Proceedings of the 23rd International Conference on Principles of Knowledge Representation and Reasoning
Main Track

A Study of Belief Revision Postulates in Multi-Agent Systems

Michael Thielscher[1] , **Tran Cao Son**[2]
[1]University of New South Wales, Sydney, Australia
[2]New Mexico State University, Las Cruces, New Mexico, USA
mit@unsw.edu.au, stran@nmsu.edu

Abstract

We investigate the belief revision problem in epistemic planning, i.e., what will be the beliefs of all agents in a multi-agent system after an agent gains the belief in some state property. Based on the standard representation in epistemic planning of agents' beliefs via a *single multi-agent Kripke model*, we generalize the classical AGM belief revision postulates to the multi-agent setting, with the aim to provide a formal framework for evaluating dynamic epistemic reasoning frameworks in which the beliefs of all agents as the result of actions are computed. As an example of a simple operator that satisfies all of the generalized AGM postulates, we present generalized *full-meet* multi-agent belief revision. We moreover define a generalization of the standard postulates for *iterated* revision, present a more sophisticated, event model based revision operator, and discuss the potential issues in defining an epistemic operator on Kripke models that can satisfy all of the generalized postulates for iterated multi-agent belief revision.

1 Introduction

Multi-agent epistemic reasoning about actions and planning has garnered much research attention recently, as a formal framework for controlling heterogeneous, collaborative—or competitive—agents and robots. An example are service robots that interact with humans (Baral et al. 2017); other emerging applications for dynamic epistemic reasoning include the combination with large language models in order to plan interactions based on user beliefs (Davila et al. 2024); model reconciliation (Vasileiou et al. 2024); and reasoning about epistemic responsibility (Chockler and Halpern 2004) applied in legal contexts.

Dynamic epistemic logic (DEL) is considered a quasi standard as the most expressive formalism for modeling these domains (Bolander and Andersen 2011; Burigana and Fabiano 2026). In this setting, the knowledge that different agents have, both of the environment and of each others' knowledge, is represented by a *single Kripke model* of the possible states that the environment could be in according to the limited information of the agents. Such a multi-agent Kripke structure encodes all the beliefs of all agents, including those of higher order, i.e., beliefs about each other's beliefs. The latter is essential in most settings for dynamic epistemic reasoning when agents need to take into account what they know of the other (cooperating or competing) agents' knowledge or beliefs. The use of a *single* Kripke model to encode all beliefs of all agents in epistemic planning is motivated by efficiency: In order to predict the effects of a sequence of actions for the purpose of planning, it suffices to update one initial Kripke model action by action (Baral et al. 2022; Burigana and Fabiano 2026).

Actions are encoded by *event models*, which similarly represent what the different agents know and can observe about the effects of an action. Combining one with the other, known as *event model update*, results in a new Kripke model representing the updated knowledge after the action (Bolander 2017). As an example, the Kripke model on the left in the figure below represents the beliefs of two agents a and b about the status of a coin (h: heads up); the event model in the middle (with colored arrows) encodes the pulic announcement of $\neg h$; and the rightmost model represents the beliefs of a, b after the announcement:

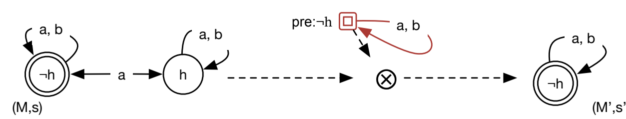

Since event models can be complex to specify directly because they represent all possible views on an action occurrence, several high-level languages for describing actions in epistemic multi-agent environments have been developed (Baral et al. 2015; Rajaratnam and Thielscher 2021). Multi-agent epistemic reasoning and planning is thus well-understood when agents have (incomplete) *knowledge* of their environment; and event model update provides a clear semantics of how to update that knowledge after an action.

In reality, however, actions and sensors are often uncertain, so that agents can only hold *beliefs*. These may be wrong, in which case they may need to be revised as a result of an action that reveals a false belief. But there is no established theory of how a Kripke structure should be updated to reflect this, and yet at the same time, multi-agent belief revision is becoming ever more important in the area of agentic AI: in collaborative plan execution, for example, when agents have different beliefs about the precondition of actions they need to execute; when a chatbot tries to persuade a user to purchase a product and must be able to reason about the mental state of the user to decide on a 'right' price; when an AI system wishes to convince a human of its trustworthiness and has to reconcile its beliefs about the human's model with its own beliefs in order to provide explanations about its actions; and for formal accounts of responsibility by reasoning about beliefs of plaintiffs and defendants.

Classical belief revision deals with the problem of identifying the beliefs of a single agent in the presence of a new piece of information. In their seminal work, Alchourrón, Gärdenfors, and Makinson (1985) proposed the foundational properties, referred to as *AGM Postulates*, that a rational belief revision operator should satisfy. Darwiche and Pearl (1997) added four properties for *iterated* belief revision, commonly referred to as *DP Postulates*, which focus on sequences of revisions. The postulates have been extensively studied in the literature, including concrete revision operators based on the distance between models before and after a revision (Dalal 1988; Satoh 1988); belief revision based on possible worlds (Katsuno and Mendelzon 1992) and in the context of answer set programming (Aravanis and Peppas 2017; Delgrande et al. 2008); and iterated belief revision and change for dynamic worlds (Baltag and Smets 2008; Hunter and Delgrande 2011; Shapiro et al. 2011; Peppas 2014; Tardivo et al. 2021).

Multi-agent revision has also been considered very early on; however, most of this work (Dragoni and Puliti 1994; Dragoni and Giorgini 1996; Liu and Williams 1999; Malheiro, Jennings, and Oliveira 1994) focuses on maintaining the consistency of the local knowledge bases of agents in a distributed manner. Proposals for changing the beliefs of multiple agents in dynamic environments mostly focus on belief *updates* (Baltag, Moss, and Solecki 1998; Herzig, Lang, and Marquis 2005; van Benthem, van Eijck, and Kooi 2006; van Benthem 2007; van Ditmarsch, van der Hoek, and Kooi 2007; Lorini and Schwarzentruber 2021; Baral et al. 2022). To the best of our knowledge, the only systematic study on the AGM and the DP postulates in belief revision in multi-agent environements is by Aucher (2010), who considers belief sets represented by *all* Kripke models consistent with a set of belief formulas. This allows for a straightforward application of the AGM postulates to a multi-agent setting but does not address the issue of revising a *single* Kripke model as used in practical approaches to epistemic reasoning about actions and planning.

In this paper, *we expand the logic of belief revision to the dynamic, multi-agent setting in order to provide a general, systematic framework for assessing multi-agent belief revision operators that are defined over a single Kripke model as used in dynamic epistemic reasoning about actions and planning*. To this end, we develop a suitable multi-agent generalization of the AGM postulates. We present a simple, generalized so-called *full-meet* multi-agent belief revision operator and prove that it satisfies the generalized AGM postulates. We also generalize the DP postulates to iterated multi-agent belief revision and show that full-meet satisfies all but one of them. We then develop a more sophisticated, event model-based belief revision operator and analyze different strategies to define such operators and discuss their consequences in satisfying the DP posulates.

All omitted proofs are available in an accompanying technical report (Thielscher and Son 2026).

2 Background

We begin with a concise outline of the formal foundations of dynamic epistemic reasoning and belief revision

Dynamic Epistemic Reasoning A *multi-agent* domain is defined by a pair $\langle \mathcal{A}, \mathcal{P} \rangle$ where $\mathcal{A}$ is a finite and non-empty set of agents and $\mathcal{P}$ a set of propositions. *Belief formulas* over $\langle \mathcal{A}, \mathcal{P} \rangle$ are defined by the BNF: $\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \mathbf{B}_i\varphi$ where $p \in \mathcal{P}$ and $i \in \mathcal{A}$. Standard connectives like $\rightarrow$ and $\leftrightarrow$ are used through their usual abbreviations. A belief formula which does not contain any occurrence of $\mathbf{B}_i$ is referred to as a proposition formula. $\mathcal{L}_\mathcal{P}$ (resp. $\mathcal{L}_\mathcal{A}$) denotes the set of proposition formulas over $\mathcal{P}$ (resp. the set of belief formulas over $\langle \mathcal{A}, \mathcal{P} \rangle$).

Satisfaction of belief formulas is defined over *pointed Kripke structures* (a.k.a. pointed Kripke models or epistemic states) (Fagin et al. 1995). A Kripke structure/model M is a tuple $\langle W, \{R_a\}_{a\in\mathcal{A}}, \pi \rangle$, where W is a set of worlds, $\pi : W \mapsto 2^\mathcal{P}$ is a function that associates an interpretation of $\mathcal{P}$ to each element of W, and for $a \in \mathcal{A}$, $R_a \subseteq W \times W$ is a binary relation over W. We write $R_a(u, w)$ and use this notation interchangeably with $(u, w) \in R_a$. For $u \in W$ and $\varphi \in \mathcal{L}_\mathcal{P}$, $M[\pi](u)$ and $M[\pi](u)(\varphi)$ denote the interpretation associated with u via π and the truth value of φ with respect to $M[\pi](u)$. For a world $s \in W$, referred to as *true state of the world*, (M, s) is a *pointed Kripke structure*.

The satisfaction relation $\models$ between a state (M, s) and belief formulas is defined as follows: (*i*) $(M, s) \models p$ if $p \in \mathcal{P}$ and $M[\pi](s) \models p$; (*ii*) $(M, s) \models \mathbf{B}_i\varphi$ if $\forall t.[(s, t) \in R_i \Rightarrow (M, t) \models \varphi]$; (*iii*) $(M, s) \models \neg\varphi$ if $(M, s) \not\models \varphi$; (*iv*) $(M, s) \models \varphi_1 \vee \varphi_2$ if $(M, s) \models \varphi_1$ or $(M, s) \models \varphi_2$; (*v*) $(M, s) \models \varphi_1 \wedge \varphi_2$ if $(M, s) \models \varphi_1$ and $(M, s) \models \varphi_2$.

Two pointed Kripke structures $\langle W, \{R_a\}_{a\in\mathcal{A}}, \pi \rangle$ and $\langle W', \{R'_a\}_{a\in\mathcal{A}}, \pi' \rangle$ are *bisimilar* if there is a relation $\mathcal{Z} \subseteq W \times W'$ such that for all $(w, w') \in \mathcal{Z}$: (*i*) $M[\pi](w) = M'[\pi'](w')$; (*ii*) for each $w_1 \in W$ such that $R_a(w, w_1)$, $R'_a(w', w'_1)$ for some $(w_1, w'_1) \in \mathcal{Z}$; (*iii*) for each $w'_1 \in W'$ such that $R'_a(w', w'_1)$, $R_a(w, w_1)$ for some $(w_1, w'_1) \in \mathcal{Z}$.

For a set of formulas A, $Cn(A)$ denotes the set of logical consequences of A. Cn is assumed to be supraclassical, i.e., if p can be derived from A by classical truth-functional logic, then $p \in Cn(A)$. Cn satisfies the following properties: (*i*) $A \subseteq Cn(A)$; (*ii*) if $A \subseteq B$ then $Cn(A) \subseteq Cn(B)$; (*ii*) $Cn(A) = Cn(Cn(A))$. We say that A is a belief set if and only if $A = Cn(A)$.

For a set K of formulas, $K \vdash p$ (resp. $K \not\vdash p$) is an alternative notation for $p \in Cn(K)$ (resp. $p \notin Cn(K)$). $Cn(\emptyset)$ is the set of tautologies. The *expansion* of K by a formula p, i.e., the operation that just adds p and removes nothing, is denoted $K + p$ and defined by: $K + p = Cn(K \cup \{p\})$.

Logic of Belief Revision and AGM Postulates We follow Darwiche and Pearl (1997) and state the standard postulates on the basis of *epistemic states* M, which implicitly determine a set of beliefs K_M. Two epistemic states are *equivalent*, written $M_1 \equiv M_2$, iff $K_{M_1} = K_{M_2}$, that is, they entail the same beliefs. The intuitive meaning of $M * p$ is to revise the beliefs so as to ensure that K_{M*p} contains p and is consistent (unless p is inconsistent). Alchourrón, Gärdenfors,

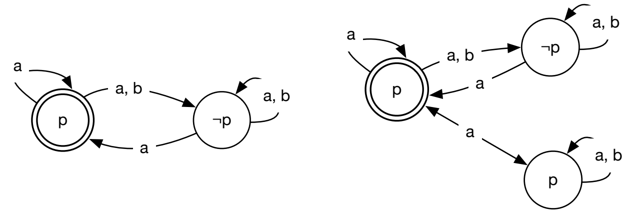

Figure 1: Two pointed Kripke structures (M_1, s) and (M_2, s) with the same "true" world s (marked with a double circle — so both structures agree that p is actually true) but which determine different multi-agent belief sets $K_1 = K_{(M_1,s)}$ and $K_2 = K_{(M_2,s)}$.

and Makinson (1985), a.k.a. AGM, proposed the following eight basic postulates for one-shot belief revision:

- *Closure*: $K_{M*p} = Cn(K_{M*p})$
- *Success*: $p \in K_{M*p}$
- *Inclusion*: $K_{M*p} \subseteq K_M + p$
- *Vacuity*: if $\neg p \notin K$ then $K_{M*p} = K_M + p$
- *Consistency*: K_{M*p} is consistent if p is consistent
- *Extensionality*: if $\vdash p \leftrightarrow q$ then $M * p \equiv M * q$
- *Superexpansion*: $K_{M*(p\wedge q)} \subseteq K_{M*p} + q$
- *Subexpansion*: $K_{M*p} + q \subseteq K_{M*(p\wedge q)}$ if $\neg q \notin K_{M*p}$

Darwiche and Pearl (1997) augmented the AGM framework by four more postulates for *iterated* revision:

(DP1) if $q \vdash p$ then $K_{(M*p)*q} = K_{M*q}$

(DP2) if $q \vdash \neg p$ then $K_{(M*p)*q} = K_{M*q}$

(DP3) if $p \in K_{M*q}$ then $p \in K_{(M*p)*q}$

(DP4) if $\neg p \notin K_{M*q}$ then $\neg p \notin K_{(M*p)*q}$

Another standard postulate, which strengthens (**DP3**) and (**DP4**), is the following (Jin and Thielscher 2007):

(IN) *Independence*: if $\neg p \notin K_{M*q}$ then $p \in K_{(M*p)*q}$

3 Multi-agent Belief Revision (MBR) — Basic Concepts

We begin by discussing in detail the foundations for multi-agent belief revision in the context of dynamic epistemic reasoning, where single Kripke models represent the beliefs of all agents. In the section that follows, we then develop generalizations of the AGM postulates to this setting.

Belief sets As customary in epistemic reasoning about actions and planning, a *belief set* in MBR shall be represented by a single, pointed Kripke structure (M, s) and is the set of all formulas entailed by (M, s):

$$K_{(M,s)} \;=\; \{\varphi \in \mathcal{L}_{\mathcal{A}} \mid (M,s) \models \varphi\}$$

From now on, we will assume that for any pointed Kripke structure (M, s) in discussion, s is the true state of the world.

As an example, Fig. 1 shows two pointed Kripke structures, with agents a and b, side by side. The egdes labelled a, b from the actual world into a world in which $\neg p$ holds indicate that both agents consider it possible that p is false. In fact, agent b (falsely) believes in $\neg p$ because this is the only b-accessible world from the true world, s. Agent a, on the other hand, considers both p and $\neg p$ possible because there is also a self-loop on s labelled a. Also, in both belief sets K_1 and K_2, b believes that a does not believe in either, that is, $\mathbf{B}_b[(\neg\mathbf{B}_a p) \wedge (\neg\mathbf{B}_a \neg p)]$, because from the only b-accessible world there is one a-accessible world with p and one with $\neg p$. The two belief sets differ in another higher-degree belief, however: Only in K_1, a believes that b believes $\neg p$. Formally, $\mathbf{B}_a\mathbf{B}_b\,\neg p \in K_1$ whereas $\neg\mathbf{B}_a\mathbf{B}_b\,\neg p \in K_2$, because in (M_2, s) there is an additional a-accessible world in which b does believe p.

Note that *bisimilar* multi-agent Kripke models define the same belief set provided the two designated worlds coincide in their interpretation. When two structures are *dissimilar*, they still induce the same belief set if they entail the same belief formulas from the respective true state of the world.

Deductive closure of a belief set For a belief set K, the deductive closure $Cn(K)$ is given by the entailment relation "$\vdash$". Any belief set therefore satisfies $Cn(K) = K$.

Consistent beliefs/belief sets Agents can have inconsistent beliefs: By definition, $\mathbf{B}_a\bot \in K_{(M,s)}$ if, and only if, $R_a(s) = \emptyset$ in M. Agents can also believe that other agents have inconsistent beliefs etc., but a belief set itself is always consistent because $M[\pi](s) \not\models \bot$ for any pointed Kripke structure. Hence, the concept of consistency in the classical AGM postulates needs to be adapted to consistency of the beliefs *of individual agents* in the multi-agent setting when representing a belief set by a single Kripke structure.

Revision With the aim to model actions that change the beliefs of agents (e.g., sensing actions, announcement actions) in a multi-agent environment, we consider revising belief sets to reflect the result of one agent making an observation or receiving some information about the environment. Formally, a belief set will be revised by *first-degree belief formulas*. These are defined as

$$\mathcal{B}_{\mathcal{A},\mathcal{P}} = \{\mathbf{B}_a\varphi \mid a \in \mathcal{A},\ \varphi \text{ proposition formula over } \mathcal{P}\}$$

The belief set after revision by a first-degree belief $\mathbf{B}_a\varphi \in \mathcal{B}_{\mathcal{A},\mathcal{P}}$ should again be represented by a single, pointed Kripke structure, denoted by $(M, s) * \mathbf{B}_a\varphi$. We will simply write $K * \mathbf{B}_a\varphi$ to refer to the belief set $K_{(M,s)*\mathbf{B}_a\varphi}$ when it is clear from the context that (M, s) is the underlying Kripke model of the belief set K.

Subset relation over belief sets The classical postulates also require us to define the concept of a subset relation among belief sets. Because every belief set is represented by a single Kripke structure, we cannot define this relation based on the set of all formulas entailed, since otherwise the subset relation would be satisfied only if the two Kripke structures entail identical sets of formulas. Therefore, and in line with the definiton of revision formulas, we consider all formulas of the form $\mathbf{B}_a\varphi$, where $a \in \mathcal{A}$ and φ is a proposition formula over $\mathcal{P}$, when comparing two belief sets:

$$K_{(M_1,s_1)} \subseteq K_{(M_2,s_2)}$$
$$\text{iff}$$
$$\forall\, \mathbf{B}_a\varphi \in \mathcal{B}_{\mathcal{A},\mathcal{P}}.[(M_1, s_1) \models \mathbf{B}_a\varphi \Rightarrow (M_2, s_2) \models \mathbf{B}_a\varphi]$$

Recall, for example, the belief sets represented in Fig. 1. Although $K_1 \neq K_2$, both of them entail the same first-degree belief formulas. This follows from the fact that for both Kripke structures (M, s), we have

$$\exists(s,w),(s,w') \in R_a.\ M[\pi](w) \models p \wedge M[\pi](w') \models \neg p \\ \wedge \forall(s,w) \in R_b.\ M[\pi](w) \models p$$

Hence, in both models it holds that $(M, s) \models \mathbf{B}_a\varphi$ iff φ is a tautological propositional formula while $(M_i, s) \models \mathbf{B}_b\varphi$ iff $p \vdash \varphi$. Consequently, $K_1 \subseteq K_2$ and $K_2 \subseteq K_1$.

Minimal belief sets A consequence of the above definition is that the "smallest" (w.r.t. the subset relation) representable belief sets are exactly those that include only first-degree belief formulas of the form $\mathbf{B}_a\varphi$ with φ a propositional tautology over $\mathcal{P}$. There are different Kripke structures that can be used to represent this set; a *generic* minimal Kripke structure can be constructed as follows: $M_\emptyset = \langle W, \{R_a\}_{a\in\mathcal{A}}, \pi\rangle$ with $W = 2^{\mathcal{P}}$, $R_i = W \times W$ for all $i \in \mathcal{A}$, and $\pi(w) = w$.

Lemma 1. *For any $s \in 2^{\mathcal{P}}$ we have that $K_{(M_\emptyset,s)} \subseteq K$ for all belief sets K.*

It is worth stressing that not all smallest belief sets are equal as they can be based on structurally different Kripke models and hence contain different nested beliefs. For example, $(M_\emptyset, s)$ always entails $\mathbf{B}_a\neg\mathbf{B}_b\varphi$ for all agents $a, b \in \mathcal{A}$ and non-tautological proposition formulas φ, that is, every agent believes that no agent believes in anything other than tautological properties about the environment. This may not be the case in other minimal belief sets.

Expansion A key concept in the classical AGM postulates is the expansion of a belief set by a new belief. The intuition behind this concept is to add a new belief while retaining the existing beliefs, together with all the logical consequences of the old and new beliefs, but in a minimal fashion. The intuition behind the following generalization of this concept to multi-agent beliefs given by a pointed Kripke structure, is that $K + \mathbf{B}_a\varphi$ is obtained by constructing a new structure that is a combination of: (1) M (the old structure); (2) a new true state of the world s'; and (3) a "replica" structure obtained from M by (i) removing all links labeled a into a world in which φ is false, and (ii) adding links labeled a going from s' to worlds in which there is a link labeled a.

Formally, let $M = (W, \{R_a\}_{a\in\mathcal{A}}, \pi)$, then *expanding* a pointed Kripke model (M, s) by a first-degree belief formula results in the pointed Kripke model $(M, s) + \mathbf{B}_a\varphi = (M', s')$ with $M' = (W', \{R'_a\}_{a\in\mathcal{A}}, \pi')$ such that

- $W' = W \cup W^r \cup \{s'\}$ with $W^r = \{s^r \mid s \in W\}$ — the *replica* of W — and where s' is a new world symbol that does not occur in $W \cup W^r$;
- for $w \in W$ and $w' \in \{w, w^r\}$, $\pi'(w') = \pi(w)$; and $\pi'(s') = \pi(s)$;
- for $x \in \mathcal{A} \setminus \{a\}$,
  - if $(u, v) \in R_x$ then (u, v) and (u^r, v^r) belong to R'_x,
  - if $(s, v) \in R_x$ then (s', v) belongs to R'_x;
- for $x = a$,
  - if $(u, v) \in R_x$ and $\pi(v) \models \varphi$ then (u, v) and (u^r, v^r) belong to R'_x,
  - if $(u, v) \in R_x$ and $\pi(v) \not\models \varphi$ then (u, v) is in R'_x,
  - if $(s, v) \in R_x$ and $\pi(v) \models \varphi$ then (s', v^r) is in R'_x.

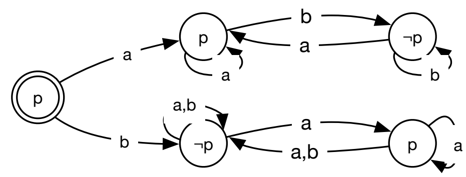

Figure 2: The Kripke model resulting from expanding (M_1, s) in Fig. 1 by $\mathbf{B}_a p$. The new designated world s' (marked with a double circle) is linked via a to the replica of the old Kripke structure (top) and via b to the old Kripke structure itself (bottom). All a-links to $\neg p$-worlds have been removed in the replicated part. As a result, a now believes in p (but also still in $\mathbf{B}_b\neg p$).

Let K be a belief set represented by the pointed Kripke model (M, s), and let $\mathbf{B}_a\varphi \in \mathcal{B}_{\mathcal{A},\mathcal{P}}$, then the *expansion* $K + \mathbf{B}_a\varphi$ is defined as the belief set $K_{(M',s')}$ where $(M, s') = (M, s) + \mathbf{B}_a\varphi$.

As an example, Fig. 2 depicts the results of expanding the Kripke structure to the left in Fig. 1 by $\mathbf{B}_a p$. Obviously, agent a now believes in p, i.e. $\mathbf{B}_a p \in K_1 + \mathbf{B}_a p$, since the only a-accessible world from s' satisifies p. We can also see that b retains exactly her old beliefs, e.g., $\mathbf{B}_b\neg p \in K_1 + \mathbf{B}_A p$ and also $\mathbf{B}_b((\neg\mathbf{B}_a p) \wedge (\neg\mathbf{B}_a\neg p)) \in K_1 + \mathbf{B}_a p$. This is so because the b-accessible worlds from s' are exactly those that were previously accessible from s, and with the same structure. Meanwhile, a also retained his belief that b believes in $\neg p$, that is, $\mathbf{B}_a\mathbf{B}_b\neg p \in K_1 + \mathbf{B}_a p$.

The next lemma shows that expanding a belief set with $\mathbf{B}_a\varphi$ does not change the first-degree beliefs of other agents.

Lemma 2. *For any formula ψ and agent $x \in \mathcal{A} \setminus \{a\}$, $(M, s) \models B_x\psi$ iff $(M', s') \models B_x\psi$.*

Proof. By construction of $(M, s) + \mathbf{B}_a\varphi$, we have that $(s', u) \in R'_x$ iff $(s, u) \in R_x$. Therefore, $(M, s) \models B_x\psi$ iff $(M', s') \models B_x\psi$. □

The definition of expansion applies to any multi-agent Kripke structure and new first-degree belief, including when no world in the model satisfies the new belief. The example depicted in Fig. 3 illustrates that in this case the resulting expanded model still does not include a world that satisfies the new belief, here: $\mathbf{B}_a\neg p$. This implies that there are no a-reachable worlds at all from the new designated state, which in turn means that agent a ends up with inconsistent beliefs (and hence, in particular, believes $\neg p$). This is very much in the spirit of expansion in the classical case, when expanding by a new belief that is inconsistent with the current ones results in an inconsistent belief set (Alchourrón, Gärdenfors, and Makinson 1985). In the multi-agent generalization, all other agents maintain their beliefs, however (cf. Lemma 2).

It is easy to prove that expanding a belief set with $\mathbf{B}_a\varphi$ always results in agent a believing φ when a ends up with consistent beliefs; and if the result is that a has inconsistent beliefs, then again, vacuously, a believes in φ.

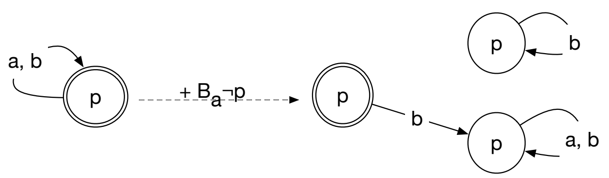

Figure 3: A Kripke structure in which p is common belief among agents a, b is expanded by the new belief $\mathbf{B}_a\neg p$. Agent b's first- and higher-order beliefs remain unchanged thanks to the retained, b-reachable copy of the old structure (bottom) whereas there is no a-reachable world in the replica (top), hence $\mathbf{B}_a\bot$ after expansion.

Lemma 3. $(M, s) + \mathbf{B}_a\varphi \models \mathbf{B}_a\varphi$.

Proof. By definion of the entailment of formulas from a pointed Kripke structure, if $R'_a(s') = \emptyset$ then $(M', s') \models \bot$ so that $(M', s') \models \mathbf{B}_a\psi$ for any formula ψ, hence specifically, $(M', s') \models \mathbf{B}_a\varphi$.

Assume that $R'_a(s') \neq \emptyset$. By construction, for every u^r such that $(s', u^r) \in R'_a$, we have that $(s, u) \in R_a$ and $\pi(u) \models \varphi$. This implies that $\pi(u^r) \models \varphi$, and thus, $(M', s') \models \mathbf{B}_a\varphi$ because $R'_a(s') \neq \emptyset$. □

The next lemma shows that the first-degree beliefs of agent a are exactly the logical consequences of its old beliefs plus the new belief.

Lemma 4. *If $R'_a(s') \neq \emptyset$ then $(M', s') \models \mathbf{B}_a\psi$ if, and only if, $\psi \in Cn\left(\{\varphi\} \cup \{\phi \in \mathcal{B}_{\mathcal{A},\mathcal{P}} \mid \mathbf{B}_a\phi \in K_{(M,s)}\}\right)$.*

Proof. The proof of this lemma is similar to the proof of Lemma 3 with the observation that for $\psi \in Cn\left(\varphi \cup \{\phi \in \mathcal{B}_{\mathcal{A},\mathcal{P}} \mid \mathbf{B}_a\phi \in K_{(M,s)}\}\right)$, $(s', u^r) \in R'_a$ and $\pi'(u^r) \models \psi$ iff $(s, u) \in R_a$ and $\pi(u) \models \psi$. □

4 Generalizing AGM Postulates

Using the basic formal concepts for multi-agent belief sets represented by a single, pointed Kripke strucure as developed in the previous section, we can now define generalizations of the standard postulates, beginning with the AGM potulates for one-shot revision.

Closure A revised belief set is deductively closed, denoted by *Cn*, under the modal logic being interpreted:

$$K * \mathbf{B}_a\varphi = Cn(K * \mathbf{B}_a\varphi)$$

Since the set of belief formulas entailed by a single Kripke structure is deductively closed, this postulate holds when a revised belief set is represented by a pointed Kripke model.

Success The result of revising a Kripke model by a first-degree belief formula should include the new belief:

$$\mathbf{B}_a\varphi \in K * \mathbf{B}_a\varphi$$

Inclusion A revised Kripke model should only contain first-degree belief formulas that would be included in the expanded Kripke model:

$$K * \mathbf{B}_a\varphi \subseteq K + \mathbf{B}_a\varphi$$

From Lemma 2 and 4 it follows that under the Inclusion principle, a revised Kripke model contains only first-degree belief formulas that follow logically from the old and new first-degree beliefs.

Worthy of note, this generalized Inclusion postulate does not stipulate any requirements about second- or higher-degree belief formulas. In particular it allows for another agent, b, to change her beliefs about agent a believing in φ. Hence, $K * \mathbf{B}_a\varphi$ may contain beliefs that are not included in $K + \mathbf{B}_a\varphi$ if these are not first-order beliefs.

Vacuity If a Kripke model is revised by a first-degree belief that is consistent with the current beliefs, then the result should contain all first-degree belief formulas that are included in the expanded Kripke model:

$$\mathbf{B}_a\neg\varphi \notin K \;\Rightarrow\; K + \mathbf{B}_a\varphi \subseteq K * \mathbf{B}_a\varphi$$

By Lemma 2 and 4 it follows that under the Vacuity principle, a revised Kripke model contains all first-degree belief formulas that follow logically from the old and new first-degree beliefs, provided the old beliefs did not include the opposite of the new belief.

Similar to the generalized Inclusion principle, the Vacuity postulate does allow for belief revision operators in which other agents change their belief about the belief of agent a in φ as a result of this revision.

It should also be noted that Vacuity and Inclusion together are a weaker requirement than stipulating that revision be identical to expansion in case a new belief is consistent with the old ones. They merely postulate that the first-degree beliefs are the same, while they do not demand anything about other belief formulas. In particular, they do not prescribe the specific structure from our definition of expansion of a pointed Kripke model with a new belief (cf. Section 3).

Consistency 1 Any revision by a logically consistent belief should result in a consistent belief for the agent:

$$\not\vdash \varphi \rightarrow \bot \;\Rightarrow\; \mathbf{B}_a\bot \notin K * \mathbf{B}_a\varphi$$

It is worth noting that this does not postulate overall consistency of beliefs as it cannot be generally assumed that one agent changing their beliefs would always mean that any other agent that may have had inconsistent beliefs would automatically end up with consistent beliefs too. However, while it is possible that other agents change their beliefs about agent a's beliefs, it is reasonable to postulate that they do not end up with inconsistent beliefs as a result if they had consistent beliefs beforehand. For this reason, we suggest the following additional postulate on consistency.

Consistency 2 If a new belief $\mathbf{B}_a\varphi$ is consistent, then any agent with consistent beliefs will have consistent beliefs after the revision:

$$\mathbf{B}_b\bot \notin K \wedge \not\vdash \varphi \rightarrow \bot \;\Rightarrow\; \mathbf{B}_b\bot \notin K * \mathbf{B}_a\varphi$$

Extensionality If two new belief formulas are logically equivalent w.r.t. the underlying modal logic, then a belief set revised by either of the two should give the same result:

$$\vdash \mathbf{B}_a\varphi \leftrightarrow \mathbf{B}_b\psi \;\Rightarrow\; K * \mathbf{B}_a\varphi = K * \mathbf{B}_b\psi$$

It is easy to see that $\mathbf{B}_a\varphi$ and $\mathbf{B}_b\psi$ are logically equivalent under *any* pointed Kripke structure if, and only if, $a = b$ and φ and ψ are logically equivalent proposition formulas.

Superexpansion Revising a belief set by a conjunction $\varphi \wedge \psi$ of two new beliefs of an agent a should not result in more first-degree beliefs than the expansion by $\mathbf{B}_a\psi$ of the result of revising by $\mathbf{B}_a\varphi$:

$$K * \mathbf{B}_a(\varphi \wedge \psi) \subseteq (K * \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$$

Similarly to Inclusion and Vacuity, this generalized postulate does not make any assumptions about second- or higher-degree belief formulas of any agent. This is also true for the following counterpart, subexpansion.

Subexpansion If the second belief $\mathbf{B}_a\psi$ is consistent with the result of revising a belief set by the first belief $\mathbf{B}_a\varphi$, then revision by the conjunction of the two should not result in fewer first-degree beliefs than the expansion by $\mathbf{B}_a\psi$ of the result of revising by $\mathbf{B}_a\varphi$:

$$\mathbf{B}_a\neg\psi \notin K * \mathbf{B}_a\varphi \Rightarrow (K * \mathbf{B}_a\varphi) + \mathbf{B}_a\psi \subseteq K * \mathbf{B}_a(\varphi \wedge \psi)$$

5 Multi-agent Belief Revision Operators

Having defined generalized AGM postulates for MBR, in this section we present a generalization of the well-known "full-meet" revision operator for classical Belief Revision (Alchourrón and Makinson 1982) and show that it satisfies all generalized AGM postulates.

We define this *multi-agent full meet revision*, denoted by the opeartor name $*_{\text{fm}}$, as follows:

$$K_{(M,s)} *_{\text{fm}} \mathbf{B}_a\varphi = \begin{cases} K_{(M,s)} + \mathbf{B}_a\varphi & \text{if } \mathbf{B}_a\neg\varphi \notin K_{(M,s)} \\ K_{(M_\emptyset,s)} + \mathbf{B}_a\varphi & \text{otherwise} \end{cases}$$

The principle behind this definition is the same as for classical full-meet revision (Alchourrón, Gärdenfors, and Makinson 1985). If a new belief is consistent with the current beliefs, the underlying Kripke structure is simply expanded by that belief. Otherwise, the new belief is incorporated in the most conservative manner by starting with the minimal belief set $K_{(M_\emptyset,s)}$ (cf. Section 3), in which it is common knowledge that no agent believes in anything but tautological proposition formulas, and then expanding the undelying Kripke model $(M_\emptyset, s)$ by the new belief.

It should be noted, however, that unlike with the classical full-meet belief revision operator, the belief set resulting from revision by a belief that is inconsistent with the old beliefs does entail more beliefs than follow logically from the new one. This is so because the generic Kripke structure $(M_\emptyset, s)$ makes strong assumptions about higher-degree beliefs. For example, full-meet revision of the Kripke structure depicted on the left-hand side in Figure 3 by $\mathbf{B}_a\neg p$ results in agent b not only *losing* both her first-degree belief in p as well as her higher-degree belief that p is common knowledge, but also *gaining* second-order beliefs of "ignorance", such as, say, $\mathbf{B}_b(\neg\mathbf{B}_a p \wedge \neg\mathbf{B}_a\neg p)$.

This notwithstanding, the generalized full-meet operator provably satisfies all of the generalized AGM postulates.

Theorem 1. $*_{\text{fm}}$ *satisfies the generalized AGM postulates.*

Proof. Due to limited space, we include below only the proof for **Superexpansion** as it is somewhat more complicated than the others. Detailed proofs for all theorems and lemmas in the paper are available in the supplementary report (Thielscher and Son 2026).

1. Suppose $\mathbf{B}_a\neg(\varphi \wedge \psi) \notin K$. Since K is deductively closed, it follows that $\mathbf{B}_a\neg\varphi \notin K$. Hence, agent a has consistent beliefs in $K + \mathbf{B}_a(\varphi \wedge \psi)$, in $K + \mathbf{B}_a\varphi$, and in $(K + \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$. By Lemma 4 it follows that $(M, s) + \mathbf{B}_a(\varphi \wedge \psi) \models \mathbf{B}_a\chi$ iff $\chi \in Cn(\{\varphi \wedge \psi\} \cup \{\phi \in \mathcal{B}_{\mathcal{A},\mathcal{P}} \mid \mathbf{B}_a\phi \in K\})$. This is equivalent to $\chi \in Cn(\{\psi\} \cup Cn(\{\varphi\} \cup \{\phi \in \mathcal{B}_{\mathcal{A},\mathcal{P}} \mid \mathbf{B}_a\phi \in K\}))$, which in turn by Lemma 4 is equivalent to $(K + \mathbf{B}_a\varphi) + \mathbf{B}_a\psi \models \mathbf{B}_a\chi$. By Lemma 2 it follows that all other agents too have the same beliefs in $K + \mathbf{B}_a(\varphi \wedge \psi)$ and in $(K + \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$. Hence, $K *_{\text{fm}} \mathbf{B}_a(\varphi \wedge \psi) \subseteq (K *_{\text{fm}} \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$.
2. Suppose $\mathbf{B}_a\neg(\varphi \wedge \psi) \in K$. By definition of $(M_\emptyset, s) + \mathbf{B}_a(\varphi \wedge \psi)$ and Lemma 4 it follows that, for any first-degree belief, $\mathbf{B}_a\phi \in K *_{\text{fm}} \mathbf{B}_a(\varphi \wedge \psi)$ iff $\phi \in Cn(\varphi \wedge \psi)$, and for $x \in \mathcal{A} \setminus \{a\}$, $\mathbf{B}_x\phi \in K *_{\text{fm}} \mathbf{B}_a(\varphi \wedge \psi)$ iff $\models \phi$.
   We distinguish two cases: If $\mathbf{B}_a\neg\varphi \notin K$, then $(K *_{\text{fm}} \mathbf{B}_a\varphi) + \mathbf{B}_a\psi = (K + \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$. From $\mathbf{B}_a\neg(\varphi \wedge \psi) \in K$ it follows that $\mathbf{B}_a\neg\psi \in K + \mathbf{B}_a\varphi$. Hence, a has inconsistent beliefs in $(K + \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$ while the beliefs of all other agents $x \in \mathcal{A} \setminus \{a\}$ are the same in $(K + \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$ and $K + \mathbf{B}_a(\varphi \wedge \psi)$. If, on the other hand, $\mathbf{B}_a\neg\varphi \in K$, then $(K *_{\text{fm}} \mathbf{B}_a\varphi) + \mathbf{B}_a\psi = (K_{(M_\emptyset,s)} + \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$. Hence, a's first-degree beliefs in $(K *_{\text{fm}} \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$ are exactly the logical consequences of $\varphi \wedge \psi$ while the beliefs of all other agents $x \in \mathcal{A} \setminus \{a\}$ are the same in $(K *_{\text{fm}} \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$ and $K + \mathbf{B}_a(\varphi \wedge \psi)$. Thus, $K *_{\text{fm}} \mathbf{B}_a(\varphi \wedge \psi) \subseteq (K *_{\text{fm}} \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$. □

While satisfying all our generalized postulates for one-shot revision, multi-agent full-meet revision makes for a very drastic revision in cases where simple expansion would leave an agent with inconsistent beliefs. We introduce a more refined alternative in Section 7; prior to this, we first turn to the question of successive multi-agent revisions.

6 Generalized Postulates for Iterated MBR

We complete our framework for belief revision for multi-agent epistemic reasoning and planning by generalizing the standard postulates for *iterated* revision.

DP1 – Successive revision respect Revision by one belief followed by revising by a stronger belief makes the first revision redundant:

$$\vdash \mathbf{B}_b\psi \rightarrow \mathbf{B}_a\varphi \;\Rightarrow\; (K * \mathbf{B}_a\varphi) * \mathbf{B}_b\psi \doteq K * \mathbf{B}_b\psi$$

Here, $K_1 \doteq K_2$ means $K_1 \subseteq K_2$ and $K_2 \subseteq K_1$.

DP2 – Irrelevance of superseded beliefs If for two successive revisions, the second belief contradicts the first one, then the resulting beliefs for all agents should be the same as from revising according to the second belief only.

$$\vdash \mathbf{B}_b\psi \rightarrow \mathbf{B}_a\neg\varphi \;\Rightarrow\; (K * \mathbf{B}_a\varphi) * \mathbf{B}_b\psi \doteq K * \mathbf{B}_b\psi$$

DP3 – Consistency preservation across revisions If one agent's first-degree belief would be contained in the belief set after revision by any other first-degree belief, then in case the former is used to revise the belief set first, that first belief should be preserved through the second revision:

$$\mathbf{B}_a\varphi \in K * \mathbf{B}_b\psi \;\Rightarrow\; \mathbf{B}_a\varphi \in (K * \mathbf{B}_a\varphi) * \mathbf{B}_b\psi$$

DP4 – Minimal change when reaffirming a belief After two consecutive revisions, the first agent should not end up believing the opposite unless they would do so if the belief set was revised by the second belief only:

$$\mathbf{B}_a\neg\varphi \notin K * \mathbf{B}_b\psi \ \Rightarrow\ \mathbf{B}_a\neg\varphi \notin (K * \mathbf{B}_a\varphi) * \mathbf{B}_b\psi$$

IN – Independence After two consecutive revisions, the belief of the first agent should be preserved unless they would believe the opposite if the belief set was revised by the second belief only:

$$\mathbf{B}_a\neg\varphi \notin K * \mathbf{B}_b\psi \ \Rightarrow\ \mathbf{B}_a\varphi \in (K * \mathbf{B}_a\varphi) * \mathbf{B}_b\psi$$

Interestingly, while in the classical, single-agent case the Independence postulate **IN** strengthens both **DP3** and **DP4** (Jin and Thielscher 2007) for any operator that satisfies the AGM postulates, in the generalized case Independence only strengthens **DP4**.

Lemma 5. *Consider a belief revision operator that satisfies the generalized AGM postulates, then the operator satisfies multi-agent DP4 if it satisfies multi-agent Independence.*

Proof. Independence obviously implies DP4 unless $\{\mathbf{B}_a\phi, \mathbf{B}_a\neg\phi\} \subseteq (K * \mathbf{B}_a\varphi) * \mathbf{B}_b\psi$. The latter would mean that a has inconsistent beliefs at the end of the two revisions, which according to Inconsistency 1 and 2 can only happen if $\varphi \vdash \bot$. This in turn implies $\mathbf{B}_a\neg\varphi \notin K * \mathbf{B}_b\psi$, thus DP4 holds vacuously in this case also. □

It is noteworthy that **IN** would not entail **DP4** without the additional **Consistency 2** postulate, which guarantees that agent a does not end up believing in both φ and $\neg\varphi$ as a result of further revision by $\mathbf{B}_b\psi$ after revising by $\mathbf{B}_a\varphi$.

IN does not imply **DP3** even if an operator satisfies all multi-agent AGM postulates, for the following reason: If $\mathbf{B}_a\bot \in K * \mathbf{B}_b\psi$ then **IN** vacuously holds while **DP3** is violated if $\mathbf{B}_a\bot \notin (K * \mathbf{B}_a\bot\varphi) * \mathbf{B}_b\psi$. This is possible if an operator allows an agent to regain consistent beliefs when revising a multi-agent belief set by another agent's belief.

The next theorem summarizes the satisfaction of the generalized DP postulates of $*_{\text{fm}}$.

Theorem 2. $*_{\text{fm}}$ *satisfies* **DP1**, **DP3**, *and* **DP4** *but not* **DP2** *nor* **IN**. *However,* $*_{\text{fm}}$ *does satisfy a weak version of* **DP2** *where* $\doteq$ *is replaced by* $\subseteq$.

Proof. **(DP2)** Generalized full-meet does not satisfy **DP2** for the following reason: If $\mathbf{B}_b\psi$ is consistent with the current belief set K then $K *_{\text{fm}} \mathbf{B}_b\psi = K + \mathbf{B}_b\psi$, hence by Lemma 2, all agents retain all their beliefs when revising K by b's new belief. But if K is revised by $\mathbf{B}_a\varphi$ first and $\models \mathbf{B}_b\psi \to \mathbf{B}_a\varphi$ holds then $(K *_{\text{fm}} \mathbf{B}_b\varphi) *_{\text{fm}} \mathbf{B}_b\psi = K_{(M_\emptyset,s)} + \mathbf{B}_b\psi$, which implies that the beliefs of all other agents have been erased.

The fact that full-meet does not satisfy **DP2** mirrors a result in classical, single-agent belief revision (Jin and Thielscher 2007). To show that full-meet multi-agent belief revision satisfies one direction of **DP2** (namely, that revising by a belief that is then superseded by a second, contradictory belief never introduces more beliefs than revision with the second belief directly) we make a case distinction.

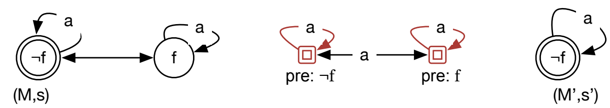

Figure 4: Revising with $\mathbf{B}_a f$ is not the same as sensing f

1. Suppose that $\models \varphi \leftrightarrow \bot$, then $\mathbf{B}_a\neg\varphi \in K$, hence $K *_{\text{fm}} \mathbf{B}_a\varphi = K_{(M_\emptyset,s)} + \mathbf{B}_a\bot$, that is, there is no a-accessible world in s while all other agents' first-degree beliefs are tautological proposition formulas. Consequently, $(K *_{\text{fm}} \mathbf{B}_a\varphi) *_{\text{fm}} \mathbf{B}_b\psi \subseteq K *_{\text{fm}} \mathbf{B}_b\psi$.
2. Otherwise, $\models \mathbf{B}_b\psi \to \mathbf{B}_a\varphi$ implies $a = b$ and $\psi \models \neg\varphi$, hence $\varphi \models \neg\psi$. It follows that $\mathbf{B}_b\psi \notin K *_{\text{fm}} \mathbf{B}_a\varphi$, hence $(K *_{\text{fm}} \mathbf{B}_a\varphi) *_{\text{fm}} \mathbf{B}_b\psi = K_{(M_\emptyset,s)} + \mathbf{B}_b\psi$, which implies $(K *_{\text{fm}} \mathbf{B}_a\varphi) *_{\text{fm}} \mathbf{B}_b\psi \subseteq K *_{\text{fm}} \mathbf{B}_b\psi$. □

7 Event Model-Based Belief Revision

While full-meet is an instructive operator to show that all generalized AGM and most DP postulates can be simultaneously satisfied, it is obviously too strong for practical purposes in that any (non-tautological) belief may be abandoned in case of a true revision. In this section, we introduce a more sophisticated operator for MBR based on event models. Event models (a.k.a. update models) have been used to describe transformations of epistemic states in multi-agent domains according to a predetermined transformation pattern (Baltag, Moss, and Solecki 1998)). This has also been used in modeling belief-altering actions in high-level action languages (Baral et al. 2022; Rajaratnam and Thielscher 2021). Nonetheless, event models proposed for the purpose of updating beliefs of agents after an action occurrence are not suitable for revising the beliefs of agents. This is illustrated in Fig. 4: In the model (M, s) (left), simplified to consider only a single agent, a does not believe f nor $\neg f$. The event model for sensing f is in the middle, whose update on (M, s) is (M', s') (right), which indicates that a believes $\neg f$ after sensing f. Such an update does not allow for the revision of (M, s) with $\mathbf{B}_a f$.

For the discussion of the generalized postulates, we need some extra notation. A *literal* is either a proposition $p \in \mathcal{P}$ or its negation $\neg p$. For a literal ℓ, $\neg\ell$ denotes its negation, with $\neg\neg p = p$ for $p \in \mathcal{P}$. An *inference rule* (or rule) among literals is of the form $\lambda \to \delta$ where λ and δ are sets of literals. Given a set of literals w and a set of rules R, $C_R(w)$ denotes the minimal set (w.r.t. $\subseteq$) of literals w' such that $w \subseteq w'$ and, for every $\lambda \to \delta$ in R, if $w \models \lambda$ then $w' \models \delta$.

In the presence of the set of rules R, the valuation function π of any Kripke model M must satisfy R as well. Therefore, we require that for every world u in M, $\pi(u)$ is consistent, i.e., $\pi(u) = C_R(\pi(u))$. For a set of literals φ and an interpretation u, let $u \star \varphi$ denote an interpretation u' such that $u' = C_R((u \cap u') \cup \varphi)$. In general, there might exist several interpretations u' satisfying the aforementioned equation. Fortunately, it is well-known that there are conditions on R such that there is a unique u' that satisfies the equation (see, e.g., Tu et al.'s (2011) work). In this paper, we will assume that R only consists of rules of the form $p \to q$

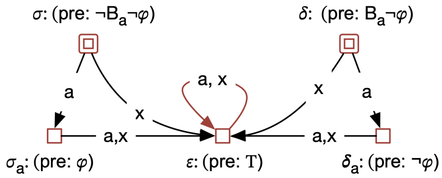

Figure 5: Revision event model for $\mathbf{B}_a\varphi$

where p, q are literals and, for each u and p, there is a unique $u' = C_R((u \cap u') \cup \{p\})$, since this is sufficient for dealing with the **DP1** and **DP2** postulates.

Consider $a \in \mathcal{A}$ and a set of literals φ. We define $\mathbf{\Sigma}^a(\varphi)$, called the *event model for revision by* $\mathbf{B}_a\varphi$, as the event model $\langle \Sigma, \{E_a\}_{a\in\mathcal{A}}, pre, eff\rangle$ where

- $\Sigma = \{\sigma, \delta, \sigma_a, \delta_a, \epsilon\}$;
- $E_a = \{(\sigma, \sigma_a), (\sigma_a, \epsilon), (\delta, \delta_a), (\delta_a, \epsilon), (\epsilon, \epsilon)\}$;
- $E_x = \{(\eta, \epsilon) \mid \eta \in \Sigma \setminus \{\sigma_a, \delta_a\}\}$ for $x \in \mathcal{A} \setminus \{a\}$;
- $pre(\sigma) = \neg\mathbf{B}_a\neg\varphi$, $pre(\delta) = \mathbf{B}_a\neg\varphi$, $pre(\sigma_a) = \varphi$, $pre(\delta_a) = \neg\varphi$, and $pre(\epsilon) = \top$.
- for $u \in W$, $eff(u, \eta) = \pi(u)$ for $\eta \in \Sigma \setminus \{\delta_a\}$, and $eff(u, \delta_a) = \pi(u) \star \varphi$.

In the above definition, Σ is the set of events representing possible views of the event "(M, s) *is revised by* $\mathbf{B}_a\varphi$". Intuitively, this revision could affect the worlds accessible by a in the following ways: If $\neg\mathbf{B}_a\neg\varphi$ is true, i.e., $\mathbf{B}_a\varphi$ or $\neg(\mathbf{B}_a\varphi \wedge \mathbf{B}_a\neg\varphi)$ holds, then some worlds accessible by a satisfy φ and a can just eliminate all worlds satisfying $\neg\varphi$ from its accessibility. This is represented by the events σ and σ_a (Fig. 5) where σ denotes a *designated event* and σ_a has the precondition φ. If, on the other hand, $\mathbf{B}_a\neg\varphi$ is true, then agent a should revise his beliefs in every world accessible by a, via the definition of eff. This is represented by the events δ and δ_a. For agent $x \in \mathcal{A}\setminus\{a\}$, nothing changes and thus, x's view is that only the event ϵ occurs.

The revision of $\mathbf{B}_a\varphi$ in (M, s), denoted by $(M, s) \ast_{\text{ev}} \mathbf{B}_a\varphi$, is a pointed Kripke structure (M^φ, s^φ) where $M^\varphi = (W^\varphi, \{R^\varphi_x\}_{x\in\mathcal{A}}, \pi^\varphi)$ and

- $W^\varphi = \{(u, \tau) \mid u \in W, \tau \in \Sigma, (M, u) \models pre(\tau)\}$;
- $((u, \tau), (u', \tau')) \in R^\varphi_x$ iff $(u, \tau), (u', \tau') \in W^\varphi$ along with $(u, u') \in R_x$ and $(\tau, \tau') \in E_x$;
- $\pi^\varphi((u, \tau)) = eff(u, \tau)$; and
- $s^\varphi = (s, \sigma)$ if $(M, s) \models \neg\mathbf{B}_a\neg\varphi$; otherwise, $s^\varphi = (s, \delta)$.

Before we discuss the properties of $\ast_{\text{ev}}$ in detail, let us observe that if $R_a(s) \neq \emptyset$ then $R^\varphi_a(s^\varphi) \neq \emptyset$, i.e., a has consistent belief in (M^φ, s^φ). This is the main difference between $\ast_{\text{ev}}$ and $\ast_{\text{fm}}$. We prove that beliefs of other agents do not change.

Lemma 6. *For $x \in \mathcal{A} \setminus \{a\}$ and a proposition formula ψ, $(M^\varphi, s^\varphi) \models \mathbf{B}_x\psi$ iff $(M, s) \models \mathbf{B}_x\psi$.*

Proof. By the construction of (M^φ, s^φ), we have that $s^\varphi = (s, \eta)$ such that $(M, s) \models pre(\eta)$. Furthermore, $(u, \epsilon) \in W^\varphi$ for every $u \in W$ and $((s, \eta), (u, \epsilon)) \in R^\varphi_x$ iff $(s, u) \in R_x$. Because $\pi^\varphi((u, \epsilon)) = \pi(u)$ for every $u \in W$, $(M^\varphi, s^\varphi) \models \mathbf{B}_x\psi$ iff $(M, s) \models \mathbf{B}_x\psi$. □

Theorem 3. $\ast_{\text{ev}}$ *satisfies the generalized AGM postulates.*

Proof. As above, proofs are given for selected interesting postulates only. We assume a consistent K that is represented by (M, s) with $M = (W, \{R_x\}_{x\in\mathcal{A}}, \pi)$. We use $(M, s) \ast_{\text{ev}} \mathbf{B}_a\varphi$ and (M^φ, s^φ) interchangeably where $M^\varphi = (W^\varphi, \{R^\varphi_x\}_{x\in\mathcal{A}}, \pi^\varphi)$. Successive revisions such as $(K \ast_{\text{ev}} \mathbf{B}_a\varphi) \ast_{\text{ev}} \mathbf{B}_a\psi$ will be denoted by $(M^{\varphi;\psi}, s^{\varphi;\psi})$ etc. Furthermore, $R_a(s|\varphi) = \{u \mid u \in R_a(s), \pi(u) \models \varphi\}$. Below, we use conjunctions of literals and sets of literals interchangeably, and we always assume that a conjunction of literals φ is consistent in the classical sense, i.e., $\varphi \not\models \bot$. For a set of literals φ, $\neg\varphi$ denotes $\{\neg\ell \mid \ell \in \varphi\}$.

(*Success*) We show $K \ast_{\text{ev}} \mathbf{B}_a\varphi \models \mathbf{B}_a\varphi$ via two cases:

1. $(M, s) \models \neg\mathbf{B}_a\neg\varphi$, i.e., a does not believe in $\neg\varphi$ before the revision. This implies that $s^\varphi = (s, \sigma)$. Furthermore, because K is consistent, we have that $R_a(s) \neq \emptyset$ and there exists some $u \in W$ such that $(s, u) \in R_a$ and $\pi(u) \models \varphi$, which implies that $((s, \sigma), (u, \sigma_a)) \in R^\varphi_a$, i.e., $R^\varphi(s^\varphi) \neq \emptyset$. In addition, if $((s, \sigma), (u, \rho)) \in R^\varphi_a$ then $\rho = \sigma_a$, and hence, $\pi(u) \models \varphi$, because E_a contains only one element related to σ, (σ, σ_a), and $(s, u) \in R_a$. Thus, we have that $(M^\varphi, s^\varphi) \models \mathbf{B}_a\varphi$.
2. $(M, s) \models \mathbf{B}_a\neg\varphi$. This implies that $s^\varphi = (s, \delta)$. Again, because K is consistent, we have that $R_a(s) \neq \emptyset$ and for each $u \in R_a(s)$, $\pi(u) \models \neg\varphi$. By the construction of (M^φ, s^φ), $R^\varphi_a(s^\varphi) \neq \emptyset$. Consider $u' \in R^\varphi_a(s^\varphi)$. We have that $u' = (u, \delta_a)$ for some $u \in R_a(s)$, and hence, $\pi^\varphi(u') = C_R((\pi(u) \cap \pi^\varphi(u')) \cup \{\varphi\})$, which implies $\varphi \in \pi^\varphi(u')$. It follows that $(M^\varphi, s^\varphi) \models \mathbf{B}_a\varphi$. □

(Superexpansion) Consider a propositional formula ϕ, then Lemma 6 and Lemma 2 imply that to prove that $\ast_{\text{ev}}$ satisfies this postulate, it suffices to show that $K \ast_{\text{ev}} \mathbf{B}_a(\varphi \wedge \psi) \models \mathbf{B}_a\phi$ implies $(K \ast_{\text{ev}} \mathbf{B}_a\varphi) + \mathbf{B}_a\psi \models \mathbf{B}_a\phi$. Let (M', s') denote $(M^\varphi, s^\varphi) + \mathbf{B}_a\psi$. We consider two cases:

- $(M, s) \models \neg\mathbf{B}_a\neg(\varphi \wedge \psi)$. So, $s^{\varphi\wedge\psi} = (s, \sigma^{\varphi\wedge\psi})$ and $(s^{\varphi\wedge\psi}, u^{\varphi\wedge\psi}) \in R^{\varphi\wedge\psi}_a$ where $u^{\varphi\wedge\psi} = (u, \sigma^{\varphi\wedge\psi}_a)$ iff $(s, u) \in R_a$ and $\pi(u) \models (\varphi \wedge \psi)$. Since $(M, s) \models \neg\mathbf{B}_a\neg(\varphi \wedge \psi)$, we have that $(M, s) \models \neg\mathbf{B}_a\neg\varphi$, i.e., $u \in R_a(s|\varphi \wedge \psi)$. Thus $s^\varphi = (s, \sigma^\varphi)$ and $(s^\varphi, u^\varphi) \in R^\varphi_a$ where $u^\varphi = (u, \sigma^\varphi_a)$ iff $(s, u) \in R_a$ and $\pi(u) \models \varphi$, i.e., $u \in R_a(s|\varphi)$. Because $R_a(s|\varphi) = R_a(s|\varphi \wedge \psi) \cup R_a(s|\varphi \wedge \neg\psi)$ and the construction of $(K \ast_{\text{ev}} \mathbf{B}_a\varphi) \models \mathbf{B}_a\phi$ we have $R'_a(s') = \{u^r \mid u \in R_a(s), \pi(u) \models \varphi \wedge \psi\}$, which proves the consequence of the postulate.
- $(M, s) \models \mathbf{B}_a\neg(\varphi \wedge \psi)$. Then, $s^{\varphi\wedge\psi} = (s, \delta^{\varphi\wedge\psi})$ and for every $u^{\varphi\wedge\psi}$ such that $(s^{\varphi\wedge\psi}, u^{\varphi\wedge\psi}) \in R^{\varphi\wedge\psi}_a$ holds, we have $u^{\varphi\wedge\psi} = (u, \delta^{\varphi\wedge\psi}_a)$ iff $(s, u) \in R_a$. Therefore, we can conclude that $K \ast_{\text{ev}} \mathbf{B}_a(\varphi \wedge \psi) \models \mathbf{B}_a\phi$ iff for every $u \in R_a(s)$, $\pi(u) \star (\varphi \wedge \psi) \models \phi$. There are two cases:
  1. $(M, s) \models \mathbf{B}_a\neg\varphi$. In this case, similar arguments to the above allow us to conclude that $(s^\varphi, u^\varphi) \in R^\varphi_a$ iff $u^\varphi = (u, \delta^\varphi_a)$ for some $(s, u) \in R_a$ and $\pi(u^\varphi) = \pi(u) \star \varphi$. Because of the construction of $(K \ast_{\text{ev}} \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$, we have two sub-cases:
     (a) $K \ast_{\text{ev}} \mathbf{B}_a\varphi \models \mathbf{B}_a\neg\psi$. In this case, a has inconsistent beliefs in $(K \ast_{\text{ev}} \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$, i.e., the postulate holds trivially.

(b) $K *_{\text{ev}} \mathbf{B}_a\varphi \models \neg\mathbf{B}_a\neg\psi$. In this case, we can conclude that $R'_a(s') \neq \emptyset$. Furthermore, $(s', u') \in R'_a$ iff $u' = u^\varphi$ and $\pi(u^\varphi) \models \psi$. Given that $\star$ satisfies the classical AGM postulate of Superexpansion, we can conclude that $K *_{\text{ev}} \mathbf{B}_a(\varphi \wedge \psi) \models \mathbf{B}_a\phi$ implies $K *_{\text{ev}} \mathbf{B}_a\varphi + \mathbf{B}_a\psi \models \mathbf{B}_a\phi$.

2. $(M, s) \models \neg\mathbf{B}_a\neg\varphi$. We have $s^\varphi = (s, \sigma^\varphi)$ and $(s^\varphi, u^\varphi) \in R^\varphi_a$ where $u^\varphi = (u, \sigma^\varphi_a)$ iff $(s, u) \in R_a$ and $\pi(u) \models \varphi$, which implies that $\pi^\varphi(u^\varphi) \models \neg\psi$ for every $(s^\varphi, u^\varphi) \in R^\varphi_a$. This means that a's belief is inconsistent in $(K *_{\text{ev}} \mathbf{B}_a\varphi) + \mathbf{B}_a\psi$, and thus, the postulate also holds in this case. □

While $*_{\text{ev}}$ satisfies all generalized AGM postulates, it does not satisfy **DP1**-**DP2**. This is because $*_{\text{ev}}$ attempts to remove uncertainty in the beliefs of an agent before revising them, which is encoded in the event σ in which the agent is uncertain about φ and, due to the precondition of σ_a, the operator only retains worlds satisfying φ. Nevertheless, these postulates, in a simpler form where φ and ψ are literals, are satisfied under restricted conditions.

Theorem 4. $*_{\text{ev}}$ *satisfies* (i) **DP1** *if* $\mathbf{B}_a\neg p \notin K *_{\text{ev}} \mathbf{B}_a q$; (ii) **DP2** *if* $\mathbf{B}_a q \vee \mathbf{B}_a p \in K$; (iii) **DP3**; (iv) **DP4** *and* **IN** *if agent a has consistent beliefs in K.*

Proof. **(DP1)** Since $\models \mathbf{B}_b p \rightarrow \mathbf{B}_a q$, we have that $p \rightarrow q$ and a and b are identical. Therefore, we will show that $*_{\text{ev}}$ satisfies this postulate by proving that if $\models \mathbf{B}_a p \rightarrow \mathbf{B}_a q$ and $K *_{\text{ev}} B_a q \not\models \mathbf{B}_a\neg p$ hold then $(K *_{\text{ev}} \mathbf{B}_a q) *_{\text{ev}} \mathbf{B}_a p \doteq K *_{\text{ev}} \mathbf{B}_a p$.

1. $(M, s) \models \neg\mathbf{B}_a\neg q$. So, if $(s^q, u^q) \in R^q_a$ then $u \in R_a(s)$ and $\pi(u) \models q$ and $\pi^q(u^q) = \pi(u)$. Since $(M, s) *_{\text{ev}} \mathbf{B}_a q \models \neg\mathbf{B}_a\neg p$, for $u^{q;p} \in W^{q;p}$ such that $(s^{q;p}, u^{q;p}) \in R^{q;p}_a$ iff $(s^q, u^q) \in R^q_a$ and $\pi^q(u^q) \models p$. This implies that $u^p \in W^p$ and $(s^p, u^p) \in R^p_a$.

   $(M, s) *_{\text{ev}} \mathbf{B}_a q \models \neg\mathbf{B}_a\neg p$ also implies that $(M, s) \models \neg\mathbf{B}_a\neg p$. Therefore, for every $u \in R_a(s)$ such that $u^p \in W^p$ and $(s^p, u^p) \in R^p_a$, we can conclude that $(s^{q;p}, u^{q;p}) \in R^{q;p}_a$ because $\models p \rightarrow q$. The above imply that $(K *_{\text{ev}} \mathbf{B}_a q) *_{\text{ev}} \mathbf{B}_a p \doteq K *_{\text{ev}} \mathbf{B}_a p$ for this case.

2. $(M, s) \models \mathbf{B}_a\neg q$. This implies that $(M, s) \models \mathbf{B}_a\neg p$ and $s^q = (s, \delta^q)$ and $(s^q, u^q) \in R^q_a$ iff $u^q = (u, \delta^q_a)$ and $\pi^q(u^q) = \pi(u) \star q$. By the definition of $\star$, we can show that $(M, s) \models \mathbf{B}_a\neg p$ implies $\pi^q(u^q) \models \neg p$. This implies that $s^{q;p} = ((s, \sigma^q), \delta^{q;p})$ and for every $u \in R_s(a)$, $(s^{q;p}, u^{q;p}) \in R^{q;p}_a$, $(s^q, u^q) \in R^q_a$, and $\pi^{q;p}(u^{q;p}) = (\pi(u) \star q) \star p$. On the other hand, by construction of $(M, s) *_{\text{ev}} p$, $s^p = (s, \delta^p)$ and, for every $u \in W$, $(s^p, u^p) \in R^p_a$ where $u^p = (u, \delta^p_a)$ and $\pi^p(u^p) = \pi(u) \star p$. Again by definition of $\star$, we have that $\pi(u) \star p = (\pi(u) \star q) \star p$ for every $u \in R_a(s)$. This proves that $(K *_{\text{ev}} \mathbf{B}_a q) *_{\text{ev}} \mathbf{B}_a p \doteq K *_{\text{ev}} \mathbf{B}_a p$ for this case. □

We conclude this section with the introduction of an event-based revision operator, denoted by $*_{\text{rb}}$, that satisfies all the generalized postulates if the underlying classical revision operator satisfies the classical postulates. The operator $*_{\text{rb}}$ differs from $*_{\text{ev}}$ in that it does not attempt to remove the agent's uncertainty before revision. It employs the event model

$$\mathbf{\Sigma}^a_b(\varphi) = \langle \Sigma, \{E_a\}_{a \in \mathcal{A}}, pre, \mathit{eff} \rangle$$

where $\Sigma = \{\sigma, \sigma_a, \epsilon\}$ and

- $E_a = \{(\sigma, \sigma_a), (\sigma_a, \epsilon), (\epsilon, \epsilon)\}$;
- $E_x = \{(\sigma, \epsilon), (\sigma_a, \epsilon), (\epsilon, \epsilon)\}$ for $x \in \mathcal{A} \setminus \{a\}$;
- $pre(\eta) = \top$ for every $\eta \in \Sigma$; and
- for $u \in W$, $\mathit{eff}(u, \eta) = \pi(u)$ for $\eta \in \Sigma \setminus \{\sigma_a\}$, and $\mathit{eff}(u, \sigma_a) = \pi(u) \star \varphi$.

(M^φ, s^φ), the result of $(M, s) *_{\text{rb}} \mathbf{B}_a\varphi$, is defined similar to $(M, s) *_{\text{ev}} \mathbf{B}_a\varphi$ with two changes: (*i*) $\mathbf{\Sigma}^a_b(\varphi)$ is used instead of $\mathbf{\Sigma}^a(\varphi)$ and (*ii*) $s^\varphi = (s, \sigma)$. $*_{\text{rb}}$ satisfies every postulate.

Theorem 5. $*_{\text{rb}}$ *satisfies all MBR postulates.*

As a short proof sketch for this theorem, we observe the following properties of $(M, s) *_{\text{rb}} \varphi$: For every $u \in W$ and $\lambda \in \Sigma$, $u^\varphi = (u, \lambda) \in W^\varphi$. Furthermore, $(s^\varphi, u^\varphi) \in R^\varphi_a$ iff $(s, u) \in R_a(s)$ and $\pi^\varphi(u^\varphi) = \pi(u) \star \varphi$. The conclusion of the theorem then relies on the following observations:

- for every u, $\pi^\varphi(u^\varphi) = \pi(u) \star \varphi \models \varphi$; and
- for every u, if $\pi(u) \star \varphi \not\models \neg\psi$ then $\pi(u) \star \varphi \wedge \psi = (\pi(u) \star \varphi) + \psi$, and if $\pi(u) \star \varphi \models \neg\psi$ then $(\pi(u) \star \varphi) + \psi$ is inconsistent.

8 Discussion and Outlook

We defined the problem of multi-agent belief revision (MBR) based on a single Kripke model as commonly used in epistemic reasoning about actions and planning, and we proposed a generalization of the AGM and DP postulates to this multi-agent case. We identified the challenges faced by the task of constructing MBR operators that adhere to all generalized postulates including those for iterated revision, and we presented results with a generalized full-meet operator and an event-based revision operator that satisfy most but not all of them.

Theorem 2 indicates that our generalized full-meet operator $*_{\text{fm}}$ does not satisfy **DP2**. This is because $K *_{\text{fm}} \mathbf{B}_a\varphi$ may include $\mathbf{B}_b\neg\psi$, and then further revision by $\mathbf{B}_b\psi$ erases all beliefs other than the new one, hence the conclusion of the postulate does not hold. For this reason, it is generally challenging to construct belief revision operators that satisfy **DP2**. Looking at the definition of $*_{\text{fm}}$, it is obvious that the first challenge is the problem of dealing with false beliefs, i.e., defining $K * \mathbf{B}_a\varphi$ given that $\mathbf{B}_a\neg\varphi \in K$. The second challenge is related to the revision of the accessibility relation of a. More precisely, let us denote with $R_a|\eta$ the set $\{(s, u) \in R_a \mid \pi(u) \models \eta\}$ for $\eta \in \mathcal{L}_\mathcal{P}$. Operator $*_{\text{fm}}$ essentially removes $R_a|\neg\varphi$ from consideration in constructing $K *_{\text{fm}} \mathbf{B}_a\varphi$. Thus, if $R_a|\neg\varphi \neq \emptyset$ then the set of worlds accessible by a in $K *_{\text{fm}} \mathbf{B}_a\varphi$ is a proper subset of R_a, the set of worlds accessible by a in K, which prevents $*_{\text{fm}}$ from satisfying **DP2**.

The above discussion implies that for a belief revision operator to satisfy the generalized DP postulates, it must deal with false beliefs as well as the accessibility relations of agents *properly*. We expect that ideas from studies of reasoning with false beliefs in dynamic epistemic logic (van

Ditmarsch, Hendriks, and Verbrugge 2020) or using update models (Son, Pham, and Pontelli 2024) can be helpful in addressing the first issue. To deal with the second issue, approaches based on a pre-order relation between Kripke models (see, e.g., the works by Aucher (2010) and Ditmarsch, van der Hoek, and Kooi (2007)) might be necessary. It is worth noting that the premise of **DP2**, $\vdash \mathbf{B}_b\psi \rightarrow \mathbf{B}_a\neg\varphi$, could also be considered as a culprit for this postualte not to be satisfied by $*_{\text{fm}}$. Modifying the premise to $\mathbf{B}_b\psi \rightarrow \mathbf{B}_a\neg\varphi \in K$, and also requiring that $a \neq b$, could be a viable alternative to generalizing **DP2** worthy of consideration.

Turning to the issue of higher-degree beliefs, the generalized revision postulates presented in this paper focus on first-degree beliefs, which are arguably the most fundamental ones, and postulating anything about the retention or revision of higher-degree beliefs seems more difficult to justify. But of course the concrete revision operators discussed in this paper all do precisely define how all higher-degree beliefs, including common beliefs, are changed upon revision, too. It is therefore an interesting direction for future work to develop further generalizations of the AGM and DP postulates to higher-order beliefs, with the aim of classifying concrete multi-agent belief revision operators according to their treatment of each nested belief.

It is worth noting that the strong relationship between MBR and *announcements*, a type of actions specific to multi-agent domains and very important for epistemic planning, raises several interesting problems worthy of consideration, too, especially since higher-degree beliefs of agents are an important subject of study in formalizing announcement actions in the literature. Several approaches to dealing with announcements have been proposed, but the majority of them place restrictions on the announcements (e.g., only considering public announcements) and the relationship between different methods is hardly understood. A systematic evaluation of these approaches under the lens of a set of generalized postulates for MBR could provide insights into the development of a general approach to dealing with announcements. In this regard, we note that Baltag and Smets (2006) also discuss belief revision in the multi-agent case, as we did in this paper. Our approach differs in two key respects, however: First, they employ epistemic plausibility models as the underlying representation, whereas we adopt Kripke structures as used in dynamic epistemtic reasoning and planning. Second, an epistemic plausibility model contains a priori plausibility relations for agents that dictate how agents would revise their beliefs. In our approach, agents do not have such policies. A detailed comparison between the two approaches will be one of our interesting research topics in the near future. Similarly, Lorini, Perrotin, and Schwarzentruber (2022) discuss belief updates for epistemic actions over belief bases, and thus could also be used for agents to revise their beliefs when an announcment is made. However, their focus is to define the updates. It would be interesting to determine whether the result of the update by a private announcement "*agent a was told that φ*" satisfies the generalized postulates discussed in this paper, similar to our $*_{\text{ev}}$ operator. Again, we leave this for future work.

Acknowledgements

This research was partially supported by the Australian Research Council (grant #DP250101822) and by the Australian Government through the CRC-P project *Urban Copilot – AI accelerating where we build Australia's future*. The second author acknowledges the partial support of the NSF grants #2139028, #2139028, #2151254, and the internal IAAM grant #139198.

AI Declaration

AI tools were used solely for automated spelling and grammar checks and to provide stylistic suggestions.

References

Alchourrón, C. E., and Makinson, D. 1982. On the logic of theory change: Contraction functions and their associated revision functions. *Theoria* 48(1):14–37.

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

Aravanis, T. I., and Peppas, P. 2017. Belief revision in answer set programming. In *Proceedings of the 21st Pan-Hellenic Conference on Informatics, Larissa, Greece, September 28-30, 2017*, 2:1–2:5.

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

Baltag, A., and Smets, S. 2006. Conditional doxastic models: A qualitative approach to dynamic belief revision. *Electronic Notes in Theoretical Computer Science* 165:5–21. Proceedings of the 13th Workshop on Logic, Language, Information and Computation (WoLLIC 2006).

Baltag, A., and Smets, S. 2008. A qualitative theory of dynamic interactive belief revision. In *Proceedings of the 7th Conference on Logic and the Foundations of Game and Decision Theory (LOFT)*, Texts in Logic and Games 3, 13–60. Amsterdam University Press.

Baltag, A.; Moss, L.; and Solecki, S. 1998. The logic of public announcements, common knowledge, and private suspicions. In *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 43–56.

Baral, C.; Gelfond, G.; Pontelli, E.; and Son, T. C. 2015. An Action Language for Mutli-Agent Domains: Foundations. *arXiv preprint https://arxiv.org/abs/1511.01960*.

Baral, C.; Bolander, T.; McIlraith, S.; and Ditmarsch, H. V. 2017. Epistemic planning. Technical report, Dagstuhl.

Baral, C.; Gelfond, G.; Pontelli, E.; and Son, T. C. 2022. An action language for multi-agent domains. *Artificial Intelligence* 302:103601.

Bolander, T., and Andersen, M. 2011. Epistemic Planning for Single and Multi-Agent Systems. *Journal of Applied Non-Classical Logics* 21(1).

Bolander, T. 2017. A gentle introduction to epistemic planning: The DEL approach. *arXiv preprint https://arxiv.org/abs/1703.02192*.

Burigana, A., and Fabiano, F. 2026. The epistemic planning domain definition language: Official guideline. *arXiv preprint https://arxiv.org/abs/2601.20969*.

Chockler, H., and Halpern, J. Y. 2004. Responsibility and blame: A structural-model approach. *Journal of Artificial Intelligence Research* 22:93–115.

Dalal, M. 1988. Investigations into theory of knowledge base revision. In *Proceedings of the 7th AAAI National Conference on Artificial Intelligence, AAAI*, 449–479.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial Intelligence* 89(1-2):1–29.

Davila, J. F.; Longin, D.; Emiliano; and Maris, F. 2024. Logic-based cognitive planning for conversational agents. *Autonomous Agents and Multi-Agent Systems* 38(1):20.

Delgrande, J. P.; Schaub, T.; Tompits, H.; and Woltran, S. 2008. Belief revision of logic programs under answer set semantics. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, 411–421.

Ditmarsch, H. v.; van der Hoek, W.; and Kooi, B. 2007. *Dynamic Epistemic Logic*. Springer Publishing Company, Incorporated, 1st edition.

Dragoni, A. F., and Giorgini, P. 1996. Belief revision through the belief-function formalism in a multi-agent environment. In Müller, J. P.; Wooldridge, M. J.; and Jennings, N. R., eds., *Proceedings of the 3rd Workshop on Agent Theories, Architectures, and Languages (ATAL)*, volume 1193 of *Lecture Notes in Computer Science*, 103–115. Springer.

Dragoni, A. F., and Puliti, P. 1994. Distributed belief revision versus distributed truth maintenance. In *Proceedings of the Sixth International Conference on Tools with Artificial Intelligence (TAI)*, 499–505. IEEE Computer Society.

Fagin, R.; Halpern, J.; Moses, Y.; and Vardi, M. 1995. *Reasoning about Knowledge*. MIT Press.

Herzig, A.; Lang, J.; and Marquis, P. 2005. Action Progression and Revision in Multiagent Belief Structures. In *Sixth Workshop on Nonmonotonic Reasoning, Action, and Change (NRAC)*.

Hunter, A., and Delgrande, J. P. 2011. Iterated belief change due to actions and observations. *Journal of Artificial Intelligence Research* 40:269–304.

Jin, Y., and Thielscher, M. 2007. Iterated belief revision, revised. *Artificial Intelligence* 171(1):1–18.

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

Liu, W., and Williams, M. 1999. A framework for multi-agent belief revision, part I: the role of ontology. In Foo, N. Y., ed., *Proceedings of the 12th Australian Joint Conference on Artificial Intelligence, AI*, volume 1747 of *Lecture Notes in Computer Science*, 168–179. Springer.

Lorini, E., and Schwarzentruber, F. 2021. Multi-agent belief base revision. In Zhou, Z., ed., *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1959–1965.

Lorini, E.; Perrotin, E.; and Schwarzentruber, F. 2022. Epistemic Actions: Comparing Multi-agent Belief Bases with Action Models. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning*, 236–246.

Malheiro, B.; Jennings, N. R.; and Oliveira, E. 1994. Belief revision in multi-agent systems. In Cohn, A. G., ed., *Proceedings of the 11th European Conference on Artificial Intelligence*, 294–298. John Wiley and Sons, Chichester.

Peppas, P. 2014. A panorama of iterated revision. In *David Makinson on Classical Methods for Non-Classical Problems. Outstanding Contributions to Logic*. Springer Netherlands. 71–94.

Rajaratnam, D., and Thielscher, M. 2021. Representing and reasoning with event models for epistemic planning. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 18(1), 519–528.

Satoh, K. 1988. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems (FGCS)*, 455–462. Springer.

Shapiro, S.; Pagnucco, M.; Lespérance, Y.; and Levesque, H. J. 2011. Iterated belief change in the situation calculus. *Artificial Intelligence* 175(1):165–192.

Son, T. C.; Pham, L.; and Pontelli, E. 2024. On dealing with false beliefs and maintaining KD45n property. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, 1763–1771.

Tardivo, F.; Pham, L.; Son, T. C.; and Pontelli, E. 2021. A logic programming approach to regression based repair of incorrect initial belief states. In *Proceedings of the 23rd International Symposium on Practical Aspects of Declarative Languages (PADL)*, volume 12548 of *Lecture Notes in Computer Science*, 73–89. Springer.

Thielscher, M., and Son, T. C. 2026. A study of belief revision postulates in multi-agent systems (extended version). *arXiv preprint https://arxiv.org/abs/2605.02249*.

Tu, P.; Son, T.; Gelfond, M.; and Morales, R. 2011. Approximation of action theories and its application to conformant planning. *Artificial Intelligence* 175(1):79–119.

van Benthem, J.; van Eijck, J.; and Kooi, B. P. 2006. Logics of communication and change. *Information and Computation* 204(11):1620–1662.

van Benthem, J. 2007. Dynamic logic of belief revision. *Journal of Applied Non-Classical Logics* 17(2):129–155.

van Ditmarsch, H.; Hendriks, P.; and Verbrugge, R. 2020. Editors' review and introduction: Lying in logic, language, and cognition. *Topics in Cognitive Science* 12(2):466–484.

van Ditmarsch, H.; van der Hoek, W.; and Kooi, B. 2007. *Dynamic Epistemic Logic*. Springer.

Vasileiou, S. L.; Kumar, A.; Yeoh, W.; Son, T. C.; and Toni, F. 2024. Dialectical reconciliation via structured argumentative dialogues. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, volume 21(1), 777–787.