1F2-4

On Opinion Control in Belief Revision Games

Nicolas Schwind[*1*2] Katsumi Inoue[*2*3*4] Gauvain Bourgne[*5] Sébastien Konieczny[*6]
Pierre Marquis[*6]

[*1] National Institute of Advanced Industrial Science and Technology, Tokyo, Japan
[*2] National Institute of Informatics, Tokyo, Japan
[*3] The Graduate University for Advanced Studies, Tokyo, Japan
[*4] Tokyo Institute of Technology, Tokyo, Japan
[*5] CNRS & Sorbonne Universités, UPMC Univ Paris 06, UMR 7606, LIP6, F-75005, Paris, France
[*6] CRIL - CNRS, Université d'Artois, Lens, France

We consider the problem of belief propagation in a network of communicating agents, modeled in the recently introduced Belief Revision Game (BRG) framework. In this setting, each agent expresses her belief through a propositional formula and revises her own belief at each step by considering the beliefs of her acquaintances, using belief change tools. We study the extent to which BRGs satisfy some monotonicity properties, i.e., whether promoting some desired piece of belief to a given set of agents is actually always useful for making it accepted by all of them. We show that some basic monotonicity properties are not satisfied by BRGs in general, even when the agents merging-based revision policies are rational (in the AGM sense), but we identify a class where they hold.

1. Introduction

We are interested in the issue of monotonicity in a multi-agent system, represented as a Belief Revision Game (BRG) [9]. The BRG setting is a framework for modeling the belief dynamics of a group of agents V, for instance agents involved in a social network. A BRG is a dynamical system where agents have their own belief bases (representing their belief states), and communicate synchronously with their acquaintances. The acquaintance relationship is given through a binary relation A over V, i.e., (V, A) is a graph. Let us introduce a motivating example:

Example 1. *Consider a group of friends Alex, Beth and Chris who are discussing on whether they should trust the quality of the food served in a given restaurant. Alex and Beth know each other, Alex and Chris as well, but Beth and Chris do not know each other. Two meals are considered by them. At the beginning, Alex believes that either the first meal or the second one is healthy, but not both of them; Beth believes that none of the two meals is healthy; whereas Chris believes that at least one of the two meals is healthy.*

At each communication step, each agent revises her beliefs by considering her acquaintances' beliefs. Several merging-based revision policies have been defined, each of them reflecting how much an agent is ready to question her current beliefs in front of her acquaintances' beliefs. Then a piece of belief φ is accepted by an agent i of V when there exists a step of the game from which φ holds in the belief bases of i at each subsequent step; and φ is unanimously accepted when it is accepted by every agent of V.

Now, given a piece of belief φ, is adding "more φ" in a BRG always beneficial to φ? More precisely, whenever a piece of belief φ is unanimously accepted in a BRG, is it always harmless to replace at the beginning some bases by φ, or more generally, by a base that is "closer" to φ, i.e., by a "promotion" of φ? This monotonicity condition is essential when one wants to investigate the potential manipulation of such sytems, in particular the control issue: consider a set of agents from a predefined subset C of V and an additional agent M who can "control" the agents from C, i.e., M can modify the initial beliefs of agents from C, then is it possible for M to make a piece of belief unanimously accepted? Such a control issue is significant for a number of multi-agent problems, including brand crisis management [2]; in such applications, it is useful to know what information agents from C should convey to their acquaintances in order to avoid the propagation of negative perceptions.

2. Belief Revision Games

Let $\mathcal{L}_\mathcal{P}$ be a propositional language built up from a finite set of propositional variables $\mathcal{P}$ and the usual connectives, including $\oplus$, the xor connective. $\bot$ (resp. $\top$) is the Boolean constant always false (resp. true). Formulae are interpreted in a standard way. $Mod(\varphi)$ denotes the set of models of the formula φ, $\models$ 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 *belief base* B denotes the set of beliefs of an agent, it is a finite set of propositional formulae interpreted conjunctively, so that B is identified with the conjunction of its elements. A *profile* $\mathcal{C} = \langle B_1, \ldots, B_n \rangle$ is a finite vector of belief bases. A Belief Revision Game (BRG for short) is formalized as follows [9]:

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

- *$V = \{1, \ldots, n\}$ is a finite set of agents;*
- *$A \subseteq V \times V$ is an irreflexive binary relation on V which represents the set of* acquaintances *between the agents;*
- *$\mathcal{L}_\mathcal{P}$ is a finite propositional language;*
- *$\mathcal{B}$ is a mapping from V to $\mathcal{L}_\mathcal{P}$ where for each $i \in V$, $B(i)$ (noted B_i) is the initial belief base of agent i;*

- $\mathcal{R} = \{R_1, \ldots, R_n\}$, *where each* R_i *is the revision policy of agent* i, *i.e., 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 if* $in(i) = 0$, *then* R_i *is the identity function.*

Let $G = (V, A, \mathcal{L}_\mathcal{P}, \mathcal{B}, \mathcal{R})$ be a BRG and let us denote $\mathcal{C}_i$ the *context* of i, defined as the profile $\mathcal{C}_i = \langle B_{i_1}, \ldots, B_{i_{in(i)}} \rangle$ where $\{i_1, \ldots, i_{in(i)}\} = \{i_j \mid (i_j, i) \in A\}$. Then $R_i(B_i, \mathcal{C}_i)$ is the belief base of agent i once revised by taking into account her own current beliefs B_i and her current context $\mathcal{C}_i$.

In a BRG, the beliefs of each agent evolve at each time step using her revision policy. This induces for each $i \in V$ a *belief sequence* $(B_i^s)_{s \in \mathbb{N}}$ where B_i^s denotes the belief base of agent i after s steps, defined as $B_i^0 = B_i$ and for each $s \geq 0$, $B_i^{s+1} = R_i(B_i^s, \mathcal{C}_i^s)$,[*1] where $\mathcal{C}_i^s$ is the context of i at step s. In [9] we showed that in any BRG, the belief sequence of each agent i is *cyclic*, i.e., in $(B_i^s)_{s \in \mathbb{N}}$ there exists a finite subsequence $B_i^b, \ldots, B_i^e$ such that for every $j > e$, we have $B_j^j \equiv B_i^{b+((j-b) mod (e-b+1))}$; the *belief cycle* $Cyc(i)$ of an agent i corresponds to the series of this subsequence of belief bases $Cyc(i) = B_i^b, B_i^{b+1}, \ldots, B_i^e$. As we are interested in determining the pieces of beliefs resulting from the interaction of the agents, we focus on the *outcome* of each agent i in G, denoted $Acc_G(i)$ and defined as:

$$Acc_G(i) = \bigvee \{B_i^s \mid B_i^s \in Cyc(i)\}.$$

We say that a formula φ is *accepted* by i in G if $Acc_G(i) \models \varphi$, which means that φ is a logical consequence of every belief base in the belief cycle of i. *G converges at step s* if for each $i \in V$, $B_i^{s+1} = B_i^s$.

The formalization of a BRG allows each agent i to consider any revision policy $R_i \in \mathcal{R}$. However, one can take advantage of theoretical tools from Belief Change Theory (see e.g. [1]), more precisely, belief revision and merging operators. A merging operator Δ associates any formula μ (the *integrity constraint*) and any profile $\mathcal{C}$ with a new formula $\Delta_\mu(\mathcal{C})$ (the *merged base*). A merging operator Δ aims at defining the merged base as the beliefs of a group of agents represented by the profile, under some integrity constraint. Standard properties (denoted **(IC0)**–**(IC8)**) are expected for merging operators, and such operators are called *IC merging operators* (see [7] for more details).

IC merging operators include some *distance-based* operators, i.e., operators that are based on the selection of some models of the integrity constraint, the "closest" ones to the given profile. These operators are characterized by a distance d between interpretations and an aggregation function f [5]. They associate with every formula μ and every profile $\mathcal{C}$ a belief base $\Delta_\mu^{d,f}(\mathcal{C})$ satisfying $Mod(\Delta_\mu^{d,f}(\mathcal{C})) = \min(Mod(\mu), \leq_\mathcal{C}^{d,f})$, where $\leq_\mathcal{C}^{d,f}$ is the total preorder over interpretations induced by $\mathcal{C}$ defined by $\omega \leq_\mathcal{C}^{d,f} \omega'$ if and only if $d^f(\omega, \mathcal{C}) \leq d^f(\omega', \mathcal{C})$, where $d^f(\omega, \mathcal{C}) = f_{B \in \mathcal{C}}\{d(\omega, B)\}$ and $d(\omega, B) = \min_{\omega' \models B} d(\omega, \omega')$. 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). Using aggregation functions such as Σ and GMax lead to IC merging operators. For instance, GMax operators consider for each profile $\mathcal{C}$ the total preorder over interpretations $\leq_\mathcal{C}^{d,\mathsf{GMax}}$ defined by $\omega \leq_\mathcal{C}^{d,\mathsf{GMax}} \omega'$ if and only if $d^{\mathsf{GMax}}(\omega, \mathcal{C}) \leq^{lex} d^{\mathsf{GMax}}(\omega', \mathcal{C})$ (where $\leq^{lex}$ is the lexicographic ordering induced by the natural order) and $d^{\mathsf{GMax}}(\omega, \mathcal{C})$ is the vector of numbers $d_1, \ldots, d_n$ obtained by sorting in a non-increasing order the vector $\langle d(\omega, B_i) \mid B_i \in \mathcal{C} \rangle$. Lastly, *belief revision operators* can be seen as belief merging operators applied to singleton profiles: indeed, if Δ is an IC merging operator then the revision operator $\circ_\Delta$ *induced by* Δ defined for all bases B_1, B_2 as $B_1 \circ_\Delta B_2 = \Delta_{B_2}(\langle B_1 \rangle)$ satisfies the standard AGM revision postulates [1, 4].

steps s	B_1^s	B_2^s	B_3^s
0	$p_1 \oplus p_2$	$\neg p_1 \wedge \neg p_2$	$p_1 \vee p_2$
$2k+1$	$\neg p_1 \vee \neg p_2$	$p_1 \oplus p_2$	$p_1 \oplus p_2$
$2k+2$	$p_1 \oplus p_2$	$\neg p_1 \vee \neg p_2$	$\neg p_1 \vee \neg p_2$

Table 1: The belief sequences of Alex, Beth and Chris.

Let us go back to BRGs. Six classes of revision policies have been proposed in [9]. Each of them, denoted R_Δ^k ($k \in \{1, \ldots, 6\}$) is parameterized by an IC merging operator Δ. Each class is defined as follows,[*2] at each step s and for any agent i such that $\mathcal{C}_i \neq \emptyset$:

- $R_\Delta^1(B_i^s, \mathcal{C}_i^s) = \Delta(\mathcal{C}_i^s)$;
- $R_\Delta^2(B_i^s, \mathcal{C}_i^s) = \Delta_{\Delta(\mathcal{C}_i^s)}(\langle B_i^s \rangle) \quad [= B_i^s \circ_\Delta \Delta(\mathcal{C}_i^s)]$;
- $R_\Delta^3(B_i^s, \mathcal{C}_i^s) = \Delta(\langle B_i^s, \mathcal{C}_i^s \rangle)$;
- $R_\Delta^4(B_i^s, \mathcal{C}_i^s) = \Delta(\langle B_i^s, \Delta(\mathcal{C}_i^s) \rangle)$;
- $R_\Delta^5(B_i^s, \mathcal{C}_i^s) = \Delta_{B_i^s}(\Delta(\mathcal{C}_i^s)) \quad [= \Delta(\mathcal{C}_i^s) \circ_\Delta B_i^s]$;
- $R_\Delta^6(B_i^s, \mathcal{C}_i^s) = \Delta_{B_i^s}(\mathcal{C}_i^s)$.

Example 1 (continued). *We consider the BRG* $G_* = (V_*, A_*, \mathcal{L}_{P_*}, \mathcal{B}_*, \mathcal{R}_*)$ *defined as follows.* $V_* = \{1, 2, 3\}$ *where* 1 *corresponds to Alex,* 2 *to Beth, and* 3 *to Chris.* $A_* = \{(1,2), (2,1), (1,3), (3,1)\}$ *expresses that Alex knows Beth and vice-versa, and Alex knows Chris and vice-versa, but Beth and Chris are not connected.* $\mathcal{L}_{\mathcal{P}_*}$ *is the propositional language defined from the variables* $\mathcal{P}_* = \{p_1, p_2\}$, *where* p_1 *stands for "the first meal is healthy" and* p_2 *means "the second meal is healthy." The initial beliefs of the group members are* $B_1 = p_1 \oplus p_2$, $B_2 = \neg p_1 \wedge \neg p_2$ *and* $B_3 = p_1 \vee p_2$. *Assume that all agents use the same revision policy,* $R_i = R^1_{\Delta^{d_H, \mathsf{GMax}}}$ *for each* $i \in V_*$. *The belief sequences associated with the three agents are given in Table 1: the belief cycle of agent* 1 *(resp.* 2, 3*) is given by* (B_1^0, B_1^1) *(resp.* (B_2^1, B_2^2), (B_3^1, B_3^2)*). We have for each* $i \in V_*$, $Cyc(i) = p_1 \oplus p_2, \neg p_1 \vee \neg p_2$ *and* $Acc_G(i) = \neg p_1 \vee \neg p_2$.

In [9] we studied the extent to which BRGs satisfy some basic logical properties depending on the class of revision policies used by the agents. We focused on the case where all agents use the same revision policy ranging over R_Δ^k, $k \in \{1, \ldots, 6\}$. It turned out that when the revision policy is induced from the merging operator $\Delta^{d_D, \Sigma}$, i.e., the merging operator based on the drastic distance and the summation function, the underlying BRGs satisfy a number of expected properties [9]. We also developed a software available online at *http://www.cril.fr/brg/brg.jar*, allowing one to play BRGs with various revision policies.

*1 Abusing notations, a context $\mathcal{C}_i = \langle B_{i_1}, \ldots, B_{i_{in(i)}} \rangle$ is here identified with the sequence $B_{i_1}, \ldots, B_{i_{in(i)}}$.

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

3. On the Notion of Promotion

Belief control in a multi-agent system can take various forms, depending on the meaning given to "control." Here, we are specifically interested in control strategies that consist in promoting a certain belief φ in the beliefs of the agents. A key issue to be addressed is then to determine what "promoting" precisely means in this context. A simple view is to consider that promoting φ in the belief base B of an agent consists in replacing B by a base equivalent to φ. While such a drastic way of promoting φ makes sense, it is not the only one. Thus, for instance, revising B by φ is another approach to do the job. Considering the whole spectrum of promotion techniques is interesting because in some scenarios it could be the case that the agent under consideration can be ready to promote φ by revising her own beliefs B with it, while she would be reluctant in questioning her whole base B and replacing it by φ. In a bribery context, she could for instance ask much more money to accept to change her base B to φ than to change it to the revision of B by φ.

We now characterize the notion of "promotion" of φ thanks to a preorder $\preceq_\varphi$ over belief bases which intuitively reflects the closeness relationship to φ; thus, $B' \preceq_\varphi B$ states that B' is at least as close to φ as B. On this ground, promoting B consists in replacing B by any B' satisfying $B' \preceq_\varphi B$, which roughly means that B' and φ are closer to each other than B and φ are:

Definition 2 (φ-promotion)**.** *For every formula $\varphi \in \mathcal{L}_\mathcal{P}$, the φ-promotion relation is the binary relation $\preceq_\varphi$ on $\mathcal{L}_\mathcal{P} \times \mathcal{L}_\mathcal{P}$ defined for all belief bases B, B' as $B' \preceq_\varphi B$ if and only if $B \wedge \varphi \models B' \models B \vee \varphi$.*

Obviously, one can check that for any $\varphi \in \mathcal{L}_\mathcal{P}$, the binary relation $\preceq_\varphi$ is reflexive and transitive. More generally:

Proposition 1. *For every formula $\varphi \in \mathcal{L}_\mathcal{P}$, $(\mathcal{L}_\mathcal{P}, \preceq_\varphi)$ is a Boolean lattice $(\mathcal{L}_\mathcal{P}, \sqcap_\varphi, \sqcup_\varphi, \dot{\neg}, \varphi, \neg\varphi)$.*

Every base B' promoting φ in B satisfies $\varphi \preceq_\varphi B' \preceq_\varphi B$. Thus B is (up to logical equivalence) the greatest formula w.r.t. $\preceq_\varphi$ promoting φ in B, and φ is (up to logical equivalence) the least formula w.r.t. $\preceq_\varphi$ promoting φ in B. Stated otherwise, the least demanding promotion of φ w.r.t. B consists in letting B unchanged, while the promotion of φ w.r.t. B leading to a formula as close as possible to φ consists in replacing B by φ.

The following model-theoretic characterization of the notion of promotion can be derived easily. $\blacktriangle$ denotes the symmetric difference between sets:

Proposition 2. *Given a formula φ, let B and B' be two belief bases such that B' promotes φ w.r.t. B. Then $\exists S \subseteq Mod(B) \blacktriangle Mod(\varphi)$ such that $Mod(B') = (Mod(B) \cap Mod(\varphi)) \cup S$.*

This proposition also illustrates the fact that B and φ plays symmetric role in the notion of promotion, i.e., $B' \preceq_\varphi B$ if and only if $B' \preceq_B \varphi$.

When a promotion of φ in B is achieved, the set of interpretations assigning different truth values to the base B and to φ may only diminish. Formally:

Proposition 3. *Given a formula φ, let B and B' be two belief bases such that B' promotes φ in B. Then $Mod(B') \blacktriangle Mod(\varphi) \subseteq Mod(B) \blacktriangle Mod(\varphi)$.*

We define a *promotion operator* $\odot$ as a mapping from $\mathcal{L}_\mathcal{P} \times \mathcal{L}_\mathcal{P}$ to $\mathcal{L}_\mathcal{P}$ such that $\psi \odot \varphi \preceq_\varphi \psi$.

We can now lift the relation of formula promotion to BRGs defined on the same set of variables V, acquaintance relation A, propositional language $\mathcal{L}_\mathcal{P}$ and revision policies R. Given two BRGs $G = (V, A, \mathcal{L}_\mathcal{P}, \mathcal{B}, \mathcal{R})$ and $G' = (V, A, \mathcal{L}_\mathcal{P}, \mathcal{B}', \mathcal{R})$ and a formula φ, we note $G' \preceq_\varphi G$ if and only if for each agent $i \in V$, $B'_i \preceq_\varphi B_i$. Finally, we note $G \odot \varphi$ any BRG G' such that $G' \preceq_\varphi G$. Observe that such a promotion operation of φ in G can be non-uniform, i.e., it is not necessarily the case that the promotion of φ in distinct bases of G' is achieved thanks to the same promotion operator. For instance, it can be the case that $B'_i = B_i$ for agent $i \in V$, while $B'_j = B_j \circ \varphi$ for agent $j \in V$, and $B'_k = \varphi$ for agent $k \in V$.

4. On Monotonicity in BRGs

In this section, we focus on the issue of monotonicity for BRGs instantiated with revision policies from the six classes defined in the previous section. Given a merging operator Δ and $E \subseteq \{1, \ldots, 6\}$ R^E_Δ denotes the set $\{R^k_\Delta \mid k \in E\}$. We use the simpler notation R^k_Δ instead of R^E_Δ when $E = \{k\}$. Given a class $\mathcal{G}$ of BRGs and $E \subseteq \{1, \ldots, 6\}$, $\mathcal{G}(R^E_\Delta)$ is the subclass of all BRGs $(V, A, \mathcal{L}_\mathcal{P}, \mathcal{B}, \mathcal{R})$ from $\mathcal{G}$ where for each $R_i \in \mathcal{R}$, $R_i \in R^E_\Delta$. Additionally, a set of revision policies R^E_Δ is said to satisfy a given property P on a given class $\mathcal{G}$ of BRGs if all BRGs from $\mathcal{G}(R^E_\Delta)$ satisfy P.

Given a BRG $G = (V, A, \mathcal{L}_\mathcal{P}, \mathcal{B}, \mathcal{R})$, a subset C of V of so-called "controllable agents" whose initial beliefs can be modified, and a formula φ, one is interested in determining how to modify the belief bases of agents in C in order to make φ unanimously accepted in the resulting game (when possible). The objective is thus to determine a successful "control strategy" to be implemented in order to reach the goal when it can be reached. A control strategy for G given C takes here the form of a mapping σ from C to $\mathcal{L}_\mathcal{P}$, stating for each $i \in C$ that B_i must be replaced by $\sigma(i)$; it is successful for φ if φ is unanimously accepted in the BRG obtained by applying σ to G. Typically, one wants to minimize the number of agents in C to be controlled, but the optimization problem under consideration can be much more complex (for instance, it may consider the cost of controlling each agent in C which is not always uniform).

Among the potential control strategies is the *basic strategy* σ_φ for G given C defined for any $i \in C$ by $\sigma(i) \equiv \varphi$: it simply amounts to promoting φ as much as possible in the belief base of each agent from C. We have the following surprising result:

Proposition 4. *Given a BRG $G = (V, A, \mathcal{L}_\mathcal{P}, \mathcal{B}, \mathcal{R})$, a set $C \subseteq V$ of controllable agents, and a formula φ, it can be the case that the basic strategy σ_φ for G given C is not successful for φ, while a control strategy for G given C which is successful for φ exists.*

Example 1 (continued)**.** *Assume that the goal of the restaurant manager is to convince all protagonists that at least one of the meals is healthy, i.e., $\varphi = p_1 \vee p_2$. Note that at the beginning, Chris' beliefs coincide with this goal (since $B_3 = \varphi$) and that φ is not unanimously accepted. So if Chris is the only controllable agent, the basic strategy σ_φ is not successful. However, when we replace Chris' beliefs by $p_1 \wedge p_2$ instead, we get that $Acc_{G_*}(i) = p_1 \oplus p_2$ for each*

steps s	B_1^s	B_2^s	B_3^s
0	$p_1 \oplus p_2$	$\neg p_1 \wedge \neg p_2$	$p_1 \wedge p_2$
$s \geq 1$	$p_1 \oplus p_2$	$p_1 \oplus p_2$	$p_1 \oplus p_2$

Table 2: An example of control strategy for G_* where $\varphi = p_1 \vee p_2$ is unanimously accepted.

$i \in V$, so $p_1 \vee p_2$ is unanimously accepted (see Table 2). This shows that control is possible here given $C = \{3\}$ as the only controllable agent, but not with the basic strategy.

This illustrates the complexity of the controllability issue for BRGs. It is therefore important to identify some conditions on BRGs for which focusing on the basic strategy would be enough to decide whether a positive answer can be given or not to the controllability question. The following strong monotonicity property, based on the promotion relation, can be used as such conditions:

Definition 3 (Strong Monotonicity **(SMon)**)**.** *A BRG $G = (V, A, \mathcal{L}_\mathcal{P}, \mathcal{B}, \mathcal{R})$ satisfies **(SMon)** if for each $i \in V$, if φ is unanimously accepted in G, then φ is unanimously accepted in any BRG $G \odot \varphi$.*

BRGs satisfying **(SMon)** are interesting in terms of strategy-proofness. Indeed, Proposition 1 tells us that φ is the least element of $(\mathcal{L}_\mathcal{P}, \preceq_\varphi)$. As a consequence, for such BRGs, determining whether it is possible to convince all the agents involved in the BRG to accept some piece of belief φ simply amounts to determining whether φ is unanimously accepted in the BRG obtained by the promotion of φ in every B_i associated with a controllable agent so that $B_i \odot \varphi = \varphi$. Stated otherwise:

Proposition 5. *Let $G = (V, A, \mathcal{L}_\mathcal{P}, \mathcal{B}, \mathcal{R})$ satisfying **(SMon)**, a set $C \subseteq V$ of controllable agents, and a formula φ. If the basic strategy σ_φ for φ is not successful, then there is no control strategy for G given C that is successful for φ.*

An interesting issue now is to know whether there are BRGs satisfying **(SMon)**. We provide a positive answer to this question in the next section.

5. The Case of Complete Graphs

We now study the extent to which **(SMon)** is satisfied by BRGs whose acquaintance graph is a *complete* graph. This simple topology is adequate to the cases when all the agents of V know each other (for instance, this is the case in meetings where all agents are around a table). We simply call the corresponding class of BRGs the complete BRGs:

Definition 4 (Complete BRG)**.** *A BRG $G = (V, A, \mathcal{L}_\mathcal{P}, \mathcal{B}, \mathcal{R})$ is* complete *if (V, A) is a complete graph, i.e., for all $i, j \in V$, $i \neq j$, $(i, j) \in A$. Given a class $\mathcal{G}$ of BRGs, $\mathcal{G}_{com}$ denotes the subclass of complete BRGs from $\mathcal{G}$.*

In the general case, **(SMon)** is not satisfied by BRGs from $\mathcal{G}_{com}(R^k_\Delta)$ with $k \in \{1, \dots, 6\}$, even when Δ is "fully" rational in the sense that it satisfies all IC postulates:

Proposition 6. *For $\Delta \in \{\Delta^{d_H,\Sigma}, \Delta^{d_H,\mathsf{GMax}}\}$, for any $k \in \{1, \dots, 6\}$, R^k_Δ does not satisfy **(SMon)** on $\mathcal{G}_{com}(R^k_\Delta)$.*

Let us now consider complete BRGs when the merging operator used for defining the revision policies is the distance-based merging operator based on the *drastic distance* (and the summation function) $\Delta^{d_D,\Sigma}$. Computing $\Delta^{d_D,\Sigma}_\mu(\mathcal{C})$ consists in selecting in the models of the integrity constraint μ those satisfying as many bases of the profile $\mathcal{C}$ as possible. Several works have proved this specific operator to satisfy a number of expected properties, e.g., some (language) independence conditions [6, 8]. In particular, $\Delta^{d_D,\Sigma}$ is robust from the point of view of strategy-proofness [3], this is why this operator appears as a good candidate for the monotonicity issue. Indeed:

Proposition 7. *Let $k \in \{1, \dots, 6\}$. Then all BRGs from the class $\mathcal{G}_{com}(R^k_{\Delta^{d_D,\Sigma}})$ satisfy **(SMon)**.*

As a consequence, to decide whether a control strategy exists for making a formula φ accepted by all agents in such a BRG, it is enough to focus on the basic strategy.

6. Conclusion

We pointed out a quite paradoxical result in belief diffusion: in a network of agents, replacing some agents' belief bases by a piece of belief φ may fail to make φ unanimously accepted, while other successful strategies exist nevertheless. However, we have identified a class of BRGs satisfying a property of strong monotonicity, and thus, avoiding this paradox, making for these BRGs the belief control issue easier to manage.

As perspectives for further research, we plan to identify additional classes of BRGs for which the strong monotonicity property holds, e.g., by considering other topologies as the line, single loop and more generally regular graph topologies. For all these BRGs, the next step will be to search for control strategies by focusing now on the "who" issue, i.e., which agents from a predefined set C should be considered. Another interesting research direction is to investigate on its own the notion of "promotion" of a piece of belief, as an attempt to provide an uniformized axiomatization of a large class of standard belief change operators.

References

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

[2] N. Dawar and Y. Pillutla. Impact of product-harm crises on brand equity: the moderating role of consumer expectations. *Journal of Marketing Research*, 37(2):215–226, 2000.

[3] P. Everaere, S. Konieczny, and P. Marquis. The strategy-proofness landscape of merging. *Journal of Artificial Intelligence Research*, 28:49–105, 2007.

[4] H. Katsuno and A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1992.

[5] S. Konieczny, J. Lang, and P. Marquis. DA^2 merging operators. *Artificial Intelligence*, 157:49–79, 2004.

[6] S. Konieczny, P. Marquis, and N. Schwind. Belief base rationalization for propositional merging. In *IJCAI'11*, pages 951–956, 2011.

[7] S. Konieczny and R. Pino Pérez. Merging information under constraints: a logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.

[8] P. Marquis and N. Schwind. Lost in translation: Language independence in propositional logic - application to belief change. *Artificial Intelligence*, 206:1–24, 2014.

[9] N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Belief revision games. In *AAAI'15*, pages 1590–1596, 2015.