



HAL
open science

Is Promoting Beliefs Useful to Make Them Accepted in Networks of Agents?

Nicolas Schwind, Katsumi Inoue, Gauvain Bourgne, Sébastien Konieczny,
Pierre Marquis

► **To cite this version:**

Nicolas Schwind, Katsumi Inoue, Gauvain Bourgne, Sébastien Konieczny, Pierre Marquis. Is Promoting Beliefs Useful to Make Them Accepted in Networks of Agents?. Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, Jul 2016, New York, NY, United States. pp.1237-1243. hal-01562250

HAL Id: hal-01562250

<https://hal.science/hal-01562250>

Submitted on 13 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Is Promoting Beliefs Useful to Make Them Accepted in Networks of Agents?

Nicolas Schwind

National Institute of Advanced Industrial
Science and Technology, Tokyo, Japan
email: nicolas-schwind@aist.go.jp

Katsumi Inoue

National Institute of Informatics / SOKENDAI,
Tokyo Institute of Technology, 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

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. In this paper, we investigate 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 formally capture such a concept of promotion by a new family of belief change operators. We show that some basic monotonicity properties are not satisfied by BRGs in general, even when the agents merging-based revision policies are fully rational (in the AGM sense). We also identify some classes 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) [Schwind *et al.*, 2015]. 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 the beliefs of her acquaintances. 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 the beliefs of her acquaintances. 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 systems, 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 [Dawar and Pillutla, 2000]; 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.

In the next section we provide some preliminaries on BRGs. We then define and investigate in Section 3 the notion of promoting a formula φ in a belief base, and show that many rational change operators (including revision and contraction operators) can be used for the promotion purpose. We then formalize in Section 4 the monotonicity property based on promotion operators and show that this condition is not always satisfied by BRGs. As a consequence, replacing some agents’ initial beliefs by a belief φ is not always the best way

to make φ unanimously accepted. In Section 5 we show that the monotonicity property holds for BRGs based on the complete graph topology, focusing on revision policies relying on the merging operator based on a specific merging operator. We discuss some related work in Section 6 before concluding. For space reasons, proofs are omitted but can be found online at <http://www.cril.fr/marquis/ijcai16-long.pdf>.

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. \perp (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* $\mathcal{C} = \langle B_1, \dots, B_n \rangle$ is a finite vector of belief bases. A Belief Revision Game (BRG for short) is formalized as follows [Schwind *et al.*, 2015]:

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*

- $V = \{1, \dots, n\}$ 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 ;
- $\mathcal{R} = \{R_1, \dots, R_n\}$, 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 $\mathcal{C}_i = \langle B_{i_1}, \dots, B_{i_{in(i)}} \rangle$ where $\{i_1, \dots, i_{in(i)}\} = \{j \mid (i_j, i) \in A\}$. 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 . Schwind *et al.* [2015] 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, \dots, B_i^e such that for every $j > e$, we have $B_i^j \equiv B_i^{b + ((j-b) \bmod (e-b+1))}$; the *belief cycle* $Cyc(i)$ of an agent i corresponds to the series of this subsequence of belief bases $Cyc(i) = B_i^b, B_i^{b+1}, \dots, B_i^e$. 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

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

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. [Alchourrón *et al.*, 1985]), 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 [Konieczny and Pino Pérez, 2002] for details on these properties).

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 [Konieczny *et al.*, 2004]. They associate with every formula μ and every profile \mathcal{C} a belief base $\Delta_{\mu}^{d,f}(\mathcal{C})$ satisfying $Mod(\Delta_{\mu}^{d,f}(\mathcal{C})) = \min(Mod(\mu), \leq_C^{d,f})$, where $\leq_C^{d,f}$ is the total preorder over interpretations induced by \mathcal{C} defined by $\omega \leq_C^{d,f} \omega'$ 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 $\leq_C^{d, GMax}$ defined by $\omega \leq_C^{d, GMax} \omega'$ if and only if $d^{GMax}(\omega, \mathcal{C}) \leq^{lex} d^{GMax}(\omega', \mathcal{C})$ (where \leq^{lex} is the lexicographic ordering induced by the natural order) and $d^{GMax}(\omega, \mathcal{C})$ is the vector of numbers d_1, \dots, 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_{Δ} 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 [Alchourrón *et al.*, 1985, Katsuno and Mendelzon, 1992].

Let us go back to BRGs. Six classes of revision policies have been proposed in [Schwind *et al.*, 2015]. Each of them, denoted R_{Δ}^k ($k \in \{1, \dots, 6\}$) 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 $\mathcal{C}_i \neq \emptyset$:

- $R_{\Delta}^1(B_i^s, \mathcal{C}_i^s) = \Delta(\mathcal{C}_i^s)$;

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

| 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 in G_* .

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

Example 1 (continued). We consider the BRG $G_* = (V_*, A_*, \mathcal{L}_{P_*}, B_*, \mathcal{R}_*)$ 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}_{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, $R_i = R_{\Delta}^{d_H, \text{GMax}}$ 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 $i \in V_*$, $\text{Cyc}(i) = p_1 \oplus p_2, \neg p_1 \vee \neg p_2$ and $\text{Acc}_G(i) = \neg p_1 \vee \neg p_2$.

[Schwind *et al.*, 2015] studied the extent to which BRGs satisfy some basic logical properties depending on the class of revision policies used by the agents. They focused on the case where all agents use the same revision policy ranging over $R_{\Delta}^k, k \in \{1, \dots, 6\}$. It turned out that when the revision policy is induced from the merging operator $\Delta^{d_H, \Sigma}$, i.e., the merging operator based on the drastic distance and the summation function, the underlying BRGs satisfy a number of expected properties [Schwind *et al.*, 2015]. They also developed a software available online at <http://www.cril.fr/brg/brg.jar>, allowing one to play BRGs with various revision policies.

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 through a preorder \preceq_{φ} over belief bases which 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 $B' \preceq_{\varphi} B$:

Definition 2 (φ -promotion). For every formula $\varphi \in \mathcal{L}_{\mathcal{P}}$, the φ -promotion relation is the binary relation \preceq_{φ} 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_{φ} is reflexive and transitive. More precisely:

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

- \sqcap_{φ} is the meet operation defined for each $B, B' \in \mathcal{L}_{\mathcal{P}}$ as $B \sqcap_{\varphi} B' = (B \wedge B') \vee ((B \vee B') \wedge \varphi)$,
- \sqcup_{φ} is the join operation defined for each $B, B' \in \mathcal{L}_{\mathcal{P}}$ as $B \sqcup_{\varphi} B' = (B \wedge B') \vee ((B \vee B') \wedge \neg\varphi)$,
- $\dot{\neg}$ is the complement corresponding to the standard logical negation \neg ,
- φ is the least element, i.e., for each $B \in \mathcal{L}_{\mathcal{P}}$, $\varphi \preceq_{\varphi} B$,
- $\neg\varphi$ is the top element, i.e., for each $B \in \mathcal{L}_{\mathcal{P}}$, $B \preceq_{\varphi} \neg\varphi$.

where formulae of $\mathcal{L}_{\mathcal{P}}$ are considered up to equivalence.

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_{φ} promoting φ in B , and φ is (up to logical equivalence) the least formula w.r.t. \preceq_{φ} 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 $\exists S \subseteq \text{Mod}(B) \blacktriangle \text{Mod}(\varphi)$ such that $\text{Mod}(B') = (\text{Mod}(B) \cap \text{Mod}(\varphi)) \cup S$.

This proposition also illustrates the fact that B and φ plays symmetric role in the notion of promotion. Indeed, making B closer to φ is precisely the same as making φ closer to B :

Proposition 3. $B' \preceq_{\varphi} B$ 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 4. Given a formula φ , let B and B' be two belief bases such that B' promotes φ in B . Then $\text{Mod}(B') \blacktriangle \text{Mod}(\varphi) \subseteq \text{Mod}(B) \blacktriangle \text{Mod}(\varphi)$.

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 $\psi \odot \varphi \preceq_{\varphi} \psi$. Interestingly, some

belief change operators of the literature are promotion operators. We recall below some of the postulates for rational revision operators \circ , and contraction operators $-$ [Katsuno and Mendelzon, 1992, Caridroit *et al.*, 2015], as well as some postulates for arbitration (alias symmetric revision) operators \diamond [Liberatore and Schaerf, 1998]:

- (R1) $\varphi \circ \mu \models \mu$.
- (R2) If $\varphi \wedge \mu$ is consistent, then $\varphi \circ \mu \equiv \varphi \wedge \mu$.
- (C1) $\varphi \models \varphi - \mu$.
- (C2) If $\varphi \not\models \mu$, then $\varphi - \mu \models \varphi$.
- (C4) If $\varphi \models \mu$, then $(\varphi - \mu) \wedge \mu \models \varphi$.
- (LS2) $\varphi \wedge \mu \models \varphi \diamond \mu$.
- (LS7) $\varphi \diamond \mu \models \varphi \vee \mu$.

Proposition 5. *Let φ be a formula and let B be a belief base. Let \circ be any revision operator satisfying (R1) and (R2), let $-$ be any KM contraction operator satisfying (C1), (C2) and (C4), and let \diamond be any arbitration operator satisfying (LS2) and (LS7). We have:*

1. (i) $B \circ \varphi \preceq_{\varphi} B \wedge \varphi \preceq_{\varphi} B$, (ii) $B \vee \varphi \preceq_{\varphi} B - \neg \varphi \preceq_{\varphi} B$, and (iii) $B \diamond \varphi \preceq_{\varphi} B$;
2. $B \circ \varphi$, $B \wedge \varphi$, $B \vee \varphi$, $B - \neg \varphi$ and $B \diamond \varphi$ are incomparable w.r.t. \preceq_{φ} in the general case.³

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 $E \subseteq \{1, \dots, 6\}$ R_{Δ}^E denotes the set $\{R_{\Delta}^k \mid k \in E\}$. We use the simpler notation R_{Δ}^k instead of R_{Δ}^E when $E = \{k\}$. Given a class \mathcal{G} of BRGs and $E \subseteq \{1, \dots, 6\}$, $\mathcal{G}(R_{\Delta}^E)$ is the subclass of all BRGs $(V, A, \mathcal{L}_{\mathcal{P}}, \mathcal{B}, \mathcal{R})$ from \mathcal{G} where for each $R_i \in \mathcal{R}$, $R_i \in R_{\Delta}^E$. Additionally, a set of revision policies R_{Δ}^E is said to satisfy a given property P on a given class \mathcal{G} of BRGs if all BRGs from $\mathcal{G}(R_{\Delta}^E)$ 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

³This is even the case for \circ and $-$ when they correspond one another thanks to Levi/Harper identities.

| 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.

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 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 also be much more complex (for instance, it may take into account 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 6. *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 $\text{Acc}_{G_*}(i) = p_1 \oplus p_2$ for each $i \in V$, 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 example illustrates the complexity of the controllability issue for BRGs in the general case. This explains why it is 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. In the following, we show that some monotonicity properties can be used as such conditions. A first property of monotonicity has been introduced in [Schwind *et al.*, 2015]. Given a BRG $G = (V, A, \mathcal{L}_{\mathcal{P}}, \mathcal{B}, \mathcal{R})$, a formula α and an agent $i \in V$, let us denote $G_{i \rightarrow \alpha}$ the BRG $(V, A, \mathcal{L}_{\mathcal{P}}, \mathcal{B}', \mathcal{R})$ where \mathcal{B}' maps V to $\{B'_1, \dots, B'_n\}$, with $B'_i = B_i \wedge \alpha$ and for every $j \in V$, $j \neq i$, $B'_j = B_j$.

Definition 3 (Monotonicity (Mon)). *A BRG $G = (V, A, \mathcal{L}_{\mathcal{P}}, \mathcal{B}, \mathcal{R})$ satisfies (Mon) if whenever φ is unanimously accepted in G , φ is unanimously accepted in $G_{i \rightarrow \varphi}$ for every $i \in V$.*

(Mon) is similar to the *monotonicity criterion* in Social Choice Theory. In the BRG context, a formula φ that is unanimously accepted should still be so if some agent’s initial

| | $B_i^s \wedge \varphi_{\max}(B^s) \models \perp$ (case (i)) | $\varphi_{\max}(B^s) \models B_i^s$ (case (ii)) | Otherwise (case (iii)) |
|------------------------------|-------------------------------------------------------------|----------------------------------------------------------------------|-----------------------------------------|
| $R_{\Delta^{d_D, \Sigma}}^1$ | $\varphi_{\max}(B^s)$ | $\varphi_{\max}(B^s) \vee (\varphi_{\max-1}(B^s) \wedge \neg B_i^s)$ | $\varphi_{\max}(B^s) \wedge \neg B_i^s$ |
| $R_{\Delta^{d_D, \Sigma}}^2$ | $\varphi_{\max}(B^s)$ | $\varphi_{\max}(B^s)$ | $\varphi_{\max}(B^s) \wedge \neg B_i^s$ |
| $R_{\Delta^{d_D, \Sigma}}^3$ | $\varphi_{\max}(B^s)$ | $\varphi_{\max}(B^s)$ | $\varphi_{\max}(B^s)$ |
| $R_{\Delta^{d_D, \Sigma}}^4$ | $\varphi_{\max}(B^s) \vee B_i^s$ | $\varphi_{\max}(B^s)$ | $\varphi_{\max}(B^s) \vee B_i^s$ |
| $R_{\Delta^{d_D, \Sigma}}^5$ | B_i^s | $\varphi_{\max}(B^s)$ | B_i^s |
| $R_{\Delta^{d_D, \Sigma}}^6$ | B_i^s | $\varphi_{\max}(B^s)$ | $\varphi_{\max}(B^s) \wedge B_i^s$ |

Table 3: B_i^{s+1} depending on the revision policy applied by agent i in any BRG from $\mathcal{G}_{com}(R_{\Delta^{d_D, \Sigma}}^{\{1, \dots, 6\}})$.

beliefs were “strengthened” by φ . Note that the strengthening of a belief base by a formula φ involved here consists in just expanding the base by φ . However, one could consider weaker versions than expansion, i.e., the promotion operators \odot we formalized in the previous section. So we introduce now a stronger version of the monotonicity property (cf. Definition 3) on BRGs based on the promotion relation:

Definition 4 (Strong Monotonicity (**SMon**)). *A BRG $G = (V, A, \mathcal{L}_P, \mathcal{B}, \mathcal{R})$ satisfies (**SMon**) if for each $i \in V$, if φ is unanimously accepted in G , then φ is unanimously accepted in any BRG $G \odot \varphi$.*

It is easy to see that (**SMon**) implies (**Mon**) and that the converse does not hold in general. BRGs satisfying (**SMon**) are interesting in terms of strategy-proofness. Indeed, Proposition 1 tells us that φ is the least element of $(\mathcal{L}_P, \preceq_\varphi)$. 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 7. *Let $G = (V, A, \mathcal{L}_P, \mathcal{B}, \mathcal{R})$ satisfying (**SMon**), a set $C \subseteq V$ of controllable agents, and a formula φ . If the basic strategy σ_φ for G given C is not successful for φ then no control strategy for G given C 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 sitting around a table). We simply call the corresponding class of BRGs the complete BRGs:

Definition 5 (Complete BRG). *A BRG $G = (V, A, \mathcal{L}_P, \mathcal{B}, \mathcal{R})$ 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 $\mathcal{G}_{com}(R_{\Delta}^k)$ with $k \in \{1, \dots, 6\}$, even when Δ is “fully” rational in the sense that it satisfies all IC postulates. Indeed:

Proposition 8. *For $\Delta \in \{\Delta^{d_H, \Sigma}, \Delta^{d_H, \text{GM}^{\max}}\}$, for any $k \in \{1, \dots, 6\}$, R_{Δ}^k does not satisfy (**SMon**) on $\mathcal{G}_{com}(R_{\Delta}^k)$.*

In the following, we investigate the monotonicity issue for 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 $\Delta_{\mu}^{d_D, \Sigma}(\mathcal{C})$ 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 [Konieczny *et al.*, 2011, Marquis and Schwind, 2014]. In particular, $\Delta^{d_D, \Sigma}$ is robust from the point of view of strategy-proofness [Everaere *et al.*, 2007], this is why this operator appears as a good candidate for the monotonicity issue.

The following notations are used in the rest of this section. Let \mathcal{C} be a profile and ω be an interpretation. $\bigwedge \mathcal{C} = \bigwedge \{B_i \mid B_i \in \mathcal{C}\}$. $\#E$ denotes the number of elements from a finite set E . $\#\max(\mathcal{C})$ is the maximal number of belief bases from \mathcal{C} that are consistent when interpreted conjunctively, i.e., $\#\max(\mathcal{C}) = \max(\#\{\mathcal{C}' \mid \mathcal{C}' \subseteq \mathcal{C}, \bigwedge \mathcal{C}' \not\models \perp\})$. $\{\mathcal{C} \mid \omega \models \mathcal{C}\} = \{B_i \in \mathcal{C} \mid \omega \models B_i\}$. $\varphi_{\max}(\mathcal{C}) = \Delta^{d_D, \Sigma}(\mathcal{C})$, and $\varphi_{\max-1}(\mathcal{C})$ is any formula φ such that $\text{Mod}(\varphi) = \{\omega \mid \#\{\mathcal{C} \mid \omega \models \mathcal{C}\} = \#\max(\mathcal{C}) - 1\}$.

We now present a number of characterization results concerning stability and acceptability for such BRGs, which are interesting by their own, and will be exploited to identify classes of BRGs offering (**SMon**). We first show that a stronger version of the agreement preservation property (**AP**) (see [Schwind *et al.*, 2015]) is satisfied by these BRGs:

Proposition 9. *Let $G = (V, A, \mathcal{L}_P, \mathcal{B}, \mathcal{R})$ be a BRG from $\mathcal{G}_{com}(R_{\Delta^{d_D, \Sigma}}^{\{1, \dots, 6\}})$. If there is a step s where $\bigwedge \{B_i^s \mid i \in V\}$ is consistent, then for each $i \in V$, $\text{Acc}_G(B_i) = \bigwedge \{B_i \mid i \in V\}$. Moreover, G converges at step $s+2$ at most. It converges at step $s+1$ at most if $G \in \mathcal{G}_{com}(R_{\Delta^{d_D, \Sigma}}^{\{2, \dots, 6\}})$.*

We now intend to characterize the outcome of a BRG from its initial state. Beforehand, let us introduce an intermediary result characterizing the belief base of any agent in a BRG from $\mathcal{G}_{com}(R_{\Delta^{d_D, \Sigma}}^k)$ where $k \in \{1, \dots, 6\}$, at some step $s+1$ given the belief bases of all agents at step s :

Lemma 1. *Let $E \subseteq \{1, \dots, 6\}$ and G be a BRG from $\mathcal{G}_{com}(R_{\Delta^{d_D, \Sigma}}^E)$. Then for each agent $i \in V$ and each step s , her base B_i^{s+1} can be characterized as shown in Table 3.*

Obviously enough, a consequence of Table 3 is that all six revision policies from $R_{\Delta^{d_D, \Sigma}}^k$, $k \in \{1, \dots, 6\}$, remain distinct. Moreover, in the case when all agents have the same revision policy among $R_{\Delta^{d_D, \Sigma}}^k$, $k \in \{1, \dots, 4\}$, the outcome of complete BRGs can be fully characterized:

Proposition 10. Let $k \in \{1, \dots, 4\}$ and $G = (V, A, \mathcal{L}_{\mathcal{P}}, \mathcal{B}, \mathcal{R})$ be a BRG from $\mathcal{G}_{com}(R_{\Delta^{d_D, \Sigma}}^k)$. Then for each $i \in V$,

$$Acc_G(i) = \Delta^{d_D, \Sigma}(\langle B_1, \dots, B_n \rangle).$$

The cases when all agents have the same revision policy among $R_{\Delta^{d_D, \Sigma}}^k$, $k \in \{5, 6\}$ is simpler to characterize:

Proposition 11. Let $k \in \{5, 6\}$ and $G = (V, A, \mathcal{L}_{\mathcal{P}}, \mathcal{B}, \mathcal{R})$ be a BRG from $\mathcal{G}_{com}(R_{\Delta^{d_D, \Sigma}}^k)$.

- If $k = 5$, then for all $i \in V$, $Acc_G(i) = B_i$ if $\Delta^{d_D, \Sigma}(\langle B_1, \dots, B_n \rangle) \not\models B_i$, otherwise $Acc_G(i) = \Delta^{d_D, \Sigma}(\langle B_1, \dots, B_n \rangle)$.
- If $k = 6$, then for all $i \in V$, $Acc_G(i) = B_i$ if $\Delta^{d_D, \Sigma}(\langle B_1, \dots, B_n \rangle) \wedge B_i \models \perp$, otherwise $Acc_G(i) = \Delta^{d_D, \Sigma}(\langle B_1, \dots, B_n \rangle) \wedge B_i$.

These results allow us now to derive monotonicity results for the class of complete BRGs we are dealing with:

Proposition 12. Let $k \in \{1, \dots, 6\}$. Then all BRGs from the class $\mathcal{G}_{com}(R_{\Delta^{d_D, \Sigma}}^k)$ satisfy (**SMon**).

A consequence of Propositions 7 and 12 is that considering the basic strategy is enough for BRGs from $\mathcal{G}_{com}(R_{\Delta^{d_D, \Sigma}}^k)$, i.e., 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. Let us illustrate how these monotonicity results apply to the following slightly modified version of our running example:

Example 1 (continued). Consider again the BRG $G_* = (V_*, A_*, \mathcal{L}_{\mathcal{P}_*}, \mathcal{B}_*, \mathcal{R}_*)$ from our running example, and let G'_* be the BRG $(V_*, A'_*, \mathcal{L}_{\mathcal{P}_*}, \mathcal{B}'_*, \mathcal{R}_*)$ where:

- A' is the total relation over V_* , i.e., Beth and Chris know each other as well;
- B'_* maps V_* to $\{B'_1, B'_2, B'_3\}$, with $B'_1 = B_1$, $B'_2 = B_2$ and $B'_3 = \neg p_1 \wedge \neg p_2$, i.e., the beliefs of Alex and Beth in G'_* remain the same as in G_* , but Chris now initially believes in G'_* that none of the two meals is healthy.

From Proposition 10, we get that $Acc_{G'_*}(i) = \Delta^{d_D, \Sigma}(\langle B'_1, B'_2, B'_3 \rangle) \equiv \neg p_1 \wedge \neg p_2$ for each $i \in V_*$. Note that $p_1 \vee p_2$ is not accepted by any agent. By applying the basic strategy σ_φ where $\varphi = p_1 \vee p_2$ and when Chris is the only controllable agent, one gets that the outcome of each agent in the resulting BRG becomes $p_1 \oplus p_2$. We have now φ being unanimously accepted, which means that σ_φ is successful.

6 Related Work

As far as we know, the control problem for multi-agent systems has not been investigated so far in settings similar to BRGs. Though there are many works on opinion dynamics where some control or influence issues have been studied, the opinion of an agent takes typically the form of a real number, as in Hegselmann-Krause's model [Hegselmann and Krause, 2005]. Let us mention among others [Tsang and Larson, 2014, Bloembergen *et al.*, 2014, Ranjbar-Sahraei *et al.*, 2014, Blondel *et al.*, 2009]. The very nature of the agents' opinions (real number vs. propositional formula) makes these studies and our own approach quite unrelated. Indeed, unlike real numbers that can be simply averaged, merging propositional formulae is a much more complex process (this is why

many works on the topic in the knowledge representation area have been done for two decades). The quite paradoxical result presented in the paper (roughly, promoting some belief is not guaranteed to be the safest way to make it unanimously accepted) echoes paradoxes pointed out in other settings for multi-agent systems, in particular social network games [Apt *et al.*, 2013]. However, the resemblance is only at a shallow level since the BRG model departs a lot from social network games. Especially, social network games are "true" games, unlike BRGs (in the BRG model, agents do not have strategies and there is no payoff based on the adequacy of an agent's beliefs with the beliefs of her acquaintances).

On the other hand, some control issues (under various forms: manipulation, bribery, etc.) have received much attention recently in voting theory, especially from the point of view of computational complexity. See among many others [Amanatidis *et al.*, 2015, Bredereck *et al.*, 2015, Faliszewski *et al.*, 2015, Loreggia *et al.*, 2015, Faliszewski *et al.*, 2014, Bouveret and Lang, 2014, Christian *et al.*, 2007]. However, the results obtained in voting theory cannot be exploited directly for BRGs since in BRGs the beliefs of the agents may evolve after each communication step (when voting is considered, the users preferences are typically considered as fixed and the agents do not influence each other). An exception is [Hassanzadeh *et al.*, 2013] which shows how to reach a consensus on rankings (encoding the agents preferences) by iterative voting. Nevertheless, the corresponding model is quite different from the BRG one since it deals with preferences, and not with beliefs, and the system dynamics is ruled by distinct mechanisms (voting vs. revision). Furthermore, belief control has not been investigated in this model (the focus is on convergence issues).

Finally, though other approaches have been considered for modeling the belief dynamics of a group of agents (see dynamic epistemic logic [Harel *et al.*, 2000, Hendricks and Symons, 2006], or more recent studies in propositional logic [Gauwin *et al.*, 2007, Delgrande *et al.*, 2007, Grandi *et al.*, 2015]), the belief control issue for such settings has not been considered as well.

7 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

- [Alchourrón *et al.*, 1985] 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.
- [Amanatidis *et al.*, 2015] G. Amanatidis, N. Barrot, J. Lang, E. Markakis, and B. Ries. Multiple referenda and multiwinner elections using hamming distances: Complexity and manipulability. In *AAMAS’15*, pages 715–723, 2015.
- [Apt *et al.*, 2013] K. R. Apt, E. Markakis, and S. Simon. Paradoxes in social networks with multiple products. *CoRR*, abs/1301.7592, 2013.
- [Bloembergen *et al.*, 2014] D. Bloembergen, B. Ranjbar Sahraei, H. Bou-Ammar, K. Tuyls, and G. Weiss. Influencing social networks: An optimal control study. In *ECAI’14*, pages 105–110, 2014.
- [Blondel *et al.*, 2009] V. D. Blondel, J. M. Hendrickx, and J. N. Tsitsiklis. On Krause’s multi-agent consensus model with state-dependent connectivity. *IEEE Transactions on Automatic Control*, 54(11):2586–2597, 2009.
- [Bouveret and Lang, 2014] S. Bouveret and J. Lang. Manipulating picking sequences. In *ECAI’14*, pages 141–146, 2014.
- [Bredereck *et al.*, 2015] R. Bredereck, J. Chen, R. Niedermeier, and T. Walsh. Parliamentary voting procedures: Agenda control, manipulation, and uncertainty. In *IJCAI’15*, pages 164–170, 2015.
- [Caridroit *et al.*, 2015] T. Caridroit, S. Konieczny, and P. Marquis. Contraction in propositional logic. In *EC-SQARU’15*, pages 186–196, 2015.
- [Christian *et al.*, 2007] R. Christian, M. Fellows, F. Rosamond, and A. Slinko. On complexity of lobbying in multiple referenda. *Review of Economic Design*, 11(3):217–224, 2007.
- [Dawar and Pillutla, 2000] 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.
- [Delgrande *et al.*, 2007] J. P. Delgrande, J. Lang, and T. Schaub. Belief change based on global minimisation. In *IJCAI’07*, pages 2468–2473, 2007.
- [Everaere *et al.*, 2007] P. Everaere, S. Konieczny, and P. Marquis. The strategy-proofness landscape of merging. *Journal of Artificial Intelligence Research*, 28:49–105, 2007.
- [Faliszewski *et al.*, 2014] P. Faliszewski, E. Hemaspaandra, and L. A. Hemaspaandra. The complexity of manipulative attacks in nearly single-peaked electorates. *Artificial Intelligence*, 207:69–99, 2014.
- [Faliszewski *et al.*, 2015] P. Faliszewski, Y. Reisch, J. Rothe, and L. Schend. Complexity of manipulation, bribery, and campaign management in bucklin and fallback voting. *Autonomous Agents and Multi-Agent Systems*, 29(6):1091–1124, 2015.
- [Gauwin *et al.*, 2007] O. Gauwin, S. Konieczny, and P. Marquis. Conciliation through iterated belief merging. *Journal of Logic and Computation*, 17(5):909–937, 2007.
- [Grandi *et al.*, 2015] U. Grandi, E. Lorini, and L. Perrussel. Propositional opinion diffusion. In *AAMAS’15*, pages 989–997, 2015.
- [Harel *et al.*, 2000] D. Harel, J. Tiuryn, and D. Kozen. *Dynamic Logic*. MIT Press, Cambridge, MA, USA, 2000.
- [Hassanzadeh *et al.*, 2013] F. Farnoud Hassanzadeh, E. Yaakobi, B. Touri, O. Milenkovic, and J. Bruck. Building consensus via iterative voting. In *ISIT’13*, pages 1082–1086, 2013.
- [Hegselmann and Krause, 2005] R. Hegselmann and U. Krause. Opinion dynamics driven by various ways of averaging. *Computational Economics*, 25:381–405, 2005.
- [Hendricks and Symons, 2006] V. Hendricks and J. Symons. *Epistemic Logic*. Edward N. Zalta, editor, The Stanford Encyclopedia of Philosophy, 2006.
- [Katsuno and Mendelzon, 1992] H. Katsuno and A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1992.
- [Konieczny and Pino Pérez, 2002] S. Konieczny and R. Pino Pérez. Merging information under constraints: a logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.
- [Konieczny *et al.*, 2004] S. Konieczny, J. Lang, and P. Marquis. DA² merging operators. *Artificial Intelligence*, 157:49–79, 2004.
- [Konieczny *et al.*, 2011] S. Konieczny, P. Marquis, and N. Schwind. Belief base rationalization for propositional merging. In *IJCAI’11*, pages 951–956, 2011.
- [Liberatore and Schaerf, 1998] P. Liberatore and M. Schaerf. Arbitration (or how to merge knowledge bases). *IEEE Transactions on Knowledge and Data Engineering*, 10(1):76–90, 1998.
- [Loreggia *et al.*, 2015] A. Loreggia, N. Narodytska, F. Rossi, K. B. Venable, and T. Walsh. Controlling elections by replacing candidates or votes. In *AAMAS’15*, pages 1737–1738, 2015.
- [Marquis and Schwind, 2014] P. Marquis and N. Schwind. Lost in translation: Language independence in propositional logic - application to belief change. *Artificial Intelligence*, 206:1–24, 2014.
- [Ranjbar-Sahraei *et al.*, 2014] B. Ranjbar-Sahraei, H. Bou Ammar, D. Bloembergen, K. Tuyls, and G. Weiss. Evolution of cooperation in arbitrary complex network. In *AAMAS’14*, pages 677–684, 2014.
- [Schwind *et al.*, 2015] N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Belief revision games. In *AAAI’15*, pages 1590–1596, 2015.
- [Tsang and Larson, 2014] A. Tsang and K. Larson. Opinion dynamics of skeptical agents. In *AAMAS’14*, pages 277–284, 2014.