Belief Revision Games (extended version including proofs of the propositions)

Nicolas Schwind
Transdisciplinary Research Integration Center
National Institute of Informatics
Tokyo, Japan
email: schwind@nii.ac.jp

Katsumi Inoue
National Institute of Informatics
The Graduate University for Advanced Studies
Tokyo, Japan
email: inoue@nii.ac.jp

Gauvain Bourgne
CNRS & Sorbonne Universités,
UPMC Univ Paris 06, UMR 7606,
LIP6, F-75005, Paris, France
email: Gauvain.Bourgne@lip6.fr

Sébastien Konieczny
CRIL - CNRS
Université d'Artois
Lens, France
email: konieczny@cril.fr

Pierre Marquis
CRIL - CNRS
Université d'Artois
Lens, France
email: marquis@cril.fr

Abstract

Belief revision games (BRGs) are concerned with the dynamics of the beliefs of a group of communicating agents. BRGs are "zero-player" games where at each step every agent revises her own beliefs by taking account for the beliefs of her acquaintances. Each agent is associated with a belief state defined on some finite propositional language. We provide a general definition for such games where each agent has her own revision policy, and show that the belief sequences of agents can always be finitely characterized. We then define a set of revision policies based on belief merging operators. We point out a set of appealing properties for BRGs and investigate the extent to which these properties are satisfied by the merging-based policies under consideration.

Introduction

In this paper, we introduce belief revision games (BRGs), that are concerned with the dynamics of the beliefs of a group of communicating agents. BRGs can be viewed as "zero-player" games: at each step of the game each agent revises her current beliefs (expressed in some finite propositional language) by taking account for the beliefs of her acquaintances. The aim is to study the dynamics of the game, i.e., the way the beliefs of a group of agents evolve depending on how agents are ready to share their beliefs. BRGs could be useful to model the evolution of beliefs in a group of agents in social networks, and to study several interesting notions such as influence, manipulation, gossip, etc. In this paper we mainly focus on the definition of BRGs, using formal tools coming from belief change theory, and investigate their behavior with respect to a set of expected logical properties. Let us introduce a motivating example of a BRG.

Example 1 *Consider a group of three undergraduate students, Alice, Bob and Charles, following the same CS curriculum. Bob is a friend of both Alice and Charles, but Alice and Charles do not know each other. Alice, Bob and Charles want to prepare the final exam of the "Basics of programming" course. Each student has some feelings about the topics which will be considered by their teacher for this exam. At start, Alice believes that "Binary search" will not be among the topics of the final exam, unlike "Bubble sort"; Bob believes that "Binary search" will be kept by the teacher, and that if "Bubble sort" is kept then "Quick sort" will be chosen as well by the teacher; finally, Charles just feels that "Binary search" will not be considered by the teacher. Each pair of friends exchange their opinions by sending e-mails in the evening. Each student is ready to make her opinion evolve by adopting the opinions of her friends when this does not conflict with hers, and by considering as most plausible any state of affairs which is as close as possible to the set of opinions at hand (her own one plus her friends' ones) in the remaining case. At the end of each day, Alice e-mails to Bob with her feelings, Bob to both Alice and Charles, and Charles to Bob. One is asked now about what can be inferred from this description. Some of the key questions are: (1) How beliefs must be updated? (2) Will agents always agree on some pieces of belief if they agree on it at the beginning of the game? (3) Will they eventually stop changing their beliefs?*

In the following, we present a formal setting for BRGs. Our very objective is to provide some answers to the questions above. Thus, we address question (1) by putting forward a set of revision policies which are based on existing belief merging operators from the literature and the induced belief revision operators. We identify a set of valuable properties for BRGs. They include unanimity preservation which models question (2) and convergence which models question (3). For each revision policy under consideration, we determine whether such properties are satisfied or not.

The proofs of propositions are given in an appendix.

Belief Revision Games

Belief sets are represented using a propositional language $\mathcal{L}_\mathcal{P}$ defined from a finite set of propositional variables $\mathcal{P}$ and the usual connectives. $\bot$ (resp. $\top$) is the Boolean constant always false (resp. true). An interpretation is a total function from $\mathcal{P}$ to $\{0, 1\}$. The set of all interpretations is denoted $\mathcal{W}$. An interpretation ω is a model of a formula $\varphi \in \mathcal{L}_\mathcal{P}$

if and only if it makes it true in the usual truth functional way. $Mod(\varphi)$ denotes the set of models of the formula φ, i.e., $Mod(\varphi) = \{\omega \in \mathcal{W} \mid \omega \models \varphi\}$. $\models$ denotes logical entailment and $\equiv$ logical equivalence, i.e., $\varphi \models \psi$ iff $Mod(\varphi) \subseteq Mod(\psi)$ and $\varphi \equiv \psi$ iff $Mod(\varphi) = Mod(\psi)$. A profile $\mathcal{K} = \langle\varphi_1, \ldots, \varphi_n\rangle$ is a finite vector of propositional formulae. Two profiles of formulae $\mathcal{K}_1 = \langle\varphi_1^1, \ldots, \varphi_n^1\rangle$ and $\mathcal{K}_2 = \langle\varphi_1^2, \ldots, \varphi_n^2\rangle$ are said to be equivalent, denoted $\mathcal{K}_1 \equiv \mathcal{K}_2$ if there is a permutation f over $\{1, \ldots, n\}$ such that for every $i \in 1, \ldots, n$, $\varphi_i^1 \equiv \varphi_{f(i)}^2$. Let us now introduce the formal definition of a Belief Revision Game.

Definition 1 (Belief Revision Game) *A* Belief Revision Game (BRG) *is a 5-tuple* $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ *where*

- $V = \{1, \ldots, n\}$ *is a finite set;*
- $A \subseteq V \times V$ *is an irreflexive binary relation on* V*;*
- $\mathcal{L}_\mathcal{P}$ *is a finite propositional language;*
- B *is a mapping from* V *to* $\mathcal{L}_\mathcal{P}$*;*
- $\mathcal{R} = \{R_1, \ldots, R_n\}$*, where each* R_i *is a mapping from* $\mathcal{L}_\mathcal{P} \times \mathcal{L}_\mathcal{P}{}^{in(i)}$ *to* $\mathcal{L}_\mathcal{P}$ *with* $in(i) = |\{j \mid (j, i) \in A\}|$ *the in-degree of* i*, such that for all formulae* $\varphi_0^1, \varphi_1^1, \ldots, \varphi_{in(i)}^1$*,* $\varphi_0^2, \varphi_1^2, \ldots, \varphi_{in(i)}^2$*, if* $\varphi_0^1 \equiv \varphi_0^2$ *and* $\langle\varphi_1^1, \ldots, \varphi_{in(i)}^1\rangle \equiv \langle\varphi_1^2, \ldots, \varphi_{in(i)}^2\rangle$*, then* $R_i(\varphi_0^1, \varphi_1^1, \ldots, \varphi_{in(i)}^1) \equiv R_i(\varphi_0^2, \varphi_1^2, \ldots, \varphi_{in(i)}^2)$*, and such that if* $in(i) = 0$*, then* R_i *is the identity function.*

Let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG. The set V represents the set of agents under consideration in G. The set A represents the set of *acquaintances* between the agents. Intuitively, if $(i, j) \in A$ then agent j is "aware" of the beliefs of agent i in the sense that agent i communicates her beliefs to agent j during the game. The set B represents each agent's beliefs expressed by a formula from $\mathcal{L}_\mathcal{P}$: for each $i \in V$, the formula $B(i)$ (noted B_i for short) is called a *belief state* and represents the initial beliefs of agent i. Lastly, each element $R_i \in \mathcal{R}$ is called the *revision policy* of agent i. Let us denote $\mathcal{C}_i$ the *context* of i, defined as the sequence $\mathcal{C}_i = B_{i_1}, \ldots, B_{i_{in(i)}}$ where $i_1 < \cdots < i_{in(i)}$ and $\{i_1, \ldots, i_{in(i)}\} = \{i_j \mid (i_j, i) \in A\}$. Then $R_i(B_i, \mathcal{C}_i)$ is the belief state of agent i once revised by taking into account her own current beliefs B_i and her current context. It is assumed by definition that all beliefs are considered up to equivalence (i.e., the syntactical form of the beliefs does not matter) and that an agent's beliefs do not evolve spontaneously when she has no neighbor.

Playing a BRG consists in determining how the beliefs of each agent evolve each time a revision step is performed. This calls for a notion of "belief sequence", which makes precise the dynamics of the game:

Definition 2 (Belief Sequence) *Given a BRG* $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ *and an agent* $i \in V$*, the* belief sequence *of* i*, denoted* $(B_i^s)_{s\in\mathbb{N}}$*, states how the beliefs of agent* i *evolve while moves take place.* $(B_i^s)_{s\in\mathbb{N}}$ *is inductively defined as follows:*

- $B_i^0 = B_i$*;*
- $B_i^{s+1} = R_i(B_i^s, \mathcal{C}_i^s)$ *for every* $s \in \mathbb{N}$*, where* $\mathcal{C}_i^s$ *is the context of* i *at step* s*.*

B_i^s denotes the belief state of agent i after s moves.

Since $\mathcal{L}_\mathcal{P}$ is a finite propositional language, there exists only finitely many formulae up to equivalence, hence only finitely many belief states can be reached. To make it formal, we need the concept of belief cycle:

Definition 3 (Belief Cycle) *A sequence* $(K^s)_{s\in\mathbb{N}}$ *of formulae from* $\mathcal{L}_\mathcal{P}$ *is* cyclic *if there exists a finite subsequence* $K^b, \ldots, K^e$ *such that for every* $j > e$*, we have* $K^j \equiv K^{b+((j-b) mod (e-b+1))}$*. In this case, the (characteristic)* belief cycle *of* $(K^s)_{s\in\mathbb{N}}$ *is defined by the subsequence* $K^b, \ldots, K^e$ *for which* b *and* e *are minimal.*

By the above argument, it is easy to prove that:

Proposition 1 *For every BRG* $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ *and every agent* $i \in V$*, the belief sequence of* i *is cyclic.*

As a consequence, each agent i is associated with a belief cycle which we simply denote $Cyc(B_i)$: the belief sequence of every agent i (which is an infinite sequence) can always be finitely described, since it is entirely characterized by its initial segment $B_i^0, B_i^1, \ldots, B_i^{b-1}$ and its belief cycle $Cyc(B_i) = B_i^b, B_i^{b+1}, \ldots, B_i^e$, which will be repeated (up to equivalence) ad infinitum in the sequence.

While there is no winner in a BRG G, such a game A BRG G can be "stopped" after a finite number of steps $stop(G) = \max_{i\in V}(\{e \mid Cyc(B_i) = B_i^b, \ldots, B_i^e\})$, since when this step $stop(G)$ is reached the belief cycles of all agents can be determined up to equivalence and the future evolution of the agents' beliefs can be predicted from the sequences of beliefs reached before $stop(G)$.

In the following, we are interested in determining the pieces of beliefs which result from the interaction of the agents in a BRG, focusing on the agents' belief cycles. A formula φ is considered accepted by an agent when it holds in every state of its belief cycle, which means that from some step s, φ will always hold. Then we define the notion of acceptability at the agent level and at the group level:

Definition 4 (Acceptability) *Let* $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ *be a BRG and* $\varphi \in \mathcal{L}_\mathcal{P}$*.* φ *is* accepted by $i \in V$ *if and only if for every* $B_i^s \in Cyc(B_i)$*, we have* $B_i^s \models \varphi$*.* φ *is* unanimously accepted *in* G *if and only if* φ *is accepted by all* $i \in V$*.*

A case of interest is when $|Cyc(B_i)| = 1$, i.e., the belief cycle of agent i has length 1. In such a case, the beliefs of agent i "stabilize" once the belief cycle is reached. A specific case is achieved by *stable* BRGs:

Definition 5 (Stability) *Let* $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ *be a BRG. A belief state* $B_i \in B$ *is said to be* stable *in* G *if* $|Cyc(B_i)| = 1$*. The BRG* G *is said to be* stable *iff each* $B_i \in B$ *is stable in* G*.*

Stability of a game is an interesting property, since it says in a sense that we reach some equilibrium point, where no agent further changes her belief. These two concepts will take part of some further properties on BRGs which we will introduce and investigate in the following.

Merging-Based Revision Policies

While all kinds of possible revision policies are allowed for BRGs, we now focus on revision policies R that are rationalized by theoretical tools from Belief Change Theory (see e.g. (Alchourrón, Gärdenfors, and Makinson 1985)), in particular belief merging and belief revision operators. Before introducing specific classes of revision policies of interest, let us introduce some necessary background on belief merging and belief revision. Formally, given a propositional language $\mathcal{L}_\mathcal{P}$ a merging operator Δ is a mapping from $\mathcal{L}_\mathcal{P} \times {\mathcal{L}_\mathcal{P}}^n$ to $\mathcal{L}_\mathcal{P}$. It associates any formula μ (the *integrity constraints*) and any profile $\mathcal{K} = \langle K_1, \ldots, K_n \rangle$ of belief states with a new formula $\Delta_\mu(\mathcal{K})$ (the *merged state*). A merging operator Δ aims at defining the merged state as the beliefs of a group of agents represented by the profile, under some integrity constraints. A set of nine standard properties denoted **(IC0)–(IC8)** are expected for merging operators (Konieczny and Pino Pérez 2002). Such operators are called *IC merging operators*. For space reasons, we just recall those used in the rest of the paper:

(IC0) $\Delta_\mu(\mathcal{K}) \models \mu$;

(IC1) If $\mu \not\models \bot$, then $\Delta_\mu(\mathcal{K}) \not\models \bot$;

(IC2) If $\bigwedge_{K \in \mathcal{K}} K \wedge \mu \not\models \bot$, then $\Delta_\mu(\mathcal{K}) \equiv \bigwedge_{K \in \mathcal{K}} K \wedge \mu$;

(IC3) If $\mathcal{K}_1 \equiv \mathcal{K}_2$ and $\mu_1 \equiv \mu_2$, then $\Delta_{\mu_1}(\mathcal{K}_1) \equiv \Delta_{\mu_2}(\mathcal{K}_2)$;

(IC4) If $K_1 \models \mu$, $K_2 \models \mu$ and $\Delta_\mu(\langle K_1, K_2 \rangle) \wedge K_1 \not\models \bot$, then $\Delta_\mu(\langle K_1, K_2 \rangle) \wedge K_2 \not\models \bot$.

A couple of additional postulates have been investigated in the literature, which are appropriate for some merging scenarios. We recall below one of them, Disjunction (Everaere, Konieczny, and Marquis 2010):

(Disj) If $\bigvee \mathcal{K} \wedge \mu$ is consistent, then $\Delta_\mu(\mathcal{K}) \models \bigvee \mathcal{K}$.

(Disj) is not satisfied by all IC merging operators but is expected in the case when it is assumed that (at least) one of the agent is right (her beliefs hold in the actual world), but we do not know which one.

Distance-based merging operators $\Delta^{d,f}$ are characterized by a pseudo-distance d (i.e., triangular inequality is not mandatory) between interpretations and an (aggregation) function f from $\mathbb{R}^+ \times \cdots \times \mathbb{R}^+$ to $\mathbb{R}^+$ (some basic conditions are required on f, including symmetry and non-decreasingness conditions, see (Konieczny, Lang, and Marquis 2004) for more details). They associate with every formula μ and every profile $\mathcal{K}$ a belief state $\Delta_\mu^{d,f}(\mathcal{K})$ which satisfies $Mod(\Delta_\mu^{d,f}(\mathcal{K})) = \min(Mod(\mu), \leq_\mathcal{K}^{d,f})$, where $\leq_\mathcal{K}^{d,f}$ is the total preorder over interpretations induced by $\mathcal{K}$ defined by $\omega \leq_\mathcal{K}^{d,f} \omega'$ if and only if $d^f(\omega, \mathcal{K}) \leq d^f(\omega', \mathcal{K})$, where $d^f(\omega, \mathcal{K}) = f_{K \in \mathcal{K}}\{d(\omega, K)\}$ and $d(\omega, K) = \min_{\omega' \models K} d(\omega, \omega')$. Usual distances are d_D, the drastic distance ($d_D(\omega, \omega') = 0$ if $\omega = \omega'$ and 1 otherwise), and d_H the Hamming distance ($d_H(\omega, \omega') = n$ if ω and ω' differ on n variables).

IC merging operators include some distance-based ones. We mention here two subclasses of them: the summation operators $\Delta^{d,\Sigma}$ (i.e., the aggregation function is the sum Σ) and the GMin operators $\Delta^{d,\mathsf{GMin}}$. GMin operators[1] associate with every formula μ and every profile $\mathcal{K}$ a belief state $\Delta_\mu^{d,f}(\mathcal{K})$ which satisfies $Mod(\Delta_\mu^{d,f}(\mathcal{K})) = \min(Mod(\mu), \leq_\mathcal{K})$, where $\leq_\mathcal{K}^{d,\mathsf{GMin}}$ is the total preorder over interpretations induced by $\mathcal{K}$ defined by $\omega \leq_\mathcal{K}^{d,\mathsf{GMin}} \omega'$ if and only if $d^{\mathsf{GMin}}(\omega, \mathcal{K}) \leq^{lex} d^{\mathsf{GMin}}(\omega', \mathcal{K})$ (where $\leq^{lex}$ is the lexicographic ordering induced by the natural order) and $d^{\mathsf{GMin}}(\omega, \mathcal{K})$ is the vector of numbers $d_1, \ldots, d_n$ obtained by sorting in a non-decreasing order the multiset $\langle d(\omega, K_i) \mid K_i \in \mathcal{K} \rangle$.

ω	K_1	K_2	K_3	$d_H^\Sigma(\omega, \mathcal{K})$	$d_H^{\mathsf{GMin}}(\omega, \mathcal{K})$
11	0	2	2	4	$(\mathbf{0}, \mathbf{2}, \mathbf{2})$
10	1	1	1	**3**	$(1, 1, 1)$
01	1	1	1	**3**	$(1, 1, 1)$

Table 1: The merging operators $\Delta^{d_H,\Sigma}$ and $\Delta^{d_H,\mathsf{GMin}}$.

Example 2 *Let $\mathcal{P} = \{a, b\}$, $\mathcal{K} = \langle K_1, K_2, K_3 \rangle$ where $K_1 = a \wedge b$, $K_2 = K_3 = \neg a \wedge \neg b$, and $\mu = a \vee b$. We consider both summation and* GMin *operators based on the Hamming distance. Table 1 shows for each interpretation $\omega \in Mod(\mu)$ the distances $d_H(\omega, K_i)$ for $i \in \{1, 2, 3\}$, and the distances $d_H^\Sigma(\omega, \mathcal{K})$ and $d_H^{\mathsf{GMin}}(\omega, \mathcal{K})$ (interpretations ω are denoted as binary sequences following the ordering $a < b$). We get that $\Delta_\mu^{d_H,\Sigma}(\mathcal{K}) \equiv (a \wedge \neg b) \vee (\neg a \wedge b)$ and $\Delta_\mu^{d_H,\mathsf{GMin}}(\mathcal{K}) \equiv a \wedge b$.*

Noteworthy, summation operators and GMin operators satisfy all **(IC0)–(IC8)** postulates (whatever the pseudo-distance under consideration), and additionally, GMin operators satisfy **(Disj)**, as well as the operator $\Delta^{d_D,\Sigma} = \Delta^{d_D,\mathsf{GMin}}$ (**(Disj)** is not satisfied by $\Delta^{d_H,\Sigma}$).

Belief revision operators can be viewed as belief merging operators restricted to singleton profiles: the revision $K_1 \circ K_2$ of a belief state K_1 by another belief state K_2 consists in "merging" the singleton profile $\langle K_1 \rangle$ under the integrity constraints K_2. Accordingly, if Δ is an IC merging operator then the revision operator $\circ_\Delta$ *induced by* Δ defined for all states K_1, K_2 as $K_1 \circ_\Delta K_2 = \Delta_{K_2}(\langle K_1 \rangle)$ satisfies the standard AGM revision postulates (Alchourrón, Gärdenfors, and Makinson 1985; Katsuno and Mendelzon 1992).

We are now ready to introduce several classes of revision policies R_i which are parameterized by an IC merging operator Δ and for some of them, by the corresponding revision operator $\circ_\Delta$.[2] Let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG. In the following, we assume for the sake of simplicity that all agents $i \in V$ apply the same revision policy, i.e., given an IC merging operator Δ, for all $R_i \in \mathcal{R}$, $R_i = R_\Delta$. Then let us consider the following revision policies, defined at each step s for any agent i who has a non-empty context $\mathcal{C}_i$:

[1]Here we give an alternative definition of $\Delta^{d,\mathsf{GMin}}$ by means of lists of numbers. However using Ordered Weighted Averages, one could fit the definition of a distance-based operator (Konieczny, Lang, and Marquis 2004).

[2]When using a merging operator without integrity constraints we just note $\Delta(\mathcal{K})$ instead of $\Delta_\top(\mathcal{K})$ for improving readibility.

Definition 6 (Merging-Based Revision Policies)

- $R^1_\Delta(B^s_i, \mathcal{C}^s_i) = \Delta(\langle \mathcal{C}^s_i \rangle)$;
- $R^2_\Delta(B^s_i, \mathcal{C}^s_i) = \Delta_{\Delta(\langle \mathcal{C}^s_i \rangle)}(\langle B^s_i \rangle) \quad [= B^s_i \circ_\Delta \Delta(\langle \mathcal{C}^s_i \rangle)]$;
- $R^3_\Delta(B^s_i, \mathcal{C}^s_i) = \Delta(\langle B^s_i, \mathcal{C}^s_i \rangle)$;
- $R^4_\Delta(B^s_i, \mathcal{C}^s_i) = \Delta(\langle B^s_i, \Delta(\langle \mathcal{C}^s_i \rangle) \rangle)$;
- $R^5_\Delta(B^s_i, \mathcal{C}^s_i) = \Delta_{B^s_i}(\Delta(\langle \mathcal{C}^s_i \rangle)) \quad [= \Delta(\langle \mathcal{C}^s_i \rangle) \circ_\Delta B^s_i]$;
- $R^6_\Delta(B^s_i, \mathcal{C}^s_i) = \Delta_{B^s_i}(\langle \mathcal{C}^s_i \rangle)$.

First of all, please note that since **(IC3)** requires Δ to be syntax-independent (i.e., profiles and integrity constraints are considered up to equivalence), these revision policies are all consistent with the conditions given in Definition 1.

Intuitively, these strategies are ranked according to the relative importance given to each agent's beliefs compared to her neighbors' opinion. For R^1_Δ, only the aggregated opinion of the neighbors is relevant. For R^2_Δ, the current opinion of the agent is revised by the aggregated opinion of the neighbors; doing so, an agent is ready to adopt the part of the merged beliefs of her neighbors which are as close as possible to her own current beliefs. For R^3_Δ the agent considers that her opinion is as important as each one of her neighbors. For R^4_Δ the agent considers that her opinion is as important as the aggregated opinion of her neighbors. For R^5_Δ and R^6_Δ, the agent does not give up her current beliefs and just accepts additional information compatible with them. Noteworthy, R^5_Δ and R^6_Δ are not equivalent: for R^5_Δ the agent first aggregates her neighbors' opinion, and then revise the merged result by her own opinion; for R^6_Δ the agent proceeds with her neighbors' opinion and her own one in a single step.[3]

Example 1 (continued) *We formalize the example presented in the introduction as the BRG $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ defined as follows. Let $V = \{1, 2, 3\}$ where 1 corresponds to Alice, 2 to Bob, and 3 to Charles. $A = \{(1,2), (2,1), (2,3), (3,2)\}$ expresses that Alice and Bob are connected, and that Bob and Charles are connected. $\mathcal{L}_\mathcal{P}$ is built up from the set of propositional variables $\mathcal{P} = \{s, b, q\}$, where s stands for "Binary Search ", b for "Bubble Sort" and q for "Quick Sort". The initial beliefs of agents are expressed as $B_1 = \neg s \wedge b$, $B_2 = s \wedge (b \Rightarrow q)$ and $B_3 = \neg s$. Since in the case of conflicting beliefs, each agent considers to merge her friends' opinions and her own one together, revision policies R^3_Δ are appropriate candidates for each agent. Let us consider the summation operator based on the Hamming distance. We have $R_1 = R_2 = R_3 = R^3_{\Delta^{d_H,\Sigma}}$. The belief sequences associated with the three agents are given in Table 2: the belief cycle of agent 1 (resp. 2, 3) is given by (B^2_1) (resp. (B^1_2), (B^2_3)). G is a stable game. Note that $\neg s \wedge b \wedge q$ is unanimously accepted in G (as well as all formulae entailed by it).*

At first, Alice believes that "Binary Search" will be considered, unlike "Bubble Sort". On the next day, she still believes that "Bubble sort" will be considered (since this does not conflict with Bob's view), but she now believes that "Quick sort" will be considered as well (she adopts the fact that if "Bubble sort" is kept then "Quick sort" will be chosen from Bob's view since this does not conflict with her opinion); finally, she cancels her opinion about the fact that "Binary search" will not be considered since Bob disagrees with it. Thus Alice's beliefs evolve from $\neg s \wedge b$ to $b \wedge q$. Similarly, Bob's beliefs evolve from $s \wedge (b \Rightarrow q)$ to $\neg s \wedge b \wedge q$. Note here that since both friends of Bob agree about $\neg s$, Bob changes his mind about it. Charles' beliefs evolve from $\neg s$ to $b \Rightarrow q$. At the end of the day, a further e-mail exchange process takes place. It makes the three friends modifying their beliefs and now sharing the same opinion about the exam topics, namely $\neg s \wedge b \wedge q$. Their opinions then do not change any longer.

The belief sequences are graphically represented in Figure 1. At each time step, blue nodes are agents accepting $\varphi = s \wedge b \wedge q$, red nodes denote agents accepting $\neg\varphi = \neg s \vee \neg b \vee \neg q$ and gray nodes stand for agents accepting neither φ nor $\neg\varphi$.

[3] Consider for instance $\mathcal{C}_i = p \wedge q, \neg p, \neg p \wedge \neg q$ and $B_i = p$. Then $R^5_{\Delta^{d_D,\Sigma}}(B_i, \mathcal{C}_i) \equiv p \wedge \neg q$ whereas $R^6_{\Delta^{d_D,\Sigma}}(B_i, \mathcal{C}_i) \equiv p \wedge q$.

step i	B^i_1	B^i_2	B^i_3
0	$\neg s \wedge b$	$s \wedge (b \Rightarrow q)$	$\neg s$
1	$b \wedge q$	$\neg s \wedge b \wedge q$	$b \Rightarrow q$
≥ 2	$\neg s \wedge b \wedge q$	$\neg s \wedge b \wedge q$	$\neg s \wedge b \wedge q$

Table 2: The belief sequences of Alice, Bob and Charles.

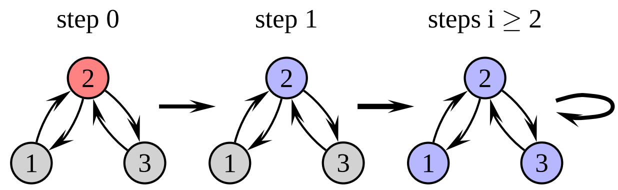

Figure 1: A graphical representation of belief sequences.

Logical Properties for Belief Revision Games

We introduce now some expected logical properties for BRGs, and investigate which BRGs satisfy them depending on the chosen revision policy. While the properties hereafter are relevant to all BRGs, we focus on BRGs which are instantiated with revision policies from the six classes defined in the previous section, and assume that the same revision policy is applied for each agent. Given a revision policy R^k_Δ, $\mathcal{G}(R^k_\Delta)$ is the set of all BRGs $(V, C, \mathcal{L}_\mathcal{P}, C, \mathcal{R})$ where for each $R_i \in \mathcal{R}$, $R_i = R^k_\Delta$. Additionally, R^k_Δ is said to satisfy a given property P on BRGs if all BRGs from $\mathcal{G}(R^k_\Delta)$ satisfy P.

We start with a set of "preservation" properties which are counterparts of some postulates on belief merging operators (cf. previous section). These properties express the idea that the interaction between agents should not lead them to "degrade" their belief states.

Definition 7 (Consistency Preservation (CP)) *A BRG $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ satisfies **(CP)** if for each $B_i \in B$, if B_i is consistent then all beliefs from $(B^s_i)_{s \in \mathbb{N}}$ are consistent.*

(CP) requires that agents with consistent initial beliefs never become self-conflicting in their belief sequence. It is the direct counterpart of **(IC1)** for merging operators:

Proposition 2 *For every $k \in \{1, \ldots, 6\}$, R^k_Δ satisfies **(CP)** if Δ satisfies **(IC1)**.*

Definition 8 (Agreement Preservation (AP)) *A BRG $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ satisfies **(AP)** if given any consistent formula $\varphi \in \mathcal{L}_\mathcal{P}$, if for each $B_i \in B$, $\varphi \models B_i$ then for each $B_i \in B$ and at every step $s \geq 0$, $\varphi \models B^s_i$.*

(AP) requires that if all agents initially agree on some alternatives, then they will not change their mind about them. It corresponds to **(IC2)** for merging operators:

Proposition 3 *For every $k \in \{1, \ldots, 6\}$, R^k_Δ satisfies **(AP)** if Δ satisfies **(IC2)**.*

Definition 9 (Unanimity Preservation (UP)) *A BRG $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ satisfies **(UP)** if given any formula $\varphi \in \mathcal{L}_\mathcal{P}$, if for each $B_i \in B$, $B_i \models \varphi$ then for each $B_i \in B$ and at every step $s \geq 0$, $B^s_i \models \varphi$.*

(UP) states that every formula which is a logical consequence of the initial agents' beliefs should remain so in their belief sequence; note that in such a case, the formula is unanimously accepted in the BRG under consideration (cf. Definition 4). It is interesting to note that the statements of **(AP)** and **(UP)** have quite a similar structure. However, **(AP)** expresses a unanimity on models whereas **(UP)** is concerned with unanimity on formulae. The corresponding properties for merging operators have been presented in (Everaere, Konieczny, and Marquis 2010), where the authors also showed that the corresponding postulate of unanimity on formulae for merging operators is equivalent to **(Disj)** (cf. previous section).

Proposition 4 *For every $k \in \{1, \ldots, 6\}$, R^k_Δ satisfies **(UP)** if Δ satisfies:*

- ***(IC0)** when $k \in \{5, 6\}$;*
- ***(Disj)** when $k \in \{1, 3, 4\}$;*
- ***(IC0)** and **(Disj)** when $k = 2$.*

In the general case, revision policies R^k_Δ with $k \in \{1, 2, 3, 4\}$ do not satisfy **(UP)** for merging operators Δ which do not satisfy **(Disj)**. This is because such merging operators may produce new beliefs absent from the states of the profile under consideration: some interpretations that do not satisfy any of the input belief states can be models of the merged state. However, for R^5_Δ and R^6_Δ, Δ is not required to satisfy **(Disj)** since in the presence of **(IC0)** alone these policies are the most change-reluctant ones: each agent who accepts φ at some step will keep accepting φ at the next step since she will only refine her own beliefs. We address precisely the behavior of all merging-based revision policies in terms of agents' responsiveness to their neighbors:

Definition 10 (Responsiveness (Resp)) *A BRG $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ satisfies **(Resp)** if for each $B_i \in B$ such that $\mathcal{C}_i$ is not empty, for every step $s \geq 0$, if (i) for every $B^s_{i_j} \in \mathcal{C}^s_i$, $B^s_{i_j} \wedge B^s_i \models \bot$, and (ii) $\bigwedge_{B^s_{i_j} \in \mathcal{C}^s_i} B^s_{i_j} \not\models \bot$, then $B^{s+1}_i \not\models B^s_i$.*

Informally, **(Resp)** demands that an agent should take into consideration the beliefs of her neighbors whenever (i) her beliefs are inconsistent with the beliefs of each one of her neighbors, and (ii) her neighbors agree on some alternatives. Accordingly, **(Resp)** is not satisfied by R^5_Δ and R^6_Δ:

Proposition 5 *If Δ satisfies **(IC0)**, then R^5_Δ and R^6_Δ do not satisfy **(Resp)**.*

But **(Resp)** is satisfied by most of the remaining revision policies R^k_Δ under some basic conditions on Δ:

Proposition 6 *For every $k \in \{1, 2, 4\}$, R^k_Δ satisfies **(Resp)** if Δ satisfies:*

- ***(IC2)** when $k = 1$;*
- ***(IC0)** and **(IC2)** when $k = 2$;*
- ***(IC2)** and **(IC4)** when $k = 4$.*

Intuitively, R^3_Δ seems to be less change-reluctant than R^4_Δ, since for R^3_Δ the agent considers her beliefs as being as important as each one of her neighbors whereas for R^4_Δ, she considers her beliefs as being as important as the aggregated beliefs of her neighbors. However, surprisingly R^3_Δ does not satisfy **(Resp)** even when some "fully rational" IC merging operators Δ are used:

Proposition 7 *$R^3_{\Delta^{d_H,\Sigma}}$ does not satisfy **(Resp)**.*

Recall that the merging operator $\Delta^{d_H,\Sigma}$ satisfies all the standard IC postulates **(IC0)–(IC8)**. Thus, the fact that Δ satisfies those postulates is not enough for R^3_Δ to satisfy **(Resp)**. However, we show below that these postulates are consistent with **(Resp)**, in the sense that there exists a merging operator Δ satisfying **(IC0)–(IC8)** (and **(Disj)**) which makes R^3_Δ a responsive policy:

Proposition 8 *For any aggregation function f, $R^3_{\Delta^{d_D,f}}$ satisfies **(Resp)**.*

In particular, the revision policy $R^3_{\Delta^{d_D,\Sigma}} = R^3_{\Delta^{d_D,GMin}}$ satisfies **(Resp)**.

Given a BRG $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$, a formula φ and an agent $i \in V$, let us denote $G_{i \to \varphi}$ the BRG $(V, A, \mathcal{L}_\mathcal{P}, B', \mathcal{R})$ defined as $B'_i = B'_i \wedge \varphi$ and for every $j \in V$, $j \neq i$, $B'_j = B_j$.

Definition 11 (Monotonicity (Mon)) *A BRG $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ satisfies **(Mon)** if whenever φ is unanimously accepted in G, φ is also unanimously accepted in $G_{i \to \varphi}$ for every $i \in V$.*

(Mon) is similar to the *monotonicity criterion* in Social Choice Theory. It is expressed in (Woodall 1997) as the condition where a candidate should not be harmed if she is raised on some ballots without changing the orders of the other candidates. In the BRG context, a formula φ which is unanimously accepted should still be unanimously accepted if some agent's initial beliefs were "strengthened" by φ.

For each revision policy R^k_Δ, $k \in \{1, \ldots, 6\}$, **(Mon)** is not guaranteed even when the merging operator under consideration satisfies the postulates **(IC0)–(IC8)**:

Proposition 9 *For every $k \in \{1, \ldots, 6\}$, $R^k_{\Delta^{d_H,\Sigma}}$ does not satisfy **(Mon)**.*

The existence of revision policies R^k_Δ which satisfy **(Mon)** remains an open issue. However, one conjectures that for every $k \in \{1, \ldots, 6\}$, $R^k_{\Delta_{d_D,\Sigma}}$ satisfies **(Mon)**. This claim is supported by some empirical evidence. We have conducted a number of tests when four propositional symbols are considered in the language $\mathcal{L}_\mathcal{P}$, for various graph topologies up to 10 agents and for $k \in \{1, \ldots, 6\}$. All the tested instances supported the claim.

The last property we provide concerns the stability issue:

Definition 12 (Convergence) *A BRG satisfies **(Conv)** if it is stable.*

Proposition 10 *The revision policies R^5_Δ and R^6_Δ satisfy **(Conv)** if Δ satisfies **(IC0)**.*

None of the remaining revision policies R^k_Δ, $k \in \{1, 2, 3, 4\}$ satisfy **(Conv)** in the general case. In fact, for these policies the stability of BRGs cannot be guaranteed as soon as the merging operator under consideration satisfies some basic IC postulates.

Proposition 11 *For every $k \in \{1, 2, 3, 4\}$, R^k_Δ does not satisfy **(Conv)** if Δ satisfies:*

- ***(IC2)** when $k = 1$;*
- ***(IC0)** and **(IC2)** when $k = 2$;*
- ***(IC1)**, **(IC2)** and **(IC4)** when $k \in \{3, 4\}$.*

All the results are summarized in Table 3. For each class R^k_Δ of revision policies and each property on revision policies, for some (set of) postulate(s) **(P)** on merging operators or directly for some merging operators, $\surd$**(P)** (resp. $\times$**(P)**) means that R^k_Δ satisfies (resp. does not satisfy) the corresponding property when Δ satisfies **(P)** or is one of the merging operators which are specified. One can observe that under some basic conditions on Δ, for $k \in \{1, 2, 4\}$ the revision policies R^k_Δ are well-behaved in terms of responsiveness but do not guarantee the stability of all BRGs, while the converse holds for the revision policies R^5_Δ and R^6_Δ.

Before closing the section, we go further in the investigation of the convergence property by considering a subclass of so-called *directed acyclic* BRGs $(V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ which require the underlying graph (V, A) not to contain any cycle:

Proposition 12 *For $k \in \{1, 2, 3, 4\}$, all directed acyclic BRGs from $\mathcal{G}(R^k_\Delta)$ satisfy **(Conv)** when $k = 1$ or if:*

- *when $k = 2$, Δ satisfies **(IC0)** and **(IC2)**;*
- *when $k = 3$, Δ is a distance-based merging operator;*
- *when $k = 4$, Δ satisfies **(IC2)**, **(IC4)** and **(Disj)**, or Δ is a distance-based merging operator.*

Related Work

Belief revision games are somehow related to many settings where some interacting ”agents” are considered, including cellular automata (Wolfram 1983), Boolean networks (Kauffman 1969; 1993; Aldana 2003), opinion dynamics (Hegselmann and Krause 2005; Riegler and Douven 2009; Tsang and Larson 2014), and many complex systems (Latane and Nowak 1997; Kacpersky and Holyst 2000; Olshevsky and Tsitsiklis 2009; Bloembergen et al. 2014; Ranjbar-Sahraei et al. 2014). We focus here on related work strongly connected to Belief Revision Games.

In (Delgrande, Lang, and Schaub 2007), the authors introduce a general framework for minimizing disagreements among beliefs associated with points connected through a graph. They define a completion operator which consists in revising the belief state of each point with respect to the belief states of its “neighbors”. This operator outputs a new graph where each belief state is strengthened and restricted to the models which are the closest ones to the neighbor states. Suitable applications include the case when points in the graph are interpreted as regions in space (Würbel, Jeansoulin, and Papini 2000). Though the idea of embedding belief states into a graph structure is similar to our approach, it differs from BRGs on several aspects. First, only undirected graphs are considered. Second, their completion operator is idempotent so it cannot be used iteratively. Third, belief states are strengthened by the operation of completion, whereas in BRGs agents can “give up” beliefs (e.g., when considering responsive policies such as R^1, R^2 and R^4).

In (Gauwin, Konieczny, and Marquis 2007), the authors introduce and study families of so-called iterated merging conciliation operators. Such operators are considered to rule the dynamics of the profile $\mathcal{K}$ of belief states associated with a group of agents. At each step the state B_i of agent i is modified, by revising the merged state $\Delta(\mathcal{K})$ by B_i (skeptical approach), or by revising B_i by the merged state $\Delta(\mathcal{K})$ (credulous approach). Such merge-then-revise change functions are closely related to our merging-based revision policies R^2 (for the credulous one) and R^5 (for the skeptical one). They do not coincide with them nevertheless since in our approach B_i does not belong to its context $\mathcal{C}_i$; clearly enough, this amounts to giving more importance to B_i when majoritarian merging operators are considered, and as a consequence the states obtained after the “revision” of B_i may differ. Notwithstanding the merging-based revision policies used, such conciliation processes correspond to specific BRGs where the topology is the clique one. One of the main issues considered in (Gauwin, Konieczny, and Marquis 2007) is the stationarity of the process (i.e., the convergence of the policies), which is proved in the skeptical approach; however, preservation issues, as well as responsiveness and monotonicity are not studied.

Our work also is relevant to the opinion dynamics problem, which raises an abundant literature in philosophy for the last two decades. One of the most influential model to opinion dynamics is Hegselmann-Krause’s one (see e.g., (Hegselmann and Krause 2005)). In the original Hegselmann-Krause’s model, a set of agents aims at determining the value of a given parameter $p \in (0, 1]$. Each agent i has some belief p_i, her estimate of the right value of p. Each agent updates her belief p_i by replacing it by the average of p_i with the beliefs of its ”neighbors”, i.e., the set of all values p_j which are sufficiently close to p_i, i.e., $\mid p_i - p_j \mid \leq \epsilon$ where ϵ is a preset constant. Available results take the form of analytical results or of empirical results achieved using computer simulations and show the existence of diverging converging groups in the basic model. Many extensions of it have been pointed out so far, a closest one to our work being

	(CP)	(AP)	(UP)	(Resp)	(Mon)	(Conv)
R^1_Δ	$\surd$**(IC1)**	$\surd$**(IC2)**	$\surd$**(Disj)**	$\surd$**(IC2)**	$\times_{(\Delta^{d_H,\Sigma})}$	$\times$**(IC2)**
R^2_Δ	$\surd$**(IC1)**	$\surd$**(IC2)**	$\surd$**(IC0) & (Disj)**	$\surd$**(IC0) & (IC2)**	$\times_{(\Delta^{d_H,\Sigma})}$	$\times$**(IC0) & (IC2)**
R^3_Δ	$\surd$**(IC1)**	$\surd$**(IC2)**	$\surd$**(Disj)**	$\surd_{(\Delta^{d_D,f})}$ / $\times_{(\Delta^{d_H,\Sigma})}$	$\times_{(\Delta^{d_H,\Sigma})}$	$\times$**(IC1) & (IC2) & (IC4)**
R^4_Δ	$\surd$**(IC1)**	$\surd$**(IC2)**	$\surd$**(Disj)**	$\surd$**(IC2) & (IC4)**	$\times_{(\Delta^{d_H,\Sigma})}$	$\times$**(IC1) & (IC2) & (IC4)**
R^5_Δ	$\surd$**(IC1)**	$\surd$**(IC2)**	$\surd$**(IC0)**	$\times$**(IC0)**	$\times_{(\Delta^{d_H,\Sigma})}$	$\surd$**(IC0)**
R^6_Δ	$\surd$**(IC1)**	$\surd$**(IC2)**	$\surd$**(IC0)**	$\times$**(IC0)**	$\times_{(\Delta^{d_H,\Sigma})}$	$\surd$**(IC0)**

Table 3: Properties satisfied by the revision policies R^k_Δ for $k \in \{1, \ldots, 6\}$.

Riegler-Douven's one (Riegler and Douven 2009). Indeed, in Riegler-Douven's model, the belief states take the form of propositional theories. Proximity between belief states is evaluated as the minimal Hamming distance between their propositional models. The objective of the agents is to track the truth, which is rendered possible by incorporating at each update step some piece of evidence e_i supposed to be true in the actual state of affairs. Update proceeds by a specific way of averaging over the "neighbors" beliefs together with the evidence. Thus, this work departs from our own one in many dimensions; mainly the way beliefs are revised, the handling of pieces of evidence, the concept of neighborhood which depends on the proximity of the belief states and the nature of the results (which mainly amounts here to determining using computer simulations for which values of the parameters used in the model the beliefs are converging to the truth).

Conclusion

In this paper, we formalized the concept of belief revision game (BRG) for modeling the dynamics of the beliefs of a group of agents. We pointed out a set of properties for BRGs which address several preservation issues, as well as responsiveness, monotonicity and convergence. As a first attempt to investigate the behavior of BRGs with respect these properties, we introduced several classes of revision policies which are based on belief merging operators. We considered the case where all agents use the same revision policy and investigated the extent to which the BRGs concerned with these policies satisfy the properties. Additionally, we developed a software available online at http://www.cril.fr/brg/brg.jar. It consists of graphical interface which allows one to play BRGs considering any of the 18 revision policies from $\{R^k_\Delta \mid k \in \{1, \ldots, 6\}, \Delta \in \{\Delta^{d_D,\Sigma}, \Delta^{d_H,\Sigma}, \Delta^{d_H,\mathsf{Gmin}}\}\}$. Some instances of BRGs are provided together with the software, including the BRG from our motivating example (Example 1) and the counter-examples used in the proofs of some propositions.

Practical applications of the BRG model are numerous. For instance, in brand crisis management, negative content regarding a brand could disseminate rapidly over social media and generate negative perceptions (Dawar and Pillutla 2000). In such a case, identifying how information is propagated within a social network and which are the influential agents (the opinion leaders) is a hot research topic. As a consequence, our general framework leaves the way open to many extensions and additional theoretical studies. Perspectives include a further investigation of the robustness of BRGs in terms of belief manipulation. For example, one could investigate how "controllable" a BRG is with respect to some piece of belief. An intuitive notion of controllability with respect to some piece of belief would consider the minimal number of controlled agents the role of which is to make this piece of belief unanimously accepted in the BRG.

References

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(2):510–530.

Aldana, M. 2003. Boolean dynamics of networks with scale-free topology. *Physica D: Nonlinear Phenomena* 185(1):45–66.

Bloembergen, D.; Sahraei, B. R.; Bou-Ammar, H.; Tuyls, K.; and Weiss, G. 2014. Influencing social networks: An optimal control study. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*, 105–110.

Dawar, N., and Pillutla, M. 2000. Impact of product-harm crises on brand equity: the moderating role of consumer expectations. *Journal of Marketing Research* 37.

Delgrande, J. P.; Lang, J.; and Schaub, T. 2007. Belief change based on global minimisation. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI'07)*, 2468–2473.

Everaere, P.; Konieczny, S.; and Marquis, P. 2010. Disjunctive merging: Quota and gmin merging operators. *Artificial Intelligence* 174(12-13):824–849.

Gauwin, O.; Konieczny, S.; and Marquis, P. 2007. Conciliation through iterated belief merging. *Journal of Logic and Compution* 17(5):909–937.

Hegselmann, R., and Krause, U. 2005. Opinion dynamics driven by various ways of averaging. *Computational Economics* 25:381–405.

Kacpersky, K., and Holyst, J. 2000. Phase transition as a persistent feature of groups with leaders in models of opinion formation. *Physica A* 287:631–643.

Katsuno, H., and Mendelzon, A. O. 1992. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3):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.

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.; Lang, J.; and Marquis, P. 2004. DA^2 merging operators. *Artificial Intelligence* 157:49–79.

Latane, B., and Nowak, A. 1997. *Progress in Communication Sciences*. Ablex Publishing Corporation. chapter "Self-organizing social systems: necessary and sufficient conditions for the emergence of clustering, consolidation, and continuing diversity".

Olshevsky, A., and Tsitsiklis, J. 2009. Convergence speed in distributed consensus and averaging. *SIAM Journal on Control and Optimization* 48(1):33–55.

Ranjbar-Sahraei, B.; Ammar, H. B.; Bloembergen, D.; Tuyls, K.; and Weiss, G. 2014. Evolution of cooperation in arbitrary complex network. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'14)*, 677–684.

Riegler, A., and Douven, I. 2009. Extending the Hegselmann-Krause model III: from single beliefs to complex belief states. *Episteme* 6:145–163.

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.

Wolfram, S. 1983. Statistical mechanics of cellular automata. *Reviews of Modern Physics* 55(3):601–644.

Woodall, D. R. 1997. Monotonicity of single-seat preferential election rules. *Discrete Applied Mathematics* 77(1):81–98.

Würbel, E.; Jeansoulin, R.; and Papini, O. 2000. Revision: an application in the framework of GIS. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning (KR'00)*, 505–515.

Appendix: proofs of propositions

Proposition 1
For every BRG $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ and every agent $i \in V$, the belief sequence of i is cyclic.

Proof: Let us first introduce the notion of equivalence between BRGs which we will exploit in this proof. Two BRGs $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$, $G' = (V', A', \mathcal{L}_\mathcal{P}', B', \mathcal{R}')$ are said to be equivalent, denoted $G \equiv G'$, if $V = V'$, $A = A'$, $\mathcal{L}_\mathcal{P} = \mathcal{L}_\mathcal{P}'$, $\mathcal{R} = \mathcal{R}'$ and for every $i \in V$, $B_i \equiv B'_i$. Let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG. Since $\mathcal{L}_\mathcal{P}$ is a propositional language defined on a finite set of propositional variables, the number of belief bases from $\mathcal{L}_\mathcal{P}$ which are distinct up to logical equivalence is finite. Therefore, there are only finitely many n-vectors of formulae from $\mathcal{L}_\mathcal{P}$ up to equivalence. According to Definition 1, each $R_i \in \mathcal{R}$ is a mapping from $\mathcal{L}_\mathcal{P} \times {\mathcal{L}_\mathcal{P}}^{in(i)}$ to $\mathcal{L}_\mathcal{P}$ which considers formulae up to equivalence. This means that there are finitely many BRGs G^s up to equivalence when s ranges over $\mathbb{N}$. So the series $(G^s)_{s\in\mathbb{N}}$ is "cyclic" in the sense that there exists a finite subsequence $(G^b, \ldots, G^e)$ such that for every $j > e$, we have $G^j \equiv G^{b+((j-b) mod (e-b+1))}$. Therefore, for such indexes b, e, for each agent $i \in V$ and for the subsequence $(B_i^b, \ldots, B_i^e)$ from her belief sequence $(B_i^s)_{s\in\mathbb{N}}$, we get that for every $j > e$, $B_i^j \equiv B_i^{b+((j-b) mod (e-b+1))}$. This means that the belief sequence of each agent $i \in V$ is cyclic. ■

Proposition 2
For every $k \in \{1, \ldots, 6\}$, R^k_Δ satisfies **(CP)** if Δ satisfies **(IC1)**.

Proof: Let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG from any of the classes $\mathcal{G}(R^k_\Delta)$ where $k \in \{1, \ldots, 6\}$ and where Δ satisfies **(IC1)**. We must prove that G satisfies **(CP)**. Assume that for each $B_i \in B$, B_i is consistent. We prove that for each step $s \geq 0$, B_i^s is consistent by recursion on s. This is trivial for each $k \in \{1, \ldots, 6\}$ when $s = 0$. Now, let $s \geq 0$ and assume that for each $B_i^s \in G^s$, B_i^s is consistent. Then:
• when $k \in \{1, 3, 4\}$, B_i^{s+1} is a belief base of the form $\Delta(\mathcal{K}) = \Delta_\top(\mathcal{K})$ for some profile $\mathcal{K}$. Since $\top$ is consistent, by **(IC1)** we get that $\Delta_\top(\mathcal{K})$ is also consistent, thus B_i^{s+1} is consistent;
• when $k = 2$, $B_i^{s+1} = R_i^2(B_i^s, \mathcal{C}_i^s) = \Delta_{\Delta(\langle \mathcal{C}_i^s \rangle)}(\langle B_i^s \rangle)$. Since by **(IC1)**, $\Delta(\langle \mathcal{C}_i^s \rangle) = \Delta_\top(\langle \mathcal{C}_i^s \rangle)$ is consistent, we get by **(IC1)** that $\Delta_{\Delta(\langle \mathcal{C}_i^s \rangle)}(\langle B_i^s \rangle)$ is also consistent. Hence, B_i^{s+1} is consistent;
• when $k \in \{5, 6\}$, B_i^{s+1} is a belief base of the form $\Delta_{B_i^s}(\mathcal{K})$ for some profile $\mathcal{K}$. Since B_i^s is consistent by the recursion hypothesis, we get by **(IC1)** that $\Delta_{B_i^s}(\mathcal{K})$ is consistent. Hence, B_i^{s+1} is consistent.

This concludes the proof. ■

Proposition 3
For every $k \in \{1, \ldots, 6\}$, R^k_Δ satisfies **(AP)** if Δ satisfies **(IC2)**.

The proof uses the following lemma:

Lemma 1 *Let Δ be a merging operator which satisfies **(IC2)**. Then for every consistent propositional formula φ, for every formula μ and every non-empty profile $\mathcal{K}$, if $\varphi \models \mu$ and $\varphi \models K$ for every $K \in \mathcal{K}$, then $\varphi \models \Delta_\mu(\mathcal{K})$.*

Proof: Let Δ be a merging operator which satisfies **(IC2)**, φ be a consistent formula, μ be a formula and $\mathcal{K}$ be a non-empty profile such that $\varphi \models \mu$ and $\varphi \models K$ for every $K \in \mathcal{K}$. Since φ is consistent, $\bigwedge_{K\in\mathcal{K}} K \wedge \mu$ is consistent and $\varphi \models \bigwedge_{K\in\mathcal{K}} K \wedge \mu$. Then by **(IC2)**, $\Delta_\mu(\mathcal{K}) \equiv \bigwedge_{K\in K} K \wedge \mu$. Therefore, $\varphi \models \Delta_\mu(\mathcal{K})$. ■

We now prove Proposition 3:

Proof: Let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG from any of the classes $\mathcal{G}(R^k_\Delta)$ where $k \in \{1, \ldots, 6\}$ and where Δ satisfies **(IC2)**. We must prove that G satisfies **(AP)**. Let φ be a consistent propositional formula from $\mathcal{L}_\mathcal{P}$ and assume that for each $B_i \in B$, $\varphi \models B_i$. We prove that for each $B_i \in B$ and at every step $s \geq 0$, $\varphi \models B_i^s$ by recursion on

s. This is trivial for each $k \in \{1, \dots, 6\}$ when $s = 0$. Now, let $s \geq 0$ and assume that for each $B_i^s \in G^s$, $\varphi \models B_i^s$. Then using Lemma 1, for each $k \in \{1, \dots, 6\}$ since Δ satisfies **(IC2)** it can be easily checked that for each agent $i \in V$, B_i^{s+1} is consistent. This concludes the proof. ■

Proposition 4
For every $k \in \{1, \dots, 6\}$, R^k_Δ satisfies **(UP)** if Δ satisfies:

- **(IC0)** when $k \in \{5, 6\}$;
- **(Disj)** when $k \in \{1, 3, 4\}$;
- **(IC0)** and **(Disj)** when $k = 2$.

Proof: Let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG from any of the classes $\mathcal{G}(R^k_\Delta)$ where $k \in \{1, \dots, 6\}$. We must prove that G satisfies **(UP)** under the conditions on Δ given within the statement of the proposition. Let φ be a propositional formula from $\mathcal{L}_\mathcal{P}$ and assume that for each $B_i \in B$, $B_i \models \varphi$. We prove that for each $B_i \in B$ and at every step $s \geq 0$, $B_i^s \models \varphi$ by recursion on s. This is trivial for each $k \in \{1, \dots, 6\}$ when $s = 0$. Now, let $s \geq 0$ and assume that for each $B_i^s \in G^s$, $B_i^s \models \varphi$. Then:
• when $k \in \{1, 3\}$, assume that Δ satisfies **(Disj)**. Now, B_i^{s+1} is a belief base of the form $\Delta(\mathcal{K}) = \Delta_\top(\mathcal{K})$ for some profile $\mathcal{K}$ such that $\bigvee \mathcal{K} \models \varphi$ (by the recursion hypothesis). Yet by **(Disj)** we get that $\Delta(\mathcal{K}) \models \bigvee \mathcal{K}$, thus $\Delta(\mathcal{K}) \models \varphi$. Hence, $B_i^{s+1} \models \varphi$;
• when $k = 4$, assume that Δ satisfies **(Disj)**. We know that $\Delta(\langle \mathcal{C}_i^s \rangle) \models \varphi$ (the proof is similar to the one given in the preceding item). So here, B_i^{s+1} is a belief base of the form $\Delta(\langle B_i, K \rangle) = \Delta_\top(\langle B_i, K \rangle)$ where $K \models \varphi$. Since we have as well $B_i \models \varphi$ by the recursion hypothesis, we get from **(Disj)** that $\Delta(\langle B_i, K \rangle) \models \varphi$. Hence, $B_i^{s+1} \models \varphi$;
• when $k = 2$, assume that Δ satisfies **(IC0)** and **(Disj)**. We have $B_i^{s+1} = \Delta_{\Delta(\langle \mathcal{C}_i^s \rangle)}(\langle B_i^s \rangle)$, and we already proved that $\Delta(\mathcal{C}_i^s) \models \varphi$. So by **(IC0)**, we get that $B_i^{s+1} \models \Delta(\mathcal{C}_i^s)$. Hence, $B_i^{s+1} \models \varphi$;
• when $k \in \{5, 6\}$, assume that Δ satisfies **(IC0)**. Here, B_i^{s+1} is a belief base of the form $\Delta_{B_i^s}(\mathcal{K})$ for some profile $\mathcal{K}$. Since $B_i^s \models \varphi$ by the recursion hypothesis and since by **(IC0)** we have $\Delta_{B_i^s}(\mathcal{K}) \models B_i^s$, we get that $\Delta_{B_i^s}(\mathcal{K}) \models \varphi$. Hence, $B_i^{s+1} \models \varphi$.
This concludes the proof. ■

Proposition 5
If Δ satisfies **(IC0)**, then R^5_Δ and R^6_Δ do not satisfy **(Resp)**.

Proof: Let R^k_Δ be any revision policy where $k \in \{5, 6\}$ and where Δ satisfies **(IC0)**. We must prove that R^k_Δ does not satisfy **(Resp)**. That is to say, we must show that there exists a BRG from $\mathcal{G}(R^k_\Delta)$ which does not satisfy **(Resp)**. Then let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG from $\mathcal{G}(R^k_\Delta)$ defined as $V = \{1, 2\}$, $A = \{(1, 2)\}$, $\mathcal{L}_\mathcal{P}$ is the propositional language defined from $\mathcal{P} = \{p\}$, $B_1 = B_1^0 = p$ and $B_2 = B_2^0 = \neg p$. Note that conditions (i) and (ii) in the statement of the responsiveness definition (cf. Definition 10) are satisfied for B_2: we have (i) $B_1^0 \wedge B_2^0 \models \bot$ and (ii) $\bigwedge \{B_1^0\} = B_1^0 \not\models \bot$. But since Δ satisfies **(IC0)**, for both revision policies R^5_Δ and R^6_Δ it is required that $B_2^1 \models B_2^0$, which contradicts $B_2^1 \not\models B_2^0$. Hence, G does not satisfy **(Resp)**, that concludes the proof. ■

Proposition 6
For every $k \in \{1, 2, 4\}$, R^k_Δ satisfies **(Resp)** if Δ satisfies:

- **(IC2)** when $k = 1$;
- **(IC0)** and **(IC2)** when $k = 2$;
- **(IC2)** and **(IC4)** when $k = 4$.

Proof: Let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG from any of the classes $\mathcal{G}(R^k_\Delta)$ where $k \in \{1, 2, 4\}$. We must prove that G satisfies **(Resp)** under the conditions on Δ given within the statement of the proposition. Let $B_i \in B$, $s \in \mathbb{N}$, and assume that (i) $\forall j \in V$, if $(j, i) \in A$ then $B_i^s \wedge B_j^s \models \bot$, and (ii) $\bigwedge_{j \in V} \{B_j^s \mid (j, i) \in A\} \not\models \bot$. We must prove that $B_i^{s+1} \not\models B_i^s$:
• when $k = 1$, assume that Δ satisfies **(IC2)**. By condition (ii), $\bigwedge_{j \in V} \{B_j^s \mid (j, i) \in A\}$ is consistent, so **(IC2)** requires that $B_i^{s+1} = \Delta(\mathcal{C}_i^s) \equiv \bigwedge_{j \in V} \{B_j^s \mid (j, i) \in A\} \equiv \bigwedge_{B_j^s \in \mathcal{C}_j^s} B_j^s$. By condition (i), we get that $B_i^s \wedge \bigwedge_{B_j^s \in \mathcal{C}_i^s} B_j^s \models \bot$, or equivalently, that $B_i^s \wedge B_i^{s+1} \models \bot$. In particular, we get that $B_i^{s+1} \not\models B_i^s$.
• when $k = 2$, assume that Δ satisfies **(IC0)** and **(IC2)**. By the preceding item, by condition (ii) and from **(IC2)** we know that $\Delta(\mathcal{C}_i^s) \equiv \bigwedge_{B_j^s \in \mathcal{C}_j^s} B_j^s$. Now, **(IC0)** requires that $B_i^{s+1} \models \Delta(\mathcal{C}_i^s)$, thus $B_i^{s+1} \models \bigwedge_{B_j^s \in \mathcal{C}_j^s} B_j^s$. Yet by condition (i) we have that $B_i^s \wedge \bigwedge_{B_j^s \in \mathcal{C}_i^s} B_j^s \models \bot$. Hence, $B_i^s \wedge B_i^{s+1} \models \bot$. In particular, we get that $B_i^{s+1} \not\models B_i^s$.
• when $k = 4$, assume that Δ satisfies **(IC2)** and **(IC4)**. We know by the preceding items, by condition (ii) and from **(IC2)** that $\Delta(\mathcal{C}_i^s) \equiv \bigwedge_{B_j^s \in \mathcal{C}_j^s} B_j^s$. So by definition of B_i^{s+1} we have that $B_i^{s+1} = \Delta(B_i^s, \bigwedge_{B_j^s \in \mathcal{C}_j^s} B_j^s)$. Toward a contradiction, assume that $B_i^{s+1} \models B_i^s$. In particular, we have that $B_i^s \wedge B_i^{s+1} \not\models \bot$. Now, **(IC4)** requires that $B_i^{s+1} \wedge \bigwedge_{B_j^s \in \mathcal{C}_i^s} B_j^s \not\models \bot$. This contradicts condition (i) which states that $B_i^s \wedge \bigwedge_{B_j^s \in \mathcal{C}_i^s} B_j^s \models \bot$.
This concludes the proof. ■

Proposition 7
$R^3_{\Delta^{d_H,\Sigma}}$ does not satisfy **(Resp)**.

Proof: We must show that there exists a BRG from $\mathcal{G}(R^3_{\Delta^{d_H,\Sigma}})$ which does not satisfy **(Resp)**. Then let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG from $\mathcal{G}(R^3_{\Delta^{d_H,\Sigma}})$ defined as $V = \{1, 2, 3\}$, $A = \{(1, 3), (2, 3)\}$, $\mathcal{L}_\mathcal{P}$ is the propositional language defined from $\mathcal{P} = \{p, q, r\}$, $B_1 = B_1^0 = p \wedge (q \Leftrightarrow r)$, $B_2 = B_2^0 = q \wedge (p \Leftrightarrow r)$ and $B_3 = B_3^0 = \neg p \wedge \neg q \wedge \neg r$. Note that conditions (i) and (ii) in the statement of the responsiveness definition (cf. Definition 10) are satisfied for B_3: we have (i) $B_1^0 \wedge B_3^0 \models \bot$ and $B_2^0 \wedge B_3^0 \models \bot$, and (ii) $B_1^0 \wedge B_2^0 \equiv p \wedge q \wedge r \not\models \bot$. Yet one can verify that at step 1, we get that $B_3^1 \equiv \Delta^{d_H,\Sigma}(\langle B_1^0, B_2^0, B_3^0\rangle) \equiv \neg p \wedge \neg q \wedge \neg r \equiv B_3^0$, which contradicts $B_3^1 \not\models B_3^0$. Hence, G does not satisfy **(Resp)**, that concludes the proof. ■

Proposition 8
For any aggregation function f, $R^3_{\Delta^{d_D,f}}$ satisfies **(Resp)**.

Proof: Let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG from the class $\mathcal{G}(R^3_{\Delta^{d_D,f}})$. We must prove that G satisfies **(Resp)**. Let $B_i \in B$, $s \in \mathbb{N}$, and assume that (i) $\forall j \in V$, if $(j,i) \in A$ then $B_i^s \wedge B_j^s \models \bot$, and (ii) $\bigwedge_{j \in V}\{B_j^s \mid (j,i) \in A\} \not\models \bot$. We must prove that $B_i^{s+1} \not\models B_i^s$, or equivalently, that there exists an interpretation $\omega \models \Delta(B_i^s, \mathcal{C}_i^s)$ such that for every interpretation $\omega' \models B_i^s$, we have $d_D^f(\omega, \langle B_i^s, \mathcal{C}_i^s\rangle) \leq d_D^f(\omega', \langle B_i^s, \mathcal{C}_i^s\rangle)$. Yet by condition (ii), $\bigwedge_{j \in V}\{B_j^s \mid (j,i) \in A\} \equiv \bigwedge_{B_j^s \in \mathcal{C}_i^s} B_j^s \not\models \bot$. So on the one hand, let $\omega \models \bigwedge_{B_j^s \in \mathcal{C}_i^s} B_j^s$; since $\omega \models B_j^s$ for every $B_j^s \in \mathcal{C}_i^s$, we have $d_D(\omega, B_j^s) = 0$; and by condition (i), $\omega \not\models B_i^s$, so $d_D(\omega, B_i^s) = 1$; thus we get that $d_D^f(\omega, \langle B_i^s, \mathcal{C}_i^s\rangle) = f\{1, \underbrace{0, \ldots, 0}_{|\mathcal{C}_i^s| \text{ times}}\}$. On the other hand, let ω' be any model of B_i^s, i.e., $\omega' \models B_i^s$; we have $d_D(\omega', B_i^s) = 0$, and by condition (i), $\omega' \not\models B_j^s$ for every $B_j^s \in \mathcal{C}_i^s$, so for every $B_j^s \in \mathcal{C}_i^s$, we have $d_D(\omega', B_j^s) = 1$; thus we get that $d_D^f(\omega', \langle B_i^s, \mathcal{C}_i^s\rangle) = f\{0, \underbrace{1, \ldots, 1}_{|\mathcal{C}_i^s| \text{ times}}\}$. From the symmetry and non-decreasingness of f (Konieczny, Lang, and Marquis 2004), we get that $d_D^f(\omega, \langle B_i^s, \mathcal{C}_i^s\rangle) \leq d_D^f(\omega', \langle B_i^s, \mathcal{C}_i^s\rangle)$. This concludes the proof. ■

Proposition 9
For every $k \in \{1, \ldots, 6\}$, $R^k_{\Delta^{d_H,\Sigma}}$ does not satisfy **(Mon)**.

Proof:
We must show that for every $k \in \{1, \ldots, 6\}$, there exists a BRG from $\mathcal{G}(R^k_{\Delta^{d_H,\Sigma}})$ which does not satisfy **(Mon)**. Let us first provide a counter-example for the case where $k \in \{1,2,3,4\}$. Let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG from one of the classes $\mathcal{G}(R^k_{\Delta^{d_H,\Sigma}})$ where $k \in \{1,2,3,4\}$, defined as $V = \{1,2,3,4,5\}$, $A = \{(1,2), (2,3), (2,4), (3,4), (3,5), (4,3), (4,5), (5,3), (5,5)\}$, $\mathcal{L}_\mathcal{P}$ is the propositional language defined from $\mathcal{P} = \{p,q,r,s\}$, and the belief bases B_i, $i \in \{1,2,3,4,5\}$ are defined as follows:

$$\begin{cases} B_1 = p \wedge q \wedge r \wedge s, \\ B_2 = (p \wedge q \wedge r \wedge s) \vee (\neg p \wedge \neg q \wedge \neg r \wedge \neg s), \\ B_3 = (p \wedge \neg q \wedge \neg r \wedge \neg s) \vee (\neg p \wedge \neg q \wedge r \wedge s), \\ B_4 = B_5 = B_3. \end{cases}$$

Lastly, let φ be the formula from $\mathcal{L}_\mathcal{P}$ defined as

$$\varphi = (p \wedge q \wedge r \wedge s) \vee (p \wedge \neg q \wedge \neg r \wedge \neg s).$$

Then one can verify that for each $k \in \{1,2,3,4\}$, whenever each revision policy $R_i \in \mathcal{R}$ is $R_i = R^k_{\Delta^{d_H,\Sigma}}$, the belief sequences associated with the five agents in G (respectively, in $G_{2\to\varphi}$) correspond to the ones given in Table 4 (respectively, Table 5). Both BRGs are stable and all agents reach their belief cycle by at most step 1 in both cases. For both tables, blue cells are associated with beliefs for which φ is a logical consequence; red cells are associated with beliefs for which $\neg\varphi$ is a logical consequence; and white cells correspond to the remaining cases. Accordingly, one can see that φ is unanimously accepted in G, whereas in $G_{2\to\varphi}$, φ is not accepted by agent 3, 4 and 5. This shows that for every $k \in \{1,2,3,4\}$, $R^k_{\Delta^{d_H,\Sigma}}$ does not satisfy **(Mon)**.

step i	B_1^i	B_2^i	B_3
0	$p \wedge q \wedge r$	$p \wedge (q \Leftrightarrow r)$	$\neg p \wedge \neg r$
≥ 1	$p \wedge q \wedge r$	$p \wedge q \wedge r$	$\neg p \wedge \neg q \wedge \neg r$

Table 6: The belief sequences of agents in G' (cf. proof of Proposition 9).

step i	B_1^i	B_2^i	B_3^i
0	$p \wedge q \wedge r$	$p \wedge q \wedge r$	$\neg p \wedge \neg r$
≥ 1	$p \wedge q \wedge r$	$p \wedge q \wedge r$	$\neg p \wedge q \wedge \neg r$

Table 7: The belief sequences of agents in $G'_{2\to\varphi'}$ (cf. proof of Proposition 9).

Now, let us provide a counter-example for the case where $k \in \{5,6\}$. Let $G' = (V', A', \mathcal{L}_{\mathcal{P}'}, B', \mathcal{R}')$ be a BRG from one of the classes $\mathcal{G}(R^k_{\Delta^{d_H,\Sigma}})$ where $k \in \{5,6\}$, defined as $V' = \{1,2,3\}$, $A' = \{(1,2), (2,3)\}$, $\mathcal{L}_{\mathcal{P}'}$ is the propositional language defined from $\mathcal{P}' = \{p,q,r\}$, and the belief bases B_i', $i \in \{1,2,3\}$ are defined as follows:

$$\begin{cases} B_1' = p \wedge q \wedge r, \\ B_2' = p \wedge (q \Leftrightarrow r), \\ B_3' = \neg p \wedge \neg r. \end{cases}$$

Lastly, let φ' be the formula from $\mathcal{L}_{\mathcal{P}'}$ defined as

$$\varphi' = (p \wedge q \wedge r) \vee (\neg p \wedge \neg q \wedge \neg r).$$

Then one can verify that for each $k \in \{5,6\}$, whenever each revision policy $R_i \in \mathcal{R}$ is $R_i = R^k_{\Delta^{d_H,\Sigma}}$, the belief sequences associated with the three agents in G' (respectively, in $G'_{2\to\varphi'}$) correspond to the ones given in Table 6 (respectively, Table 7). Both BRGs are stable and all agents reach their belief cycle by at most step 1 in both cases. The cell colors of tables 6 and 7 follow similar rules as for tables 4 and 5: blue cells are associated with beliefs for which φ' is a logical consequence; red cells are associated with beliefs for which $\neg\varphi'$ is a logical consequence; and white cells correspond to the remaining cases. Accordingly, one can see that φ' is unanimously accepted in G', whereas in $G'_{2\to\varphi'}$, φ is not accepted by agent 3. This shows that for every $k \in \{5,6\}$, $R^k_{\Delta^{d_H,\Sigma}}$ does not satisfy **(Mon)**.
This concludes the proof. ■

Proposition 10
The revision policies R^5_Δ and R^6_Δ satisfy **(Conv)** if Δ satisfies **(IC0)**.

step i	B_1^i	B_2^i	$B_3^i \equiv B_4^i \equiv B_5^i$
0	$p \wedge q \wedge r \wedge s$	$(p \wedge q \wedge r \wedge s) \vee (\neg p \wedge \neg q \wedge \neg r \wedge \neg s)$	$(p \wedge \neg q \wedge \neg r \wedge \neg s) \vee (\neg p \wedge \neg q \wedge r \wedge s)$
≥ 1	$p \wedge q \wedge r \wedge s$	$p \wedge q \wedge r \wedge s$	$p \wedge \neg q \wedge \neg r \wedge \neg s$

Table 4: The belief sequences of agents in G (cf. proof of Proposition 9).

step i	B_1^i	B_2^i	$B_3^i \equiv B_4^i \equiv B_5^i$
0	$p \wedge q \wedge r \wedge s$	$p \wedge q \wedge r \wedge s$	$(p \wedge \neg q \wedge \neg r \wedge \neg s) \vee (\neg p \wedge \neg q \wedge r \wedge s)$
≥ 1	$p \wedge q \wedge r \wedge s$	$p \wedge q \wedge r \wedge s$	$\neg p \wedge \neg q \wedge r \wedge s$

Table 5: The belief sequences of agents in $G_{2\to\varphi}$ (cf. proof of Proposition 9).

Proof: Let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG from any of the classes $\mathcal{G}(R^k_\Delta)$ where $k \in \{5, 6\}$ and where Δ satisfies **(IC0)**. We must prove that G is stable. Toward a contradiction, assume that for some agent $i \in V$, $|Cyc(B_i)| \geq 2$. Then there exist two belief bases B_i^s, $B_i^{s'}$ in $Cyc(B_i)$ such that $B_i^s \not\models B_i^{s'}$. This contradicts **(IC0)** which requires that at every step $s \in \mathbb{N}$, we have $B_i^{s+1} \models B_i^s$. This concludes the proof. ■

Proposition 11
For every $k \in \{1, 2, 3, 4\}$, R^k_Δ does not satisfy **(Conv)** if Δ satisfies:

- **(IC2)** when $k = 1$;
- **(IC0)** and **(IC2)** when $k = 2$;
- **(IC1)**, **(IC2)** and **(IC4)** when $k \in \{3, 4\}$.

Proof: For every $k \in \{1, 2, 3, 4\}$, we must prove that R^k_Δ does not satisfy **(Conv)** under the conditions on Δ given within the statement of the proposition. That is to say, for each revision policy we must show that there exists a BRG from $\mathcal{G}(R^k_\Delta)$ which is not stable. Then let $G = (V, A, \mathcal{L}_\mathcal{P}, B, \mathcal{R})$ be a BRG from $\mathcal{G}(R^k_\Delta)$ defined as $V = \{1, 2, 3\}$, $A = \{(1, 2), (2, 3), (3, 1)\}$, $\mathcal{L}_\mathcal{P}$ is the propositional language defined from $\mathcal{P} = \{p\}$, $B_1 = B_1^0 = p$, $B_2 = B_2^0 = \neg p$ and $B_3 = B_3^0 = \top$. Then:
• when $k = 1$, assume that Δ satisfies **(IC2)**. Then one can verify that for every step $s \in \mathbb{N}$, we have $B_1^{s+1} \equiv B_3^s$, $B_2^{s+1} \equiv B_1^s$ and $B_3^{s+1} \equiv B_2^s$. This means that for each agent $i \in V$, $|Cyc(B_i)| = 3$. Hence, G is not stable;
• when $k = 2$, assume that Δ satisfies **(IC0)** and **(IC2)**. One can verify that at step 1, $B_1^1 \equiv B_2^1 \equiv p$ and $B_3^1 \equiv \neg p$. From then, for every step $s \geq 1$, we have $B_1^{s+1} \equiv B_3^s$, $B_2^{s+1} \equiv B_1^s$ and $B_3^{s+1} \equiv B_2^s$. This means that for each agent $i \in V$, $|Cyc(B_i)| = 3$. Hence, G is not stable;
• when $k \in \{3, 4\}$, assume that Δ satisfies **(IC1)**, **(IC2)** and **(IC4)**. At step 1, one can verify that for $k \in \{3, 4\}$, $B_1^1 \equiv p$, $B_2^1 \equiv \top$ and $B_3^1 \equiv \neg p$. Similarly, the evolution of beliefs for steps $m \geq 2$ can be completely determined Doing so, one can verify that for $k \in \{3, 4\}$, $G^6 = G^0$ and that $G^s \neq G^0$ for every $s \in \{1, \ldots, 5\}$. This means that for each agent $i \in V$, $|Cyc(B_i)| = 6$. Hence, G is not stable.
This concludes the proof. ■

Proposition 12
For $k \in \{1, 2, 3, 4\}$, all directed acyclic BRGs from $\mathcal{G}(R^k_\Delta)$ satisfy **(Conv)** when $k = 1$ or if:

- when $k = 2$, Δ satisfies **(IC0)** and **(IC2)**;
- when $k = 3$, Δ is a distance-based merging operator;
- when $k = 4$, Δ satisfies **(IC2)**, **(IC4)** and **(Disj)**, or Δ is a distance-based merging operator.

Proof: Let $G = (V, A, L, B, R^k_\Delta)$ be a directed acyclic BRG where $k \in \{1, 2, 3, 4\}$. We must prove that G satisfies **(Conv)** under the conditions on Δ given within the statement of the proposition. Beforehand, since the graph (V, A) does not contain any cycle, there must exist a non-empty subset of agents $V_{root} \subseteq V$ such that for each $i \in V_{root}$, $in(i) = 0$, that is, such that there is no $j \in V$ such that $(j, i) \in A$. Then one can associate with each agent $i \in V$ a number $depth(i) \in \mathbb{N}$ which corresponds to the highest number r such that the series $(j_0, j_1), (j_1, j_2), \ldots, (j_{r-1}, j_r = i)$ satisfies for each $x \in \{0, \ldots, r-1\}$, $(j_x, j_{x+1}) \in A$ and $j_0 \in V_{root}$. For each $d \in \mathbb{N}$, let us denote $V_d = \{i \in V \mid depth(i) = d\}$. Let $d_{max} = \max\{d \in \mathbb{N} \mid V_d \neq \emptyset\}$. Note that such a number d_{max} exists since each agent $i \in V$ can be associated with a number $depth(i)$.

Now, it is enough to show that for every $d \in \{0, \ldots, d_{max}\}$ and for each agent $i \in V_d$, we have $|Cyc(B_i)| = 1$. We prove it by recursion on d. For each agent $i \in V_0$, we have $in(i) = 0$. By Definition 1, R_i is the identity function, thus we trivially get that $|Cyc(B_i)| = 1$. Now, let $d \geq 0$ and assume that for each agent $i \in V_d$, we have $|Cyc(B_i)| = 1$. Let $i \in V_{d+1}$. We know that for each $B_j \in \mathcal{C}_i$, $j \in V_{d'}$ for some $d' < d$, so that $|Cyc(B_j)| = 1$ by the recursion hypothesis. So let us denote $s_{max} \in \mathbb{N}$ the smaller (step) number which satisfies $B_j^{s_{max}} = B_j^{s_{max}-1}$ for every $B_j \in \mathcal{C}_i$. We have that for each step $s \geq s_{max}$, $\mathcal{C}_i^{s+1} = \mathcal{C}_i^s$. Then:
• when $k = 1$, for each step $s \geq s_{max}$, we directly get that $B_i^{s+1} = \Delta(\mathcal{C}^{s_{max}})$. Hence, $|Cyc(B_i)| = 1$;
• when $k = 2$, assume that Δ satisfies **(IC0)** and **(IC2)**. We have $B_i^{s_{max}+1} = \Delta_{\Delta(\mathcal{C}^{s_{max}})}(\langle B_i^{s_{max}} \rangle)$. By **(IC0)** we get that $B_i^{s_{max}+1} \models \Delta(\mathcal{C}^{s_{max}})$. This means that $B_i^{s_{max}+1} \wedge \Delta(\mathcal{C}^{s_{max}}) \not\models \perp$. In this case, we have that $B_i^{s_{max}+2} \equiv \Delta_{\Delta(\mathcal{C}^{s_{max}})}(\langle B_i^{s_{max}+1} \rangle)$ which is equivalent to $B_i^{s_{max}+1} \wedge \Delta(\mathcal{C}^{s_{max}})$ by **(IC2)**. Then for each $s \geq s_{max} + 2$, we have that $B_i^{s+1} = B_i^{s_{max}+2} \wedge \Delta(\mathcal{C}^{s_{max}})$. Thus $|Cyc(B_i)| = 1$;
• let $k = 4$ and assume that Δ satisfies **(IC4)** and **(Disj)**. We have $B_i^{s_{max}+1} = \Delta(B_i^{s_{max}}, \Delta(\mathcal{C}^{s_{max}}))$. By **(Disj)**

we get that $B_i^{s_{max}+1} \models B_i^{s_{max}} \vee \Delta(\mathcal{C}^{s_{max}})$. Then by **(IC4)** we have that $B_i^{s_{max}+1} \wedge \Delta(\mathcal{C}^{s_{max}}) \not\models \bot$. Now, by **(IC2)**, $B_i^{s_{max}+2} \equiv B_i^{s_{max}+1} \wedge \Delta(\mathcal{C}^{s_{max}})$. Then for each $s \geq s_{max}+2$, we have that $B_i^{s+1} = B_i^{s_{max}+2} \wedge \Delta(\mathcal{C}^{s_{max}})$. Thus $|Cyc(B_i)| = 1$;

• let $k \in \{3,4\}$ and assume that Δ is a distance-based merging operator. That is to say, $\Delta = \Delta^{d,f}$ for some pseudo-distance d between interpretations and some aggregation function f. We prove that $|Cyc(B_i)| = 1$ in the case where $k = 3$ (the proof is similar for the case where $k = 4$). Let us show that for each step $s \geq s_{max}$, $B_i^{s+1} \models B_i^s$. Let $\omega \models B_i^{s+1}$. We have $\omega \models \Delta^{d,f}(\langle B_i^s, \mathcal{C}^{s_{max}}\rangle)$, which means that for every interpretation ω', $d^f(\omega, \langle B_i^s, \mathcal{C}^{s_{max}}\rangle) \leq d^f(\omega', \langle B_i^s, \mathcal{C}^{s_{max}}\rangle)$. In particular, for every $\omega' \models B_i^s$, we have that $d^f(\omega, \langle B_i^s, \mathcal{C}^{s_{max}}\rangle) \leq d^f(\omega', \langle B_i^s, \mathcal{C}^{s_{max}}\rangle)$, or equivalently that $f\{d(\omega, B_i^s), d(\omega, B_{i_1}^{s_{max}}), \ldots, d(\omega, B_{i_{in(i)}}^{s_{max}}\} \leq f\{d(\omega', B_i^s), d(\omega, B_{i_1}^{s_{max}}), \ldots, d(\omega, B_{i_{in(i)}}^{s_{max}}\}$ (we denote $\mathcal{C}_i^{d_{max}} = \langle B_{i_1}^{s_{max}}, \ldots, B_{i_{in(i)}}^{s_{max}}\rangle$). On the one hand, since $\omega' \models B_i^s$, by the identity of indiscernibles property of the pseudo-distance d, we have $d(\omega', B_i^s) = 0$. Thus $f\{d(\omega, B_i^s), d(\omega, B_{i_1}^{s_{max}}), \ldots, d(\omega, B_{i_{in(i)}}^{s_{max}}\} \leq f\{0, d(\omega, B_{i_1}^{s_{max}}), \ldots, d(\omega, B_{i_{in(i)}}^{s_{max}}\}$. Since the aggregation function f is required to satisfy the non-decreasingness condition, we get that $d(\omega, B_i^s) = 0$. So by the identity of indiscernibles property of the pseudo-distance d, we get that $\omega \models B_i^s$. We just proved that for each step $s \geq s_{max}$, $B_i^{s+1} \models B_i^s$. Then assume toward a contradiction that $|Cyc(B_i)| \geq 2$. Then there must exist two belief bases $B_i^s, B_i^{s'} \in Cyc(B_i)$ such that $B_i^s \not\models B_i^{s'}$. This contradicts the fact that for each step $s \geq s_{max}$, $B_i^{s+1} \models B_i^s$. Hence, $|Cyc(B_i)| = 1$.

This concludes the proof. ■