



THE UNIVERSITY  
OF AUCKLAND

LIBRARIES AND LEARNING SERVICES

## ResearchSpace@Auckland

### **Suggested Reference**

Liu, FR., Seligman, J., & Girard, P. (2013). Logical dynamics of belief change in the community. 28 pages.

### **Copyright**

Items in ResearchSpace are protected by copyright, with all rights reserved, unless otherwise indicated. Previously published items are made available in accordance with the copyright policy of the publisher.

<https://researchspace.auckland.ac.nz/docs/uoa-docs/rights.htm>

## Logical dynamics of belief change in the community

Fenrong Liu · Jeremy Seligman · Patrick Girard

the date of receipt and acceptance should be inserted later

Wherever we are, it is our friends that make our world.  
–Henry Drummond (1851-1897)

**Abstract:** In this paper we explore the relationship between norms of belief revision that may be adopted by members of a community and the resulting dynamic properties of the distribution of beliefs across that community. We show that at a qualitative level many aspects of social belief change can be obtained from a very simple model, which we call ‘threshold influence’. In particular, we focus on the question of what makes the beliefs of a community stable under various dynamical situations. We also consider refinements and alternatives to the ‘threshold’ model, the most significant of which is to consider changes to plausibility judgements rather than mere beliefs. We show first that some such change is mandated by difficult problems with belief-based dynamics related to the need to decide on an order in which different beliefs are considered. Secondly, we show that the resulting plausibility-based account results in a deterministic dynamical system that is *non*-deterministic at the level of beliefs.

**Keywords** Belief revision · belief influence · community · plausibility judgement

When one moves from reasoning about what one believes or knows to what other people believe or know, certain conceptual distinctions become quickly necessary. Perhaps the most well-known of these is the distinction between common knowledge (everyone knows and knows that they know, etc.) and distributed knowledge

---

Fenrong Liu  
Department of Philosophy, Tsinghua University  
E-mail: fenrong@tsinghua.edu.cn

Jeremy Seligman  
Department of Philosophy, The University of Auckland  
E-mail: j.seligman@auckland.ac.nz

Patrick Girard  
Department of Philosophy, The University of Auckland  
E-mail: p.girard@auckland.ac.nz

(together, we would know). The multiplicity of knowers also encourages a shift of topic, from procedures for revising one's own beliefs to those for determining the opinion of a group. When one considers, in addition, social relationships, other topics related to the transmission of knowledge become relevant. This paper considers one of these in particular: the influence on one's beliefs of other agents to whom one is socially related.

Explicit modeling of social relationships plays a significant role in social psychology, artificial intelligence and economics. The seminal [16], which studies the patterns of interpersonal relations for groups in terms of social power, has led to many more recent mathematical and also computational models. For example, motivated by applications to marketing, the question of how to find an influential member of a network and how to maximize social influence have been recently studied in [15] and [21]. Social networks have been extensively studied also in economics (see, e.g., [20]) and there is a new field of 'social simulation' which uses computational models to predict and explain social dynamics (see, e.g., [17]). Yet little attention has been paid to the norms that guide people's reasoning about social relationships. Perhaps this is because of the individual-centered history of epistemology. Other people are inherently unreliable, and so one cannot have purely logical grounds for changing one's belief in response to the opinions of one's peers. Nonetheless, the fact that people do change their beliefs in this way allows information to spread along social channels, and this is of epistemological significance and benefit (perhaps) to the community. Moreover, an understanding of how we can reason about this spread is of interest. Here we will build on our previous work ([28] and especially [23]) to explore the consequences for logic of some simple assumptions about social belief change and propagation.

In Section 1, we introduce an account of how standard models of belief revision, in the tradition of [1], can be extended to models of social influence on one's beliefs. In particular, in Section 1.2, we introduce a specific model based on some fairly conservative assumptions about the thresholds required for us to change our beliefs when influenced by others. This is used to demonstrate a variety of dynamic phenomena, that can be analysed logically using the method of automata (from [23]). In Section 2, we examine more closely the question of what makes a community's beliefs stable with respect to social influence, including consideration of de-stabilizing changes such as when an individual changes her beliefs unilaterally (Section 2.1) and when new social relationships are formed or dissolved (Section 2.3). We also consider the effect of these changes on aggregations of belief across the community (Section 2.2). In Section 3, we consider various alternatives to the simple model of threshold influence, in which those with whom we are socially connected are ranked in some way as more or less reliable. Finally, in Section 4, we note some problems with the reliance on the single proposition revision/contraction model we have inherited. An alternative based on plausibility relations (from [12] and [30]) is explored. Although the procedure is completely determined by the distribution of plausibility judgements across the community, if one looks only at the distribution of *beliefs*, it is non-deterministic, in an interesting way. We also note, in a manner similar to [23], the conflict between these mechanisms for social influence based on plausibility and the requirement that plausibility is transitive.

A few words on the methodology of this paper might be useful. Firstly, as written is a technically ‘light’ style so as to address a broad range of readers and also to allow us space to cover a wider range of topics, and so demonstrate the fruitfulness of taking social relationships seriously when studying the logic of belief change. For those readers who expect a very technical approach, we refer to our other works, e.g. [28], [18] and [29]. Secondly, our approach is squarely within the tradition of dynamic epistemic logic ([9]), and not in any way in competition with it. While other researchers in this community search for models for belief revision ([8], [6], and [14]) our interest is to take social relationships into consideration. Consequently, few details of belief revision as such will be given. The work presented here is compatible with many such accounts. And finally, we are also not in competition with the extensive literature on models of social belief change using probabilistic and other quantitative computational methods. Our focus is different in that we aim to study the logic of belief change, that is, how to reason about the way in which people change their beliefs when influenced by others. While the quantitative approaches are certainly relevant to this project, it is not clear exactly how to adapt them to this purpose. Building further bridges between the logical and the quantitative methodologies is certainly a worthy project, and we have made some small steps in that direction in an Appendix.

## 1 Doxastic influence

To be influenced by my friends is to change my beliefs so that they correspond better to theirs. To begin with, we will consider influence regarding a single proposition  $p$ . If I do not believe  $p$  and some significant number or proportion of my friends do believe it, there are several ways I could respond. I could, of course, ignore their opinions and remain doxastically unperturbed. But if I am influenced to change my beliefs there are at least two ways of doing so: I may *revise* so that I too believe  $p$  or (more cautiously) merely *contract*, removing my belief in its negation  $\neg p$ . We will write  $Rp$  for the action of revision and  $Cp$  for the action of contraction.<sup>1</sup> The only assumptions we will make about revision and contraction is that they are ‘successful.’ This means that after I perform the action  $Rp$ , I will believe that  $p$ , and after I perform the action  $Cp$ , I will not disbelieve  $p$  (i.e., I will not believe  $\neg p$ ).<sup>2</sup> In logical terms, this means accepting the following as axioms:<sup>3</sup>

$$\begin{array}{l} [R\varphi]B\varphi \\ [C\varphi]\neg B\neg\varphi \end{array}$$

Now whether or not I change my beliefs in response to my friend’s opinions and if I change them, whether I revise or merely contract depends on at least three things:

<sup>1</sup> Proponents of any one theory of belief change may read  $Rp$  and  $Cp$  according to their favorite theory. The AGM account of [1] is certainly good enough for our purposes, but nothing we say here will depend too much on the details.

<sup>2</sup> Although success is accepted as a postulate of many accounts of belief change, including AGM, it does impose some limitations. In particular, many higher-order propositions such as the Moore-like propositional form “ $p$  but I do not believe  $p$ ” are problematic.

<sup>3</sup> The general framework here is dynamic logic, in which an expression of the form  $[\pi]\varphi$  is a formula that means ‘after performing action  $\pi$ ,  $\varphi$  is the case’.

1) my own attitude regarding  $p$ , and 2) the cohesiveness of my friends' beliefs concerning  $p$ , and 3) the extent to which I regard any particular friend as an authority on  $p$ . To simplify matters initially, we will assume that everyone regards their friends to be equally authoritative concerning each proposition we consider. This assumption will be relaxed in Sections 3.1 and 3.2.<sup>4</sup> The more cohesive an opposition I face, the more pressure I have to change. But also, if I start with an open mind about  $p$ , I may be more easily influenced than if I hold a strong contrary opinion. The particular balance of these factors varies from person to person and even from belief to belief. It may also be partly determined by higher-level beliefs about the reliability of one's friends or different matters. So instead of committing ourselves to a particular theory of influence, we will merely draw a distinction between two kinds of influence: that which leads, respectively, to revision and to contraction. In the case that I am influenced to revise my beliefs positively in favour of  $p$ , we will say that I am *strongly influenced to believe  $p$* , and write this as  $Sp$ . There may be other reasons to revise my beliefs or to keep them the same, but with regard to social influence alone, *ceteris paribus*, the condition  $Sp$  is necessary and sufficient for me to revise. Likewise, when I am influenced merely to contract my belief in  $\neg p$  (if I had one), without necessarily coming to believe  $p$ , we will say that I am *weakly influenced to believe  $p$* , and write this as  $Wp$ . We will also refer to the corresponding negative conditions of being strongly or weakly influenced to believe  $\neg p$ , written as  $S\neg p$  and  $W\neg p$ , respectively; and we will assume that it is not possible to be simultaneously (strongly or weakly) influenced to believe both  $p$  and to believe  $\neg p$ .

With this terminology and notation, we can define a general operation of social influence regarding  $p$ . My being *influenced* regarding  $p$ , written  $Ip$ , is for me to revise my beliefs so as to believe  $p$  when I am strongly influenced to do so; otherwise, to contract my belief in  $\neg p$  (if I have one) when I am weakly influenced, and similarly for  $\neg p$ . If I am not even weakly influenced, my beliefs will remain unchanged. More concisely,  $Ip$  can be defined as the program<sup>5</sup>

$$\begin{aligned} &\text{if } Sp \text{ then } Rp \text{ else if } Wp \text{ then } C\neg p; \\ &\text{if } S\neg p \text{ then } R\neg p \text{ else if } W\neg p \text{ then } Cp \end{aligned}$$

From a logical point of view, this means that a logical system for  $R$ ,  $C$ ,  $S$  and  $W$  can easily be extended to a system for  $I$  also, using a standard treatment of PDL (propositional dynamic logic). Although we will not be exploring technical logical issues in this paper, this property will allow us to provide a reduction of statements about the dynamic properties of social belief change in the underlying modal language, as a result of the known reduction of iteration-free PDL expressions.<sup>6</sup>

<sup>4</sup> Thanks to Zoé Christoff for bringing to our attention a very nice class of propositions showing why one would not always want to assume that one's friends are authorities. Suppose I do not believe  $p$  and believe that I don't believe it, but all my friends believe that I do, i.e. they believe that I believe  $p$ . It would be odd (to say the least) if I were to be influenced to revise my belief so as to come to believe that I do believe  $p$ !

<sup>5</sup> The order of the positive and negative clauses is unimportant under our assumption that one cannot be both influenced to believe  $p$  and influenced to believe  $\neg p$ .


<sup>6</sup> This is an obvious consequence of standard axiomatic presentations of PDL, such as Definition 4.78 in [11]. For further details, see our [18], which is a generalisation of other systems of dynamic epistemic logic such as [5] and [10].

In the sequel, we will be offering various specific accounts of strong and weak influence, which enable a further reduction to a language that only contains operators for belief and the social relationships that structure our communities. Readers unconcerned with such technical niceties can cheerfully ignore this talk of reduction, which will not play a role in what follows.

### 1.1 The dynamics of influence

There are three possible doxastic states of an agent with respect to the proposition  $p$ : belief ( $Bp$ ), disbelief ( $B\neg p$ ) and no belief ( $\neg Bp \wedge \neg B\neg p$ ), which we abbreviate as  $Up$ . To discuss the distribution of these states among friends, we will use the framework of ‘logic in the community,’ introduced in [28], in which friendship is taken to be a symmetric and irreflexive relation. That is to say, I am a friend of any friend of mine (symmetry), and I am not one of my own friends (irreflexivity). We do not assume that friendship is transitive, so it is quite possible that my friends have friends who are not my friends. A set of agents related by friendship will be called a *social network*. A subset of agents that are connected by friendship, in the sense that for any two agents, there is a chain of friends that connect them, is said to be a *community*.

The distribution of doxastic states within a network can therefore be depicted by diagrams such as the following:




This represents a network of six agents, clustered into two communities. Agent  $a$  believes  $p$  and has friends  $b$  and  $c$ ; agent  $b$  disbelieves  $p$  and has friends  $a$  and  $d$ ; and so on. As well as describing the doxastic states of agents (as  $Bp$ ,  $B\neg p$ , or  $Up$ ), we will describe their position in the social network using the symbol  $F$  to mean ‘all my friends’. Thus  $FBp$  means that all my friends believe  $p$ , which in the example above is a true description of agents  $b, c$  and  $e$  but not of agents  $a, d$  and  $f$ . The dual operator  $\langle F \rangle$  means ‘some of my friends’, so that, for example,  $\langle F \rangle Up$  means that at least one of my friends is undecided about  $p$ . This is true only of agents  $a$  and  $d$ .<sup>7</sup>

Now, given the success axioms for revision and contraction, each agent’s doxastic state will change in a deterministic way under those operations. Revision with  $\neg p$  will take her from state  $Bp$  to state  $B\neg p$ ; mere contraction, from  $Bp$  to  $Up$ , and so on. Moreover, if we assume that the triggering conditions of strong and weak influence depend only on the distribution of doxastic states among agents in the network, this distribution will change under operation  $Ip$  in an entirely deterministic and local fashion. Careful analysis of the definitions given above

<sup>7</sup> The technical details of this language will not be relevant to our present purposes, so we will not go into them here, referring the reader to [28] for further details.

shows that the dynamics of influence is characterised (for each agent, locally) by the following finite state automaton:<sup>8</sup>



The states of the automaton are the possible doxastic states of the agent. The transitions are labelled by mutually exclusive influence conditions. For example, if an agent believes that  $p$ , represented by the state labelled  $Bp$ , and is weakly (but not strongly) influenced by her friends to believe that  $\neg p$ , as represented by the label  $(W\neg p \wedge \neg S\neg p)$ , then after she contracts her belief that  $p$ , she will be in state  $Up$ , undecided about  $p$ .

The reader may like to refer to this automaton to check the numerous examples that follow.

## 1.2 Threshold influence

To give examples of influence in action, we must provide an explicit account of strong and weak influence. A first guess is that strong influence requires cohesion among some threshold proportion of one's friends. Conservatively, we'll assume a threshold of 100%, meaning that I am strongly influenced to believe  $p$  iff all of my friends believe  $p$  (and at least one of my friends believes  $p$ ). The parenthetical clause makes social hermits immune to strong influence, which is surely correct. Strong influence to believe  $\neg p$  works similarly. For weak influence, we'll suppose that one must have at least one friend who believes and that the number of friends who disbelieve is no greater than some threshold. Again, conservatively, we'll assume that to be 0%. This means that if I am not strongly influenced to believe that  $p$  then I am weakly influenced iff none of my friends believe  $\neg p$  (and at least one of my friends believes  $p$ ).<sup>9</sup> This account of strong and weak influence is captured

<sup>8</sup> Analysis of social logical dynamics by finite state automata was used in [23] to show that some interesting dynamic properties (such as the eventual convergence to a stable distribution) can be expressed in terms of operators similar to those we are considering here, but in the domain of preference rather than belief. Here we are providing a slightly more general characterisation for belief change, which does not depend on any particular account of revision, contraction, strong or weak influence. As in [23], it is important to realise that the machine is not the definition of a dynamical system but a tool to analyse what is already implicit in the definition of the logical operators.


<sup>9</sup> This account of strong and weak influence is more-or-less parallel to that given for preference dynamics in [23].

with the following axioms:

$$S\varphi \leftrightarrow (FB\varphi \wedge \langle F \rangle B\varphi)$$

$$W\varphi \leftrightarrow (F\neg B\neg\varphi \wedge \langle F \rangle B\varphi)$$

**Example 1:**




In this example, agents  $a$  and  $d$  both believe  $p$  and are weakly influenced to drop this belief, since all of their friends do not believe  $p$  ( $b$  disbelieves and  $c$  is undecided) and one of their friends,  $b$ , believes  $\neg p$ . Thus  $W\neg p$  is true of  $a$  and  $d$ , and under the operation of social influence,  $I_p$ , they both contract their belief in  $p$  and become undecided. By contrast, agents  $b$  and  $c$  are strongly influenced to believe  $p$  by their friends,  $a$  and  $c$ . Thus  $S_p$  is true of  $b$  and  $c$ , and under the operation  $I_p$ , they both revise so as to believe  $p$ .<sup>10</sup>

## 2 Stability and flux

Example 1 shows what happens after one application of the  $I_p$  operator, but as can be seen from the resulting configuration, further changes would occur. In the configuration on the right,  $a$  and  $d$  are both strongly influenced to believe  $p$ , and so a further application of  $I_p$  would cause them to revise their beliefs, returning them to their previous doxastic states. What happens if we continue in this way? Well, in this case, there is no further change to agents  $b$  and  $c$ , so after only one more application of  $I_p$ , all four agents unanimously believe that  $p$ . Social influence will cause no further changes.

We'll say that a community is *stable* if the operator  $I_p$  has no effect on the doxastic states of any agent in the community. Unanimity within the community is sufficient for stability but not necessary, as is shown below:

**Example 2:**




<sup>10</sup> Here we assume that the agents revise their belief simultaneously. An alternative perspective is to let each agent revise her belief in certain sequential order, corresponding to the scenario in which each agent acts upon others' beliefs at a different time.



In this network, which consists of just one community, no agent is subject to either strong or weak influence. For example, agent  $b$ , who is undecided about  $p$ , has three friends who believe  $p$  but one that disbelieves. On the very weak assumptions we are making about threshold influence, this is not enough to get her to change her mind. Also note that agent  $c$ 's friends are unanimous in being undecided about  $p$ , but this has no effect on  $c$ 's belief in  $p$ .

The community of Example 1 is not stable but becomes stable after one application of  $I_p$ . We will say that such communities are *becoming stable*. Not all communities are becoming stable. Those that never become stable will be said to be *in flux*. Here is an example:

**Example 3:**



In Example 3, each agent is strongly influenced by the other at each stage, and so revises her belief regarding  $p$ , alternately believing and disbelieving  $p$ .

Thus, a brief examination of the dynamics of influence show that all three possibilities can be realised: communities that are stable, those that are in flux, and those that are not stable but are becoming stable. Moreover, since distinct communities in a social network have no influence over each other, it is possible to have a network with several communities of a different dynamic type.

This observation raises some questions for the logic of of friendship and belief. Under what conditions is an agent part of a community that is of each of these types? Characterising stability is fairly easy, because it can be done locally. An agent who is under neither strong nor weak influence to believe  $p$ , and is also under neither strong nor weak influence to believe  $\neg p$ , will not change her mind. But also, an agent who already believes  $p$ , and is under strong or weak influence to believe  $p$ , will also remain unchanged. Reflecting on the automaton, we can see that the following condition is necessary and sufficient for an agent not to change her mind (assuming that strong influence implies weak influence):<sup>11</sup>

$$\neg(B\neg p \wedge Wp) \wedge \neg(U p \wedge S p) \wedge \neg(U p \wedge S\neg p) \wedge \neg(B p \wedge W\neg p)$$

Under the assumption of threshold influence, this is equivalent to

$$\begin{aligned} & \neg(B\neg p \wedge F\neg B\neg p \wedge \langle F \rangle B p) \wedge \\ & \neg(\neg B p \wedge \neg B\neg p \wedge F B p \wedge \langle F \rangle B p) \wedge \\ & \neg((\neg B p \wedge \neg B\neg p \wedge F B\neg p \wedge \langle F \rangle B\neg p) \wedge \\ & \neg(B p \wedge F\neg B p \wedge \langle F \rangle B\neg p) \end{aligned}$$

which is expressed in terms of only friendship and belief. A community is stable when every agent in the community satisfies this condition. Characterising the

<sup>11</sup> The four clauses describe the four cases in which an agent would stay in an automaton state, and so not change the status of her belief.

state of being in flux (or, equivalently, becoming stable) is a little harder. [23] contains a theorem that shows how to do this for the preference dynamics studied there. Here, we will conjecture that a similar property holds for belief, namely, that a community (of at least two agents) is in flux if and only if every agent in the community satisfies the condition

$$(FBp \wedge FFB\text{-}p) \vee (FB\text{-}p \wedge FFBp)$$


In particular, if there is any agent in the community in state  $Up$ , then the community is becoming stable, if not stable already. To describe it in the formal language, we need a bit of global view. Eventually, after a finite number of update (we can introduce  $[Ip]^n$ ), the stable condition holds for each agent.

This highlights the ease of attaining stability. Next, we will study ways to break or introduce stability in a community, and how beliefs may propagate within a community or from one community to another as a result of influence.

## 2.1 Unilateral belief change


Agents may change their minds for many reasons other than the influence of their friends' opinions. This raises the question of if and how such changes are propagated to other members of the community. A very coherent community may resist all such changes, ensuring that any agent who changes her mind unilaterally, will soon be brought back into conformity. On the other hand, a less coherent community, may be highly affected by the change, going into flux or even following the agent who changed her mind into a new stable configuration. We will examine some of the possibilities here starting with a single agent deciding to believe that  $p$  in a unanimously undecided community:

**Example 4:**




In the initial stable configuration, there are three agents, all of whom are friends and all undecided. Now, agent  $a$  changes her belief to  $Bp$ , which we denote in the diagram by ' $a \uparrow Bp$ '. The change is very limited, however, because this new configuration is also stable. After the change, agent  $a$  has no pressure to drop her belief, as all her friends are merely undecided. Those friends,  $b$  and  $c$ , are under weak influence to believe  $p$ , but under threshold influence this is not strong enough for them to change from their undecided state. Agent  $a$ 's unilateral belief change is therefore completely isolated.


Unanimous belief within a community can be strong enough to resist unilateral belief changes even further, as the next example shows.

**Example 5:**

Here agent  $a$  first unilaterally revises her belief in  $p$  to believe  $\neg p$  but immediately reverses this change under strong influence from her friends (indicated by ' $a : Sp$ ' in the diagram). These two examples illustrate resistance to change in communities. Example 4 shows a passive resistance: the other agents are not influenced but tolerate the change. Example 5 shows more active resistance: the agent who changed her belief is forced back into conformity. More radical consequences of unilateral belief change are possible, even with threshold influence, if we change the geometry of the social network.

**Example 6:**

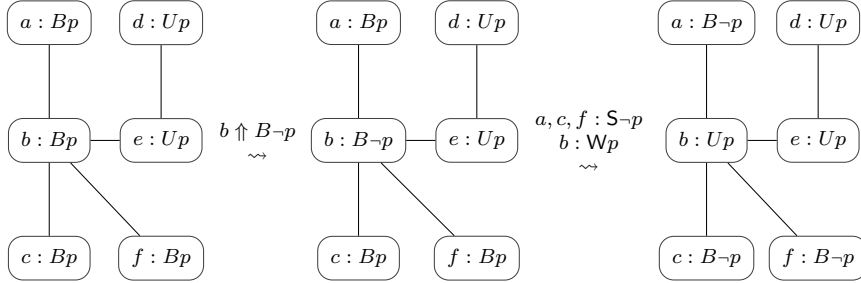
For a community of undecided agents to be influenced by a unilateral belief change, as in Example 6, the location of the agent who comes to believe  $p$  is critical. A peripheral agent will not succeed.

**Example 7:****2.2 Aggregate belief change**

We can get more of a sense of how unilateral changes in belief affect the community by calculating an aggregate opinion. This is a notoriously difficult thing to do well but here we will assume only that a sufficient condition for a community to have an aggregate belief in some proposition is that at least half of the members believe it and no member of the community disbelieves it, although some may be undecided. In such cases, we will say that the community has a *near unanimous opinion*. With

even one agent unilaterally changing her belief, even a near unanimous opinion can be overturned, as the following example shows:

**Example 8:**

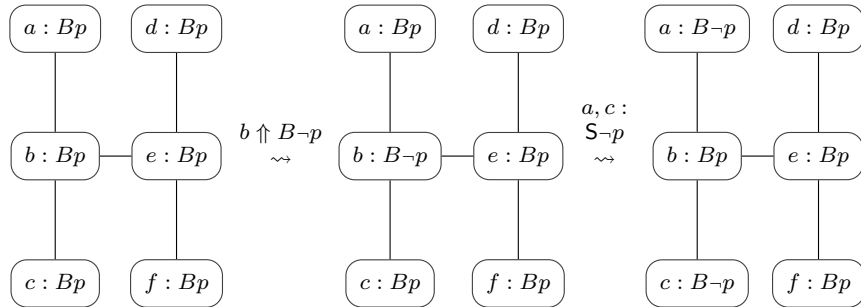


Here, the influence of centrally located agent  $b$  overturns the near unanimous opinion for  $p$ , achieving a group revision. Note that the presence of undecided agents  $d$  and  $e$  is crucial here. Without their stabilizing influence, the rest of the group would oscillate between believing and disbelieving.

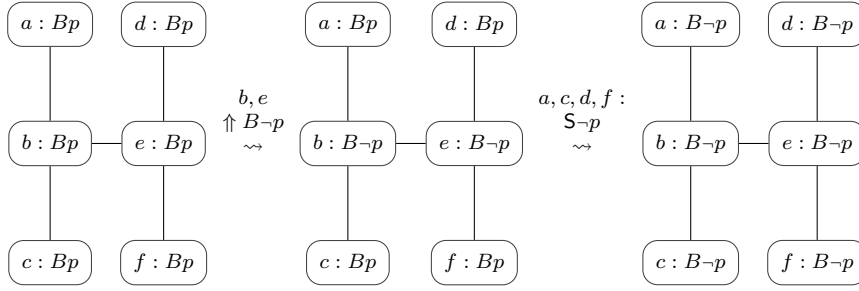
In fact, whenever any two friends are of the same opinion, they will not be influenced to change their beliefs. Take, for instance, two friends  $a$  and  $b$  who both believe that  $p$ . The only way  $a$  will change his mind under strong influence is if all his friends believe  $\neg p$ , which will not happen as long as  $b$  continues to believe  $p$ . Even weak influence on  $a$  requires, at least, that his friends all do not believe  $p$ , for which  $b$  is again a counterexample. So long as  $b$  retains her conviction in  $p$ ,  $a$  will be unaffected by social influence (of the ‘100% threshold’ kind). But the situation is entirely symmetric, and so  $b$  will also be unaffected. The only way in which either will be influenced to change his or her mind is if one of them changes for some other reason (as a ‘unilateral’ belief change.)

The stability of pairs makes it difficult for a single individual to affect large networks.

**Example 9:**




After one more iteration, the community will return to its initial unanimous belief. Yet, if two friends change their minds unilaterally but simultaneously, we can get a total reversal of unanimous opinion:

**Example 10:**

By allowing for unilateral belief change, we can subject stable social opinions to a kind of ‘stress test’. Very stable configurations will not be affected by an individual change of mind but less stable configurations will. Those configurations that are unaffected by two simultaneous unilateral belief changes are even more stable. One could use this to define a measure of the ‘resistance’ of a community to changes of opinion. Given some method of aggregating the beliefs of members of a community (with respect to  $p$ ), we can say that it is  $n$ -resistant if the aggregate opinion does not change as a result of any  $n$  members simultaneously changing their minds (in any way) with respect to  $p$ . Then the community of Example 8 is 0-resistant but not 1-resistant (with any reasonable aggregation mechanism) and that of Example 9 and 10 is 1-resistant but not 2-resistant.


## 2.3 Gaining and losing friends

Changes to the social network can also lead to changes in aggregate opinion and so can be used to distinguish between more stable and less stable communities. The simplest of these occur with a single gain or loss of a friend.

**Example 11:**

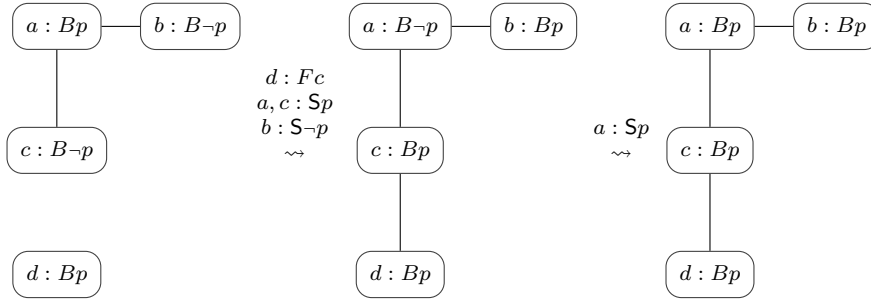
We start with a stable distribution of opinions among three mutual friends, with only one believer. One friendship is broken, putting the mutual friend into a position of greater influence. As well as changing the aggregate opinion, changes in the network can also change the dynamic status of communities, from stable to flux and vice versa

**Example 12:**



The oscillating pair of friends at the top is calmed when an indifferent agent joins their circle. In one more step, they will all become undecided. When the newcomer to a community in doxastic flux is a believer, the influence may be sufficient to convert the whole community:

**Example 13:**



Timing for newcomers is important, however. If the first agent that the newcomer met was of opposite opinion, he may be absorbed into the flux, unless, of course, he is part of another community that provides some stability.<sup>12</sup>

**3 Alternatives to threshold influence**

Our model of threshold influence is thoroughly egalitarian: when it comes to doxastic influence, all friends have equal power over us. One consequence of this assumption, as we saw above, is that pairs of friends of the same state of mind (believers, disbelievers or undecided) will never be influenced to change their beliefs. We will now consider a slightly different approach, whereby we think of the conditions for strong and weak influence as arising from ways of aggregating the opinions of our friends, who are ordered according to their relative power over us.

So, let's assume that our friends are (partially) ordered by a relation we will call 'better friend'. It is both irreflexive (no one is a better friend of mine than herself) and transitive (if  $a$  is a better friend than  $b$ , who is a better friend than  $c$ , then  $a$  is a better friend than  $c$ ) but not necessarily linear: I may have two friends  $a$  and  $b$  neither of whom is a better friend than the other. This talk of 'better friend' is only a *façon de parler*; what we really mean by it is a kind of social power. To say

<sup>12</sup> For example, if  $d$  is initially friends with another believer,  $e$ , who is not connected to  $a$ ,  $b$  or  $c$ , then  $d$  will be immune to change.

that  $a$  is a better friend (of mine) than  $b$  means only that  $a$  has greater power to influence me than  $b$ .

For strong influence to revise our beliefs, we will suppose that only our best friends are consulted, but they must believe unanimously. This amounts to aggregating our best friends' opinions with a very strong requirement for aggregation, namely, unanimity. It is still conservative but more liberal than the threshold condition we have been using, since we do not require anything of our wider circle of friends.

For weak influence to contract a belief, we will suppose that all our friends are consulted but that we aggregate their opinions in a way that gives priority to better friends. Specifically, we say that, on aggregate, our friends believe if for every disbelieving friend, we have a better friend who believes. This is a 'defeasibility' model: the opinions of disbelieving friends are defeated by their betters.<sup>13</sup>

### 3.1 Ranked Influence

To capture this new model of social influence axiomatically, in full generality, we would need to distinguish between friends in terms of their power over us, introducing a binary operator  $P$ , such that  $P\varphi\psi$  means 'for all my friends who  $\varphi$ , I have a better friend who  $\psi$ '. This presents some technical difficulties because  $P$  is not a normal modal operator.<sup>14</sup> Rather than tackle these difficulties here, we will introduce a simplifying assumption using the concept of *rank*. Our best friends are of rank 1. Those who are the best of the remainder (when we remove our best friends) are of rank 2. Those who are the best of the remainder (when we also remove our friends of rank 2) are of rank 3. And so on. The assumption is that having a higher rank is also sufficient for being a better friend: that if  $a$  is a higher ranked friend (for me) than  $b$ , then  $a$  is a better friend of mine than  $b$ .<sup>15</sup>

We assume that there are only a finite number  $N$  of agents in the social network, and so the lowest possible rank of friends is  $N - 1$ .<sup>16</sup> We will therefore introduced new symbols  $F_1, \dots, F_{N-1}$ , with  $F_i$  meaning 'all of my friends of rank  $i$ '. Thus we

<sup>13</sup> This method of aggregation has been well-studied, although mainly with regard to preference rather than belief. Our ordering of friends is what is known as a 'priority graph' in [2], and the method itself is known as 'lexicographic aggregation'. To see the connection with dictionaries, think of a pair of words (of equal length) and the order they are listed. Word  $X$  comes before word  $Y$  just in case for every letter in  $Y$  that comes before the corresponding letter in word  $X$  (in alphabetic order), there is an earlier letter in  $Y$  that comes after the corresponding letter in word  $X$ . If the words are not of equal length, this definition can still be made to work by padding the shorter word with extra 'space' characters, which are considered to come before all the letters of the alphabet.

<sup>14</sup> The normal modal operator defined over the 'better friend' relation,  $Q\varphi\psi$ , means 'for all my friends who  $\varphi$ , every better friend  $\psi$ '. But there is no way of defining  $P$  in terms of  $Q$ .

<sup>15</sup> To see that this additional assumption is non-trivial, suppose I have one best friend,  $a$ , three other friends,  $b$ ,  $c$  and  $d$ , with  $b$  a better friend than  $d$ . If  $b$  and  $c$  are incomparable, then neither is a better friend of mine than the other. But then  $a$  has rank 1,  $b$  and  $c$  have rank 2, and  $d$  has rank 3. This implies that  $c$  is a better friend of mine than  $d$ , which is an inference we could not make without the ranking assumption.

<sup>16</sup> Since friendship is assumed to be irreflexive, there must be at least two agents in order for there to be any friends at all.

have the following equivalence


$$F\varphi \leftrightarrow \bigwedge_{i < N} F_i\varphi$$

We can now express ranked strong and weak influence with the following axioms:<sup>17</sup>

$$\begin{aligned} S\varphi &\leftrightarrow F_1 B\varphi \wedge \langle F_1 \rangle B\varphi \\ W\varphi &\leftrightarrow \bigwedge_{i < N} F_i (B\neg\varphi \rightarrow \bigvee_{j < i} \langle F_j \rangle B\varphi) \wedge \langle F \rangle B\varphi \end{aligned}$$

Consider the following social network:

**Example 14:**



Each node in the diagram now represents both what the agent believes and how she ranks her friends, represented by a simple list of the ranks. So, for example,  $a$ 's best friends are  $c$  and  $d$ , with  $b$  in the second rank, and so  $a$  is weakly influenced to contract her disbelief in  $p$ . This is because her only friend who fails to believe  $p$ , namely  $b$ , is of second rank, and she has a best friend,  $c$ , who believes  $p$ . Thus under ranked influence,  $a$  will become undecided, whereas under threshold influence,  $a$  would not change her beliefs because her friend  $b$  shares her disbelief in  $p$ .

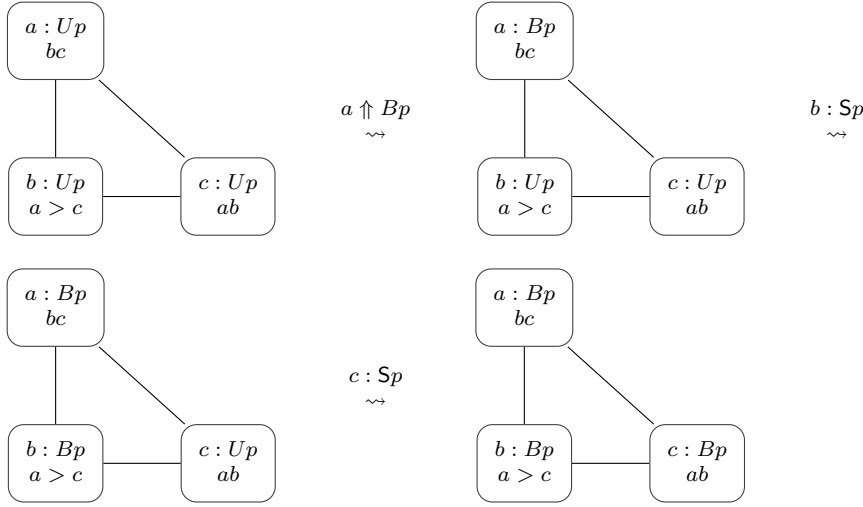
Ranked influence allows changes to spread within a community more easily, if conditions are right. In the next example, a single agent's new belief spreads to his community of previously undecided friends.

<sup>17</sup> The two axioms look superficially very different, but the first has an equivalent form that displays the difference more clearly:

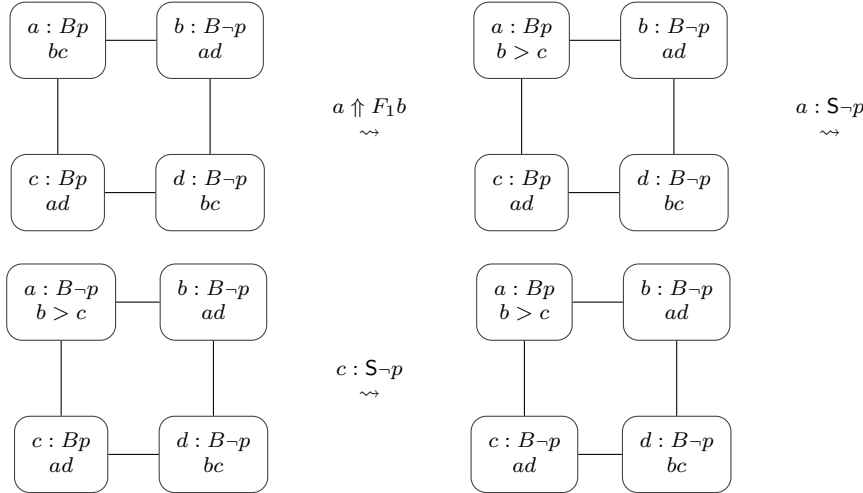
$$S\varphi \leftrightarrow \bigwedge_{i \leq N} F_i (\neg B\varphi \rightarrow \bigvee_{j < i} \langle F_j \rangle B\varphi) \wedge \langle F \rangle B\varphi$$

This expresses the apparently weaker condition that for every friend who does not believe  $\varphi$ , I have a better friend who does. But for this to be false, I must have a friend who doesn't believe  $\varphi$  and no better friend who does. But then either that friend is a best friend, or I have a best friend (and so a better friend) who does not believe  $\varphi$ , preventing strong influence.



**Example 15:**

Although the absolute stability of pairs of likeminded friends under threshold influence is disrupted when we move to ranked influence, paired best friends are still immune to change. But now we can also consider changes to the ranking of a given agent's friends, as a refinement of the actions of gaining and losing friends considered earlier.

**Example 16:**

Here, we represent the operation of promoting  $b$  to (sole) best friend by  $\uparrow F_1b$ . The new asymmetry allows  $b$ 's disbelief to spread to the rest of the community. One could think of many other operations on rankings that could be studied. For example, one could demote a friend, or promote/demote to a specific rank.

### 3.2 Believed Reliability

When the ‘better friend’ relation is re-interpreted in other more doxastically relevant ways, such as the relation of ‘regarded as having more expertise than’ or ‘taken as a better authority than’, the subjective component of an agent’s rankings becomes evident. This raises the possibility of another way of modelling the relative power of other agents to affect us: to consider a binary relation of ‘being more reliable than’ between agents and then express the conditions for social influence in terms of which agents within one’s community one *believes* to be more reliable. As in ranked influence, strong influence would require unanimity between those one takes to be most reliable, and weak influence to believe  $p$  would require any agent who disbelieves  $p$  to be defeated by an agent who believes  $p$  and whom one believes to be more reliable. Writing  $L$  for the operator ‘every more reliable friend’, we could then axiomatise this new notion of social influence, which we dub ‘reliable influence’, as follows:<sup>18</sup>


$$\begin{aligned} S\varphi &\leftrightarrow BF(\neg B\varphi \rightarrow \langle L \rangle B\varphi) \\ W\varphi &\leftrightarrow B(F(B\neg\varphi \rightarrow \langle L \rangle B\varphi) \wedge \langle F \rangle B\varphi) \end{aligned}$$

We will not go into the details of this approach here, merely noting that it offer a shift in focus from merely describing doxastic influence to proposing a normative theory of doxastic influence, according to which it is rational to change one’s beliefs when one believes that they are coherently opposed by people who one regards as reliable. Further constraints on when it is and is not rational to revise would be expected, and our aim here is only to point the way to a fruitful line of research.

### 4 Plausibility influence

The focus of the preceding sections has been the dynamics of belief change in a community. In the interest of providing clear examples we have considered social influence with respect to a single proposition. But in doing so, we have ignored a very important aspect of a person’s beliefs, namely their interdependence. Changing one belief may well affect other beliefs. So a natural question to ask is whether the order matters when it come to calculating social influence. The answer is yes.

**Example 17:**



In the depicted situation, how is  $d$  to change her beliefs? We will see that the outcome depends on whether we consider influence with respect to  $p$  and then  $q$  or vice versa. First, considering the proposition  $p$ , because all of her friends either

<sup>18</sup> The form of the axiom for strong influence adapts the alternative axiom of the ranked version given in Footnote 17.

believe  $\neg p$  or are undecided on  $p$ , and assuming threshold influence, we conclude that she will contract and so come to be undecided about  $p$ . But since she also believes  $q$  and  $q \rightarrow p$ , merely removing  $p$  from her belief set would not be successful; she would also have to remove either her beliefs that  $q$  or her belief that  $q \rightarrow p$ . Now, suppose that her belief that  $q \rightarrow p$  is more entrenched than her belief that  $q$ , so she removes  $q$  from her belief set.<sup>19</sup> Yet all of  $d$ 's friends also believe  $q$  and so, under strong influence, she will revise her beliefs so that she believes  $q$ , and then, believing  $q \rightarrow p$ , she will again believe  $p$ . The result of considering influence in the order  $p, q$  is therefore that  $d$ 's beliefs are unchanged. But in the opposite order, she will cease to believe both  $p$  and  $q$ .

This is an uncomfortably strange result, and to address it, we will move away from those approaches to belief revision that take it to be an operation defined on the propositional contents of beliefs. This is the dominant tradition in the literature but there are alternatives. The one we consider here is the tradition of [12] and [30] in which an agent's beliefs are taken to be supervenient on her judgements regarding the plausibility of specific outcomes. As we will see, this provides us with a solution to the problem of multiple issues, and enables us to make some interesting distinctions in the social setting, but comes with its own challenges.

Given a fixed domain  $W$  of possible outcomes, we will consider each agent's judgements regarding the relative plausibility of elements of  $W$ . For  $u$  and  $v$  in  $W$ , we write  $u \leq_a v$  to mean that  $a$  judges  $v$  to be at least as plausible as  $u$ . Importantly, it is possible for this relation to fail to be antisymmetric: two outcomes may be regarded as equally plausible. Also, the relation may fail to be total: there may be outcomes  $u$  and  $v$  about which the agent has no judgement regarding their relative plausibility:  $u \not\leq_a v$  and  $v \not\leq_a u$ . Thus, one fundamental change from our previous model is that there are now four (rather than three) relevant possible states of an agent: agent  $a$  may find  $v$  strictly more plausible than  $u$  ( $u \leq_a v$  and  $v \not\leq_a u$ ) or vice versa, or may regard them as equally plausible ( $u \leq_a v$  and  $v \leq_a u$ ) or have no view at all ( $u \not\leq_a v$  and  $v \not\leq_a u$ ).

There are various ways in which one might think of plausibility judgements as determining beliefs but the dominant idea (from [12] and [30]) is that an agent  $a$  believes that  $p$  just in case  $p$  holds in all the outcomes that are maximally plausible for  $a$ . An outcome is 'maximally plausible' just in case there is no other outcome that the agent judges to be strictly more plausible. Although some suggestions have been made about how to model belief revision in this framework, we will consider the separate but related question of how to model revision of plausibility judgements themselves. In the first instance, this is much more straightforward than revising her beliefs. If she wishes to revise so that she judges  $w_2$  to be at least as plausible as  $w_1$ , she should revise to

$$\leq_a \cup \{\langle w_1, w_2 \rangle\}$$

In other words, she should regard  $v$  as at least as plausible as  $u$  iff she previously regarded  $v$  as at least as plausible as  $u$  or  $u = w_1$  and  $v = w_2$ .<sup>20</sup> Likewise, if she


<sup>19</sup> For more on entrenched belief change, see [26] and [27].

<sup>20</sup> This proposal raises certain problems, especially concerning the transitivity of plausibility judgements. We will address these below.


wishes to revise so that she judges  $w_2$  to be not at least as plausible as  $w_1$ , she should revise to

$$\leq_a \setminus \{\langle w_1, w_2 \rangle\}$$

Just as in our earlier model, we must also specify the conditions under which these revisions are made. This time, we will consider only the simplest possible proposal: that an agent revises her plausibility judgements (in a positive or negative direction) iff all her friends are unanimous. So, if they all take  $v$  to be at least as plausible as  $u$ , so does she, and if they all take  $v$  not to be at least as plausible as  $u$ , nor does she. We will call this *plausibility influence* and write the corresponding operator as  $\mathcal{I}$ .<sup>21</sup> Now, let's focus on a particular couple of outcomes,  $u$  (the 'left' one) and  $v$  (the 'right' one). The action of plausibility influence is characterised by the following automaton:



(where  $\dot{F}\varphi$  is an abbreviation for  $(F\varphi \wedge \langle F \rangle \varphi)$ , i.e. the version of the universal quantifier that takes it to have existential import.) Of course, this is only half the story, the other half of which is given by swapping  $u$  and  $v$ . As mentioned above, an agent can have one of four attitudes with respect to  $u$  and  $v$ , which we will label as  $R$  (for 'right')  $v$  is strictly more plausible than  $u$ ,  $L$  (for 'left')  $u$  is strictly more plausible than  $v$ ,  $I$  (for 'impartial')  $u$  and  $v$  are equally plausible, and  $O$  (for 'no opinion'). In these terms  $u \leq v$  is  $(I \vee R)$ ,  $u \not\leq v$  is  $(O \vee L)$ ,  $v \leq u$  is  $(I \vee L)$ , and  $v \not\leq u$  is  $(O \vee R)$ . The dynamics of the two parts of the comparison can therefore be represented as the two automata




whose product completely describes the dynamics of plausibility influence with respect to the four states  $L$ ,  $R$ ,  $I$  and  $O$ :

<sup>21</sup> More precisely, plausibility influence is the operation that transforms the plausibility judgements of all agents in such a way that agent  $a$  deems  $v$  to be at least as plausible as  $u$  iff the pair  $\langle u, v \rangle$  is in the set

$$\leq_a \cup \bigcap_{a \succ b} \leq_b \setminus \bigcap_{a \succ b} \not\leq_b$$

where  $x \succ y$  means that  $x$  is friends with  $y$ . Note that the order in which the operations of adding and subtracting from the relation are performed is not important because, with at least one friend, it can never be that all my friends both do and do not regard  $v$  as at least as plausible as  $u$ .




where  $W(\alpha, \beta)$  is the condition  $(\dot{F}(\alpha \vee \beta) \wedge \neg \dot{F}\beta)$ .<sup>22</sup> This analysis of the dynamics shows that little has changed from our earlier models, in the sense that friends of the same type ( $L, R, I, O$ ) will be immune to influence from others; there are the typical unstable alternations of the form  $X \rightarrow Y \rightarrow X$  where  $X$  and  $Y$  are agent ‘types’; and a characterisation of stability could be obtained using the methods of [23].

#### 4.1 A Finer-grained Dynamics


The novelty of the plausibility approach to social influence is its fine-grained analysis of belief dynamics. Although it is still deterministic at the level of plausibility, unlike the revision/contraction approach, it is *not* deterministic at the level of beliefs. For example, suppose agents  $b$  and  $c$  are friends who both believe  $p$ . In our previous model, this makes them invulnerable to influence with respect to this belief, and if there is a third agent  $a$  who is friends with both of them (and no one else), she will be strongly influenced to believe  $p$  also, no matter what her initial view. But with plausibility influence this is no longer the case.

Suppose there are three possible outcomes,  $u_1, u_2$  and  $v$  with  $p$  true at each of  $u_1$  and  $u_2$  but not at  $v$ . And suppose that agent  $b$  regards  $u_1$  as strictly more plausible than  $v$ , agent  $c$  regards  $u_2$  as strictly more plausible than  $v$  but that these are the only judgements they make. In particular, agent  $a$  makes no judgements at all. This is depicted below:




<sup>22</sup> As before, we will save cluttering our diagrams by assuming that if none of the conditions for a transition apply then the automaton stays in its current state.

Agents  $b$  and  $c$  agree that  $u_1$  and  $u_2$  are the only maximally plausible outcomes, and so believe that  $p$ . Their friend  $a$  also allows that outcome  $v$  is maximally plausible and so is undecided about  $p$ . Thus we have the following configuration:




But after plausibility influence, agent  $b$  drops her judgement that  $u_1$  is more plausible than  $v$  because it is not supported by either friend. Likewise, agent  $c$  drops her judgement that  $u_2$  is more plausible than  $v$ , making all three agents converge to  $a$ 's initial view, and so the new configuration is




We can interpret this as capturing, to some extent, the influence of reasons rather than mere beliefs. On this interpretation, the agent's reasons for believing are given by her plausibility judgements; it is those that are influenced by her social environment. Influence on her beliefs is a secondary matter.

#### 4.2 Transitivity

One problem with plausibility influence is that it does not work very nicely with the rather natural requirement that plausibility judgements are transitive. For example, suppose instead that the three friends make the following plausibility judgements about outcomes  $u$ ,  $v$  and  $w$ :



Each of these satisfies transitivity. But, under plausibility influence, each friend will change to the paradigmatically intransitive set of plausibility judgements given by the following



Put briefly, the operation of plausibility influence fails to preserve transitivity. There are several ways to respond to this. One is to take the transitive closure of each agent's plausibility relation after calculating the effect of social influence, perhaps interpreted as an act of self-critical evaluation. This would make it marginally more difficult for an agent to be influenced by his peers. A slight drawback is that it becomes more difficult to interpret the dynamics as operating on plausibility judgements as reasons for beliefs: taking  $w$  to be more plausible than  $u$  is not a reason *in addition* to taking  $w$  to be more plausible than  $v$  and  $v$  to be more plausible than  $u$ . A finer-grained dynamics of plausibility judgements as reasons for belief, of the kind suggested above, would have to make a distinction between primary judgements and those that are inferred by transitivity. A second response is to modify the definition of belief slightly, so that it is the transitive closure of the plausibility relation, for each agent, that determines the agent's beliefs.

## 5 Conclusion and future work

We have conducted a fairly open-ended survey of some of the possibilities for modelling the way in which people's beliefs are influenced by their social relations. Unlike models offered by social psychologists and sociologists, our aim is not descriptive but normative. We are interested in the relationship between norms of belief revision that may be adopted by members of a community and the resulting dynamic properties of the distribution of beliefs across that community. Nonetheless, we have seen, at a qualitative level, that many aspects of social belief change that we see in real communities can be obtained from even a very simplistic model, which we called 'threshold influence'. In particular, we focussed on the question of what makes the beliefs of a community stable under the dynamics of influence itself and various 'stress tests' such as unilateral changes of belief by individuals within the community and changes to the social network itself. We also considered refinements and alternatives to the 'threshold' model. The most significant alternative was to move to consideration of plausibility judgements rather than mere beliefs. We showed first that some such change is mandated by difficult problems with belief-based dynamics related to the need to decide on an order in which different beliefs are considered. Secondly, we showed that the resulting plausibility-based account results in a dynamical system that is non-deterministic at the level of beliefs. Nonetheless, the plausibility-based account we considered lacks certain intuitively desirable features, such as the preservation of the transitivity of plausibility judgements.

For future work, there are many of interesting open issues to consider. The ideas charted here each point to specific formal systems whose static and dynamic prop-

erties could be studied in much more detail. The ranked model, in particular, is natural and powerful, especially when given the normative interpretation of degrees of reliability from Section 3.2. Finding axiomatisations of these logics would be challenging and may reveal interesting candidates for norms of reasoning about belief change. Exploring the consequences of threshold influence when the threshold is strictly between 0 and 1 would take us closer to more quantitative theories. In each case, theorems characterising stability and flux are clearly within reach, using the automata technique mentioned here. On a more conceptual level, a natural question to ask is whether similar techniques can be used for attributing beliefs to groups, and for exploring the dynamic interaction between belief change at the individual and group level. And in all these areas, a fuller study of the effect of the network topology on belief change is needed.

## Appendix

In this appendix, we compare our model of belief influence with established research on social networks. We base our discussion on chapters 7-9 of [20], which present the studies we think are most relevant to ours. We divide the discussion in three parts, looking at models of diffusion (or contagion), learning and games on networks. One obvious aspect of our approach, in contrast to all those presented in [20], is our reliance on a doxastic logical language of social interactions. Crucially, our aim is a logical one: to study logical patterns of reasoning about belief influence, not to provide descriptively adequate models of social phenomena. In this respect, our use of a formal language is a significant formal departure from the more established results in the literature. Nevertheless, there are clear similarities between our models and those studied in [20], and this deserves comment.

### Diffusion

In our models, we are interested in the effect of unilateral belief change and how they can propagate to a community. One typical way of analysing diffusion of information in a network is to model it as the spread of a disease in a community. Models have networks of agents that can be in one of two states: healthy or infected. Agents can contract a disease when they are exposed with a non-negative probability of infection.<sup>23</sup>

A natural question for us is whether belief influence can display the viral patterns commonly observed in diffusion through network. The analogy between a disease and belief would have to be framed so that a belief is like a virus. But to model that we would have to bias influence in favour of the viral belief, whereas our approach is more symmetric between believing something and its negation. Furthermore, our third belief state of being undecided is still left out in the analogy. One option would be to link it to immunity, but this would not be right, as strong influence may still change an undecided agent to a believer. Setting these concerns aside, let us look a bit closer at models of viral diffusion.

---

<sup>23</sup> Some models have an additional immune state, but agents that are immune are no longer relevant to the spread, as they cannot infect anybody else, nor can they contract the disease again, and so they are simply taken off the equation.



An early influential model of diffusion is known as the Bass Model [7], but doesn't have any explicit social network component in its modelling. The models of diffusion most relevant to belief influence are the models known as SIR ("susceptible, infected, removed", Kermack and McKendrick, cf., [22]) and SIS ("susceptible, infected, susceptible", cf., Bailey [3]). In each model, agents can be in one of two states, "susceptible" in which they may become infected, and "infected" when they contract the disease. In SIR models, agents that have been infected are eventually removed from the network, when they are cured from the disease, at which point they are no longer contagious, nor can they infect others. This account for the "removed" part in SIR, modelling diseases such as chickenpox. But this doesn't match belief influence. Consider for instance a community in flux, in which each agent keeps changing their beliefs at any state. In SIS models, however, agents become susceptible again after having been infected, modelling diseases such as the flu. Now, in SIS models, edges of the network represent the possibility of physical contact between agents and so the chance of an agent being infected is increased by the number of its network neighbours (its "degree"). The higher the degree, the higher the chance of infection.

One drawback of adapting SIS to model belief propagation is that, in any finite system, the infection will eventually entirely stop (see [20, p.199 and Exercise 7.5]). Thus to get flux going in SIS models, an infinite number of agents are required, which is an unacceptable assumption for us. Furthermore, as we already noted above, our models have three belief states that react to influence differently, with the undecided state only changeable in cases of strong influence. But none of the states is favoured over the others. Contagion is a better metaphor of the propagation of information than it is of belief. Once an agent receives the information, it cannot be taken from her (except perhaps by her forgetting, which is an additional complexity). This marks an important distinction between the diffusion of information in a network and the propagation of a belief state in a network via influence. In summary, diffusion concerns to spread of a disease, or a piece of information, or even a behaviour throughout a network, which has some superficial similarities with belief influence, but also very different dynamic properties.

## Learning

Another common theme of social network research is the influence of network structure on learning and consensus of beliefs. The kind of question asked here is whether a community will come to share a belief, and under which conditions, or who has most influence in such propagation. A recent influential model is the Bayesian model of observational learning by Bala and Goyal [4], in which agents learn from their experience and that of their neighbours in decision making. The main result from Bala and Goyal is that "in a connected society, local learning ensures that all agents obtain the same payoffs in the long run. Thus, if actions have different payoffs, then all agents choose the same action, and social conformism obtains." The intuition behind this result is that when agents have access to outcomes acquired by the actions of their neighbours, they will adopt the action that produced the best outcomes, until every agent converges on the same action. A question that naturally arise is whether the agents will converge on the optimal action. In doxastic terms, if the actions are to choose amongst alternatives to believe, then the prediction is that eventually a connected society would come to

a consensus on what to believe. With learning, one could ask if the society has come to learn the “right belief”. This approach is hardly consistent with belief influence, even though patterns of influence from neighbours are also an important component of models. One obvious difference is that influence models can be stable without unanimity.

An earlier model of imitation and social influence which seems more natural with respect to belief influence is the DeGroot model [13]. In this model, the initial configuration is that of a group of agents who each have their own opinion on a topic, expressed as a vector of probabilities  $p(0) = (p_1(0), \dots, p_n(0))$ . On top of that, each agent provides a hierarchy of the agents as a weighting that sums up to 1. For instance, with three agents  $\{1, 2, 3\}$ , agent 1 may weigh the agents as a vector  $(2/3, 0, 1/3)$ , meaning that she has no confidence in agent 2, and is twice as confident in herself as she is in agent 3. The dynamics in the system is thus captured by a weighted and directed  $n \times n$  non-negative matrix  $T$ . In the matrix, entry  $T_{ij}$  is the weight attributed to agent  $j$  by agent  $i$  operating on the vectors of probability:  $p(t) = Tp(t-1) = T^t p(0)$ . So  $p_i(t)$  represents the degree of belief of agent  $i$  at time  $t$ . In DeGroot models, unlike those of Bala and Goyal, it is not the case that initial conditions will always lead to consensus. As in our approach, it is possible to characterise whether or not the distribution of opinions across the network will stabilise (cf., Golub and Jackson [19]).

An adaption of this model that is similar to our threshold version of strong influence, called “majority dynamics”, is investigated in [25]. Majority dynamics works over a set of two issues, say 0 and 1, and at each iteration, each agent adopts the opinion of the majority of her friends. One can see this as the discretization of DeGroot model in which agents can only choose two alternatives instead of the whole range between 0 and 1. As a result “the probability of choosing the correct alternative approaches one as the size of the smallest social type approaches infinity, with a polynomial dependence.”

Although some similarities may be found with what we call “strong belief influence”, there are important dissimilarities. One general issue is that all these models use some sort of transformation of the degrees of belief of agents according to weighting. We parametrise our account with any version of propositional belief revision (and contraction). Furthermore, it is not obvious how one would differentiate between weak and strong influence, nor how one would give interpretations of  $Bp$ ,  $B\neg p$  and  $Up$  in those models. Furthermore, an indication that our update rules operate differently is that our notion of convergence, namely stability, does not always yield consensus. In the other direction, however, our approach raises interesting issues for learning models. The distinction between weak and strong influence implements the idea of degrees of influence using conventional distinctions in the theory of belief revision. How would one integrate this distinction, formalise it at the level of beliefs and introduce dynamic rules in learning models? Would network still converge under the same conditions as before, or would we need different necessary and sufficient conditions for convergence?

### Games on Networks

Another topic that has been studied extensively is the influence networks may have on individual behaviour, or actions, for instance in buying or selling products, or

engaging in some communal activities. One way of formalising and analysing this is by interpreting them as games, in the sense of Game Theory, and looking for equilibria. Since the actions we take often depend on what others do, game analysis of network interaction can reveal intricate patterns of strategies on networks. And in contrast to what we've seen in learning, it may sometimes be more profitable for agents to adopt opposite actions to those of their neighbours, and here games are again useful in making these distinctions. A general setting for games on networks is a probabilistic analysis of how agents react to their neighbours' actions. Game theoretic notions such as best-response then naturally give rise to studies of equilibria.

Could belief influence be analysed as a strategic interpretation of doxastic actions? For instance, we may say that if all my friends believe  $p$ , then I'm better off adopting a belief in  $p$ . Some pragmatic interpretations might motivate such an interpretation, for instance in argumentative scenarios, in which idiosyncratic beliefs are deemed implausible and thus unusable. More concretely, one could give the following game interpretation to our models: assign a utility of 1 if all my friends agree with me (they and I believe  $p$  or believe  $\neg p$ ), 0 if some of my friends agree with me and some disagree, and -1 if all of them disagree with me. One could then look at equilibria in this game and see whether they coincide with our stable configurations. We leave this for future research.

It is no surprise then that we can find game models on networks that come very close to our idea of influence. We focus on one such model, that of Morris [24]. Morris is interested in the case of contagion of actions in populations with emphasis on local interaction games, that is games in which agents react to the actions of a selected number of agents from the population.<sup>24</sup> Now, each agent has a choice between two actions 0 and 1, and, there is a threshold such that at each tick of a clock, an agent will adopt the action of her neighbours if the number of neighbours is greater than the threshold. This is the same definition as our strong influence with threshold. Say that a subset  $S$  of a network is  $r$ -cohesive if each agent has at least  $r$  of her neighbours in  $S$ . The main theorem Morris proved is that both actions 0 and 1 are played by different subsets of the networks in equilibria iff there is some non-empty and strict subset of players  $S$  that is  $q$ -cohesive and such that its complement is  $(1-q)$ -cohesive. These conditions for us would also yield stable networks - ignoring weak influence. Now, for the same reasons as above, our approach raises interesting issues for Morris's model. What would happen to the characterisation of stability if one could distinguished between weak and strong influence? We leave this question open for future research.

In conclusion, our belief influence model is novel in applying the distinction introduced in [23] between weak and strong doxastic influence to the case of doxastic influence in social networks. It would be an interesting future project to see how this distinction would impact on famous results like the ones we've considered in this brief discussion.

---

<sup>24</sup> For his results, Morris only considers infinite populations. This is unacceptable from our point of view, but necessary for his mathematical analysis. But let's ignore this for the sake of comparison.

## References

1. Alchourrón, C., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* **50**, 510–530 (1985)
2. Andréka, H., Ryan, M., Schobbens, P.Y.: Operators and laws for combining preferential relations. *Journal of Logic and Computation* **12**, 12–53 (2002)
3. Bailey, N., et al.: The mathematical theory of infectious diseases and its applications. Charles Griffin & Company Ltd, 5a Crendon Street, High Wycombe, Bucks HP13 6LE. (1975)
4. Bala, V., Goyal, S.: Learning from Neighbors. *Econometric Institute Reports*. Erasmus University (1996). URL <http://books.google.co.nz/books?id=3XrESgAACAAJ>
5. Baltag, A., Moss, L., Solecki, S.: The logic of public announcements, common knowledge and private suspicions. In: TARK 98 (1998)
6. Baltag, A., Smets, S.: A qualitative theory of dynamic interactive belief revision. In: M.W. G. Bonanno W. van der Hoek (ed.) *Logic and the Foundations of Game and Decision Theory, Texts in Logic and Games*, vol. 3. Amsterdam: Amsterdam University Press (2008)
7. Bass, F.: A new product growth for model consumer durables. *Management Science* **15**(5) (1969)
8. van Benthem, J.: Dynamic logic for belief revision. *Journal of Applied Non-Classical Logic* **17**, 129–156 (2007)
9. van Benthem, J.: *Logical Dynamics of Information and Interaction*. Cambridge: Cambridge University Press (2011)
10. van Benthem, J., van Eijck, J., Kooi, B.: Logics of communication and change. *Information and computation* **204**(11), 1620–1662 (2006)
11. Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. Cambridge University Press, Cambridge, Mass. (2001)
12. Burgess, J.: Basic Tense Logic, *Handbook of Philosophical Logic*, vol. 2, chap. 2, pp. 89–133. Dordrecht: D. Reidel (1984)
13. DeGroot, M.: Reaching a consensus. *Journal of the American Statistical Association* **69**(345), 118–121 (1974)
14. van Ditmarsch, H.: Prolegomena to dynamic logic for belief revision. *Synthese(Knowledge, Rationality & Action)* **147**, 229–275 (2005)
15. Domingos, P., Richardson, M.: Mining the network value of customers. In: *Proceedings of the Seventh International Conference on Knowledge Discovery and Data Mining*, pp. 57–66. ACM Press (2002)
16. French, J.: A formal theory of social power. *Psychological Review* **63**(3), 181–194 (1956)
17. Gilbert, N., Troitzsch, K.: *Simulation for the Social Scientist*, 2 edn. Open University Press (2005). URL <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&#38;path=ASIN/0335216005>
18. Girard, P., Seligman, J., Liu, F.: General dynamic dynamic logic. In: T. Bolander, T. Braüner, S. Ghilardi, L.S. Moss (eds.) *Advances in Modal Logic*, pp. 239–260. College Publications (2012)
19. Golub, B., Jackson, M.: Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics* **2**(1), 112–149 (2010)
20. Jackson, M.O.: *Social and Economic Networks*. Princeton University Press (2008). URL <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0691134405>
21. Kempe, D., Kleinberg, J., Tardos, E.: Maximizing the spread of influence through a social network. In: *Proceedings the 9th Acm International Conference*, pp. 137–146. ACM Press (2003)
22. Kermack, W.O., McKendrick, A.G.: A contribution to the mathematical theory of epidemics. In: *Proceedings of the Royal Society of London, A*, vol. 115, pp. 700–721 (1927)
23. Liang, Z., Seligman, J.: A logical model of the dynamics of peer pressure. *Electronic Notes in Theoretical Computer Science* **278**, 275–288 (2011)
24. Morris, S.: Contagion. *The Review of Economic Studies* **67**(1), 57–78 (2000)
25. Mossel, E., Neeman, J., Tamuz, O.: Majority dynamics and aggregation of information in social networks. *arXiv preprint arXiv:1207.0893* (2012)
26. Nayak, A., Nelson, P., Polansky, H.: Belief change as change in epistemic entrenchment. *Synthese* **109**(2), 143–174 (1996)
27. Rott, H.: Basic entrenchment. *Studia Logica* **73**, 257–280 (2003)

28. Seligman, J., Liu, F., Girard, P.: Logic in the community. In: M. Banerjee, A. Seth (eds.) Logic and Its Applications - 4th Indian Conference, ICLA 2011, Delhi, India, January 5-11, 2011. Proceedings, *Lecture Notes in Computer Science*, vol. 6521, pp. 178–188 (2011)
29. Seligman, J., Liu, F., Girard, P.: Facebook and the epistemic logic of friendship. In: B.C. Schipper (ed.) Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge, pp. 229–238 (Chennai, India, 2013)
30. Veltman, F.: Logics for conditionals. Ph.D. thesis, University of Amsterdam (1985)

**Acknowledgement** We would like to thank Frank Zenker and Carlo Proietti for their efforts in putting together this volum. A previous version of the paper was presented at the LOGICIC kick-off workshop: Belief Change in Social Context in Amsterdam, at the Centre for Mathematical Social Science seminar series in Auckland, New Zealand, and at the Workshop on Knowledge Representation and Reasoning in Guiyang, China. We would like to thank the participants of each event, and in particular Christian List, Zoé Christoff, Chenwei Shi, Shaun White and Mark C. Wilson for valuable comments and discussions. Finally, we would like to thank the anonymous referees for their useful comments. Fenrong Liu is supported by the Project of National Social Science Foundation of China (NO.13AZX018).