Private Revision in a Multi-Agent Setting

(Extended Abstract)

Thomas Caridroit
CRIL, Université d'Artois and CNRS, France
caridroit@cril.fr

Sébastien Konieczny
CRIL and CNRS, France
konieczny@cril.fr

Tiago de Lima
CRIL, Université d'Artois and CNRS, France
delima@cril.fr

Pierre Marquis
CRIL, Université d'Artois and CNRS, France
marquis@cril.fr

ABSTRACT

AGM belief change aims at modeling the evolution of an agent's beliefs about her environment. In many applications though, a set of agents sharing the same environment must be considered. For such scenarios, beliefs about other agents' beliefs must be taken into account. In this work, we study the private revision issue in a multi-agent setting represented by a KD45$_n$ model. More precisely, we investigate the changes induced by a new piece of objective information made available to one agent in the set. We point out an adaptation of AGM revision postulates to this setting, and present some specific revision operators.

1. INTRODUCTION

The main theoretical framework for belief change is AGM (Alchourrón-Gärdenfors-Makinson) and its developments [7, 1].

In most works on belief revision, the belief set of the agent consists of beliefs about the environment (the world), and is represented by a set of formulas in classical logic. However, in many applications, an agent is not alone in her environment, but shares it with other agents, who also have beliefs. Beliefs about the beliefs of other agents is an important piece of information to be considered for making the best decisions and performing the best actions. Using beliefs on beliefs of other agents is for instance crucial in game theory [9].

Here, our objective is to design operators that change the beliefs of the agents in KD45$_n$ models. This task is more complicated than in the standard AGM framework, because, in a multi-agent context, the new pieces of evidence can take different forms. In particular, a new piece of evidence can be either observed/transmitted/available to every agent or only to some of them. This kind of issue has already been studied in dynamic epistemic logic, where public and private announcements lead to distinct belief changes [10, 4].

2. PRELIMINARIES

We focus on a multi-agent setting where each agent has her own beliefs about the state of the world. Formally, let $A = \{1, \ldots, n\}$ be a finite set of agents. We consider the language L containing a propositional language L_0 plus one belief operator B_i for each agent $i \in A$. We use B_i^k to abbreviate a sequence of k operators B_i (i.e., $\mathrm{B}_i^0\varphi$ abbreviates φ and $\mathrm{B}_i^{k+1}\varphi$ abbreviates $\mathrm{B}_i\mathrm{B}_i^k\varphi$, for $k \geq 0$.) A formula of the form $\mathrm{B}_i\varphi$ is read 'agent i believes that φ is true'. Formulas in L_0 are called objective formulas, while subjective formulas are formulas of the form $\mathrm{B}_i\varphi$ and $\neg\mathrm{B}_i\varphi$, for $\varphi \in L$.

To capture agent beliefs, the semantics of L is supposed to be ruled by the standard system KD45$_n$ [6]. This system can be characterized using a specific class of Kripke models.

Definition 1 (Kripke Model). *A Kripke model is a tuple $\langle W, R, V\rangle$ where W is a non-empty set of worlds, $R = \{R_i \mid i \in A\}$, with R_i a binary accessibility relation for agent i, and $V : W \to 2^P$ is a valuation function. For each possible world $w \in W$, $V(w)$ is the set of propositional variables which are true at w. A pointed Kripke model is a pair (M, w), where $M = \langle W, R, V\rangle$ is a Kripke model and $w \in W$ is the real world.*

$R_i(w)$ denotes the set of possible worlds that are accessible from w for agent i, that is, $R_i(w) = \{w' \mid (w, w') \in R_i\}$. We note $(M, w) \models \varphi$ the fact that the formula φ is satisfied at the world w in the model M. This notion is defined using the usual satisfaction relation such that $(M, w) \models \mathrm{B}_i\varphi$ iff $\forall w' \in W$ if $(w, w') \in R_i$ then $(M, w') \models \varphi$. We use $\|\varphi\|_M$ to denote the set of possible worlds of M that satisfy φ, that is, $\|\varphi\|_M = \{w : w \in W \text{ and } (M, w) \models \varphi\}$.

Two pointed Kripke models may satisfy the same set of formulas. Such a pair of models is then considered equivalent. It is known that if two pointed Kripke models are bisimilar (noted $(M, w) \underline{\leftrightarrow} (M', w')$) (for the definition, please see [5]), then they are equivalent.

A pointed KD45$_n$ model (M, w) represents a set of n belief sets $K_i^{(M,w)}$, one for each agent $i \in A$, where $K_i^{(M,w)} = \{\varphi \mid (M, w) \models B_i\varphi\}$. We also define the objective belief set of agent i (i.e., what i believes about the state of the world). This is the set $O_i^{(M,w)} = K_i^{(M,w)} \cap L_0$.

3. PRIVATE REVISION

Let the result of the private revision of the KD45$_n$ pointed model (M, w) by the objective formula φ for agent a be denoted by $(M, w) \star_a \varphi = (\langle W', R', V'\rangle, w')$. The AGM postulates for revision can be adapted as follows to this case:

(R_n0) $V'(w') = V(w)$

(R_n1) $(M, w) \star_a \varphi \in$ KD45$_n$

(R_n2) $(M, w) \star_a \varphi \models \mathrm{B}_a\varphi$

(R_n3) $(M, w) \models \mathrm{B}_i\psi$ iff $(M, w) \star_a \varphi \models \mathrm{B}_i\psi$, for $i \neq a$

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

(R_n4) $(M, w) \models \mathrm{B}_a^k \mathrm{B}_i \psi$ iff $(M, w) \star_a \varphi \models \mathrm{B}_a^k \mathrm{B}_i \psi$, for $i \neq a$

(R_n5) If $(M, w) \star_a \varphi \models \mathrm{B}_i \psi$ then $(M, w) +_a \varphi \models \mathrm{B}_i \psi$

(R_n6) If $(M, w) \not\models \mathrm{B}_a \neg\varphi$, then $(M, w) +_a \varphi \,\underline{\leftrightarrow}\, (M, w) \star_a \varphi$

(R_n7) If $(M^1, w^1) \,\underline{\leftrightarrow}\, (M^2, w^2)$ and $\models \varphi \equiv \psi$,
then $(M^1, w^1) \star_a \varphi \,\underline{\leftrightarrow}\, (M^2, w^2) \star_a \psi$

(R_n8) If $(M, w) \star_a (\varphi \wedge \psi) \models \mathrm{B}_i \chi$,
then $((M, w) \star_a \varphi) +_a \psi \models \mathrm{B}_i \chi$

(R_n9) If $(M, w) \star_a \varphi \not\models \mathrm{B}_a \neg\psi$, then $((M, w) \star_a \varphi) +_a \psi \models \mathrm{B}_i \chi$
implies $(M, w) \star_a (\varphi \wedge \psi) \models \mathrm{B}_i \chi$.

(R_n1) ensures that the model obtained after a revision is still a KD45$_n$ model. (R_n2) is the success postulate, it states that φ is believed by a after the revision. (R_n3) states that the beliefs of any agent except a do not change. (R_n4) states that the beliefs of agent a about the other agents do not change. (R_n5) and (R_n6) state that when the new piece of evidence is consistent with the beliefs of the agent, revision is just expansion in AGM sense. (R_n7) asks for irrelevance of syntax, stating that if two formulas are logically equivalent, then they lead to the same revision results. (R_n8) and (R_n9) make precise when the revision by a conjunction can be obtained by a revision followed by an expansion.

One can show that the revision operators characterized by (R_n0)–(R_n9) are conservative extensions of usual AGM belief revision operators:

PROPOSITION 2. *Let $\star_i$ be an revision operator satisfying (R_n0)–(R_n9). The $\star$ operator defined as $O_i^{(M,w)} \star \varphi = O_i^{(M,w) \star_i \varphi}$ is an AGM revision operator (i.e., it satisfies (K*1)–(K*8) [1]).*

Let us now define a family of private revision operators parameterized by AGM belief revision operators $\circ$:

DEFINITION 3. ***Revision of** (M, w_0) **by** φ **for agent** a.*
Let $(M, w_0) = (\langle W, R, V \rangle, w_0)$ be pointed a KD45$_n$ model, let φ be a consistent objective formula (i.e., $\varphi \in L_0$), and let $\circ$ be an AGM revision operator. We define the private revision of (M, w_0) by φ for agent a (with revision operator $\circ$) as $(M, w_0) \star_a^\circ \varphi = (\langle W', R', V' \rangle, w_0')$, such that:

- *if $R_a(w_0) \cap \|\varphi\|_M \neq \emptyset$*
  - *then $E = \{V(w) \mid w \in R_a(w_0) \cap \|\varphi\|_M\}$*
  - *else $E = \{e \mid e \subseteq P$ and $e \models O_a^{(M,w_0)} \circ \varphi\}$*
- *$W' = W \cup W^\varphi \cup \{w_0'\}$ where*
  - *$W^\varphi = \bigcup_{w \in R_a(w_0)} W_w^\varphi$ and $W_w^\varphi = \bigcup_{e \in E} \{v_w^e\}$*
- *$R_a' = R_a \cup R_a^\varphi \cup R_a^0$ where*
  - *$R_a^\varphi = \{(w_1^\varphi, w_2^\varphi) | w_1^\varphi, w_2^\varphi \in W^\varphi\}$*
  - *$R_a^0 = \{(w_0', w^\varphi) | w^\varphi \in W^\varphi\}$*
- *$R_i' = R_i \cup R_i^{\overrightarrow{\varphi}} \cup R_i^0$ for $i \neq a$, where*
  - *$R_i^{\overrightarrow{\varphi}} = \{(v_w^e, w') | w R_i w', v_w^e \in W^\varphi\}$ for $i \neq a$*
  - *$R_i^0 = \{(w_0', w) | (w_0, w) \in R_i\}$ for $i \neq a$*
- *$V'(w) = V(w)$ for $w \in W$*
- *$V'(v_w^e) = e$ for $v_w^e \in W^\varphi$*
- *$V'(w_0') = V(w_0)$*

If the revision formula φ is considered possible by agent a, she performs an expansion, otherwise, each of the worlds of the new set W^φ has as valuation a (propositional) model of the new information φ. Interestingly, such operators exhibit good logical properties:

PROPOSITION 4. *The operators $\star_a^\circ$ satisfy (R_n0)–(R_n9).*

Let us finally illustrate the behaviour of our private revision operators on a simple example.

EXAMPLE 5. *We consider the model (M, w_0) of Figure 1, where agent 1 believes $\neg x \wedge \neg y$ and believes that agent 2 believes $x \wedge y$. Agent 2 believes $x \wedge y$ and believes that agent 1 believes $x \leftrightarrow y$. After the revision by $x \wedge y$, agent 1 must believe $x \wedge y$. The beliefs of agent 2 remain unchanged. The obtained model (M', w_0') is reported as well at Figure 1. In this example, agent 1 uses Dalal's AGM revision operator $\circ_D$ [8]. We can observed that the revised model obtained using Definition 3 may be non-minimal. Nevertheless, a minimal model can be obtained via a bisimulation contraction. Here, this leads to the model (M'', w_0').*

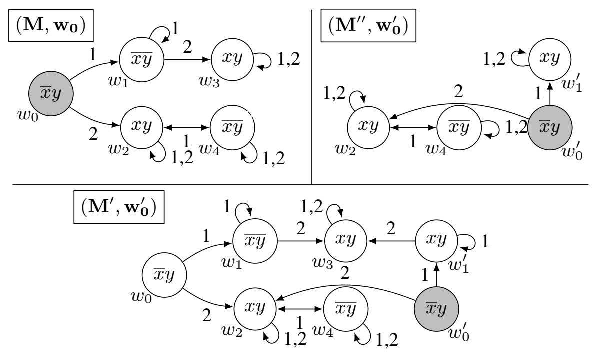

Figure 1: $(M'', w_0') \,\underline{\leftrightarrow}\, (M, w_0) \star_1^{\circ_D} (x \wedge y)$

4. CONCLUSION AND RELATED WORK

The closest work to our own one is the study of private expansion and revision made by Aucher [3, 2]. Aucher allows revision by subjective formulas and compute distances between the corresponding (epistemic) models. Aucher's revision does not allow the agent concerned by the private revision to choose, among the models of the objective formulas, the most plausible ones. We can do that thanks to the underlying AGM revision operators in the definition of the private revision operator. So our private revision result implies (sometimes strictly) the result given by Aucher's revision.

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*, 50:510–530, 1985.

[2] G. Aucher. *Perspectives on belief and change*. PhD thesis, Université Paul Sabatier; University of Otago, 2008.

[3] G. Aucher. Generalizing AGM to a multi-agent setting. *Logic Journal of the IGPL*, 18(4):530–558, 2010.

[4] G. Aucher. Private announcement and belief expansion: an internal perspective. *Journal of Logic and Computation*, 22(3):451–479, 2012.

[5] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.

[6] R. Fagin, J. Y. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. The MIT Press, 1995.

[7] P. Gärdenfors. *Knowledge in flux*. MIT Press, 1988.

[8] H. Katsuno and A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.

[9] J. van Benthem. *Logic in Games*. MIT Press, 2014.

[10] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, 2008.