

What Has Been Said? Identifying the Change Formula in a Belief Revision Scenario

Nicolas Schwind^{1*}, Katsumi Inoue^{2,3}, Sébastien Konieczny⁴,
Jean-Marie Lagniez⁴ and Pierre Marquis^{4,5}

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

²National Institute of Informatics, Tokyo, Japan

³Tokyo Institute of Technology, Tokyo, Japan

⁴CRIL-CNRS, Université d'Artois, Lens, France

⁵Institut Universitaire de France

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

Abstract

We consider the problem of identifying the change formula in a belief revision scenario: given that an unknown announcement (a formula μ) led a set of agents to revise their beliefs and given the prior beliefs and the revised beliefs of the agents, what can be said about μ ? We show that under weak conditions about the rationality of the revision operators used by the agents, the set of candidate formulae has the form of a logical interval. We explain how the bounds of this interval can be tightened when the revision operators used by the agents are known and/or when μ is known to be independent from a given set of variables. We also investigate the completeness issue, i.e., whether μ can be exactly identified. We present some sufficient conditions for it, identify its computational complexity, and report the results of some experiments about it.

1 Introduction

Announcements have been widely studied in the Dynamic Epistemic Logic framework [van Ditmarsch *et al.*, 2008; Baltag and Moss, 2004]. Three main types of announcements have been identified in this setting: private announcements, where some agents receive new information but the other agents are not aware of this communication, semi-private announcements where the other agents know that the agents receive some information about a formula φ (but they do not know whether it was φ or $\neg\varphi$), and finally, public announcements where the exact message is known by all agents.

We want to address another kind of announcement, that is in-between private and semi-private, where the other agents are aware that something has been said to the agents but do not know what. Let us call these announcements *quasi-private announcements*. This corresponds to very common situations when one can see a conversation between several persons but cannot hear what is said. Thus, in this work, we focus on the case when a formula μ (announcement) has

been said to a group of agents, and each person of this group has rationally incorporated this new piece of information into her beliefs, using a belief revision operator [Alchourrón *et al.*, 1985; Katsuno and Mendelzon, 1991]. We want to know what can be said about this formula from the prior and revised beliefs of the agents.

Technically, this issue can be considered as an inverse problem for belief revision. Usually, in belief revision, one knows the prior beliefs φ_i of an agent i and the new piece of information μ , and the aim is to compute her revised beliefs γ_i so that $\gamma_i = \varphi_i \circ \mu$. Here, our goal is to find the new piece of information μ from the prior beliefs φ_i and the revised beliefs γ_i of a group of agents. In formal terms, the problem amounts to “solving” the system of equations $\varphi_i \circ \mu = \gamma_i$, where the unknown is μ .

For the sake of illustration, let us present a scenario that will be used as a running example in the paper:

Example 1. *We know that Amy and Beniko received some announcement about a park they plan to visit. The prior beliefs of Amy were that the park is far to the north and that the park's entrance is free, whereas Beniko believes at start that the park is near to the south and that the entrance is subject to a fee. Then we saw that a third person told them something, but we do not know what. Nevertheless, we know that, after the occurrence of this event, the beliefs of Amy have evolved in such a way that she now believes that the park's entrance is free, but that the park is located to the east, and is either near to the north or far to the south. We also know that Beniko still believes that the park is near and that entrance is subject to a fee, and additionally believes that it is located south east of their current location. The issue is to determine the announcement made to Amy and Beniko.*

In this work the agents are supposed to be rational in the sense that each one of them uses a belief revision operator to revise her beliefs. This is enough to show that the set of candidate formulae for μ has the form of a logical interval. We explain how the bounds of this interval can be tightened when μ is known to be independent from a given set of variables (this captures a situation when the topic of the change formula is known), and/or when the revision operators used by the agents are known. We also study the completeness issue,

*Contact Author

i.e., whether the input allows one to identify μ (up to logical equivalence). We present some sufficient conditions for it, identify its computational complexity, and report the results of some experiments about the completeness issue when the revision operators used by the agents are Dalal's.

For space reasons, most proofs are omitted. An extended version of the paper containing all the proofs is available from <https://www.cril.univ-artois.fr/~marquis/ijcai19-IP-extended.pdf>.

2 Preliminaries

Let $\mathcal{L}_{\mathcal{P}}$ be a propositional language built up from a finite set of propositional variables \mathcal{P} and the usual connectives. \perp (resp. \top) is the Boolean constant always false (resp. true). An interpretation (or world) is a mapping from \mathcal{P} to $\{0, 1\}$. The set of all interpretations is denoted \mathcal{W} . $[\varphi]$ denotes the set of models of the formula φ , i.e., $[\varphi] = \{\omega \in \mathcal{W} \mid \omega \models \varphi\}$. \models denotes logical entailment and \equiv logical equivalence, i.e., $\varphi \models \psi$ iff $[\varphi] \subseteq [\psi]$ and $\varphi \equiv \psi$ iff $[\varphi] = [\psi]$. $Var(\alpha)$ denotes the set of propositional variables appearing in α .

A *belief base* (base for short) is a formula (or equivalently a finite set of propositional formulae considered conjunctively) that represents the set of beliefs of an agent.

Definition 1 (KM revision operator [Katsuno and Mendelzon, 1991]). *A KM revision operator \circ is a revision operator associating a base φ and formula μ with a new base $\varphi \circ \mu$, such that for all formulae $\mu, \mu', \varphi, \varphi'$:*

- (R1) $\varphi \circ \mu \models \mu$;
- (R2) If $\varphi \wedge \mu$ is consistent, then $\varphi \circ \mu \equiv \varphi \wedge \mu$;
- (R3) If μ is consistent, then $\varphi \circ \mu$ is consistent;
- (R4) If $\varphi \equiv \varphi'$ and $\mu \equiv \mu'$, then $\varphi \circ \mu \equiv \varphi' \circ \mu'$;
- (R5) $(\varphi \circ \mu) \wedge \mu' \models \varphi \circ (\mu \wedge \mu')$;
- (R6) If $(\varphi \circ \mu) \wedge \mu'$ is consistent, then $\varphi \circ (\mu \wedge \mu') \models (\varphi \circ \mu) \wedge \mu'$.

Postulates (R1 - R3) are the most essential ones. Postulate (R4) is the syntax-independence postulate. Postulates (R5 - R6) require that the minimality of change is induced by a notion of proximity between worlds (see [Katsuno and Mendelzon, 1991] for more details about the rationale of these postulates).

KM revision operators can be characterized in terms of total preorders over interpretations. Given a preorder \preceq, \simeq (resp. \prec) denotes the indifference (resp. strict) part of \preceq : a *faithful assignment* is a mapping associating every base φ with a preorder \preceq_{φ} such that for all bases $\varphi, \varphi_1, \varphi_2$ and all interpretations $\omega, \omega' \in \mathcal{W}$, the following conditions are satisfied:

1. If $\omega, \omega' \models \varphi$, then $\omega \simeq_{\varphi} \omega'$;
2. If $\omega \models \varphi$ and $\omega' \not\models \varphi$, then $\omega \prec_{\varphi} \omega'$;
3. If $\varphi_1 \equiv \varphi_2$, then $\preceq_{\varphi_1} = \preceq_{\varphi_2}$.

Proposition 1. [Katsuno and Mendelzon, 1991] *A revision operator \circ is a KM revision operator if and only if there exists a faithful assignment associating each base φ with a total preorder \preceq_{φ} such that for all $\varphi, \mu \in \mathcal{L}_{\mathcal{P}}$, $[\varphi \circ \mu] = \min([\mu], \preceq_{\varphi})$.*

An interesting class of KM revision operators is the class of distance-based revision operators \circ_d characterized by a distance between interpretations¹ $d: \mathcal{W} \times \mathcal{W} \mapsto \mathbb{N}$ [Lehmann *et al.*, 2001; Konieczny *et al.*, 2004]. For all $\varphi, \mu \in \mathcal{L}_{\mathcal{P}}$, $\varphi \circ_d \mu$ is defined as any formula whose models are the most plausible models of μ w.r.t. \preceq_{φ}^d , i.e.,

$$[\varphi \circ_d \mu] = \min([\mu], \preceq_{\varphi}^d),$$

where for all $\omega_1, \omega_2 \in \mathcal{W}$, $\omega_1 \preceq_{\varphi}^d \omega_2$ if and only if $d(\omega_1, \varphi) \leq d(\omega_2, \varphi)$, and for every $\omega \in \mathcal{W}$, $d(\omega, \varphi) = \min_{\omega' \models \varphi} d(\omega, \omega')$.

Noteworthy, KM revision operators include Dalal's revision operator \circ_{Dal} [Dalal, 1988] which corresponds to the distance-based revision operator based on the Hamming distance d_H defined for all $\omega, \omega' \in \mathcal{W}$ as $d_H(\omega, \omega') = |\{x \in \mathcal{P} \mid \omega(x) \neq \omega'(x)\}|$.

3 Identifying the Change Formula

We consider a finite set of agents $\bar{n} = \{1, \dots, n\}$, $n \geq 1$. Each agent $i \in \bar{n}$ is associated with a prior belief base φ_i . A formula, called *change formula* and denoted by μ , is assumed to be said to all agents but it is not known. Each agent $i \in \bar{n}$ revises her prior belief base φ_i by μ , resulting in the revised belief base γ_i . Lastly, we assume that each agent performs rational revision operations. More precisely, each agent $i \in \bar{n}$ uses a revision operator \circ_i which is not known, but which satisfies all KM rationality postulates :

$$\begin{aligned} \varphi_1 \circ_1 \mu &= \gamma_1 \\ &\vdots \\ \varphi_n \circ_n \mu &= \gamma_n. \end{aligned}$$

Overall, we consider an “inverse revision problem” whose input is characterized by an *IP pair* defined as follows:

Definition 2 (IP pair). *Given a prior profile $\Phi = \langle \varphi_1, \dots, \varphi_n \rangle$ and a revised profile $\Gamma = \langle \gamma_1, \dots, \gamma_n \rangle$, where for each $i \in \bar{n}$, φ_i and γ_i correspond respectively to the prior belief base and the revised belief base of agent i , the pair $IP = \langle \Phi, \Gamma \rangle$ is called an Inverse revision Problem description, or IP pair for short.*

Our goal is to characterize the change formula μ when possible, or “approximate” it as much as possible given the information provided by the IP pair.

Definition 3 (Frame). *The frame of an IP pair $IP = \langle \Phi, \Gamma \rangle$, denoted by $fr(IP)$, is the set of formulae δ from $\mathcal{L}_{\mathcal{P}}$ such that for each agent $i \in \bar{n}$, there exists a revision operator \circ_i satisfying all KM postulates and such that $\varphi_i \circ_i \delta \equiv \gamma_i$.*

From now on, we assume that (i) the change formula μ is consistent, and (ii) no agent lies on her actual prior or revised belief base. Obviously enough, under these assumptions, we always have that $fr(IP) \neq \emptyset$, since for each agent $i \in \bar{n}$, the revised base γ_i results as a matter of fact from the revision of the prior base φ_i by μ . We say that a frame $fr(IP)$ is

¹Actually a pseudo-distance is enough, i.e., the triangular inequality is not needed.

complete when for all $\delta, \delta' \in fr(IP)$, we have that $\delta \equiv \delta'$. Intuitively, this means that the change formula of an IP pair can be characterized (up to logical equivalence) precisely when $fr(IP)$ is complete.

Example 1 (continued). *Let us now model the scenario provided in the introduction. We are given two agents $\bar{2} = \{1, 2\}$, where 1 corresponds to Amy, 2 to Beniko. $\mathcal{L}_{\mathcal{P}}$ is the propositional language defined from the set of variables $\mathcal{P} = \{s, e, n, f\}$, where s stands for “the park is located to the south”, e stands for “the park is located to the east”, n means “the park is near”, and f means that “the park’s entrance is free”. Amy’s (resp. Beniko’s) prior beliefs are then represented as $\varphi_1 = \neg s \wedge \neg n \wedge f$ (resp. $\varphi_2 = s \wedge n \wedge \neg f$). After getting information from the third person, Amy’s (resp. Beniko’s) revised beliefs are represented as $\gamma_1 = e \wedge (n \Leftrightarrow \neg s) \wedge f$ (resp. $\gamma_2 = s \wedge e \wedge n \wedge \neg f$). Thus, we consider the IP pair $IP = \{(\varphi_1, \varphi_2), (\gamma_1, \gamma_2)\}$. We assume that there are two KM revision operators \circ_1, \circ_2 such that $\varphi_1 \circ_1 \mu = \gamma_1, \varphi_2 \circ_2 \mu = \gamma_2$, although nothing is known about \circ_1 and \circ_2 apart from the fact that they satisfy all KM postulates. The issue of interest is to identify, or at least approximate as much as possible, the frame $fr(IP)$, i.e., the set of possible candidates for μ .*

Let us first observe that:

$$fr(\langle \Phi \sqcup \langle \varphi_{n+1} \rangle, \Gamma \sqcup \langle \gamma_{n+1} \rangle \rangle) \subseteq fr(IP),$$

where \sqcup denotes the concatenation of vectors. This means that the more agents in an IP pair, the better the approximation of the change formula.

We show below that the frame of an IP pair can be simply characterized:

Proposition 2. *Let IP be an IP pair. Then $fr(IP)$ is precisely the set of formulae given by*

$$fr(IP) = \{ \delta \in \mathcal{L}_{\mathcal{P}} \mid \bigvee_{\gamma_i \in \Gamma} \gamma_i \models \delta \models \bigwedge_{\varphi_i \in \Phi, \gamma_i \in \Gamma} \neg \varphi_i \vee \gamma_i \}.$$

Proof. (\subseteq) Let $\delta \in fr(IP)$. By definition of $fr(IP)$, we know that for each agent $i \in \bar{n}$, there is a revision operator \circ_i satisfying (R1-R6) such that $\varphi_i \circ_i \delta \equiv \gamma_i$. Since for each $i \in \bar{n}$, \circ_i satisfies (R1), we get that for each $\gamma_i \in \Gamma$, $\gamma_i \models \delta$. Hence, $\bigvee_{\gamma_i \in \Gamma} \gamma_i \models \delta$. Let us now prove that $\delta \models \bigwedge_{\varphi_i \in \Phi, \gamma_i \in \Gamma} \neg \varphi_i \vee \gamma_i$. Let $i \in \bar{n}$. In the case where $\varphi_i \wedge \delta$ is inconsistent, we get that $\delta \models \neg \varphi_i$, so we trivially have that $\delta \models \neg \varphi_i \vee \gamma_i$. Otherwise $\varphi_i \wedge \delta$ is consistent, and since \circ_i satisfies (R2) we know that $\gamma_i \equiv \varphi_i \wedge \delta$, thus we also have that $\delta \models \neg \varphi_i \vee \gamma_i$. Therefore, $\delta \models \bigwedge_{\varphi_i \in \Phi, \gamma_i \in \Gamma} \neg \varphi_i \vee \gamma_i$.

(\supseteq) Let $\delta \in \mathcal{L}_{\mathcal{P}}$ such that $\bigvee_{\gamma_i \in \Gamma} \gamma_i \models \delta$ and $\delta \models \bigwedge_{\varphi_i \in \Phi, \gamma_i \in \Gamma} \neg \varphi_i \vee \gamma_i$. To prove that $\delta \in fr(IP)$, we need to show that one can associate with each agent $i \in \bar{n}$ a revision operator \circ_i satisfying (R1-R6) and such that $\varphi_i \circ_i \delta \equiv \gamma_i$. We proceed as follows. For each agent $i \in \bar{n}$ such that $\varphi_i \wedge \gamma_i$ is consistent, let \circ_i be the drastic revision operator, i.e., $\circ_i = \circ_D$, where \circ_D is defined for all $\alpha, \beta \in \mathcal{L}_{\mathcal{P}}$ as $\alpha \circ_D \beta = \alpha \wedge \beta$ if $\alpha \wedge \beta$ is consistent, otherwise $\alpha \circ_D \beta = \beta$. And for each remaining agent $i \in \bar{n}$, since we know that \circ_i can be represented by a faithful assignment, we associate

with each \circ_i the faithful assignment associating each formula φ with the total preorder $\preceq_{\varphi_i}^i$, where $\preceq_{\varphi_i}^i$ is defined for all interpretations $\omega, \omega' \in \mathcal{W}$ as $\omega \preceq_{\varphi_i}^i \omega'$ if and only if $(\omega \models \varphi_i$ or $\omega' \models \neg(\varphi_i \vee \gamma_i))$. Under these constructions, it can easily be verified that for each $i \in \bar{n}$, $\varphi_i \circ_i \delta \equiv \gamma_i$. Therefore, $\delta \in fr(IP)$. This concludes the proof. \square

Proposition 2 above shows that the frame of an IP pair has the form of a logical interval that can be characterized succinctly and in polynomial time. A frame $fr(IP)$ can simply be identified by a pair of propositional formulae, so that it can be denoted by $(fr^-(IP), fr^+(IP))$, where $fr^-(IP) = \bigvee_{\gamma_i \in \Gamma} \gamma_i$ and $fr^+(IP) = \bigwedge_{\varphi_i \in \Phi, \gamma_i \in \Gamma} \neg \varphi_i \vee \gamma_i$. On the one hand, $fr^+(IP)$ represents what has been said “at least” in the change formula μ given the information provided by the IP pair, i.e., for every formula α such that $fr^+(IP) \models \alpha$, we know that $\mu \models \alpha$. On the other hand, $fr^-(IP)$ characterizes what has been said “at most” in μ , i.e., for every formula α such that $\alpha \wedge fr^-(IP)$ is consistent, we know that $\mu \not\models \neg \alpha$.

A simple observation is that under the assumptions (i) (the change formula μ is consistent), and (ii) (no agent lies on her actual prior or revised belief base), we necessarily have $fr^-(IP) \models fr^+(IP)$, which is equivalent to state that $fr(IP) \neq \emptyset$. The set $fr(IP)$ contains all the interpolants of $fr^-(IP)$ and $fr^+(IP)$ in the sense of Craig / Lyndon theorem [Craig, 1957], but it is not restricted to them in general since one does not ask the admissible δ to be such that $Var(\delta) \subseteq \bigcup_{i=1}^n Var(\gamma_i)$.

Clearly enough, Proposition 2 has an interesting consequence on the computational viewpoint. Indeed, deciding whether a given formula belongs to the frame of an IP pair is not harder than the inference problem in propositional logic:

Proposition 3. *Given an IP pair IP and a formula $\alpha \in \mathcal{L}_{\mathcal{P}}$, the problem of deciding whether $\alpha \in fr(IP)$ is coNP-complete.*

Example 1 (continued). *From Proposition 2, we get that $fr(IP) = (fr^-(IP), fr^+(IP))$, where*

$$fr^-(IP) = e \wedge (s \Rightarrow (n \Leftrightarrow \neg f)) \wedge (\neg s \Rightarrow (n \wedge f))$$

and

$$fr^+(IP) = (s \vee n \vee \neg f) \wedge (\neg s \vee e \vee \neg n \vee f).$$

Hence, among what has been said by the person met by Amy and Beniko is that (cf. $fr^+(IP)$) if the park’s entrance is free, then it is not located far to the north; and if the park’s entrance is not free, then it is not located near to the south west.

Note that in the above example, little is known about μ : since $fr^-(IP)$ admits three models and $fr^+(IP)$ admits 13 models, there is significant uncertainty about μ , which can be assessed by the ratio $\frac{10}{16} = \frac{5}{8}$ given that for 10 interpretations out of 16 interpretations (the total number of interpretations over $\mathcal{P} = \{s, e, n, f\}$) one ignores whether it is a model or a counter-model of μ .

Thus, our greatest interest lies in complete frames. Taking advantage of Proposition 2, we know that a frame $fr(IP)$ is complete when $fr^-(IP) \equiv fr^+(IP)$, and in such a case, $fr(IP)$ is denoted by (δ) , where δ is any formula from the frame. Interestingly, one can easily identify the cases when a frame is complete:

Proposition 4. *Given an IP pair IP, $fr(IP)$ is complete if and only if $\bigvee_{\varphi_i \in \Phi, \gamma_i \in \Gamma} \varphi_i \vee \gamma_i$ is a valid formula.*

So this result means that if we have a varied enough set of bases, so that the disjunction of all these formulae is a valid formula, then we know that the announcement can be identified up to logical equivalence.

This result also tells us that the computational complexity of the frame completeness issue is at the first level of the polynomial hierarchy:

Proposition 5. *The problem of deciding whether the frame of an IP pair is complete is coNP-complete.*

4 Independence Assumptions

We now intend to show how the bounds of the logical interval representing a frame can be tightened in the case when the change formula μ is known to be independent from a given set of variables. We say that μ is *independent* of a set of propositional variables $X \subseteq \mathcal{P}$ if there exists a formula μ' logically equivalent to μ and such that $X \cap Var(\mu') = \emptyset$ [Lang et al., 2003]. Equivalently, μ is independent of $X \subseteq \mathcal{P}$ if μ is equivalent to $\exists X \cdot \mu$, i.e., the formula μ obtained after forgetting X :

Definition 4 (Forgetting [Lang et al., 2003]). *Let $X \subseteq \mathcal{P}$ and $\alpha \in \mathcal{L}_{\mathcal{P}}$. The forgetting of X in α is the formula $\exists X \cdot \alpha$ defined inductively as*

$$\exists X \cdot \alpha = \begin{cases} \alpha & \text{if } X = \emptyset, \\ \alpha_{x \leftarrow 0} \vee \alpha_{x \leftarrow 1} & \text{if } X = \{x\}, \\ \exists(X \setminus \{x\}) \cdot (\exists\{x\} \cdot \alpha) & \text{if } |X| > 1, x \in X, \end{cases}$$

where $\alpha_{x \leftarrow 0}$ (resp. $\alpha_{x \leftarrow 1}$) is the formula α obtained by replacing each occurrence of x in α by the constant \perp (resp. \top).

Given an IP pair IP and $X \subseteq \mathcal{P}$, we are then interested in the frame of IP projected onto $\overline{X} = \mathcal{P} \setminus X$, denoted by $fr(IP, \overline{X})$, defined as the subset of formulae from $fr(IP)$ that are independent of X . For any formula α , let us denote by $\forall X \cdot \alpha$ the formula $\neg \exists X \cdot \neg \alpha$:

Proposition 6. *Let IP be an IP pair and $X \subseteq \mathcal{P}$. Then*

$$fr(IP, \overline{X}) = (\exists X \cdot fr^-(IP), \forall X \cdot fr^+(IP)) \cap \mathcal{L}_{\overline{X}}.$$

Proof. (\subseteq) By definition of $fr(IP, \overline{X})$, we know that $\mu \in fr(IP, \overline{X})$. And by assumption, $\mu \equiv \exists X \cdot \mu$. So by Proposition 2, we get that $fr^-(IP) \models \exists X \cdot \mu$. Yet we know that for all formulae α, β , if $\alpha \models \exists X \cdot \beta$, then $\exists X \cdot \alpha \models \exists X \cdot \beta$. So we can write that $\exists X \cdot fr^-(IP) \models \exists X \cdot \mu$. Hence, $\exists X \cdot fr^-(IP) \models \mu$. Similarly, we can prove that $\mu \models \forall X \cdot fr^+(IP)$. This shows that $fr(IP, \overline{X}) \subseteq (\exists X \cdot fr^-(IP), \forall X \cdot fr^+(IP)) \cap \mathcal{L}_{\overline{X}}$.

(\supseteq) Let $\delta \in (\exists X \cdot fr^-(IP), \forall X \cdot fr^+(IP)) \cap \mathcal{L}_{\overline{X}}$. Yet we know that for any formula α , we have that $\alpha \models \exists X \cdot \alpha$ and $\forall X \cdot \alpha \models \alpha$. This means that $(\exists X \cdot fr^-(IP), \forall X \cdot fr^+(IP)) \subseteq (fr^-(IP), fr^+(IP))$, so $\delta \in fr(IP)$. Since we also have that $\delta \in \mathcal{L}_{\overline{X}}$, we get that $\delta \in fr(IP, \overline{X})$ by definition of $fr(IP, \overline{X})$. This concludes the proof. \square

Accordingly, the frame of an IP pair IP projected onto \overline{X} is also represented as a logical interval (restricted to the propositional language $\mathcal{L}_{\overline{X}}$), denoted by $fr(IP, \overline{X}) = (fr^-(IP, \overline{X}), fr^+(IP, \overline{X})) \cap \mathcal{L}_{\overline{X}}$, where $fr^-(IP, \overline{X}) = \exists X \cdot fr^-(IP)$ and $fr^+(IP, \overline{X}) = \forall X \cdot fr^+(IP)$.

Example 1 (continued). *Let us assume that in addition to their revised beliefs, Amy and Beniko have declared that the person they met did not mention anything about whether the park has a paid entrance or not. In this case, μ is independent of $\{f\}$. From Proposition 6, we get that $fr(IP, \overline{\{f\}}) = (fr^-(IP, \overline{\{f\}}), fr^+(IP, \overline{\{f\}})) \cap \mathcal{L}_{\overline{\{f\}}}$, where*

$$\begin{aligned} fr^-(IP, \overline{\{f\}}) &= \exists\{f\} \cdot (e \wedge (s \Rightarrow (n \Leftrightarrow \neg f)) \\ &\quad \wedge (\neg s \Rightarrow (n \wedge f))) \\ &\equiv e \wedge (s \vee n) \\ fr^+(IP, \overline{\{f\}}) &= \forall\{f\} \cdot ((s \vee n \vee \neg f) \\ &\quad \wedge (\neg s \vee e \vee \neg n \vee f)) \\ &\equiv (s \vee n) \wedge (\neg s \vee \neg n \vee e) \end{aligned}$$

Accordingly, f does not appear in $fr^-(IP, \overline{X})$ and $fr^+(IP, \overline{X})$. Note that here, the logical interval obtained by $fr(IP, \overline{\{f\}})$ is more restrained: among the total of eight interpretations, $fr^-(IP, \overline{\{f\}})$ admits three models and $fr^+(IP, \overline{\{f\}})$ admits five models, so there exist only two interpretations over $\{s, e, n\}$ for which one ignores whether they are models or counter-models of μ . Thus the uncertainty on μ can be quantified this time as $\frac{2}{8} = \frac{1}{4}$.

With such additional independence assumptions, the frame completeness issue can be characterized as follows:

Proposition 7. *Let IP be an IP pair and $X \subseteq \mathcal{P}$. Then $fr(IP, \overline{X})$ is complete if and only if $\exists X \cdot \bigvee_{\varphi_i \in \Phi, \gamma_i \in \Gamma} \varphi_i \vee \gamma_i$ is a valid formula.*

And as a consequence:

Proposition 8. *The problem of deciding whether the frame of an IP pair is complete whenever μ is known to be independent of $X \subseteq \mathcal{P}$ is Π_2^p -complete.*

Tu sum up, using additional independence information, a better approximation of μ can be obtained, but identifying whether the resulting frame is complete is more demanding from a computational point of view. Interestingly, when the bases from the prior and revised profiles are all represented as terms (i.e., conjunctions of literals), the complexity of the frame completeness problem drops to coNP:

Proposition 9. *The problem of deciding whether the frame of an IP pair is complete whenever μ is known to be independent of $X \subseteq \mathcal{P}$ is coNP-complete when all φ_i and γ_i are terms.*

5 When the Revision Operators are Known

So far we have considered a setting when the prior and revised bases are known for each agent, and we have assumed that each agent uses a revision operator satisfying the KM postulates. We now intend to show how the bounds of the logical interval representing a frame of the input IP pair can be tightened in the case when the revision operators used by the agents are known. One considers successively two scenarios: the case when each \circ_i is known, and its subcase when each \circ_i is a distance-based revision operator.

Each KM Revision Operator \circ_i is Provided Formally, we are now given an IP pair IP and a set $\Theta = \langle \circ_1, \dots, \circ_n \rangle$, where each \circ_i is the KM revision operator used by agent i . We are then interested in the frame of IP given Θ , denoted by $fr(IP, \Theta)$, which is defined as the subset of formulae δ from $fr(IP)$ such that for each agent $i \in \bar{n}$, $\varphi_i \circ_i \delta \equiv \gamma_i$, where \circ_i is specified by Θ . This frame can be characterized as follows:

Proposition 10. *Let IP be an IP pair and $\Theta = \langle \circ_1, \dots, \circ_n \rangle$. Then*

$$fr(IP, \Theta) = (fr^-(IP), \bigwedge_{i \in \bar{n}} (\alpha_i \vee \gamma_i)),$$

where for each $i \in \bar{n}$, α_i is any formula such that

$$[\alpha_i] = \{\omega \in \mathcal{W} \mid \exists \omega' \models \gamma_i, \omega' \prec_{\varphi_i}^i \omega\},$$

and $\preceq_{\varphi_i}^i$ is the total preorder corresponding to \circ_i and associated with φ_i through its faithful assignment.

So we also obtain a better approximation in this case, and the frame of an IP pair IP given Θ is also represented as a logical interval, denoted by $fr(IP, \Theta) = (fr^-(IP, \Theta), fr^+(IP, \Theta))$, where $fr^-(IP, \Theta) = fr^-(IP)$ and $fr^+(IP, \Theta) = fr^+(IP) \wedge \bigwedge_{i \in \bar{n}} \alpha_i$.

Propositions 6 and 10 state that the frame of an IP pair IP can be tightened in two different ways. Let us denote by $fr(IP, \bar{X}, \Theta)$ the frame of an IP pair IP projected onto \bar{X} given Θ , i.e., $fr(IP, \bar{X}, \Theta)$ is the subset of formulae from $fr(IP)$ that are independent of X and such that for each agent $i \in \bar{n}$, $\varphi_i \circ_i \delta \equiv \gamma_i$, where \circ_i is given by Θ . Interestingly, one can combine both results to strengthen the frame of IP:

Corollary 1. *Let IP be an IP pair, $\Theta = \langle \circ_1, \dots, \circ_n \rangle$, and assume that μ is independent of $X \subseteq \mathcal{P}$. Then*

$$fr(IP, \bar{X}, \Theta) = (fr^-(IP, \bar{X}, \Theta), fr^+(IP, \bar{X}, \Theta)),$$

where $fr^-(IP, \bar{X}, \Theta) = fr^-(IP, \bar{X}) \vee fr^-(IP, \Theta)$ and $fr^+(IP, \bar{X}, \Theta) = fr^+(IP, \bar{X}) \wedge fr^+(IP, \Theta)$.

Each \circ_i is a Distance-Based Revision Operator Exploiting Proposition 10 requires to consider a representation of each $\preceq_{\varphi_i}^i$ ($i \in \bar{n}$) in order to compute a corresponding α_i . Such representations are not always available. Things are easier when each $\circ_i \in \langle \circ_1, \dots, \circ_n \rangle$ is a distance-based revision operator characterized by a distance d_i since in this case, whatever φ_i , the corresponding preorder $\preceq_{\varphi_i}^i$ is induced from the underlying distance d_i . Indeed, with $k = \min_{\omega \in [\gamma_i]} d_i(\omega, \varphi_i)$, let $D_i(\varphi_i, > k)$ be any formula which is satisfied by the interpretations ω satisfying $d_i(\omega, \varphi_i) > k$. Then Proposition 10 can be specialized by replacing α_i by $D_i(\varphi_i, > k)$.

Example 1 (continued). *Assume that in addition to the assumption that μ is independent of $\{f\}$, one knows that $\Theta = \{\circ_1, \circ_2\}$, where $\circ_1 = \circ_2 = \circ_{Dal}$, i.e., both Amy and Beniko use Dalal's operator for revising their beliefs. From Corollary 1, we get that $fr(IP, \overline{\{f\}}, \Theta) = (fr^-(IP, \overline{\{f\}}, \Theta), fr^+(IP, \overline{\{f\}}, \Theta))$. On the one hand, we have that*

$$fr^-(IP, \overline{\{f\}}, \Theta) = fr^-(IP, \overline{\{f\}}) \equiv e \wedge (s \vee n).$$

On the other hand, we have that $fr^+(IP, \overline{\{f\}}, \Theta) = fr^+(IP, \overline{\{f\}}) \wedge fr^+(IP, \Theta)$; yet $fr^+(IP, \overline{\{f\}}) = (s \vee n) \wedge (\neg s \vee \neg n \vee e)$, and it can be verified that $fr^+(IP, \Theta) = fr^+(IP) \wedge \bigwedge_{i \in \bar{n}} \alpha_i$ (from Proposition 10), which is equivalent to $(s \vee n) \wedge (s \vee e \vee \neg n \vee \neg f) \wedge (\neg s \vee e \vee n \vee \neg f) \wedge (\neg s \vee e \vee \neg n \vee f)$. Overall, we get that

$$fr^+(IP, \overline{\{f\}}, \Theta) \equiv e \wedge (s \vee n).$$

Therefore, we get that

$$fr^-(IP, \overline{\{f\}}, \Theta) = fr^+(IP, \overline{\{f\}}, \Theta) \equiv e \wedge (s \vee n),$$

which means that $fr(IP, \overline{\{f\}}, \Theta)$ is complete, with $fr(IP, \overline{\{f\}}, \Theta) \equiv (e \wedge (s \vee n))$. We can deduce that what has been said to Amy and Beniko is precisely that the park is located to the south east, or near to the north east.

6 Some Experimental Results

This section reports experimental results about the number of agents needed to identify the change formula μ (up to logical equivalence) when each agent takes advantage of Dalal's operator for revising her beliefs.

Computing such a number requires to generate instances consisting of a change formula μ , and IP pairs for this μ . To this end, we have taken advantage of a random CNF generator available at <https://www.csc.kth.se/~giraldez/>. For a given number $p \geq 10$, one creates a CNF instance containing $\lceil 0.5 \times p \rceil$ variables and $\lceil 1.75 \times p \rceil$ clauses; if this formula is unsatisfiable, then one removes clauses from it in a greedy fashion until getting a satisfiable formula μ . The sizes of the clauses that are generated depend on the number of communities which is expected (this number varies between 1 and $\lceil \frac{p}{3} \rceil$). When $\lceil \frac{p}{3} \rceil < 3$, clauses of size 2 have been considered.

The bases φ_i (under the form of CNF formulae) are generated in the same way as μ considering $\lceil 0.8 \times p \rceil$ variables and $\lceil 2.8 \times p \rceil$ clauses. For each IP pair associated with μ , μ is identified whenever the upper bound of the logical interval associated with its frame implies the corresponding lower bound. Those implication tests are at the second level of the polynomial hierarchy (they correspond to instances of a Π_2^P -complete problem). Indeed, in the CNF encodings associated with the upper bounds, some extra-variables (existentially quantified) are introduced. In our approach, those variables are eliminated (using the resolution principle) so that the complexity of the resulting implication tests falls down to coNP, and a SAT solver is used to solve them. Obviously enough, the variable elimination step is computationally expensive (the output can be exponential in the size of the input). This explains why the number of variables considered in the experiments has been kept small.

The following figures synthesize the results we got by considering a set of 100 change formulae μ over the same set of 10 propositional variables.

Figure 1 (a-b) makes precise for each interval of number of agents (on the x -axis) the number of additional change formulae out of 100 (on the y -axis) that have been characterized in comparison with the previous interval. Figure 1 (a) corresponds to the case one just knows that these operators are KM

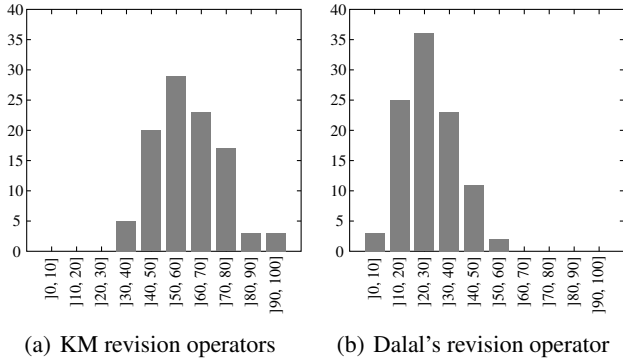


Figure 1: Number of agents required for identifying μ .

ones, and in Figure 1 (b), the reported results correspond to the case when we know that each agent uses Dalal’s revision operator (leading to tighter logical intervals).

Figure 2 presents a scatter plot enabling to compare the corresponding numbers of agents at the instance level: each point corresponds to an instance, its x -coordinate indicates the number of agents required to identify μ in case (b), and its y -coordinate indicates the number of agents required to identify μ in case (a).

One can observe on the two figures that exploiting the information given by the revision operators used by the agents (here, Dalal’s) leads to diminish significantly the number of agents required for the identification task. It is also interesting to note that the number of agents required to identify the announcement becomes quite reasonable when the revision operators are known (mostly around 20-30 agents).

7 Other Related Work

Several inverse problems of belief revision have been considered in the literature. [Booth and Nittka, 2008] considered a framework where one observes a sequence of belief states from a single agent that evolves over time. At each step we are given some formulae that hold and some formulae that do not hold in the belief state of the agent. The change formulae μ_i are also given, and a lexicographic revision (by μ_i) is assumed to be performed at each step. The main issue is to find the initial belief state. [Liberatore, 2015] considered a similar sequence-based framework, except that the belief base of the agent is fully known at each step. The focus is then given on the computational complexity of deciding whether the sequence allows for an initial epistemic state, given some assumption on the iterated revision operator that is used at each step. In both works, the unobservable component is on the agent’s side. This clearly departs from our approach where the unknown is the change formula μ .

More closely related to our setting is [Hunter *et al.*, 2017]: n agents are considered, and each agent i provides her belief base φ_i together with her revision operator \circ_i . In addition, a goal ψ_i is given for each agent i . The issue of interest is to determine whether there exists a formula μ such that $\varphi_i \circ_i \mu$, the revised base of each agent, entails the goal ψ_i associated with it. Clearly, our contribution departs from [Hunter *et al.*,

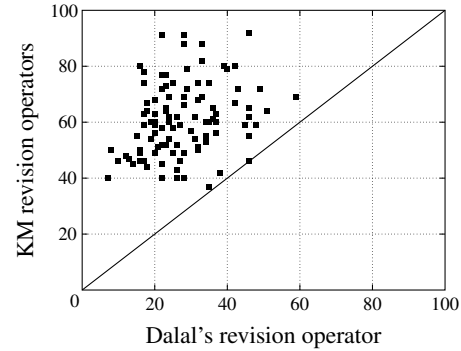


Figure 2: Comparison of the number of agents required for identifying μ in cases (a) and (b).

2017], both in terms of setting and objective. Rather than manipulating the beliefs of a set of agents, one wants to find what has been said to them, so we are not given goals to be achieved for the agents, but their *actual* revised bases. But, on the same vain, it could be interesting to study if our results can be exported for multi-agent epistemic planning problems (see, e.g., [Muise *et al.*, 2015]).

8 Conclusion and Perspectives

We have considered a belief revision scenario where an announcement (a propositional formula) μ is made to a group of agents. We assumed that we do not know μ , but we are given access for each agent to her prior beliefs and her beliefs after revision by μ . We proved that the set of μ satisfying those requirements has the form of a logical interval which can be characterized by compact bounds (i.e., formulae of size polynomial in the input size). We explained how this interval can be tightened under variable-independence assumptions or when the revision operator used by each agent is known. We also provided some experimental results about the identification issue, showing the benefits of taking into account such extra-information.

One of our findings is that the announcement μ can be completely characterized in the case when the disjunction of the agents’ prior beliefs and revised beliefs is a valid formula. A direct consequence of this result is that whenever one can identify a subset of agents such that the disjunction of their prior beliefs is valid, we know that μ *could* be characterized, were the revised beliefs of these agents available. Thus, when the access to those revised (and prior) beliefs is not “for free”, i.e., some underlying cost has to be paid to get them, the identification issue can be recast as a cost-minimization problem. Studying this optimization problem and designing algorithms can be interesting for studying causes for changes of opinion in social networks, allowing to complete the study of social networks opinion change using belief change tools [Schwind *et al.*, 2015; 2016]. This will be a perspective for further research.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number JP17H00763 and JP17K12746.

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.
- [Baltag and Moss, 2004] A. Baltag and L. S. Moss. Logics for epistemic programs. *Synthese*, 139(2):165–224, 2004.
- [Booth and Nittka, 2008] R. Booth and A. Nittka. Reconstructing an agent’s epistemic state from observations about its beliefs and non-beliefs. *Journal of Logic and Computation*, 18(5):755–782, 2008.
- [Craig, 1957] W. Craig. Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *Journal of Symbolic Logic*, 22:269–285, 1957.
- [Dalal, 1988] M. Dalal. Investigations into a theory of knowledge base revision: Preliminary report. In *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI’88)*, pages 475–479, 1988.
- [Hunter *et al.*, 2017] A. Hunter, F. Schwarzenruber, and E. Tsang. Belief manipulation through propositional announcements. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI’17)*, pages 1109–1115, 2017.
- [Katsuno and Mendelzon, 1991] H. Katsuno and A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.
- [Konieczny *et al.*, 2004] S. Konieczny, J. Lang, and P. Marquis. DA² merging operators. *Artificial Intelligence*, 157(1-2):49–79, 2004.
- [Lang *et al.*, 2003] J. Lang, P. Liberatore, and P. Marquis. Propositional independence: Formula-variable independence and forgetting. *Journal of Artificial Intelligence Research*, 18:391–443, 2003.
- [Lehmann *et al.*, 2001] D. J. Lehmann, M. Magidor, and K. Schlechta. Distance semantics for belief revision. *Journal of Symbolic Logic*, 66(1):295–317, 2001.
- [Liberatore, 2015] P. Liberatore. Revision by history. *Journal of Artificial Intelligence Research*, 52:287–329, 2015.
- [Muise *et al.*, 2015] C. Muise, V. Belle, P. Felli, S. McIlraith, T. Miller, A. Pearce, and L. Sonenberg. Planning over multi-agent epistemic states: A classical planning approach, 2015.
- [Schwind *et al.*, 2015] N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Belief revision games. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI’15)*, pages 1590–1596, 2015.
- [Schwind *et al.*, 2016] N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Is promoting beliefs useful to make them accepted in networks of agents? In *Proceedings of the Twenty Fifth International Joint Conference on Artificial Intelligence (IJCAI’16)*, pages 1237–1243, 2016.
- [van Ditmarsch *et al.*, 2008] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, 2008.