

# Iterative Social Consolidations: Forming Beliefs from Many-Valued Evidence and Peers' Opinions

Yuri David Santos<sup>1</sup> and Barteld Kooi<sup>2</sup> and Rineke Verbrugge<sup>3</sup>

**Abstract.** Recently several logics modelling evidence have been proposed in the literature. These logics often also feature beliefs. We call the process or function that maps evidence to beliefs *consolidation*. In this paper, we use a four-valued modal logic of evidence as a basis. In the models for this logic, agents are represented by nodes, peerhood connections by edges, and the *private* evidence that each agent has by a four-valued valuation. From this basis, we propose methods of consolidating the beliefs of the agents, taking into account both their private evidence as well as their peers' opinions. To this end, beliefs are computed iteratively. The final consolidated beliefs are the ones in the point of stabilisation of the model. However, it turns out that some consolidation policies will not stabilise for certain models. Finding the conditions for stabilisation is one of the main problems studied here, along with other properties of such consolidations. Our main contributions are twofold: we offer a new dynamic perspective on the process of forming evidence-based beliefs, in the context of evidence logics; and we set up and address some mathematically challenging problems, which are related to graph theory and practical subject areas such as belief/opinion diffusion and contagion in multi-agent networks.

## 1 Introduction

In artificial intelligence, epistemic and doxastic logics are used as tools to model the knowledge and belief of agents [20, 36]. In practical, real-world scenarios, however, these intelligent agents often have to rely on inconsistent or incomplete data to build up their representation of the world. We can think of this data as *evidence*, a looser and more general concept than that of *justification* as featured in justification logics [2, 3, 4, 15, 21].

Recently, a series of logics have emerged with the purpose of modelling agents who possess evidence [35, 34, 33, 32, 16, 10, 22, 27, 30, 29]. Given this setting, then, we pose the following problem: how to *consolidate* this evidence into beliefs? We highlighted the relevance of this problem in [26], where we used the term *consolidation*<sup>4</sup> to refer to the process of forming beliefs from evidence – formally represented by functions from evidence to doxastic models. The complexity of certain epistemic tasks has been studied (e.g. in [12]), and, in the same vein, looking at consolidations as processes enables us to ask questions about the complexity of such operations.

As in our previous papers [26, 28], we use four-valued epistemic logic (FVEL) [25, 27] as a base. The resulting system is reminiscent of [6, 5]: agents are represented as nodes, peerhood relations

as edges, while belief is decided iteratively via modal operators  $B_0, B_1, \dots$ . In a first moment,  $B_0\varphi$  is decided for each formula  $\varphi$ , based solely on the agent's own evidence. These are the agent's initial beliefs or  $B_0$ -beliefs. Next,  $B_1$ -beliefs are decided based on the agent's evidence again plus the  $B_0$ -beliefs of its neighbors. Then,  $B_2$ -beliefs are decided similarly but also taking into account  $B_1$ -beliefs of the neighbors, and so on. The reason for this choice is to make evidence *private* to each agent. So an agent can only access its own evidence plus its neighbors' *opinions* (or beliefs), but not their evidence as in [28], which allowed for belief consolidation in a single iteration.

## 2 Logical Language

In this section we explore a variant of *four-valued epistemic logic* (FVEL) [27] proposed in [28]. The only difference here is our new definitions for belief.

### 2.1 Syntax

Let  $At$  be a countable set of atoms. Below,  $p \in At$ ; the classical part of the language is given by  $\mathcal{L}_0$ ; the propositional part is given by  $\mathcal{L}_1$ ; and the complete language is given by  $\mathcal{L}$ :

$$\begin{aligned} \mathcal{L}_0 \quad \psi &::= p \mid \sim\psi \mid (\psi \wedge \psi) \\ \mathcal{L}_1 \quad \chi &::= \psi \mid \sim\chi \mid (\chi \wedge \chi) \mid \neg\chi \\ \mathcal{L} \quad \varphi &::= \chi \mid \sim\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi \mid B_i\psi \end{aligned}$$

where  $i \in \mathbb{N}$ . We abbreviate  $\varphi \vee \psi \stackrel{\text{def}}{=} \sim(\sim\varphi \wedge \sim\psi)$  and  $\Diamond\varphi \stackrel{\text{def}}{=} \sim\Box\sim\varphi$ . We restrict belief to classical propositional formulas ( $\mathcal{L}_0$ ) because formulas with  $\neg$  refer to evidence, and we are not interested here in expressing agents holding beliefs about evidence, only about facts. Formulas such as  $p$  are read as *the agent has evidence for  $p$* , whereas  $\neg p$  is read as *the agent has evidence against  $p$* , and  $\sim\varphi$  as *it is not the case that  $\varphi$* . We read  $\Box\varphi$  as  $\varphi$  holds for all peers and  $B_i\varphi$  as *the agent believes  $\varphi$  in iteration  $i$* .<sup>5</sup>

### 2.2 Semantics

Models are tuples  $M = (S, R, V)$ , where  $S$  is a finite non-empty set of agents,  $R$  is a binary irreflexive<sup>6</sup> relation on  $S$  representing "peerhood" and  $V : At \times S \rightarrow \mathcal{P}(\{0, 1\})$  is a four-valued valuation representing agents' evidence:  $\{1\}$  is *true* ( $t$ ),  $\{0\}$  is *false* ( $f$ ),  $\{0, 1\}$

<sup>1</sup> University of Groningen, The Netherlands, email: yurids@ime.usp.br

<sup>2</sup> University of Groningen, The Netherlands, email: b.p.kooi@rug.nl

<sup>3</sup> University of Groningen, The Netherlands, email: L.C.Verbrugge@rug.nl

<sup>4</sup> Borrowed from belief revision [17, 18], where it has the meaning of transforming a potentially inconsistent belief base into a consistent one.

<sup>5</sup> Our reading of belief formulas is non-standard:  $B_i p$  is *not the agent believes she has evidence for  $p$*  (at iteration  $i$ ), but simply *the agent believes  $p$*  (at iteration  $i$ ).

<sup>6</sup> In [6] the authors work with symmetric, serial and irreflexive relations. Irreflexivity here means that the agents are not peers of themselves.

is both (b) and  $\emptyset$  is none (n).<sup>7</sup> A satisfaction relation is defined as follows:

$M, s \models p$	iff $1 \in V(p, s)$
$M, s \models \neg p$	iff $0 \in V(p, s)$
$M, s \models \sim \varphi$	iff $M, s \not\models \varphi$
$M, s \models (\varphi \wedge \psi)$	iff $M, s \models \varphi$ and $M, s \models \psi$
$M, s \models \neg(\varphi \wedge \psi)$	iff $M, s \models \neg\varphi$ or $M, s \models \neg\psi$
$M, s \models \Box \varphi$	iff for all $t \in S$ s.t. $sRt$ , it holds that $M, t \models \varphi$
$M, s \models \neg \sim \varphi$	iff $M, s \models \varphi$
$M, s \models \neg \neg \varphi$	iff $M, s \models \varphi$

An extended valuation function can be defined differently for each type of formula. If  $\varphi \in \mathcal{L}_1$ , then:  $1 \in \bar{V}(\varphi, s)$  iff  $M, s \models \varphi$ ;  $0 \in \bar{V}(\varphi, s)$  iff  $M, s \models \neg\varphi$ . Otherwise:  $1 \in \bar{V}(\varphi, s)$  iff  $M, s \models \varphi$  iff  $0 \notin \bar{V}(\varphi, s)$ . We can also define formulas discriminating which of the four truth values formula  $\varphi \in \mathcal{L}_1$  has:  $\varphi^n \stackrel{\text{def}}{=} (\sim\varphi \wedge \sim\neg\varphi)$ ;  $\varphi^f \stackrel{\text{def}}{=} \sim\sim(\sim\varphi \wedge \neg\varphi)$ ;  $\varphi^t \stackrel{\text{def}}{=} \sim\sim(\varphi \wedge \sim\neg\varphi)$ ;  $\varphi^b \stackrel{\text{def}}{=} \sim\sim(\varphi \wedge \neg\varphi)$ . If  $\varphi$  has value  $x \in \{t, f, b, n\}$ , i.e.  $\bar{V}(\varphi, s) = x$ , then  $\varphi^x$  has value  $t$ , otherwise  $\varphi^x$  has value  $f$ . If a formula  $\varphi \in \mathcal{L}_1$  has valuation  $t$  or  $f$  we say the evidence for  $\varphi$  is *unambiguous*, otherwise we say it is *ambiguous*. If  $V(p, s) = \{1\}$  we say  $s$  is a *t-agent* (w.r.t.  $p$ , but this will usually be omitted as we will mostly be thinking of a fixed atom  $p$ ), if  $V(p, s) = \{0\}$  we say  $s$  is an *f-agent*, otherwise we say  $s$  is a *b/n-agent*.<sup>8</sup>

We say that  $\Sigma \models \varphi$  ( $\Sigma$  entails  $\varphi$ ) when for all models  $M$  and states  $s$ , if  $M, s \models \sigma$  for all  $\sigma \in \Sigma$ , then  $M, s \models \varphi$ . We say that  $M \models \varphi$  if  $M, s \models \varphi$  for all states  $s$  of  $M$ . And  $\models \varphi$  ( $\varphi$  is valid) if  $M \models \varphi$  for all  $M$ ; otherwise  $\varphi$  is *invalid*. If  $\models \sim\varphi$ , we say  $\varphi$  is *contradictory*, and if  $\varphi$  is neither contradictory nor valid, it is *contingent*. If a formula is valid or contingent, it is *satisfiable*.

Notice that the semantics for belief was left open. Our goal in this paper is to discuss a number of possible definitions for the semantics of belief, taking into account that evidence is private to each agent, therefore belief can only be defined from each agent's own evidence plus their neighbors' beliefs.

### 3 Iterative Social Consolidations: Preliminaries

The following definition will be employed throughout the paper:

**Definition 1 (Attitude)** Let  $\text{Att}_i : \mathcal{L}_0 \times S \rightarrow \{1, 0, -1\}$  be a function such that:  $\text{Att}_i(\varphi, s) = 1$  iff  $M, s \models B_i\varphi$ ;  $\text{Att}_i(\varphi, s) = -1$  iff  $M, s \models B_i\sim\varphi$ ; otherwise  $\text{Att}_i(\varphi, s) = 0$ .

The function  $\text{Att}_i$  also depends on a model  $M$ , but this will be left implicit. We will write  $\text{Att}'_i$  if we are referring to a modified model  $M'$ . Moreover, as it will become clear later,  $M, s \models B_i\varphi$  and  $M, s \models B_i\sim\varphi$  are mutually exclusive.

How to define beliefs from evidence, i.e., how to *consolidate*? Before defining any consolidations, we will present the following notion, which is similar to bisimulation:

<sup>7</sup> We stick to the standard [8] in the naming of truth values. In our context, however,  $n$  is better understood as *no evidence about  $\varphi$ ,  $t$  as only evidence for  $\varphi$  (or positive evidence),  $f$  as only evidence against  $\varphi$  (or negative evidence), and  $b$  as evidence both for and against  $\varphi$ .*

<sup>8</sup> For this paper we could have used only three values ( $t$ ,  $f$  and  $b/n$ ), but since this is part of a larger project involving FVEL we chose to keep the four values.

**Definition 2 (n-Equivalence)** We say that  $(M, s) \rightleftharpoons_n (M', s')$  iff  $(M, s)$  and  $(M', s')$  satisfy exactly the same  $\mathcal{L}_1$  formulas and exactly the same formulas of the form  $\Box B_i\varphi, \Diamond B_i\varphi, \Box \sim B_i\varphi, \Diamond \sim B_i\varphi$ , where  $i \leq n$ .

Now we employ this equivalence to limit the space of possibilities. All consolidations have to conform to the following condition:

**Definition 3 (Consolidation Definability Condition (CDC))** If  $(M, s) \rightleftharpoons_n (M', s')$ , then, for all  $\varphi \in \mathcal{L}_0$ ,  $M, s \models B_{n+1}\varphi$  iff  $M', s' \models B_{n+1}\varphi$ .

What the CDC does is to make consolidations behave as functions whose input is the initial evidence and the *belief history* of peers. It will be clear later why we want to consider the history instead of just the last iteration of peers' beliefs. Even in this limited space, there are many possibilities, so in this paper we will limit ourselves to consolidations that extend the following definition:

**Definition 4** Call regular consolidations the policies respecting, for all  $i \in \mathbb{N}$ :

$M, s \models B_0p$	iff	$M, s \models p^t$
$M, s \models B_0\sim p$	iff	$M, s \models p^f$
$M, s \models B_i\sim\sim\varphi$	iff	$M, s \models B_i\varphi$
$M, s \models B_i(\varphi \wedge \psi)$	iff	$M, s \models B_i\varphi$ and $M, s \models B_i\psi$
$M, s \models B_i\sim(\varphi \wedge \psi)$	iff	$M, s \models B_i\sim\varphi$ or $M, s \models B_i\sim\psi$

Behind Def. 4 is the idea that only beliefs in literals have to be consolidated, and from those basic beliefs others can be built by simple propositional reasoning. Moreover, the first two clauses say that if the evidence for an atom  $p$  is only positive (there is only evidence for  $p$  but not against  $p$ ) or only negative, then the agent will initially believe  $p$  or  $\sim p$ , respectively.

Before the first iteration, the agents have not formed any beliefs, so each agent can only use their own private evidence. In the next iterations, however, every agent may have formed beliefs, and therefore, in order to use all information they have available, the agents can now combine their own evidence with the opinion of their peers to form more robust beliefs.<sup>9</sup> We remark that here the iterations are not intended to model the passage of time, but are only a necessary technical device used to circumvent the lack of peers' opinions in the beginning. If the goal were to realistically model time, it would make more sense to have asynchronous updates, where one agent updates in each iteration, but we will leave this variant for future work. The beliefs under  $B_0, B_1, \dots$  here do not really mean that the agent is convinced about such beliefs at any moment; these are just steps towards the agent's actual beliefs, which we will denote by the operator  $B$  (without index), which represents the beliefs of the agent in her *point of stabilisation*. So if the agent does not stabilise its beliefs with respect to a formula  $\varphi$ , we cannot say that it has actually formed any (stable) beliefs on  $\varphi$ .

**Definition 5 (Stabilisation)** An agent  $s$  in a model  $M$  is said to be stable at iteration  $i \in \mathbb{N}$  with respect to  $\varphi \in \mathcal{L}_0$  if, for all  $j \geq$

<sup>9</sup> This iterative process might remind one of Google's famous PageRank algorithm [23].

$i$ :  $\text{Att}_i(\varphi, s) = \text{Att}_j(\varphi, s)$ . A model  $M = (S, R, V)$  is said to be stable at iteration  $i \in \mathbb{N}$  w.r.t.  $\varphi \in \mathcal{L}_0$  if for all  $s \in S$ ,  $s$  is stable at iteration  $i$  w.r.t.  $\varphi$ . If a model/agent is not stable (w.r.t. a formula) it is unstable. The smallest  $i$  such that agent  $s$  is stable at iteration  $i$  w.r.t.  $\varphi$  is called the stabilisation point of agent  $s$  w.r.t.  $\varphi$ . The largest stabilisation point among all agents in  $M$  w.r.t.  $\varphi$  is called the stabilisation point of  $M$  w.r.t.  $\varphi$ . If the stabilisation point of a model/agent (w.r.t.  $\varphi$ ) is 0, it is called static (w.r.t.  $\varphi$ ).

## 4 Consolidation Policies

In this section we will study three regular consolidation policies. Good policies follow some general principles such as not wasting information, being neither too gullible nor too skeptical, etc. We will highlight the qualities and flaws of each policy as we discuss them.

### 4.1 Policy I: Monotonic Belief Diffusion

Below we define our first consolidation (a complete definition of belief):

**Definition 6 (Policy I)** Policy I is the regular consolidation with  $M, s \models B_{n+1}p$  iff:  $M, s \models p^t$  or  $(M, s \models p^b \vee p^n$  and  $M, s \models \diamond B_n p$  and  $M, s \models \square B_n \sim p$ ). And analogously for  $B_{n+1} \sim p$ .

Now, similarly to [6], we are faced with the question of what are the conditions under which this specific policy eventually stabilises. In most cases we will only talk about stability referring to some arbitrary atom  $p$ , as the dynamics are similar for all formulas.

**Lemma 7** For regular consolidations, if a model/agent is stable at iteration  $i$  w.r.t. all atoms  $p \in \text{At}$ , then this model/agent is stable at iteration  $i$  w.r.t. all formulas  $\varphi \in \mathcal{L}_0$ .

**Proof:** Follows directly from Def. 4. ■

Actually, Policy I is guaranteed to stabilise, as we prove below.

**Proposition 8** Under Policy I, for any model  $M$  and  $\varphi \in \mathcal{L}_0$ , the stabilisation point of  $M$  w.r.t.  $\varphi$  is at most  $k$ , where  $k$  is the length of the longest directed path in  $M$  without repeated edges.

**Proof:** First we prove the proposition for an arbitrary atom  $p$ , which implies the general proposition due to Lemma 7. Let  $k$  be the length of the longest directed path without repetition, and suppose  $s$  is an unstable agent at iteration  $k$ . If  $k = 0$ , it is immediately obvious that this cannot be the case, so let us assume that  $k > 0$ . If all peers of  $s$  were stable at iteration  $k - 1$ ,  $s$  would be stable by iteration  $k$  (Def. 6). So there is an agent  $s_1$  such that  $sRs_1$  and  $s_1$  is unstable at iteration  $k - 1$ . Similar reasoning applies to  $s_1$ : she has a neighbor  $s_2$  who is unstable at iteration  $k - 2$ , and so on, until we reach agent  $s_k$  who is unstable at iteration 0. But if agent  $s_k$  is unstable, there must be an agent  $s_{k+1}$  such that  $s_kRs_{k+1}$  (which could make the beliefs of  $s_k$  change in the next iterations). But, from  $s$  to  $s_{k+1}$  there is a path of length  $k + 1$ , which by our assumption (regarding  $k$ ) means that there is at least one repeated edge in this path, and therefore one repeated agent. This, in turn, implies that we have a cycle with at most  $k$  agents (otherwise the length of the longest path without repetition would exceed  $k$ ). If  $s_{k+1}$  is one of the repeated agents, then  $s_{k+1} \in \{s, s_1, \dots, s_k\}$ ; otherwise, the repeated agents are all in

$\{s, s_1, \dots, s_k\}$ . In any case, there is a cycle whose members  $s_i$  are all in  $\{s, s_1, \dots, s_k\}$ . But, since all  $s_i \in \{s, s_1, \dots, s_k\}$  are unstable, they are all  $b/n$ -agents. But, if that is the case, then it is not hard to see that  $\text{Att}_j(p, s_i) = 0$ , for all  $j \in \mathbb{N}$ . But this means that all  $s_i \in \{s, s_1, \dots, s_k\}$  are static. Contradiction. ■

Notice that for consolidations in general, due to the CDC, we cannot talk about fixpoints in the traditional sense, i.e. an iteration  $i$  where the beliefs (the output) are the same as in iteration  $i - 1$ . In Policy I, though, it is the case that if all beliefs are the same in iteration  $i$  and  $i + 1$ , then the model is stable at  $i$ .

In this policy, if the evidence is unambiguous, the agent immediately forms belief or disbelief, and never changes. Stabilisation is explained by the following:

**Proposition 9** In Policy I, the spread of belief is monotonic: let  $l \in \{p, \sim p\}$  for some  $p \in \text{At}$ ; and for all  $i \in \mathbb{N}$ , let  $A_{i,l} = \{s \in S \mid M, s \models B_{i,l}\}$ ; then for all  $i \in \mathbb{N}$ ,  $A_{i,l} \subseteq A_{i+1,l}$ .<sup>10</sup>

**Proof:** Informally: once an agent adopts belief/disbelief, it means that all her peers have also adopted such attitude (or that she had belief/disbelief from the start, due to unambiguous evidence), which in turn implies that all *their* peers have also done so, and so on... ■

Policy I is also very restrictive: the only possible change in attitude for an agent (“in time”, or relative to the progression of iterations) is from abstention to belief/disbelief.<sup>11</sup> This leads us to our next definition: a more flexible consolidation.

### 4.2 Policy II: Unstable Consolidations

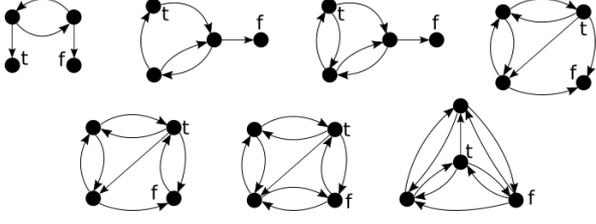
**Definition 10 (Policy II)** Policy II is the regular consolidation with  $M, s \models B_{n+1}p$  iff:  $M, s \models p^t$  or  $(M, s \models p^b \vee p^n$  and  $M, s \models \diamond B_n p$  and  $M, s \models \square \sim B_n \sim p$ ). And analogously for  $B_{n+1} \sim p$ .

What changes now is that peers that abstain are ignored (unless all of them abstain), i.e. the agents are less cautious about forming belief/disbelief when their evidence is ambiguous. Policy II is not guaranteed to stabilise. For example, the models of Fig. 1 do not stabilise – in that figure, agents where  $p$  has value  $t$  are marked with a  $t$ , and similarly for  $f$ ; for the other agents, the evidence for  $p$  is ambiguous. For the first model of Fig. 1 (top left), note that the agents with unambiguous evidence adopt belief and disbelief immediately, but  $a$  (agent on the top left of the model) and  $b$  (top right) keep changing between belief (disbelief) and abstention. First,  $B_0p$  holds for the agent marked with a  $t$ . Then, since  $b$  abstains (neither  $B_0p$  nor  $B_0 \sim p$  hold), the only non-abstaining neighbor of  $a$  believes  $p$ , therefore we get  $B_1p$  for  $a$ , and similarly  $B_1 \sim p$  for  $b$ . In the next iteration, however,  $a$  has a neighbor with  $B_1p$  (the one marked with  $t$ ) and one with  $B_1 \sim p$  (agent  $b$ ), so she abstains – similarly for  $b$ . The cycle repeats indefinitely (thence Policy II is not monotonic as in Prop. 9).

Instability is undesirable for consolidations. Even though it might be rational to be always open to changing our minds, specially upon the discovery of new evidence, our models are finite and they receive no

<sup>10</sup> A similar monotonicity holds for the Threshold Model Update in [6, Def. 2.4].

<sup>11</sup> Similar in spirit to our work, but different in many technical aspects, [19] names a change from belief to disbelief (or vice-versa) *revision* and from belief/disbelief to abstention *contraction*, adopting the classical terms from belief revision [1]. Likewise, the change from abstention to belief/disbelief could be named *expansion*.



**Figure 1.** All unstable models of size 4 under Policy II have between 4 and 10 edges. This figure shows only some of them.

new information input during the consolidation process. Therefore, rational agents are expected to decide, in a finite amount of time (or number of iterations) what are their final belief states.

Now, however, all possible attitude changes between belief, disbelief and abstention are possible.

**Proposition 11** *Stability in Policy II is decidable.*

**Proof:** For  $n$  agents, a model has  $3^n$  possible belief states (sets of attitudes of the agents) at each iteration. Since belief in one iteration depends only on the fixed evidence and on the belief state in the previous iteration, if the belief state in iteration  $3^n$  differs from that of iteration  $3^n - 1$ , the model is unstable. ■

**Proposition 12**<sup>12</sup> *Let  $R^+$  be the transitive closure of  $R$ . Under Policy II, for every  $s \in S$  that remains unstable there is a  $t$  such that  $sR^+t$  and  $t$  is in a cycle.*

**Proof:** Consider an agent  $s$  such that: (\*) there is no  $t$  such that  $sR^+t$  and  $t$  is in a cycle. Then all directed paths starting from this  $s$  are finite (forming a rooted DAG). We will prove the proposition by induction on the length  $k$  of the longest path starting at  $s$ . I.H: For each agent  $s$  such that (\*) holds and whose longest path starting from it has size  $k \leq n - 1$ ,  $s$  is stable. Base:  $k = 0$ , then obviously  $s$  is stable. Step:  $k = n$ . Consider any of the longest paths from  $s$ :  $(s, s_1, \dots, s_{n-1})$ . Then,  $(s_1, \dots, s_{n-1})$  has length  $n - 1$  and (\*) holds for  $s_1$ , which by I.H. gives us that  $s_1$  is stable. The other peers of  $s$  that belong to smaller paths are also stable, due to the I.H. Therefore, all peers of  $s$  are stable and thus so is  $s$ . ■

**Proposition 13** *A model with only one  $b/n$ -agent  $s$  is stable under Policy II.*

**Proof:** If  $s$  is the only  $b/n$ -agent, all peers of  $s$  are static, and therefore  $s$  is stable. ■

**Proposition 14** *In any model without a  $t$ -agent (or without an  $f$ -agent), the spread of belief is monotonic (under Policy II).*

**Proof:** Let  $M$  be a model without any  $f$ -agents. By Def. 10, it is impossible for any agent to have attitude  $-1$  (disbelief). Moreover, the only way that an agent  $s$  who “adopted” (changed attitude from 0 to 1) can unadopt (revert back to 0) is when:

- i. all its peers who previously had attitude 1 also unadopted;

<sup>12</sup> Here it might be possible to draw some connection to abstract argumentation frameworks [13], as in that theory odd cycles result in the inexistence of stable extensions.

- ii. one of its peers changed attitude to  $-1$ .

Since item (ii) is impossible, the only way is via item (i), but then for that to happen it is necessary that all the peers of the peers of  $s$  unadopted. This recursion cannot go on forever for our models are finite, and since we do not have an  $f$ -agent there cannot be a first unadopter. Therefore unadoption is impossible and thus belief spread is monotonic. Similar reasoning applies for the case of no  $t$ -agent. ■

**Definition 15 (Submodel)** *We say  $M' = (S', R', V')$  is a submodel of  $M = (S, R, V)$  if  $S' \subseteq S$ ,  $R' \subseteq R$ , and for all  $p \in At$  and  $s \in S'$ :  $V'(p, s) = V(p, s)$ .*

**Definition 16 (Model Restriction)** *Let  $M = (S, R, V)$ . A restriction  $M_Z$  of  $M$  to  $Z \subseteq S$  is the submodel  $M_Z = (Z, R', V')$  of  $M$  with  $R' = R \cap (Z \times Z)$ .*

**Proposition 17** *Let  $M = (S, R, V)$ ,  $s \in S$ ,  $R^*$  be the reflexive and transitive closure of  $R$  and  $R^*(s) = \{t \in S \mid sR^*t\}$ . Then, for all  $t \in R^*(s)$ , all  $\varphi \in \mathcal{L}_0$  and all  $i \in \mathbb{N}$ :  $M_{R^*(s)}, t \models B_i\varphi$  iff  $M, t \models B_i\varphi$ .*

**Proof:** Note that  $M_{R^*(s)}, t \models B_i\varphi$ , due to Def. 10, ultimately boils down to  $M_{R^*(s)}, t \models \psi$ , for some  $\psi \in \mathcal{L}$ , where  $\psi$  does not have  $B_i$  operators. By our modal semantics, it is clear that  $M_{R^*(s)}, t \models \psi$  cannot possibly be affected by any  $r \in S \setminus R^*(s)$ . ■

**Corollary 18** *Any unstable model (under Policy II) has at least one cycle  $(s_1, \dots, s_n)$ , with  $n \geq 2$  and such that for all  $s_i \in \{s_1, \dots, s_n\}$ :  $s_i$  is a  $b/n$ -agent; there are  $a, b \in S$  such that  $s_iR^+a$  and  $s_iR^+b$ ,  $a$  is a  $t$ -agent and  $b$  is an  $f$ -agent.*

**Proof:** First we modify the proof of Prop. 12 to show that: for any unstable agent  $s$  there is a  $t$  such that  $sR^+t$  and  $t$  is in a cycle consisting only of unstable  $b/n$ -agents. We prove the contrapositive by induction as before, by assuming that no such cycle exists, and therefore any path from  $s$  to any unstable  $b/n$ -agent is finite. The rest of the induction is similar. In the base case, if the agent has no unstable  $b/n$ -peer, then all its peers are stable and therefore it is stable.

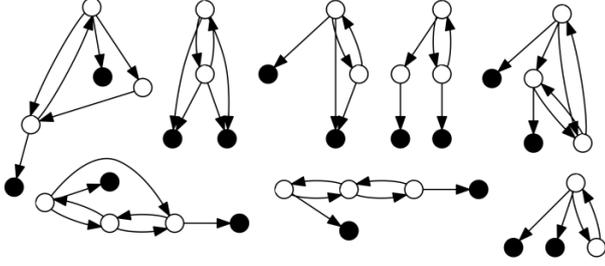
Now we just have to show that for all members  $s_i$  of this cycle, there are agents  $a$  and  $b$  as in the corollary statement. Note that for any two agents  $s_i, s_j$  in a cycle  $R^+(s_i) = R^+(s_j)$ . Now assume there is no  $t$ -agent  $a$  such that  $s_iR^+a$  is true. We know by Prop. 17 that the beliefs of  $s_i$  are the same as in  $M_{R^*(s_i)}$ , which has no  $t$ -agent. But by Prop. 14 we know that in such model the spread of belief is monotonic and therefore the model stabilises. So there has to be a  $t$ -agent  $a$  with  $s_iR^+a$ . Analogous reasoning applies for  $f$ -agent  $b$ . ■

**Corollary 19** *The first model of Fig. 1 (top left) is the smallest (in number of agents and edges) unstable model under Policy II.*

Corollary 18 gives necessary but not sufficient conditions for (see Fig. 2):<sup>13</sup>

**Open Problem 1** *What is the set of unstable models under Policy II?*

<sup>13</sup> For example, [11] solves this problem for a different logic. A more abstract study of oscillations (which we called instability) in logics is found in [31].



**Figure 2.** Some stable models satisfying the conditions of Corollary 18. In each model here, the  $b/n$ -agents are white, any one of the black agents can be taken as a  $t$ -agent and the other as an  $f$ -agent.

The set of such models is obviously infinite, but enumerable. For each  $k \in \mathbb{N}$ , we just need to generate all models of size  $k$  (which is a finite number of models), with each possible valuation (assuming here only one atom:  $At = \{p\}$ ) and combination of edges, and compute whether the model is stable or not (decidable, by Prop. 11). This algorithm works as a finite description of the set of unstable models (under Policy II). It would be more interesting, however, to have a more “structural” description, such as the one of Corollary 18.

Unfortunately, we have not managed to find such a simple structural characterisation of unstable models (and actually we do not know if such a characterisation is even possible), but the following is our attempt at finding “simplifications” that could hopefully yield models that capture the “essence” of instability.

**Definition 20 (Reduction)** Let  $\mathbb{M}$  be the class of all models. A relation  $T \subseteq \mathbb{M} \times \mathbb{M}$  is called a semi-reduction if for all models  $M = (S, R, V)$ ,  $M' = (S', R', V')$ ,  $MTM'$  iff:  $M'$  is stable iff  $M$  is stable; and  $S' \subseteq S$ . Moreover, if  $M'$  is a submodel of  $M$ , then  $T$  is called a reduction.

**Definition 21 (Faithful Reduction)** A semi-reduction  $T$  is called faithful if for all models  $M = (S, R, V)$ ,  $M' = (S', R', V')$ ,  $MTM'$  iff: for all  $i \in \mathbb{N}$ , all  $\varphi \in \mathcal{L}_0$  and all  $s \in S'$ ,  $M', s \models B_i \varphi$  iff  $M, s \models B_i \varphi$ .

Note that if for all  $M, M'$ ,  $MTM'$  only if  $M'$  is a restriction of  $M$ , then  $T$  is a faithful reduction. Below, let arbitrary models  $M = (S, R, V)$  and  $M' = (S', R', V')$ .

**Definition 22** Below we define  $T_1, T_2, T_3, T_4, T_5, T_6 \subseteq \mathbb{M} \times \mathbb{M}$  such that:

- $MT_1M'$  iff:  $M'$  is the submodel of  $M$  such that  $R \setminus R' = \{(s, t)\}$ , where  $s, t \in S$  and  $s$  is a  $t$ -agent or an  $f$ -agent, and  $S' = S$ .
- $MT_2M'$  iff: there is an  $s \in S$  such that there is no  $t \in S$  with  $sRt$  or  $tRs$ , and  $M'$  is the restriction of  $M$  to  $S \setminus \{s\}$ .
- $MT_3M'$  iff: there is a  $b/n$ -agent  $s$ , a  $t$ -agent  $a$  and an  $f$ -agent  $b$  in  $S$  such that  $sRa$  and  $sRb$ , and  $M'$  is a restriction of  $M$  to  $S \setminus \{s\}$ .
- $MT_4M'$  iff: there is a  $b/n$ -agent  $s \in S$  for which there is no  $t$ -agent  $a \in S$  with  $sR^+a$ , and no  $f$ -agent  $b \in S$  with  $sR^+b$  and  $M'$  is a restriction of  $M$  to  $S \setminus \{s\}$ .
- $MT_5M'$  iff: there are at least two distinct  $t$ -agents (or  $f$ -agents)  $a, b \in S$ ;  $S' = S$ ,  $V' = V$  and  $R' = R \cap ((S \setminus \{b\}) \times (S \setminus \{b\})) \cup Q$ , with  $Q = \{(a, s) \mid (b, s) \in S\} \cup \{(s, a) \mid (s, b) \in S\}$ .

- $MT_6M'$  iff: there is a  $b/n$ -agent  $s$  for which there is no cycle  $(s_1, \dots, s_n)$  consisting only of  $b/n$ -agents in  $M$  such that for an  $s_i$  in  $(s_1, \dots, s_n)$ , a  $t$ -agent  $a \in S$  and an  $f$ -agent  $b \in S$ , it holds that  $s_iR^+a$ ,  $s_iR^+b$  and  $s_iR^+s$ ; and  $M'$  is a restriction of  $M$  to  $S \setminus \{s\}$ .

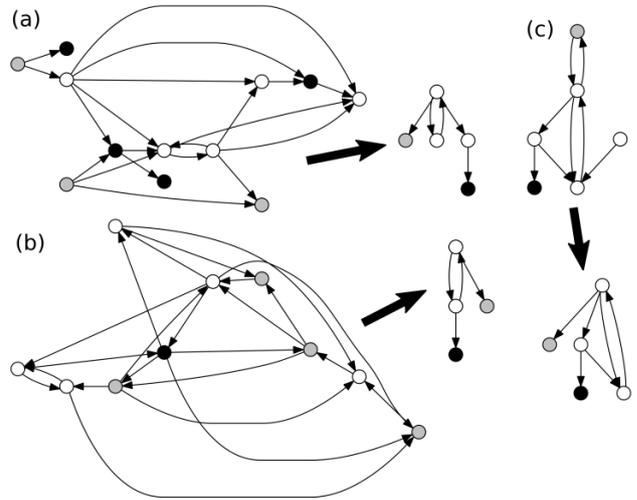
Note that  $T_5$  is the only of the relations  $T_i$  above in which  $MT_iM'$  does not require  $M'$  to be a submodel of  $M$ , which means that one has to apply it wisely if one wants to actually simplify a model (basically, one  $t$ -agent and one  $f$ -agent have to be chosen to concentrate all incoming arrows). It is called a semi-reduction because it does not necessarily yield simpler models.

**Proposition 23** The relations  $T_1, T_2, T_3, T_4$  of Def. 22 are faithful reductions,  $T_5$  is a faithful semi-reduction and  $T_6$  is a (non-faithful) reduction.

**Proof:** This proof is straightforward. In some cases, one just has to use Prop. 17 and have in mind that peers with attitude 0 do not affect any agent’s beliefs. ■

Now one can apply arbitrary sequences of the reductions above to obtain, from an arbitrary model, less cluttered counterparts that are stable if and only if the original was (see Fig. 3). One can also use only faithful reductions to obtain a simplification where all agents have exactly the same belief history.

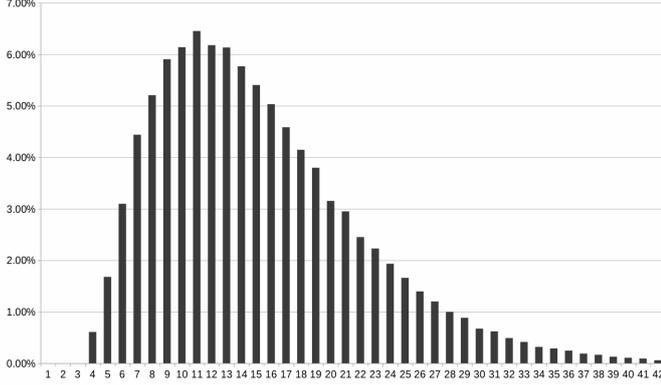
Using reductions might be a more efficient way of checking whether a model is stable, and it certainly is an easier method for humans in many cases. A more formal comparison of the complexity between checking stabilisation in the standard way versus using reductions is left for future work.



**Figure 3.** Here  $b/n$ -agents are white,  $t$ -agents are gray and  $f$ -agents are black. Many stable models reduce to a single agent model (after applying  $T_1 - T_6$  as much as possible), but there are cases like (a) above where it is not reduced completely. Likewise, many unstable models reduce to the smallest unstable model, like case (b), but some do not, as in case (c). These reduced models highlight essential features behind a model’s stability or instability.

Just as a side note, we randomly generated (using the Erdős-Rényi method [14]) and tested 100,000 models of size 4 to 60 and found

the percentages of unstable ones, as shown in Fig. 4. It is clear that we can expect a 0/1 law here (e.g. as in [37]), with the percentage going to zero in the limit when the size tends to infinity. This mitigates the problem of instability for Policy II: in large enough models (such as a big network of scientists, for example), “almost never” will the agents incur an irrational consolidation infinite loop. A possible informal explanation is that, as the size increases, the number of possible structures that can be built grows much faster than the number of possibilities for “instability-inducing” structures, which are very specific: they have to respect the conditions of Corollary 18 plus other unknown conditions (Open Problem 1). Another interesting question is: why does the percentage of unstable models peak at size 11?



**Figure 4.** Percentage of unstable models per model size, for Policy II. The curve continues as expected until size 60, but we have cropped the chart due to space.

### 4.3 Policy III: Ignoring Unstable Peers

Our next consolidation will try to tackle the instability problem by temporarily ignoring agents who have not been stable for the last  $\lambda$  iterations. Formally, we define the following abbreviation:

$$M, s \models \text{stable}_{\lambda,p}^n \text{ (with } n \geq \lambda \geq 1 \text{ and } p \in At)$$

with the meaning:  $Att_{n-1}(p, s) = Att_{n-2}(p, s) = \dots = Att_{n-\lambda}(p, s)$ . (Agent  $s$  has been stable about  $p$  in the last  $\lambda$  iterations preceding iteration  $n$ ), and set that if  $\lambda \leq 1$ , then  $M, s \models \text{stable}_{\lambda,p}^n$ ; and if  $\lambda > n$ , then  $M, s \models \text{stable}_{\lambda,p}^n$  is defined as  $M, s \models \text{stable}_{n,p}^n$ . This abbreviation does not increase expressivity, because it can always be defined by a finite propositional combination of conditions. For example,  $M, s \models \text{stable}_{2,p}^{10}$  is defined as the following disjunctive condition:  $(M, s \models B_9p \text{ and } M, s \models B_8p)$  or  $(M, s \models B_9\sim p \text{ and } M, s \models B_8\sim p)$  or  $(M, s \models \sim B_9p \wedge \sim B_9\sim p \text{ and } M, s \models \sim B_8p \wedge \sim B_8\sim p)$ .

From the above, we conclude that  $M, s \models \Box \text{stable}_{\lambda,p}^n$  means that for all  $t \in S$  such that  $sRt$ ,  $M, t \models \text{stable}_{\lambda,p}^n$ , and  $M, s \models \Diamond \text{stable}_{\lambda,p}^n$  means that there is a  $t \in S$  such that  $sRt$  and  $M, t \models \text{stable}_{\lambda,p}^n$ . To restrict the modal operators only to stable peers, we can define  $M, s \models \Box_{\lambda,p}^n \varphi$  as  $M, s \models \Box(\sim \text{stable}_{\lambda,p}^n \vee \varphi)$ , and  $M, s \models \Diamond_{\lambda,p}^n \varphi$  as  $M, s \models \Diamond(\text{stable}_{\lambda,p}^n \wedge \varphi)$ . Now we are ready for:

**Definition 24 (Policy III- $\lambda$ )** Let  $1 \leq \lambda \in \mathbb{N}$ . Policy III- $\lambda$  is the regular consolidation with  $M, s \models B_{n+1}p$  iff:  $M, s \models p^t$  or  $(M, s \models p^b \vee p^n$  and  $M, s \models \Diamond_{\lambda,p}^{n+1} B_n p$  and  $M, s \models \Box_{\lambda,p}^{n+1} \sim B_n \sim p)$ . And analogously for  $B_{n+1} \sim p$ .

It is not hard to see that Def. 24 is compliant with the CDC (Def. 3), which is also the reason why we defined the CDC based on the history of peers’ beliefs and not only on the last iteration. Note also that if the parameter  $\lambda = 1$ , Policy III- $\lambda$  coincides with Policy II, so the former is a generalisation of the latter. Fig. 5 show the evolution of belief using Policy III- $\lambda$  on the model of Fig. 1, with different values of  $\lambda$ .

$\lambda = 1$		$\lambda = 2$		$\lambda = 3$		$\lambda = 4$	
a	b	a	b	a	b	a	b
0	0	0	0	0	0	0	0
1	-1	1	-1	1	-1	1	-1
0	0	1	-1	0	0	0	0
1	-1	0	0	1	-1	1	-1
0	0	1	-1	1	-1	1	-1
	$\vdots$	1	-1	1	-1	1	-1
		0	0	0	0	1	-1
			$\vdots$	1	-1	0	0
				1	-1	1	-1
				1	-1	1	-1
				0	0	1	-1
				$\vdots$		1	-1
						0	0
						$\vdots$	

**Figure 5.** Iterations of belief for agents  $a$  and  $b$  in the first model of Fig. 1 (top left). Abstention is represented by 0, belief by 1 and disbelief by  $-1$  (as in Def. 1).

The question that immediately surfaces is whether larger values of  $\lambda$  “improve” stability. Looking at Fig. 5, we notice that larger values of  $\lambda$  make iterations of abstention less frequent. We have to formally define what “improving stability” means here.

**Definition 25 (Stability Measure)** A stability measure  $<_{M,\varphi}$  for consolidations (w.r.t. a fixed model  $M$  and an arbitrary  $\varphi \in \mathcal{L}_0$ ), where  $C <_{M,\varphi} C'$  is read as  $C'$  is more stable than  $C$  for  $\varphi$  in  $M$ , has to respect the following principles:

- i. If  $C_1$  makes  $M$  static but  $C_2$  does not, then  $C_2 <_{M,\varphi} C_1$ ;
- ii. if  $C_1$  makes  $M$  stable but  $C_2$  does not, then  $C_2 <_{M,\varphi} C_1$ ;
- iii. if  $M$  stabilises with  $C_1$  at iteration  $i$  and with  $C_2$  at  $j > i$ , then  $C_2 <_{M,\varphi} C_1$ ;

If for all models  $M$  and all  $\varphi \in \mathcal{L}_0$  it is the case that  $C <_{M,\varphi} C'$ , then we say that  $C < C'$ , i.e.  $C'$  is more stable than  $C$ .

For an arbitrary measure of stability, our initial hypothesis is:

**Hypothesis 1** Let  $1 \leq \lambda, \kappa \in \mathbb{N}$ . If  $\lambda < \kappa$ , then Policy III- $\lambda <$  Policy III- $\kappa$ .

We can start by asking whether Policy III with  $\lambda = 1$  can, in some case, be more stable than with  $\lambda = 2$ . Def. 25-i cannot be used to violate Hypothesis 1, when changing  $\lambda$  from 1 to 2. If a model is

static, changing  $\lambda$  will not produce any changes in belief. Perhaps surprisingly, though, our Hypothesis 1 can be violated by Def. 25-ii, that is, there is a model that stabilises when  $\lambda = 1$  but does not when  $\lambda = 2$ , namely the model of Fig. 6 (right). For clarity, the value of  $p$  is not shown when it is ambiguous.



**Figure 6.** Left: Unstable with  $\lambda = 1$ , stable with  $\lambda = 2$ . Right: stable with  $\lambda = 1$ , unstable with  $\lambda = 2$ .

$\lambda = 1$			$\lambda = 2$			$\lambda = 1$			$\lambda = 2$			
a	b	c	a	b	c	a'	b'	c'	a'	b'	c'	
0	0	0	0	0	0	0	0	0	0	0	0	
-1	1	0	-1	1	0	-1	1	0	-1	1	0	
0	1	-1	-1	1	0	0	0	-1	-1	1	0	
0	0	0	0	1	-1	-1	0	0	0	0	-1	
$\vdots$			0	1	0	-1	0	-1	-1	1	0	
			<i>stabilised</i>			<i>stabilised</i>				-1	1	0
									0	0	-1	
									$\vdots$			

**Figure 7.** Iterations of belief for the models of Fig. 6.

The models of Fig. 6 (left) and Fig. 6 (right) are the smallest models (considering number of agents and arrows) that feature, respectively: (a) a change from unstable with  $\lambda = 1$  to stable with  $\lambda = 2$  and (b) the opposite. We tested computationally and verified that phenomena (a) and (b) do not happen in models with 4 or less agents. Another surprising result is that, among models of size 5, there are exactly the same number of models where (a) and (b) occur. Moreover, we tested with  $\lambda = 1, \dots, 10$ , running for at most 1000 iterations, and all models fitting (a) were stable with  $\lambda = 2, \dots, 10$ , and all models fitting (b) were unstable with  $\lambda = 2, 3$  and stable otherwise. We suspect that phenomena (a) and (b) occur due to the qualitative difference between Policy III- $\lambda$  with  $\lambda = 1$ , which equals Policy II, and with  $\lambda \geq 2$ . Moreover, we can conjecture so far that increasing  $\lambda$  does have a positive effect in terms of stability in general (although Hypothesis 1 is false), as the (b)-type models became stable with larger values of  $\lambda$ , despite becoming unstable with  $\lambda = 2, 3$ .

#### 4.4 Other Policies

Other policies that have not yet been explored include the following ideas:

- i. Stopping the consolidation at a fixed iteration defined by a parameter  $\lambda$ . This policy would guarantee stability in a forceful manner. A drawback is that it would be *too sensitive* to the parameter  $\lambda$ , specially in the case of unstable models (under Policy III);
- ii. Limiting the number of times an agent can change its attitude, e.g. after going from abstention to belief/disbelief, it cannot go back to abstention. We probably will not have problems to define this consolidation respecting the CDC, for even though the agents *do not* take into account their own belief histories directly, these are definable from their previous peers' beliefs;

- iii. Defining belief based on the number of peers holding a certain attitude. In [28], we show that by introducing a dynamic operator (which increases expressivity), we can count peers with certain attitudes. This would probably be a more realistic way of consolidating beliefs, but demands a language richer than the modal logic used here.
- iv. Allowing  $t$  and  $f$ -agents to change attitudes. For this, item (iii) above might be helpful. Or,  $b/n$ -agents could have a policy similar to Policy II, whereas  $t$  and  $f$ -agents could be more resistant to change, adopting a strategy in line with Policy I.

## 5 Conclusions and Future Work

In this paper we used a many-valued modal logic (FVEL) to represent a multi-agent network of peers and their evidence, and defined belief based on this network and on the evidence. This paper is an alternative view to our previous paper [28], where the agents can access one another's evidence, therefore allowing for consolidations done in one step. By making the evidence private, we triggered an iterative process, through which the agents always (in Policy I) or almost always (Policy II) approximate a final belief. The exception to this is in the problematic cases of unstable models, which were one of the main topics explored here.

From a practical or realistic point of view, the consolidations presented here might not look very rational. But, within the abstractions and limitations of our modelling (i.e. respecting the CDC), these approaches might be among the most rational possibilities. However, a more in-depth, formal and argumentative defense of why they are rational in this context is needed, and left for future work (but [26, footnote 4] underlies some of our design choices here). One philosophical point that has to be better defended is the non-temporal aspect of these iterations. We are not trying to represent agents who are updating their beliefs as the days and years go by. This time can be seen as just a "processing time". It can also be viewed as passage of time in a very restricted situation where agents cannot do anything besides communicating with their peers – they do not have the opportunity to consult other sources of information or to make deep reflections – as if they are "deliberating" in an isolated room. This deliberation, of course, would be of a very simple kind, in which all they are allowed to do is to ask the opinions of their peers, at discrete moments of time, or "turns".

In the problem of informational cascades [7, 9], rational individual behaviour might lead to bad doxastic outcomes. In that case, this happens due to the evidence being private to each agent, who can only access the others' final judgements, which is also a feature of our models. The instability, but possibly other irrational behaviours in our system, is partially explained by that feature. A comparison with the system in [28], where others' evidence is public, would help to elucidate the impact of private evidence.

There is still much work to be done on these consolidations, especially Policy III, which has not yet been explored in great depth. Other policies such as the ones in Sec. 4.4 can also be studied. Moreover, as remarked earlier, here a three-valued logic (such as the ones in [24, Ch. 7]) of evidence would suffice, but consolidations that distinguish  $n$  and  $b$  could be developed in future work.

## REFERENCES

- [1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson, 'On the logic of theory change: Partial meet contraction and revision functions', *Journal of Symbolic Logic*, **50**(2), 510–530, (1985).

- [2] Sergei Artemov, 'Logic of proofs', *Annals of Pure and Applied Logic*, **67**(1-3), 29–59, (1994).
- [3] Sergei Artemov, 'Operational modal logic', Technical Report MSI 95-29, Cornell University, (December 1995).
- [4] Sergei Artemov, 'Explicit provability and constructive semantics', *Bulletin of Symbolic Logic*, **7**(1), 1–36, (2001).
- [5] Alexandru Baltag, Zoé Christoff, Rasmus Rendsvig, and Sonja Smets, 'Dynamic epistemic logics of diffusion and prediction in social networks', in *Proceedings of the 12th Conference on Logic and the Foundations of Game and Decision Theory*, (2016).
- [6] Alexandru Baltag, Zoé Christoff, Rasmus Rendsvig, and Sonja Smets, 'Dynamic epistemic logics of diffusion and prediction in social networks', *Studia Logica*, **107**(3), 489–531, (2019).
- [7] Abhijit V. Banerjee, 'A simple model of herd behavior', *The Quarterly Journal of Economics*, **107**(3), 797–817, (1992).
- [8] Nuel Belnap, 'A useful four-valued logic', in *Modern Uses of Multiple-valued Logic*, 5–37, Springer, (1977).
- [9] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch, 'A theory of fads, fashion, custom, and cultural change as informational cascades', *Journal of Political Economy*, **100**(5), 992–1026, (1992).
- [10] Walter Carnielli and Abilio Rodrigues, 'An epistemic approach to paraconsistency: a logic of evidence and truth', *Synthese*, 1–25, (2017).
- [11] Zoé Christoff and Davide Grossi, 'Stability in binary opinion diffusion', in *International Workshop on Logic, Rationality and Interaction (LORI)*, pp. 166–180. Springer, (2017).
- [12] Cédric Dégremond, Lena Kurzen, and Jakub Szymanik, 'Exploring the tractability border in epistemic tasks', *Synthese*, **191**(3), 371–408, (2014).
- [13] Phan Minh Dung, 'On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games', *Artificial intelligence*, **77**(2), 321–357, (1995).
- [14] Paul Erdős and Alfréd Rényi, 'On random graphs I', *Publicationes Mathematicae*, (1959).
- [15] Melvin Fitting, 'The logic of proofs, semantically', *Annals of Pure and Applied Logic*, **132**(1), 1–25, (2005).
- [16] Melvin Fitting, 'Paraconsistent logic, evidence and justification', *Studia Logica*, **105**(6), 1149–1166, (2017).
- [17] Sven Ove Hansson, *Belief base dynamics*, Ph.D. dissertation, Uppsala University, 1991.
- [18] Sven Ove Hansson, 'Semi-revision', *Journal of Applied Non-Classical Logics*, **7**(1-2), 151–175, (1997).
- [19] Fenrong Liu, Jeremy Seligman, and Patrick Girard, 'Logical dynamics of belief change in the community', *Synthese*, **191**(11), 2403–2431, (2014).
- [20] John-Jules Meyer and Wiebe van der Hoek, *Epistemic Logic for AI and Computer Science*, Cambridge University Press, 1995.
- [21] Alexey Mkrtychev, 'Models for the logic of proofs', in *International Symposium on Logical Foundations of Computer Science*, pp. 266–275. Springer, (1997).
- [22] Aybüke Özgün, *Evidence in epistemic logic: a topological perspective*, Ph.D. dissertation, University of Amsterdam, 2017.
- [23] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd, 'The pagerank citation ranking: Bringing order to the web', Technical Report 1999-66, Stanford InfoLab, (November 1999). Previous number = SIDL-WP-1999-0120.
- [24] Graham Priest, *An Introduction to Non-classical Logic: From If to Is*, Cambridge University Press, 2 edn., 2008.
- [25] Yuri David Santos, 'A dynamic informational-epistemic logic', in *Dynamic Logic. New Trends and Applications*, eds., A. Madeira and M. Benevides, volume 10669 of *Lecture Notes in Computer Science*, pp. 64–81. Springer, (2018).
- [26] Yuri David Santos, 'Consolidation of belief in two logics of evidence', in *International Conference on Logic, Rationality and Interaction (LORI)*, eds., Patrick Blackburn, Emiliano Lorini, and Meiyun Guo, volume 11813 of *Lecture Notes in Computer Science*, pp. 57–70. Springer, (2019).
- [27] Yuri David Santos, 'A four-valued dynamic epistemic logic', *Journal of Logic, Language and Information*, (2020).
- [28] Yuri David Santos, 'Social consolidations: Rational belief in a many-valued logic of evidence and peerhood', in *11th International Symposium on Foundations of Information and Knowledge Systems (FoIKS 2020)*, eds., Andreas Herzig and Juha Kontinen, pp. 58–78. Springer, (2020).
- [29] Chenwei Shi, Sonja Smets, and Fernando R. Velázquez-Quesada, 'Beliefs based on evidence and argumentation', in *International Workshop on Logic, Language, Information, and Computation (WoLLIC)*, pp. 289–306. Springer, (2018).
- [30] Chenwei Shi, Sonja Smets, and Fernando R. Velázquez-Quesada, 'Beliefs supported by binary arguments', *Journal of Applied Non-Classical Logics*, **28**(2-3), 165–188, (2018).
- [31] Johan van Benthem, 'Oscillations, logic, and dynamical systems', in *The Facts Matter. Essays on Logic and Cognition in Honour of Rineke Verbrugge*, eds., S. Ghosh and J. Szymanik, chapter 1, 9–22, College Publications, (2015).
- [32] Johan van Benthem, David Fernandez-Duque, and Eric Pacuit, 'Evidence and plausibility in neighborhood structures', *Annals of Pure and Applied Logic*, **165**(1), 106–133, (2014).
- [33] Johan van Benthem, David Fernández-Duque, Eric Pacuit, et al., 'Evidence logic: A new look at neighborhood structures.', in *Advances in Modal Logic*, eds., Thomas Bolander, Torben Braüner, Silvio Ghilardi, and Lawrence Moss, volume 9, pp. 97–118. College Publications, (2012).
- [34] Johan van Benthem and Eric Pacuit, 'Dynamic logics of evidence-based belief', Technical Report PP-2011-19, University of Amsterdam, ILLC, (2011).
- [35] Johan van Benthem and Eric Pacuit, 'Dynamic logics of evidence-based beliefs', *Studia Logica*, **99**(1), 61–92, (2011).
- [36] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi, *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*, Springer Science & Business Media, 2007.
- [37] Rineke Verbrugge, 'Zero-one laws with respect to models of provability logic and two Grzegorzczak logics', in *Advances in Modal Logic (AiML): Accepted Short Papers*, eds., Guram Bezhanishvili, Giovanna D'Agostino, George Metcalfe, and Thomas Studer, pp. 115–120, (2018).