text{iff} & \nu(a) = P \\
\mathcal{M}, a \models i & \text{iff} & g(i) = a \\
\mathcal{M}, a \models \neg\varphi & \text{iff} & \text{it is not the case that } \mathcal{M}, a \models \varphi \\
\mathcal{M}, a \models \varphi \wedge \psi & \text{iff} & \mathcal{M}, a \models \varphi \text{ and } \mathcal{M}, a \models \psi \\
\mathcal{M}, a \models G\varphi & \text{iff} & \text{for all } b \in A; \mathcal{M}, b \models \varphi \\
\mathcal{M}, a \models F\varphi & \text{iff} & \text{for all } b \in A; a \asymp b \text{ implies } \mathcal{M}, b \models \varphi \\
\mathcal{M}, a \models @_i\varphi & \text{iff} & \mathcal{M}, g(i) \models \varphi \\
\mathcal{M}, a \models [\mathcal{I}]\varphi & \text{iff} & \mathcal{M}^{\mathcal{I}}, a \models \varphi,
\end{array}
$$

where $\mathcal{M}^{\mathcal{I}}$ is the new model obtained by updating ν in accordance to our previous definition of the influence operator.

Examples

- $FI_B p$; “all of my friends privately believe that p”
- $@_i\langle F\rangle j$; “j is a friend of i” (here $\langle F\rangle$ is the dual of F).
- $[\mathcal{I}][\mathcal{I}]@_i E_U p$; “after two steps of social influence i expresses undecidedness.”

Expressing Pluralistic Ignorance

Everybody privately believes p but expresses a belief in $\neg p$:

$$PIp := G(I_B p \wedge E_B \neg p)$$

If $PI\varphi$ is true in $\mathcal{M}$ we will say that $\mathcal{M}$ is in a state of pluralistic ignorance with respect to p.

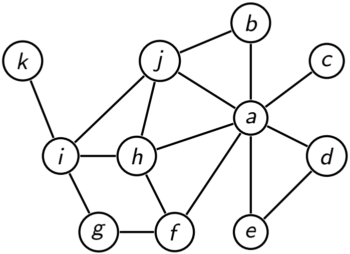

Expressing Pluralistic Ignorance

Everybody privately believes p but expresses a belief in $\neg p$:

$$PIp := G(I_B p \wedge E_B \neg p)$$

If $PI\varphi$ is true in $\mathcal{M}$ we will say that $\mathcal{M}$ is in a state of pluralistic ignorance with respect to p.

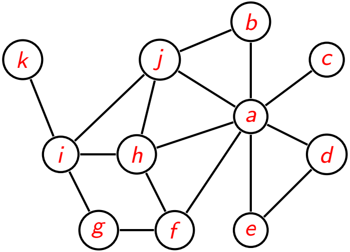

Robustness

Stability of PI

A connected network model in a state of pluralistic ignorance is stable, i.e the following is valid:

$$PIp \to [\mathcal{I}]PIp$$

Robustness

Stability of PI

A connected network model in a state of pluralistic ignorance is stable, i.e the following is valid:

$$PIp \rightarrow [\mathcal{I}]PIp$$

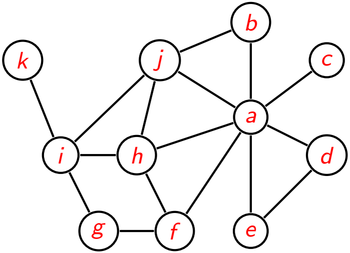

One brave/crazy agent!

Unstable state

Assume one agent i starts being sincere, for some reason, i.e. i expresses his true private belief that p:

$$UPIp := @_i(I_B p \wedge E_B p) \wedge G\big(\neg i \to (I_B p \wedge E_B \neg p)\big).$$

If $UPIp$ is true in $\mathcal{M}$ we will say that $\mathcal{M}$ *is in a state of unstable pluralistic ignorance*.

What happens next? Do all the others start being sincere too?

Breaking PI

Initially, only i is sincere.

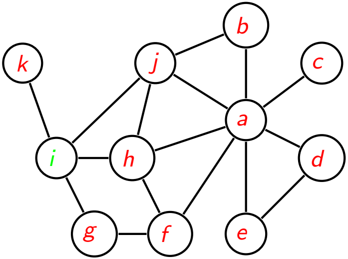

Breaking PI

Initially, only i is sincere.

- After one step, all i's friends are sincere.

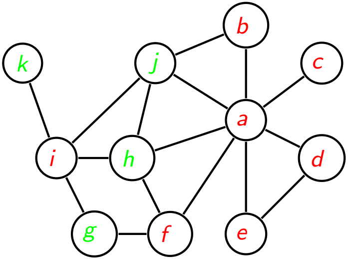

Breaking PI

Initially, only i is sincere.

- After one step, all i's friends are sincere.
- After two steps, all i's friends' friends (including i himself) are sincere.

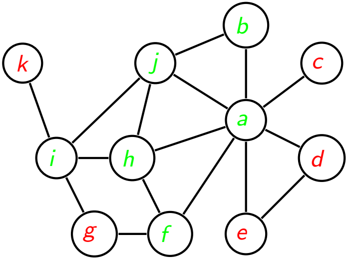

Breaking PI

Initially, only i is sincere.

- After one step, all i's friends are sincere.
- After two steps, all i's friends' friends (including i himself) are sincere.
- After three steps, everybody is sincere.

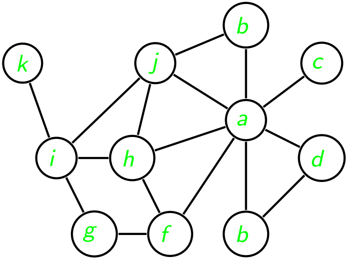

But *UPI* doesn't always dissolve!

- Initially, only *i* is sincere.

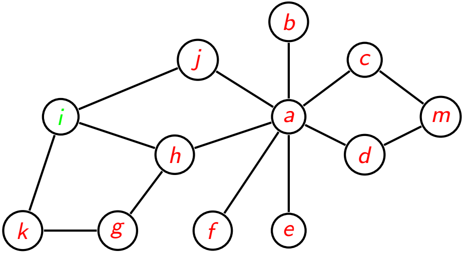

But *UPI* doesn't always dissolve!

- Initially, only i is sincere.
- After 1 step, only i's friends are sincere.

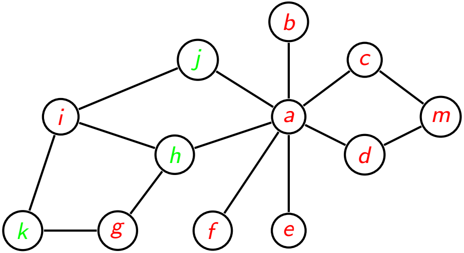

But *UPI* doesn't always dissolve!

- Initially, only i is sincere.
- After 1 step, only i's friends are sincere.
- After 2 steps, only i's friends' friends are sincere.

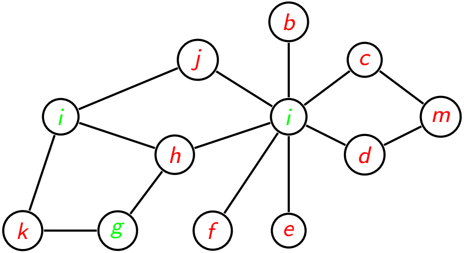

But *UPI* doesn't always dissolve!

- Initially, only i is sincere.
- After 1 step, only i's friends are sincere.
- After 2 steps, only i's friends' friends are sincere.
- After 3 steps, only i's friends' friends' friends are sincere.

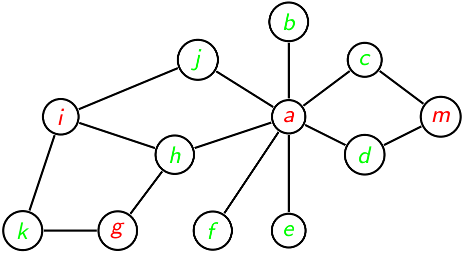

But *UPI* doesn't always dissolve!

- Initially, only *i* is sincere.
- After 1 step, only *i*'s friends are sincere.
- After 2 steps, only *i*'s friends' friends are sincere.
- After 3 steps, only *i*'s friends' friends' friends are sincere.
- ...

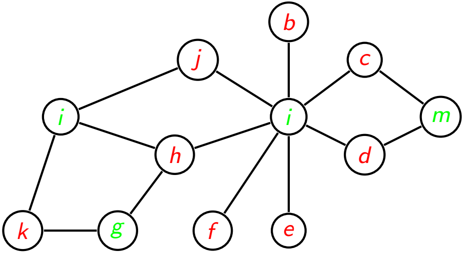

Question

When does a model satisfying *UPI* dissolve in a state of global sincerity, i.e such that $G(I_B p \wedge E_B p)$?

Let $\mathcal{M} = (A, \asymp, g, \nu)$ be a finite, connected, symmetric network model such that $\mathcal{M} \models UPI$. Then the following 6 conditions are equivalent:

Question

When does a model satisfying *UPI* dissolve in a state of global sincerity, i.e such that $G(I_B p \wedge E_B p)$?

Let $\mathcal{M} = (A, \asymp, g, \nu)$ be a finite, connected, symmetric network model such that $\mathcal{M} \models UPI$. Then the following 6 conditions are equivalent:

(i) After a finite number of updates by the influence operator $\mathcal{I}$, $\mathcal{M}$ will end up in a stable state where pluralistic ignorance is dissolved, i.e. there is a $k \in \mathbb{N}$ such that $\mathcal{M}^{\mathcal{I}^k} \models G(I_B p \wedge E_B p)$ and $\mathcal{M}^{\mathcal{I}^k} = \mathcal{M}^{\mathcal{I}^{k+1}}$.

Characterization

iff (ii) There is an agent who expresses her true private belief in p for two rounds in a row, i.e. there is an $a \in A$ and a $k \in \mathbb{N}$ such that $\mathcal{M}^{\mathcal{I}^k}, a \models E_B p$ and $\mathcal{M}^{\mathcal{I}^{k+1}}, a \models E_B p$.

Characterization

iff (ii) There is an agent who expresses her true private belief in p for two rounds in a row, i.e. there is an $a \in A$ and a $k \in \mathbb{N}$ such that $\mathcal{M}^{\mathcal{I}^k}, a \models E_B p$ and $\mathcal{M}^{\mathcal{I}^{k+1}}, a \models E_B p$.

iff (iii) There are two agents that are friends and both express their true beliefs in p in the same round, i.e. there are $a, b \in A$ and a $k \in \mathbb{N}$ such that $a \asymp b$, $\mathcal{M}^{\mathcal{I}^k}, a \models E_B p$, and $\mathcal{M}^{\mathcal{I}^k}, b \models E_B p$.

Characterization

iff (ii) There is an agent who expresses her true private belief in p for two rounds in a row, i.e. there is an $a \in A$ and a $k \in \mathbb{N}$ such that $\mathcal{M}^{\mathcal{I}^k}, a \models E_B p$ and $\mathcal{M}^{\mathcal{I}^{k+1}}, a \models E_B p$.

iff (iii) There are two agents that are friends and both express their true beliefs in p in the same round, i.e. there are $a, b \in A$ and a $k \in \mathbb{N}$ such that $a \asymp b$, $\mathcal{M}^{\mathcal{I}^k}, a \models E_B p$, and $\mathcal{M}^{\mathcal{I}^k}, b \models E_B p$.

iff (iv) There are two agents that are friends and have paths of the same length to the agent named by i, i.e. there are agents $a, b \in A$ and a $k \in \mathbb{N}$ such that $a \asymp b$, $\mathcal{M}, a \models \langle F \rangle^k i$, and $\mathcal{M}, b \models \langle F \rangle^k i$.

Characterization

iff (ii) There is an agent who expresses her true private belief in p for two rounds in a row, i.e. there is an $a \in A$ and a $k \in \mathbb{N}$ such that $\mathcal{M}^{\mathcal{I}^k}, a \models E_B p$ and $\mathcal{M}^{\mathcal{I}^{k+1}}, a \models E_B p$.

iff (iii) There are two agents that are friends and both express their true beliefs in p in the same round, i.e. there are $a, b \in A$ and a $k \in \mathbb{N}$ such that $a \asymp b$, $\mathcal{M}^{\mathcal{I}^k}, a \models E_B p$, and $\mathcal{M}^{\mathcal{I}^k}, b \models E_B p$.

iff (iv) There are two agents that are friends and have paths of the same length to the agent named by i, i.e. there are agents $a, b \in A$ and a $k \in \mathbb{N}$ such that $a \asymp b$, $\mathcal{M}, a \models \langle F \rangle^k i$, and $\mathcal{M}, b \models \langle F \rangle^k i$.

iff (v) There is a cycle in $\mathcal{M}$ of odd length starting at the agent named by i, i.e. there is a $k \in \mathbb{N}$ such that $\mathcal{M} \models @_i \langle F \rangle^{2k-1} i$.

Characterization

iff (ii) There is an agent who expresses her true private belief in p for two rounds in a row, i.e. there is an $a \in A$ and a $k \in \mathbb{N}$ such that $\mathcal{M}^{\mathcal{I}^k}, a \models E_B p$ and $\mathcal{M}^{\mathcal{I}^{k+1}}, a \models E_B p$.

iff (iii) There are two agents that are friends and both express their true beliefs in p in the same round, i.e. there are $a, b \in A$ and a $k \in \mathbb{N}$ such that $a \asymp b$, $\mathcal{M}^{\mathcal{I}^k}, a \models E_B p$, and $\mathcal{M}^{\mathcal{I}^k}, b \models E_B p$.

iff (iv) There are two agents that are friends and have paths of the same length to the agent named by i, i.e. there are agents $a, b \in A$ and a $k \in \mathbb{N}$ such that $a \asymp b$, $\mathcal{M}, a \models \langle F \rangle^k i$, and $\mathcal{M}, b \models \langle F \rangle^k i$.

iff (v) There is a cycle in $\mathcal{M}$ of odd length starting at the agent named by i, i.e. there is a $k \in \mathbb{N}$ such that $\mathcal{M} \models @_i \langle F \rangle^{2k-1} i$.

iff (vi) There is a cycle in $\mathcal{M}$ of odd length, i.e. there is a $k \in \mathbb{N}$ and $a_1, a_2, \ldots, a_{2k-1} \in A$ such that $a_1 \asymp a_2, a_2 \asymp a_3, \ldots, a_{2k-2} \asymp a_{2k-1}, a_{2k-1} \asymp a_1$.

Example

We have seen an example of network structure which doesn't dissolve from *UPI*. Why? What would need to change to make this possible?

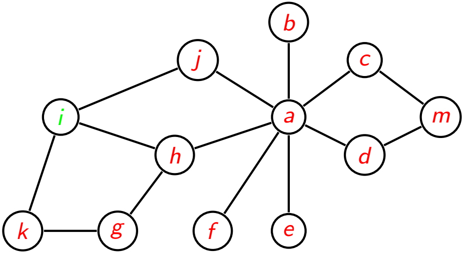

How to change the network

How do we need to change the example to force UPI to dissolve?

- Create a cycle of odd length in the graph (vi)
- Create an cycle of odd length from *i* (v)
- Make two friends have a path of the same length to *i* (iv)

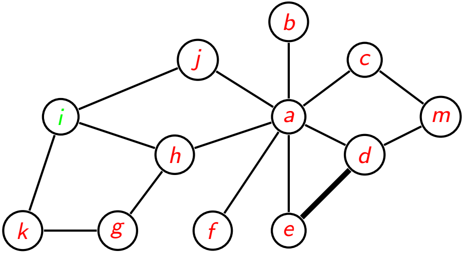

How to change the network

How do we need to change the example to force UPI to dissolve?

- Create a cycle of odd length in the graph (vi)
- Create an cycle of odd length from *i* (v)
- Make two friends have a path of the same length to *i* (iv)

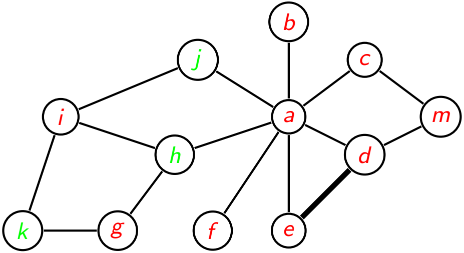

How to change the network

How do we need to change the example to force UPI to dissolve?

- Create a cycle of odd length in the graph (vi)
- Create an cycle of odd length from *i* (v)
- Make two friends have a path of the same length to *i* (iv)

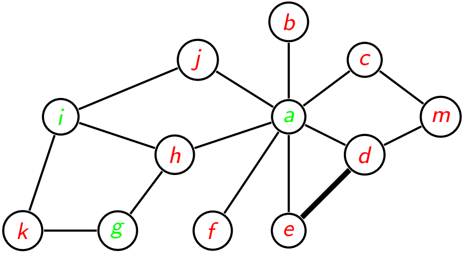

How to change the network

How do we need to change the example to force UPI to dissolve?

- Create a cycle of odd length in the graph (vi)
- Create an cycle of odd length from i (v)
- Make two friends have a path of the same length to i (iv)
- Make two friends express their true belief at the same round (iii)

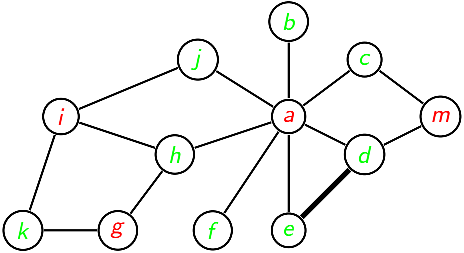

How to change the network

How do we need to change the example to force UPI to dissolve?

- Create a cycle of odd length in the graph (vi)
- Create an cycle of odd length from i (v)
- Make two friends have a path of the same length to i (iv)
- Make two friends express their true belief at the same round (iii)
- Make one agent express his true belief for two consecutive rounds (ii)

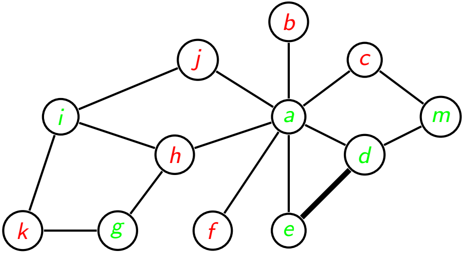

How to change the network

How do we need to change the example to force UPI to dissolve?

- Create a cycle of odd length in the graph (vi)
- Create an cycle of odd length from i (v)
- Make two friends have a path of the same length to i (iv)
- Make two friends express their true belief at the same round (iii)
- Make one agent express his true belief for two consecutive rounds (ii)

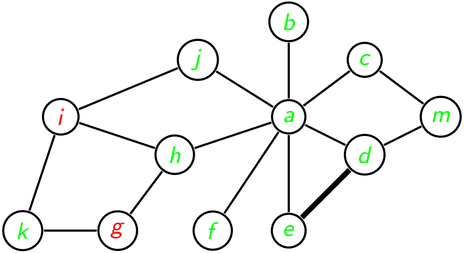

How to change the network

How do we need to change the example to force UPI to dissolve?

- Create a cycle of odd length in the graph (vi)
- Create an cycle of odd length from i (v)
- Make two friends have a path of the same length to i (iv)
- Make two friends express their true belief at the same round (iii)
- Make one agent express his true belief for two consecutive rounds (ii)
- UPI is dissolved, after 6 applications of $\mathcal{I}$ (i).

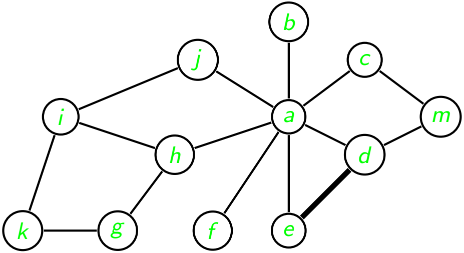

A general framework: A hybrid dynamic network logic

Generalisation

The 2-layer framework can be generalized as follows:

- Assume a finite set of variables (layers) $\{V_1, V_2, ..., V_n\}$
- Each variable V_l takes a value from a finite set R_l, for each $l \in \{1, ..., n\}$.
- Now, all atomic propositions has the form: $V_l = r$, for an $l \in \{1, \ldots, n\}$ and an $r \in R_l$.

General network models

A general network model $\mathcal{M} = (A, \asymp, g, \nu)$ is

- A is a (non-empty) set of agents
- $\asymp \subseteq A \times A$ interpreted as the network structure
- g : NOM $\rightarrow A$ is a function assigning an agent to each nominal
- $\nu : A \rightarrow \mathcal{V}$ is a valuation assigning a value to each variable for each agent.

  $\mathcal{V}$ denotes the set of all *assignments*.

  An assignment $s : \{1, ..., n\} \rightarrow R_1 \times ... \times R_n$ gives a value $\in R_l$ to each variable V_l and given an agent $a \in A$, $\nu(a)$ is an assignment assigning characteristics to a for all variables $V_1, ..., V_n$.

General network models

A general network model $\mathcal{M} = (A, \asymp, g, \nu)$ is

- A is a (non-empty) set of agents
- $\asymp \subseteq A \times A$ interpreted as the network structure
- g : NOM $\to A$ is a function assigning an agent to each nominal
- $\nu : A \to \mathcal{V}$ is a valuation assigning a value to each variable for each agent.

  $\mathcal{V}$ denotes the set of all *assignments*.

  An assignment $s : \{1, ..., n\} \to R_1 \times ... \times R_n$ gives a value $\in R_l$ to each variable V_l and given an agent $a \in A$, $\nu(a)$ is an assignment assigning characteristics to a for all variables $V_1, ..., V_n$.

Truth

Given $\mathcal{M} = (A, \asymp, g, \nu)$, $a \in A$, and a formula φ:

$$\mathcal{M}, a \models V_l = r \qquad \text{iff} \qquad \nu(a)(l) = r$$

An operator for network dynamics

A network dynamic operator $\mathcal{I} = (\Phi, \text{post})$ is

- Φ a finite set of pairwise inconsistent formulas (called "preconditions")
- post : $\Phi \to \mathcal{V}$ is a post-condition function.
  - The intuition is: if an agent satisfies $\varphi \in \Phi$ (in which case, φ is necessarily unique), then after the dynamic event $\mathcal{I}$, a will have the characteristics specified by the post condition post(φ).

Dynamic updating of network models

Given a model $\mathcal{M} = (A, \asymp, g, \nu)$ and a network dynamic operator $\mathcal{I} = (\Phi, \text{post})$, the updated model $\mathcal{M}^{\mathcal{I}} = (A, \asymp, g, \nu')$ is defined by:

$$\nu'(a) = \begin{cases} \text{post}(\varphi) & \text{if there is a } \varphi \in \Phi \text{ such that } \mathcal{M}, a \models \varphi \\ \nu(a) & \text{otherwise} \end{cases} \tag{1}$$

What is the general framework for?

A language for specifying simple network dynamics

- The logic allows us to talk about simple network dynamics in a very general way
- For instance, we can easily talk about networks with a mixture of different types of agents

A language for verifying simple dynamic network properties

- The logic allows us to prove and verify properties of networks and simple dynamics on them

What is the general framework for?

A language for specifying simple network dynamics

- The logic allows us to talk about simple network dynamics in a very general way
- For instance, we can easily talk about networks with a mixture of different types of agents

A language for verifying simple dynamic network properties

- The logic allows us to prove and verify properties of networks and simple dynamics on them

Notes

- We have a complete Hilbert style proof system for the logic
- We conjecture that a terminating tableau system for the logic is also easily obtainable
- Thus, we expect the logic to be decidable

What we have done so far

- Given some motivation for a multi-layer framework to represent social influence in social networks.
- Treated an initial example of two-layer case dynamics.
- Given a characterization of the class of frames on which *UPI* dissolves.
- Designed a logic to deal with the general n-layer cases, allowing us to model any type of complex properties distribution change within a network.
- Proved completeness of this logic.

Further research

- Prove some more general results within our new framework.
- For instance, show how modifying the network structure allows/prevents spreading of some combination of properties.
- Characterize stabilization and speed of stabilization.
- Compare different types of agents.
- Implement the framework in agent-based simulations.

References

Jeremy Seligman, Patrick Girard, and Fenrong Liu.
Logical dynamics of belief change in the community.
under submission, 2013.

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