

BeliefFlow: A Framework for Logic-Based Belief Diffusion via Iterated Belief Change

Nicolas Schwind¹, Katsumi Inoue^{2,3}, Sébastien Konieczny⁴, Pierre Marquis^{4,5}

¹ National Institute of Advanced Industrial Science and Technology, Tokyo, Japan

² National Institute of Informatics, Tokyo, Japan

³ Graduate Institute for Advanced Studies, SOKENDAI, Tokyo, Japan

⁴ Univ. Artois, CNRS, CRIL, Lens, France

⁵ Institut Universitaire de France

nicolas-schwind@aist.go.jp, inoue@nii.ac.jp, konieczny@cril.fr, marquis@cril.fr

Abstract

This paper presents BeliefFlow, a novel framework for representing how logical beliefs spread among interacting agents within a network. In a Belief Flow Network (BFN), agents communicate asynchronously. The agents' beliefs are represented using epistemic states, which encompass their current beliefs and conditional beliefs guiding future changes. When communication occurs between two connected agents, the receiving agent changes its epistemic state using an improvement operator, a well-known type of rational iterated belief change operator that generalizes belief revision operators. We show that BFNs satisfy appealing properties, leading to two significant outcomes. First, in any BFN with strong network connectivity, the beliefs of all agents converge towards a global consensus. Second, within any BFN, we show that it is possible to compute an optimal strategy for influencing the global beliefs. This strategy, which involves controlling the beliefs of a least number of agents through bribery, can be identified from the topology of the network and can be computed in polynomial time.

Introduction

Consider the following scenario with three individuals: Alice, Bob and Charles. Each initially holds distinct beliefs. Alice got a memo about a lockdown happening tomorrow. Bob lacks any information regarding the impending lockdown, yet he assumes that all aspects of daily life, especially transportation services, will continue without interruption. Charles, in contrast to Alice's viewpoint, does not believe a lockdown will take place. However, he does hold a conditional belief: if a lockdown were indeed implemented, he expects that transportation services would stop. Alice, Bob, and Charles frequently engage in private one-on-one conversations, and they are about to dive into a discussion about the upcoming lockdown and transportation situation.

Our focus lies in predicting the outcome of their forthcoming discussion. More precisely, we are interested in the following issue: will the group of agents collectively attain a stable state of beliefs, thus achieving consensus? This is a challenging question, given that multiple variable parameters come into play. Each of the three individuals, when presented with new information from their two friends, may

have varying levels of skepticism, differing immediate or delayed responses to conflicting information. Moreover, while it is possible for friends to communicate with each other, we do not make the assumption that the order in which each communication pair occurs is known.

In this paper, we introduce a new framework to model these interactions, named Belief Flow Networks (BFNs). In a BFN, the communication protocol within a set of agents is characterized by an acquaintance graph and a stochastic process that is used to trigger communication pairs randomly, in a sequence of steps. At each step, an agent sends their currently held beliefs (a propositional formula) to a connected agent, who receives it and modifies their belief state in light of this new information using an improvement operator (Konieczny and Pino Pérez 2008; Konieczny, Medina Grespan, and Pino Pérez 2010), a general form of iterated belief change operator (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988; Darwiche and Pearl 1997).

BFNs share similarities with many settings for modelling the belief dynamics of a group of agents, including Boolean networks (Kauffman 1969, 1993; Aldana 2003), opinion dynamics (Hegselmann and Krause 2005; Riegler and Douven 2009; Tsang and Larson 2014; Grandi, Lorini, and Perrussel 2015; Novaro et al. 2018, 2019), and many other complex systems. But the most closely related work to our proposal is the framework of *Belief Revision Games (BRGs)* (Schwind et al. 2015, 2016), a comparable logical framework involving a network of interacting agents. However, our approach deviates from BRGs in several significant aspects. The first noteworthy distinction between BFNs and BRGs is that BRGs adopt a *synchronous* update scheme. At each step, every agent receives the beliefs of all their acquaintances, processes these beliefs using a belief merging operator, and then revises their own beliefs based on the merged result. In contrast, our BFN approach employs an *asynchronous* update scheme, and there is no overseeing agent dictating the order with which communication pairs are activated. Another significant difference pertains to the revision policy employed by each agent in BRGs, that is characterized by a classical revision operator (Katsuno and Mendelzon 1991) and a merging operator (Konieczny and Pino Pérez 2002), both of which align with the principles of the AGM theory that delineates the behavior of *one-step* re-

vision. But the AGM approaches suffers from limitations in handling sequences of revisions, potentially leading to undesirable outcomes, which has prompted the exploration of iteration principles in belief change, addressing those limitations (Nayak 1994; Nayak et al. 1994, 1996; Darwiche and Pearl 1997; Jin and Thielscher 2007; Rott 2009). One of the most general classes of iterated belief change operators is called *improvement operators* (Konieczny and Pino Pérez 2008; Konieczny, Medina Grespan, and Pino Pérez 2010), that generalize iterated revision operators. They allow for the incorporation of varying degrees of change reluctance for an agent. So, in contrast to BRGs, in a BFN each agent uses an improvement operator to modify their beliefs. This requires each agent to be assigned with an *epistemic state* instead of a mere propositional formula (Darwiche and Pearl 1997). This state encodes both the agent’s beliefs at the current step and conditional information guiding subsequent change steps. BFNs are thus flexible enough to capture a more realistic spectrum of belief change scenarios compared to BRGs.

Although BRGs have been shown to satisfy a number of appealing properties, they have also shown vulnerability to “belief cycles” (Schwind et al. 2015): even when agents agree on certain options, they may continuously change their beliefs *ad infinitum*. BRGs have also been demonstrated to exhibit paradoxical results in terms of manipulability (Schwind et al. 2016). The behavior of BFNs is quite different. First, under strong network connectivity, agents within BFNs converge towards a unanimous global consensus. Second, we establish the feasibility of a straightforward bribery technique for influencing collective beliefs. Indeed, identifying a minimal set of agents to bribe to make a given piece of beliefs unanimously accepted across the network can be achieved in polynomial time.

The proofs of propositions are available online.¹

Preliminaries on Iterated Belief Change

Let $\mathcal{L}_{\mathcal{P}}$ be a propositional language built up from a finite set of propositional variables \mathcal{P} and the usual connectives. \perp (resp. \top) is the Boolean constant always false (resp. true). A world is a mapping from \mathcal{P} to $\{0, 1\}$. The set of all worlds is denoted by Ω . A world ω is a model of a formula φ , denoted by $\omega \models \varphi$, if it makes φ true. $[\varphi]$ denotes the set of models of φ . \models also denotes logical entailment and \equiv logical equivalence between formulae, i.e., $\varphi \models \psi$ iff $[\varphi] \subseteq [\psi]$ and $\varphi \equiv \psi$ iff $[\varphi] = [\psi]$. Given a preorder² \preceq over worlds, we define $\min([\varphi], \preceq) = \{\omega \models \varphi \mid \nexists \omega' \models \varphi : \omega' \prec \omega\}$.

We assume that each agent processes incoming pieces of information using an iterated belief change operator called *improvement operator* (Konieczny and Pino Pérez 2008; Konieczny, Medina Grespan, and Pino Pérez 2010; Medina Grespan and Pino Pérez 2013). Considering this setting requires the belief state of each agent to be characterized by an epistemic state, a more general and complex object than a simple propositional formula (Darwiche and Pearl 1997;

Schwind, Konieczny, and Pino Pérez 2022). Indeed, an epistemic state Ψ allows one to represent the current beliefs of an agent, denoted by $Bel(\Psi)$ and some conditional information guiding the change process for future changes. Given Ψ , the beliefs of the agent can be extracted through a mapping Bel , so that $Bel(\Psi)$ is a propositional formula from $\mathcal{L}_{\mathcal{P}}$. Formally, let \mathcal{E} be the set of all epistemic states, which is considered fixed. Then Bel is a mapping from \mathcal{E} to $\mathcal{L}_{\mathcal{P}}$. An (iterated) change operator \circ associates an epistemic state and a change formula with an epistemic state, i.e., it is a mapping $\circ : \mathcal{E} \times \mathcal{L}_{\mathcal{P}} \rightarrow \mathcal{E}$. An improvement operator is an iterated change operator satisfying a set of nine rationality principles (I1)–(I9) (see (Konieczny and Pino Pérez 2008; Konieczny, Medina Grespan, and Pino Pérez 2010) for a detailed justification of these postulates). In the following, given a change operator \circ , $\Psi \circ^k \varphi$ is inductively defined as $\Psi \circ^1 \varphi = \Psi \circ \varphi$ and for each $k > 1$, $\Psi \circ^k \varphi = (\Psi \circ^{k-1} \varphi) \circ \varphi$. Then $\Psi \star \varphi$ is defined as $\Psi \circ^k \varphi$, where k is the least integer such that $Bel(\Psi \circ^k \varphi) \models \varphi$ (such an integer is assumed to exist from postulate (I1) below):

- (I1) $\exists k \in \mathbb{N}_* \text{ s.t. } Bel(\Psi \circ^k \varphi) \models \varphi$
- (I2) If $Bel(\Psi) \wedge \varphi \not\models \perp$, then $Bel(\Psi \star \varphi) \equiv Bel(\Psi) \wedge \varphi$
- (I3) If $\varphi \not\models \perp$, then $Bel(\Psi \circ \varphi) \not\models \perp$
- (I4) if $\varphi_i \equiv \beta_i$ for all $i \in \{1, \dots, m\}$, then $Bel(\Psi \circ \varphi_1 \circ \dots \circ \varphi_m) \equiv Bel(\Psi \circ \beta_1 \circ \dots \circ \beta_m)$
- (I5) $Bel(\Psi \star \varphi) \wedge \beta \models Bel(\Psi \star (\varphi \wedge \beta))$
- (I6) If $Bel(\Psi \star \varphi) \wedge \beta \not\models \perp$, then $Bel(\Psi \star (\varphi \wedge \beta)) \models Bel(\Psi \star \varphi) \wedge \beta$
- (I7) If $\varphi \models \beta$, then $Bel((\Psi \circ \beta) \star \varphi) \equiv Bel(\Psi \star \varphi)$
- (I8) If $\varphi \models \neg \beta$, then $Bel((\Psi \circ \beta) \star \varphi) \equiv Bel(\Psi \star \varphi)$
- (I9) If $Bel(\Psi \star \varphi) \not\models \neg \beta$, then $Bel((\Psi \circ \beta) \star \varphi) \models \beta$

Among the nine set of postulates, (I1-I6) are the most basic ones, and (I7-I9) are principles ruling iteration.

Improvement operators can be characterized by associating each epistemic state with a total preorder over worlds. Formally, a function $\Psi \mapsto \preceq_{\Psi}$ that maps each epistemic state Ψ to a total preorder over worlds \preceq_{Ψ} is called a *gradual assignment* iff:


1. If $\omega, \omega' \models Bel(\Psi)$, then $\omega \simeq_{\Psi} \omega'$
2. If $\omega \models Bel(\Psi)$ and $\omega' \not\models Bel(\Psi)$, then $\omega \prec_{\Psi} \omega'$
3. For any positive integer n , if $\varphi_i \equiv \beta_i$ for any $i \leq n$, then $\preceq_{\Psi \circ \varphi_1 \circ \dots \circ \varphi_n} = \preceq_{\Psi \circ \beta_1 \circ \dots \circ \beta_n}$
4. If $\omega, \omega' \models \varphi$, then $\omega \preceq_{\Psi} \omega' \Leftrightarrow \omega \preceq_{\Psi \circ \varphi} \omega'$
5. If $\omega, \omega' \models \neg \varphi$, then $\omega \preceq_{\Psi} \omega' \Leftrightarrow \omega \preceq_{\Psi \circ \varphi} \omega'$
6. If $\omega \models \varphi$ and $\omega' \models \neg \varphi$, then $\omega \preceq_{\Psi} \omega' \Rightarrow \omega \prec_{\Psi \circ \varphi} \omega'$

Proposition 1 ((Konieczny, Medina Grespan, and Pino Pérez 2010)). *An operator \circ is an improvement operator iff there exists a gradual assignment that maps each epistemic state Ψ to a total preorder over worlds \preceq_{Ψ} such that for every formula φ , $[Bel(\Psi \star \varphi)] = \min([\varphi], \preceq_{\Psi})$.*

Notably, improvement operators satisfying the original success postulate (R*1) ($Bel(\Psi \circ \varphi) \models \varphi$) (Darwiche and Pearl 1997) are AGM/DP iterated *revision* operators (Alchourrón, Gärdenfors, and Makinson 1985; Darwiche and Pearl 1997).

¹<https://nicolas-schwind.github.io/SIKM-AAAI24-proofs.pdf>

²For each preorder \preceq , \simeq denotes the corresponding indifference relation, and \prec the strict part of \preceq .


 Figure 1: Restrained revision \circ_R and one-improvement \circ_I .

Let us now give two examples of improvement operators, the one-improvement operator (Konieczny, Medina Grespán, and Pino Pérez 2010) and the restrained revision operator (Booth and Meyer 2006), and illustrate their behavior through an example. For both of these operators, the set of epistemic states \mathcal{E} corresponds to the set of all total preorders over worlds.

The *restrained revision operator* \circ_R is an improvement operator that is also a revision operator, i.e., it satisfies the success postulate (R*1). That is, \circ_R is such that for each epistemic state Ψ and each formula φ , $[Bel(\Psi \circ_R \varphi)] = \min([\varphi], \preceq_\Psi)$. It also requires for all worlds $\omega, \omega' \in \Omega$ that if $\omega \models \varphi$, $\omega \not\models Bel(\Psi \circ_R \varphi)$ and $\omega' \not\models \varphi$, then $\omega' \prec_\Psi \omega \Rightarrow \omega' \prec_{\Psi \circ_R \varphi} \omega$. Thus, this operator ensures that the minimal models of φ for \preceq_Ψ become the minimal models of $\Psi \circ_R \varphi$, and that the models of $\neg\varphi$ that were strictly more plausible than the models of φ before the revision step remain so after the revision.

The *one-improvement operator* \circ_I is the improvement operator satisfying the following property:³ if $\omega \models \varphi$ and $\omega' \not\models \varphi$, then $(\omega' \prec_\Psi \omega \Rightarrow \omega' \preceq_{\Psi \circ_I \varphi} \omega)$ and $(\omega' \ll_\Psi \omega \Rightarrow \omega' \preceq_{\Psi \circ_I \varphi} \omega')$. Roughly speaking, this operator uniformly “shifts” all models of φ one level lower in the preorder \preceq_Ψ .

Figure 1 depicts a total preorder Ψ over worlds on two propositional variables ($\mathcal{P} = \{p, q\}$),⁴ and the total preorders $\Psi \circ_R \neg q$ and $\Psi \circ_I \neg q$. We have that $Bel(\Psi) \equiv p \wedge q$, $Bel(\Psi \circ_R \neg q) \equiv p \wedge \neg q$, and $Bel(\Psi \circ_I \neg q) \equiv p$.

There are many other improvement operators (see, e.g., (Rott 2009)) some of which being not representable as transitions between total preorders, i.e., they require more complex structures, such as OCFs (Ordinal Conditional Functions) (Spohn 1988), for being formally defined (Schwind, Konieczny, and Pino Pérez 2022).

Belief Flow Networks

We are now ready to define the setup of Belief Flow Networks (BFNs). Formally, a BFN is a tuple $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$ consisting of four components.

First, G is an acquaintance graph $G = (V, A)$, where $V = \{1, \dots, n\}$ represents the set of agents and $A \subseteq V \times V$ is an acquaintance relationship between them. We require A to be irreflexive, i.e., if $(i, j) \in A$ then $i \neq j$. Roughly speaking, a pair (i, j) belongs to A if and only if communication from i to j is “possible.”

³ $\omega \ll \omega'$ is a shortcut for $(\omega \prec \omega' \text{ and } \nexists \omega'' \text{ s.t. } \omega \prec \omega'' \prec \omega')$

⁴A world ω is at the same or at a lower level than a world ω' iff $\omega \preceq_\Psi \omega'$. So minimal (i.e., most plausible) worlds are at the lowest level.

Second, $\vec{\Psi}$ is an initial epistemic state profile that captures the epistemic state of each agent at the beginning of the game. It is represented as an n -vector $\vec{\Psi} = \langle \Psi^1, \dots, \Psi^n \rangle$, where each Ψ^i corresponds to the initial epistemic state of agent i . We assume that all agents start with consistent beliefs, i.e., $Bel(\Psi^i)$ is consistent for each agent $i \in V$.

Third, $\vec{\sigma}$ is a change policy profile which defines the change policy of each agent. It is represented as an n -vector $\vec{\sigma} = \langle \sigma_1, \dots, \sigma_n \rangle$, where each element σ_i is an improvement operator. For the purpose of our results in this paper, we require each operator σ_i to satisfy a slight strengthening of (I1):

$$(I1^*) \exists k \in \mathbb{N}_* \text{ s.t. } \forall \Psi \in \mathcal{E}, \forall \alpha \in \mathcal{L}, Bel(\Psi \circ^k \alpha) \models \alpha$$

(I1*) requires that the number of change steps required for a formula to be eventually entailed in the beliefs of an agent is fixed for the agent, i.e., it does not depend on the epistemic state. This additional restriction is quite light: it is satisfied, e.g., by every revision operator and every improvement operator defined on a finite epistemic space \mathcal{E} , which includes the operators \circ_R and \circ_I presented in the previous section, among many others (Rott 2009).

Fourth, \mathcal{S} is a stochastic process on A which governs the communication protocol. It is a series of random variables $\mathcal{S} = (A_s)_{s \in \mathbb{N}}$, where each random variable A_s has the domain set A . We require this series to be a *chain with complete connections* (Iosifescu and Grigorescu 1990), i.e., there exists a real number $\delta > 0$ such that for each $s \in \mathbb{N}$ and for all edges e_0, \dots, e_s , $Pr(A_s = e_s \mid A_{s-1} = e_{s-1}, \dots, A_0 = e_0) \geq \delta$. Roughly speaking, \mathcal{S} generates some scenarios, or “runs” $\sigma = (\sigma_s)_{s \in \mathbb{N}}$. At each time step $s \geq 0$, σ_s is the specific edge to be triggered, and its value (i.e., the value for A_s) is governed by \mathcal{S} . Each agent i is associated with its epistemic state at step s in a run σ , denoted by $\Psi_{\sigma_s}^i$. At each step s , some pair $(i, j) \in A$ is selected at random with a certain positive probability which may depend on the previously selected edges, i.e., on the values of A_0, \dots, A_{s-1} . Upon triggering an edge (i, j) at each step s , agent j receives the current beliefs $Bel(\Psi_{\sigma_s}^i)$ of agent i and j modifies their current epistemic state $\Psi_{\sigma_s}^j$ accordingly using their change policy σ_j , i.e., $\Psi_{\sigma_{s+1}}^j = \Psi_{\sigma_s}^j \circ_j Bel(\Psi_{\sigma_s}^i)$. Note that the definition of \mathcal{S} as a chain with complete connections is general enough to include standard Markov chains (Iosifescu and Grigorescu 1990) when the probability of each edge to be selected at each step depends on the one selected at the previous step, and Bernoulli schemes (Shields 1973) when all random variables are independent.

Let us now introduce the various notions useful to determine how BFNs are interpreted. In what follows, we assume given a BFN $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$.

A-sequences and B-runs. An A -sequence is a (possibly infinite) sequence $\sigma = (\sigma_s)_{s \in \{0, \dots\}}$ where $\sigma_s \in A$ for each $s \in \{0, \dots\}$. A B -run σ is an infinite A -sequence generated by \mathcal{S} . Note the distinction between the notions of infinite A -sequence and B -run. Indeed, not every infinite A -sequence can be yielded by a B -run. It can be easily seen for instance that for any given $A' \subsetneq A$, an infinite A' -sequence is an infinite A -sequence but it cannot be generated by \mathcal{S} , since this

would otherwise imply $Pr(A_s = (i, j)) = 0$ for each $k \in \mathbb{N}$ and for some $(i, j) \in A \setminus A'$, which is forbidden by \mathcal{S} being a chain with complete connections. We will make precise this distinction in the next section through a characterization of the subclass of infinite A -sequences corresponding to \mathcal{B} -runs (cf. Proposition 3).

Epistemic profile sequences. BFNs \mathcal{B} are interpreted in terms of *epistemic profile sequences* in some specific scenarios characterized by \mathcal{B} -runs. The epistemic profile sequence of a BFN $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$ given a \mathcal{B} -run σ , denoted by $\vec{\Psi}_\sigma$, is an infinite series of epistemic profiles $\vec{\Psi}_\sigma = \vec{\Psi}_{\sigma_0}, \vec{\Psi}_{\sigma_1}, \dots$ where for each $s \in \mathbb{N}$, $\vec{\Psi}_{\sigma_s}$ is inductively defined as follows:

- $\vec{\Psi}_{\sigma_0} = \langle \Psi_{\sigma_0}^1, \dots, \Psi_{\sigma_0}^n \rangle$, where for each $i \in V$, $\Psi_{\sigma_0}^i = \Psi^i$
- $\vec{\Psi}_{\sigma_{s+1}} = \langle \Psi_{\sigma_{s+1}}^1, \dots, \Psi_{\sigma_{s+1}}^n \rangle$, where for each $i \in V$,

$$\Psi_{\sigma_{s+1}}^i = \begin{cases} \Psi_{\sigma_s}^i \circ_i Bel(\Psi_{\sigma_s}^j), & \text{if } \sigma_s = (j, i) \\ \Psi_{\sigma_s}^i, & \text{otherwise} \end{cases}$$

So, given a \mathcal{B} -run σ , an agent $i \in V$, and a step $s \in \mathbb{N}$, $\Psi_{\sigma_s}^i$ and $Bel(\Psi_{\sigma_s}^i)$ represent respectively the epistemic state and the belief state of agent i at step s in the specific communication scenario characterized by σ .

Belief sequences, outcomes, and stability. Given an epistemic profile sequence $\vec{\Psi}_\sigma$ as defined above, an agent i and a \mathcal{B} -run σ , the sequence $Seq_\sigma(i) = (Bel(\Psi_{\sigma_s}^i))_{s \in \mathbb{N}}$ is called the σ -belief sequence of i . We are specially interested in the notion of σ -outcome sequence of i in σ , denoted by $Seq_\sigma^*(i)$ and defined as the sequence $Seq_\sigma^*(i) = \arg \min_s \{ (Bel(\Psi_{\sigma_s}^i), Bel(\Psi_{\sigma_{s+1}}^i), \dots) \mid \forall s_1 \geq s, \exists E \subseteq \mathbb{N}, |E| = \aleph_0, \forall s_2 \in E, Bel(\Psi_{\sigma_{s_2}}^i) \equiv Bel(\Psi_{\sigma_{s_1}}^i) \}$ ⁵. In informal terms, the σ -outcome sequence of i is the earliest sequence among all subsequences of $Seq_\sigma(i)$ where every formula continues to appear infinitely often up to logical equivalence. The notion of σ -outcome sequence bears similarities to the notion of *belief cycle* introduced for Belief Revision Games (BRGs) (Schwind et al. 2015). Outcome sequences shed light on what an agent eventually comes to accept or reject, and whether these beliefs remain stable. So, a key notion based on σ -outcome sequences is the one of *acceptance*, also adapted from (Schwind et al. 2015). A formula φ is *accepted* by i in σ if for each $\alpha \in Seq_\sigma^*(i)$, $\alpha \models \varphi$; it is *unanimously accepted* in σ if it is accepted by each $i \in V$.

Since we work on a propositional language generated from a finite set of variables, each σ -outcome sequence can be characterized finitely: the σ -outcome of i is a formula $Out_\sigma(i)$ such that $[Out_\sigma(i)] = \{ \omega \in \Omega \mid \exists \alpha \in Seq_\sigma^*(i), \omega \models \alpha \}$. That is, $Out_\sigma(i)$ is a formula that characterizes what is accepted by i in σ : it is the logically weakest formula that entails each formula accepted by i in σ .

When every formula from $Seq_\sigma^*(i)$ is equivalent to $Out_\sigma(i)$, we say that i is *stable* in σ . Stated otherwise, i is stable in σ whenever the beliefs of the agent i remain eventually unchanged (up to logical equivalence) in σ . The notion of stability can be lifted to (sets of) agent(s) in a straightforward way: a set of agents $V' \subseteq V$ is said to be stable in σ

if every $i \in V'$ is stable in σ . An agent $i \in V$ (resp. a set of agents $V' \subseteq V$) is said to be *stable in \mathcal{B}* if i (resp. V') is stable in every \mathcal{B} -run σ . In the specific case when $V' = V$, we simply say that \mathcal{B} is stable.

Lastly, a set $V' \subseteq V$ is said to *reach a consensus* in a \mathcal{B} -run σ if V' is stable in σ and there is a formula α such that for each $i \in V'$, $Out_\sigma(i) \equiv \alpha$. And V' is said to be *strongly consensual* in \mathcal{B} if V' reaches a consensus in every \mathcal{B} -run.

Let us formalize and develop the example sketched in the introduction to illustrate all notions introduced thus far.

Example 1. Let $\mathcal{P} = \{l, t\}$, where l stands for “lockdown will take place”, and t stands for “transportation will not be interrupted”. We consider a BFN $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$ as follows. Let $G = (V, A)$, with $V = \{1, 2, 3\}$ where 1 (resp. 2, 3) corresponds to Alice (resp. Bob, Charles), and $A = \{(1, 2), (2, 3), (1, 3), (3, 1)\}$, i.e., Alice is not influenced by (or does not receive messages from) Bob, and Bob is not influenced by Charles. Let \mathcal{E} be the set of all total preorders over worlds, $\vec{\Psi} = \langle \Psi^1, \Psi^2, \Psi^3 \rangle$ and $\vec{\sigma} = \langle \sigma_1, \sigma_2, \sigma_3 \rangle$. Bob uses a restrained revision operator $\circ_3 = \circ_R$, whereas Alice and Charles are a bit more change-reluctant and use a one-improvement operator $\circ_1 = \circ_2 = \circ_I$. Let Alice’s initial epistemic state Ψ^1 be the two-level preorder over worlds $lt \simeq_{\Psi^1} \bar{l}\bar{t} \prec_{\Psi^1} \bar{l}t \simeq_{\Psi^1} \bar{l}\bar{l}$. As to Bob and Charles, let $\Psi^2 = lt \simeq_{\Psi^2} \bar{l}\bar{t} \prec_{\Psi^2} \bar{l}t \simeq_{\Psi^2} \bar{l}\bar{l}$, and $\Psi^3 = \bar{l}\bar{t} \simeq_{\Psi^3} \bar{l}t \prec_{\Psi^3} \bar{l}\bar{l} \prec_{\Psi^3} lt$. We have for instance $Bel(\Psi^3) \equiv \neg l$, which captures Charles’ initial belief that a lockdown won’t occur tomorrow; and $\bar{l}\bar{t} \prec_{\Psi^3} lt$ represents his conditional belief that if he were to believe in a lockdown occurring, the world where transportation is disrupted is more plausible than the one where transportation remains unaffected. Consider the \mathcal{B} -run σ starting with the A -sequence $((1, 2), (1, 3), (3, 1), (1, 2), (2, 3))$, which is interpreted as the scenario where Alice sends a message to Bob first ($\sigma_0 = (1, 2)$), then to Charles ($\sigma_1 = (1, 3)$), then Charles replies to Alice ($\sigma_2 = (3, 1)$), and so on. In this scenario, at the first step Alice sends her initial beliefs that she hears about the upcoming lockdown to Bob, who revises his epistemic state accordingly: $\Psi_{\sigma_1}^2 = \Psi_{\sigma_0}^2 \circ_2 Bel(\Psi_{\sigma_0}^1) = \Psi^2 \circ_R l$, which corresponds to the preorder $\Psi_{\sigma_1}^2 = lt \prec_{\Psi^2} \bar{l}\bar{t} \prec_{\Psi^2} \bar{l}t \prec_{\Psi^2} \bar{l}\bar{l}$ and $Bel(\Psi_{\sigma_1}^2) \equiv l \wedge t$. After this first step, the epistemic states of both Alice and Charles remain unchanged $\Psi_{\sigma_1}^1 = \Psi_{\sigma_0}^1 = \Psi^1$ and $\Psi_{\sigma_1}^3 = \Psi_{\sigma_0}^3 = \Psi^3$. Then, by building the epistemic profile sequence iteratively, it is not difficult to verify in this scenario that at step 5, the three agents reach a consensus: $Bel(\Psi_{\sigma_5}^1) \equiv Bel(\Psi_{\sigma_5}^2) \equiv Bel(\Psi_{\sigma_5}^3) \equiv l \wedge \neg t$. In particular, each agent is stable in σ , and the formula $l \wedge \neg t$ and any weaker formula is unanimously accepted in σ .

Basic Properties

Let us first show that a set of basic properties, introduced and discussed in (Schwind et al. 2015) and satisfied by Belief Revision Games (BRGs), are also satisfied by BFNs. A BFN $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$ satisfies **(CP)** (resp. **(AP)**, **(UP)**) if for every \mathcal{B} -run σ , every step $s \in \mathbb{N}$ and every formula φ :

(CP) $\forall i \in V Bel(\Psi_{\sigma_s}^i) \not\models \perp$

(AP) $(\forall i \in V \varphi \models Bel(\Psi^i)) \Rightarrow (\forall i \in V \varphi \models Bel(\Psi_{\sigma_s}^i))$

⁵ $|E| = \aleph_0$ simply means that E is countably infinite.

(UP) $(\forall i \in V \text{ Bel}(\Psi^i) \models \varphi) \Rightarrow (\forall i \in V \text{ Bel}(\Psi_{\sigma_s}^i) \models \varphi)$

(CP) (Consistency Preservation) requires that agents (who initially have consistent beliefs as required by BFNs) never become self-conflicting in any \mathcal{B} -run. **(AP)** (Agreement Preservation) asks that if all agents initially agree on some alternatives, they will not change their mind about them. And **(UP)** (Unanimity Preservation) states that every piece of beliefs which is initially entailed by every agent's beliefs should remain so in their belief sequence.

We can show that:

Proposition 2. *Every BFN satisfies (CP), (AP) and (UP).*

Another property called *monotonicity* was considered in (Schwind et al. 2015) and discussed further in (Schwind et al. 2016). This property was used as a basis for defining further notions related to control issues and bribery, but it was shown not to be satisfied by BRGs in the general case. Our BFNs do not satisfy this property either, but as opposite to BRGs, this property is not necessary to derive interesting results on bribery, which will be shown in the next sections.

Responsiveness and Stability

Two other properties were considered in (Schwind et al. 2015): responsiveness and stability. Before investigating them in the context of BFNs, let us introduce a key property of \mathcal{B} -runs. An infinite A -sequence $\sigma = (\sigma_s)_{s \in \mathbb{N}}$ is said to be *disjunctive* if σ contains all possible finite A -sequences, i.e., for each $m \in \mathbb{N}$ and every finite A -sequence $\sigma' = (\sigma'_0, \dots, \sigma'_m)$, there exists $p \in \mathbb{N}$ such that for each $q \in \{0, \dots, m\}$, $\sigma_{p+q} = \sigma'_q$. We can show that:

Proposition 3. *For each BFN \mathcal{B} and each \mathcal{B} -run σ , σ is a disjunctive A -sequence.*

This observation is pivotal in our following results.

Responsiveness. In (Schwind et al. 2015) a property of responsiveness was introduced and shown to be satisfied by BRGs. The property states that at any given step, every agent should modify their beliefs at the next step whenever their beliefs are inconsistent with the beliefs of each one of their acquaintances, and the conjunction of the beliefs of the set of her acquaintances is consistent. It is easy to see that this property is not satisfied by BFNs: BFNs evolve asynchronously, i.e., only one agent modifies their epistemic state at each step, and since the revision policy of each agent is an improvement operator and not necessarily a *revision* one, their modified belief $\text{Bel}(\Psi_{\sigma_{s+1}}^i)$ may be equivalent to their previous beliefs $\text{Bel}(\Psi_{\sigma_s}^i)$ even if those beliefs were inconsistent with the newly received ones. That being said, the core idea underlying the notion of responsiveness is that if an agent i is influenced by another agent j (i.e., there exists a link $(j, i) \in A$) and that j does not change their beliefs, then there must exist a future step when i agrees with j . We call this property **(DR)** (for Delayed Responsiveness):

(DR) $\forall (j, i) \in A, \forall s \in \mathbb{N}, \exists s' \geq s$ such that $\text{Bel}(\Psi_{\sigma_{s'}}^i) \wedge \text{Bel}(\Psi_{\sigma_{s'}}^j) \not\models \perp$

We intend to show that **(DR)** is satisfied by every BFN. Recall that since for each agent $i \in V$, each change policy

\circ_i satisfies (Π^*) . This means that in any \mathcal{B} -run σ and for any pair of agents $(j, i) \in A$, there exists a finite A -sequence σ^{ji} defined as $\sigma^{ji} = (e_s)_{s \in \{1, \dots, k_i\}}$, where k_i is the least integer such that for each epistemic state Ψ and each formula φ , $\text{Bel}(\Psi \circ_i^{k_i} \varphi) \models \varphi$ and for each $e_s \in \sigma^{ji}$, $e_s = (j, i)$. We call the sequence σ^{ji} the *control sequence from j to i* . An interesting consequence of the fact that every \mathcal{B} -run σ is disjunctive (cf. Proposition 3) is that for all agents $i, j \in V$ such that $(j, i) \in A$, the control sequence from j to i appears in σ infinitely many times. This means that:

Lemma 1. *Let σ be any \mathcal{B} -run, $(j, i) \in A$, $s \in \mathbb{N}$ and $s_* \geq s$ be any step such that $(\sigma_{s_*}, \dots, \sigma_{s_*+k_i}) = \sigma^{ji}$. Then $\text{Bel}(\Psi_{\sigma_{s_*+k_i}}^i) \models \text{Bel}(\Psi_{\sigma_{s_*+k_i}}^j)$.*

And as a direct consequence of Lemma 1, we get that:

Corollary 1. *Every BFN satisfies (DR).*

Stability. In the context of BRGs, (Schwind et al. 2015) gave sufficient, reasonable conditions on the revision policies used by the agents under which stability is satisfied, in the case when G is a directed acyclic graph. We exhibit a simple example showing that this is not the case for BFNs:

Proposition 4. *Let $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$ be any BFN such that $V = \{1, 2, 3\}$, $A = \{(1, 3), (2, 3)\}$ and $\text{Bel}(\Psi^1) \wedge \text{Bel}(\Psi^2) \models \perp$. Then for each \mathcal{B} -run σ , the agent 3 is not stable in σ .*

We intend now to characterize stable BFNs based on the topology of their acquaintance graph. Beforehand, let us introduce some basic notions on graphs. A $G = (V, A)$ is said to be *strongly connected* if for all distinct $x, y \in V$, there is a path from x to y in G . A *strongly connected component (SCC)* of G is a maximal subgraph $G' = (V', A')$ of G that is strongly connected. In the following, we simply identify an SCC $G' = (V', A')$ with the set of vertices of the corresponding subgraph, i.e., with V' . The set of all SCCs V_{Cond} of a graph G forms a partition V_1, \dots, V_m of V . When each SCC V_i of G is contracted into a single vertex, the resulting graph $\text{Cond}(G) = (V_{\text{Cond}}, A_{\text{Cond}})$ is a directed acyclic graph (DAG) and is called the *condensation* of G . For each $V_i \in V_{\text{Cond}}$, we define $\text{Parents}(V_i) = \{V_j \in V_{\text{Cond}} \mid (V_j, V_i) \in A_{\text{Cond}}\}$. An SCC $V_i \in V_{\text{Cond}}$ is called a *source SCC* if $\text{Parents}(V_i) = \emptyset$.

Let $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$ be any BFN, which will be referred to in the rest of this section. Recall that a set of agents $V' \subseteq V$ is *strongly consensual* in \mathcal{B} when for every \mathcal{B} -run σ , V' is stable in σ and there is a formula α such that for each agent $i \in V'$, $\text{Out}_\sigma(i) \equiv \alpha$. So, a set V' is strongly consensual only if V' is stable. We intend to show that the converse is also true for SCCs. Our result is based on the following strengthening of Lemma 1 on control sequences (see Lemma 2 below). Let $V' \subseteq V$, and $p = (i_1, \dots, i_m)$ be any V' -path, i.e., a path in G such that for each $i_t \in p$, i_t is an agent from V' . Define σ^p as the concatenation of all control sequences $\sigma^{i_l i_{l+1}}$ for each $l \in \{1, \dots, m-1\}$, i.e., $\sigma^p = \sigma^{i_1 i_2} \dots \sigma^{i_{m-1} i_m}$. The sequence σ^p is called a *control path* from i_1 to i_m . Note that the same agent may appear several times in the path p , and in particular one may have that $i_1 = i_m$. Based on a similar argument as the one used to prove Lemma 1, we get that:

Lemma 2. Let $V' \subseteq V$, i be an agent from V' , $p = (i_1, \dots, i_m)$ be any V' -path such that $i_m = i$, σ be any \mathcal{B} -run, $s \in \mathbb{N}$ and $s_* \geq s$ be any step such that $(\sigma_{s_*}, \dots, \sigma_{s_*^p}) = \sigma^p$, where $s_*^p = s_* + |\sigma^p| - 1$. Then for each agent $j \in V'$, $Bel(\Psi_{\sigma_{s_*^p}}^i) \models Bel(\Psi_{\sigma_{s_*^p}}^j)$.

Proposition 5. If V' is an SCC of G , then V' is stable if and only if V' is strongly consensual.

Another interesting consequence of Lemma 2 is that:

Proposition 6. If V' is a source SCC of G , then V' is stable.

And as a direct consequence of Propositions 5 and 6:

Corollary 2. If G is strongly connected, then \mathcal{B} is stable and strongly consensual.

It is interesting to notice that this result strengthens the one obtained for BRGs (Schwind et al. 2016), which required very specific conditions to conclude stability and strong consensus, i.e., for G be a complete graph and for every agent i to use a revision policy parameterized by a specific (belief merging) operator, i.e., the drastic distance-based IC merging operator (Konieczny and Pino Pérez 2002).

We now intend to characterize the class of stable BFNs, based on the topology of the acquaintance graph.

Proposition 7. If V' is an SCC of G , $Parents(V') = \{V''\}$ and V'' is stable, then V' is stable.

We are now ready to give a characterization of the class of stable BFNs, based on the topology of the acquaintance graph only. Given a graph $G = (V, A)$, we denote by $\mathcal{B}(G)$ the class of BFNs $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$, i.e., the class of BFNs having G as acquaintance graph. We say that a class of BFNs is stable whenever all BFNs from the class are stable. Now, Proposition 4 tells us that with an SCC V' with $Parents(V') \geq 2$, it is possible to build a BFN instance such that V' is not stable. Taking this observation together with Propositions 6 and 7, we get that:

Corollary 3. The class $\mathcal{B}(G)$ is stable if and only if $Cond(G)$ is a tree.

In other words, when one has no information about the beliefs of the agents in a given graph, one can guarantee that a given BFN is stable precisely when the condensation of its acquaintance graph has a tree-like structure.

Bribery

In this section, we are interested in whether the agents in a BFN can be influenced, i.e., whether a given formula can be made unanimously accepted by modifying a BFN, and whether this can be done through a “bribery” of a specific subset of influential agents.

Let us call an *influence scheme* IS a pair $IS = (V', \alpha)$, where V' is a set of agents and α is a propositional formula. Given a BFN $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$, an influence scheme $in \mathcal{B}$ is simply an influence scheme $IS = (V', \alpha)$ such that $V' \subseteq V$. An influence scheme IS is interpreted as a set of agents that one intends to influence (bribe) by sending them

α . This implies the modification of a given BFN, which is made precise through the notion of *bribed BFN*.

Given a BFN $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$ and an influence scheme $IS = (V', \alpha)$ in \mathcal{B} , a *bribed BFN* \mathcal{B} by IS is a BFN $\mathcal{B}_{IS} = \langle G_{IS}, \vec{\Psi}_{IS}, \vec{\sigma}_{IS}, \mathcal{S} \rangle$, where:

- $G_{IS} = (V_{IS}, A_{IS})$, with $V_{IS} = V \cup \{n+1\}$ and $A_{IS} = A \cup \{(n+1, i) \mid i \in V'\}$
- $\vec{\Psi}_{IS} = \langle \Psi^1, \dots, \Psi^n, \Psi^{n+1} \rangle$, with $Bel(\Psi^{n+1}) \equiv \alpha$
- $\vec{\sigma}_{IS} = \langle \sigma_1, \dots, \sigma_n, \sigma_{n+1} \rangle$

So, a bribed BFN \mathcal{B}_{IS} is interpreted as an extension of \mathcal{B} by “applying” an influence scheme $IS = (V', \alpha)$ on it, which consists in adding a new agent $n+1$ to V with beliefs equivalent to α and adding acquaintance relations from the newly added agent $n+1$ to all agents in V' .

Now, we are interested in the extent to which an influence scheme can make certain beliefs φ unanimously accepted in a BFN. Given a propositional formula φ , an influence scheme IS on \mathcal{B} is said to be *successful for φ in \mathcal{B}* if φ is unanimously accepted in every BFN \mathcal{B}_{IS} . We extend the notion to classes of BFNs as follows: recall that $\mathcal{B}(G)$ denotes the class of BFNs $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$, i.e., the class of BFNs having G as acquaintance graph. Then we say that IS is successful for φ in $\mathcal{B}(G)$ if IS is successful for φ in every BFN $\mathcal{B} \in \mathcal{B}(G)$.

We have the following characterization result of the successful influence schemes in classes $\mathcal{B}(G)$.

Proposition 8. Let $IS = (V', \alpha)$ be an influence scheme and φ be a propositional formula such that $[\varphi] \subsetneq \Omega$. Then IS is successful for φ in $\mathcal{B}(G)$ if and only if $(\alpha \models \varphi)$ and for each source SCC V_{source} of G , $V_{source} \cap V' \neq \emptyset$.

This result means that, to make a certain formula φ unanimously accepted in a BFN, it is enough to bribe one agent per source SCC by sending them φ or a formula that is logically stronger. This has an interesting consequence on the computational cost of applying an “optimal” influence scheme on a BFN. Let us assume that all agents in a graph have the same bribery “cost”, which is reflected by the fact that the cost of an influence scheme $IS = (V', \alpha)$ is simply defined as $Cost(IS) = |V'|$. Then a direct consequence of Proposition 8 and of the fact the set of all source SCCs in any graph can be computed in polynomial time (Tarjan 1972) is that:

Corollary 4. An influence IS of minimal cost and that is successful for any formula φ in $\mathcal{B}(G)$ can be computed in time polynomial in $|G|$.

Empirical Study

This section aims to experimentally emphasize the tendency of BFNs to converge to a stable state (Proposition 3) and unanimous acceptance of some pieces of beliefs in bribed BFNs (Proposition 3) within a reasonable number of steps.

We considered various BFN instances $\mathcal{B} = \langle G, \vec{\Psi}, \vec{\sigma}, \mathcal{S} \rangle$, altering the structure of the acquaintance graph while maintaining the other parameters constant: $n = |V| = 20$ agents, $|\mathcal{P}| = 4$ propositional variables (i.e., $|\Omega| = 16$ worlds), \mathcal{E} was the set of all total preorders over worlds, and each




Figure 2: Number of steps necessary to reach a stable state for each BFN instance class.

agent’s belief change policy was set to $\circ_i = \circ_I$, the one-improvement operator. Our focus was on a stochastic process \mathcal{S} activating communication pairs $(i, j) \in A$ following a Bernoulli scheme with a uniform probability distribution on $A \times A$. In other words, at each step $s \in \mathbb{N}$, $Pr(A_s = (i, j)) = 1/|A|$. The epistemic state Ψ^i of each agent was generated as follows: we defined a set of worlds W , where each world $\omega \in \Omega$ had a probability p of belonging to W equal to 0.2. Each Ψ^i was then defined as a two-level total preorder with $[Bel(\Psi^i)] = W$.

We considered the following acquaintance graph types:

- R7, R10, R20, C: four classes of random graphs, where each potential edge from $V \times V$ appears in A independently with a specific probability p . The values of p were chosen from the set 7%, 10%, 20%, 100%. When $p = 100$, it results in complete graphs, which are denoted by “C”.
- BA: Barabasi-Albert preferential attachment graphs (Albert and Barabasi 2002), generated with an initial sample vertex size of 3 and the addition of two vertices at each iteration step in the generation procedure.
- K: Kleinberg small-world graphs (Kleinberg 2000), characterized by 4 rows and 5 columns, with a clustering exponent of 2.
- P: graphs consisting of a single path $(1, \dots, n)$.
- L: graphs consisting of a single loop $(1, \dots, n, 1)$.
- P+ (resp. L+): the symmetric closure of P (resp. L).

For all instances except P and L, the graphs were symmetrically closed (i.e., an edge (j, i) was added to the acquaintance graph whenever (i, j) belonged to it). This ensured that all instances had a strongly connected acquaintance graph, i.e., all corresponding BFNs were indeed stable (cf. Corollary 2).

We have conducted two kinds of experiments, performed on 10,000 instances for each instance type. For each instance \mathcal{B} , we simulated a \mathcal{B} -run σ . In the first experiment, we computed the number of steps required to reach a stable state with global consensus (cf. Corollary 2). In the second one, we considered a bribed BFN \mathcal{B}_{IS} with $IS = (\{n + 1\}, \alpha)$, α being a formula generated at random similarly to each agent’s beliefs, with probability $p = 0.5$; for each bribed BFN, we computed the number of steps necessary to ensure




Figure 3: Number of steps necessary to reach a unanimous acceptance of a beliefs held by an additional agent.

that α is unanimously accepted in \mathcal{B}_{IS} , which is guaranteed by Proposition 8. The results are reported through Tukey’s box plots in Figures 2 and 3, corresponding respectively to the first and the second experiment. Outliers were ignored for the sake of readability.

Our findings indicate a correlation between steps needed for stability in BFNs and unanimous acceptance of a piece of beliefs, except for non-symmetric loops (L). While a thorough empirical analysis of BFNs awaits future exploration, our current focus centers on demonstrating proof of concept: our theoretical reachability results demonstrate practical applicability with observable scaling in a few thousand steps. Notably, for complete graphs, Barabasi-Albert preferential influence graphs, and Kleinberg small-world graphs — representing artificial models of real-world social networks — the target state was reached within 500 steps in half of the instances, a reasonable outcome given the inherent stochasticity of the process.

We have created a software application, available online.⁶ It includes a user-friendly graphical interface that allows users to run multiple instances of BFNs, covering all the instance types used in our empirical study.

Conclusion

We introduced Belief Flow Networks (BFNs), a framework marking a step forward in accurately modeling how beliefs change within social networks. BFNs are arguably more realistic than previous logical-based approaches for agents communicating within a network: they allow asynchronous communication among agents, consider the iterative nature of belief change, and we have demonstrated the ability to predict when a group’s beliefs will reach a stable consensus. These practical outcomes underscore the usefulness of BFNs as a powerful tool for understanding how beliefs spread in complex interconnected systems. We have also demonstrated that identifying an optimal bribery policy is achievable in polynomial time. So, as a next step, it will be essential to explore strategies that increase the complexity of bribery and other manipulative actions, in an effort to counteract and mitigate manipulations and the spread of fake news within social networks.

⁶<https://github.com/nicolas-schwind/BFN-gui>

Acknowledgments

This work has benefited from the support of the JSPS KAKENHI Grant Number JP21H04905, the JST CREST Grant Number JPMJCR22D3, the AI Chair BE4musIA (ANR-20-CHIA-0028), and the THEMIS project (ANR-20-CE23-0018) of the French National Research Agency (ANR).

References

- Albert, R.; and Barabasi, A.-L. 2002. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1): 47.
- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50: 510–530.
- Aldana, M. 2003. Boolean dynamics of networks with scale-free topology. *Physica D: Nonlinear Phenomena*, 185(1): 45–66.
- Booth, R.; and Meyer, T. A. 2006. Admissible and Restrained Revision. *Journal of Artificial Intelligence Research*, 26: 127–151.
- Darwiche, A.; and Pearl, J. 1997. On the Logic of Iterated Belief Revision. *Artificial Intelligence*, 89(1-2): 1–29.
- Gärdenfors, P. 1988. *Knowledge in flux*. MIT Press.
- Grandi, U.; Lorini, E.; and Perrussel, L. 2015. Propositional Opinion Diffusion. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS’15)*, 989–997.
- Hegselmann, R.; and Krause, U. 2005. Opinion dynamics driven by various ways of averaging. *Computational Economics*, 25: 381–405.
- Iosifescu, M.; and Grigorescu, S. 1990. *Dependence with Complete Connections and Its Applications*. Cambridge University Press.
- Jin, Y.; and Thielscher, M. 2007. Iterated belief revision, revised. *Artificial Intelligence*, 171(1): 1–18.
- Katsuno, H.; and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52: 263–294.
- Kauffman, S. A. 1969. Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, 22(3): 437–467.
- Kauffman, S. A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, USA.
- Kleinberg, J. M. 2000. Navigation in a small world. *Nature*, 406(6798): 845–845.
- Konieczny, S.; Medina Grespan, M.; and Pino Pérez, R. 2010. Taxonomy of Improvement Operators and the Problem of Minimal Change. In *Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning (KR’10)*, 161–170.
- Konieczny, S.; and Pino Pérez, R. 2002. Merging Information Under Constraints: A Logical Framework. *Journal of Logic and Computation*, 12(5): 773–808.
- Konieczny, S.; and Pino Pérez, R. 2008. Improvement Operators. In *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR’08)*, 177–187.
- Medina Grespan, M.; and Pino Pérez, R. 2013. Representation of basic improvement operators. In *Trends in Belief Revision and Argumentation Dynamics*, 195–227. College Publications.
- Nayak, A. 1994. Iterated Belief Change Based on Epistemic Entrenchment. *Erkenntnis*, 41: 353–390.
- Nayak, A. C.; Foo, N. Y.; Pagnucco, M.; and Sattar, A. 1994. Entrenchment kinematics 101. In *Proceedings of the 7th Australian Joint Conference on Artificial Intelligence (AI’94)*, 157–164.
- Nayak, A. C.; Foo, N. Y.; Pagnucco, M.; and Sattar, A. 1996. Changing conditional beliefs unconditionally. In *Proceedings of the 6th Conference of Theoretical Aspects of Rationality and Knowledge (TARK’96)*, 119–135. Morgan Kaufmann.
- Novaro, A.; Grandi, U.; Longin, D.; and Lorini, E. 2018. Goal-Based Collective Decisions: Axiomatics and Computational Complexity. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI’18)*, 468–474.
- Novaro, A.; Grandi, U.; Longin, D.; and Lorini, E. 2019. Strategic Majoritarian Voting with Propositional Goals. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’19)*, 2144–2146.
- Riegler, A.; and Douven, I. 2009. Extending the Hegselmann-Krause model III: from single beliefs to complex belief states. *Episteme*, 6: 145–163.
- Rott, H. 2009. Shifting Priorities: Simple Representations for Twenty-Seven Iterated Theory Change Operators. In Makinson, D.; Malinowski, J.; and Wansing, H., eds., *Towards Mathematical Philosophy*, 269–296. Springer.
- Schwind, N.; Inoue, K.; Bourgne, G.; Konieczny, S.; and Marquis, P. 2015. Belief Revision Games. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI’15)*, 1590–1596.
- Schwind, N.; Inoue, K.; Bourgne, G.; Konieczny, S.; and Marquis, P. 2016. Is Promoting Beliefs Useful to Make Them Accepted in Networks of Agents? In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI’16)*, 1237–1243.
- Schwind, N.; Konieczny, S.; and Pino Pérez, R. 2022. On the Representation of Darwiche and Pearl’s Epistemic States for Iterated Belief Revision. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning (KR’22)*.
- Shields, P. 1973. *The theory of Bernoulli shifts*. Univ. Chicago Press.
- Spohn, W. 1988. *Ordinal Conditional Functions: A Dynamic Theory of Epistemic States*, 105–134. Springer Netherlands.

- Tarjan, R. 1972. Depth-First Search and Linear Graph Algorithms. *SIAM Journal on Computing*, 1(2): 146–160.
- Tsang, A.; and Larson, K. 2014. Opinion Dynamics of Skeptical Agents. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'14)*, 277–284.