

On Dealing with False Beliefs and Maintaining KD45_n Property

Tran Cao Son

Computer Science Department
New Mexico State University
Las Cruces, NM, United States
stran@nmsu.edu

Loc Pham

Computer Science Department
New Mexico State University
Las Cruces, NM, United States
locpham@nmsu.edu

Enrico Pontelli

Computer Science Department
New Mexico State University
Las Cruces, NM, United States
epontell@nmsu.edu

ABSTRACT

Motivated by counterintuitive results of the update of a Kripke structure using an update model, such as the inability to correct false beliefs of agents or the creation of incoherent state of beliefs of agents, this paper explores a novel methodology for updating a Kripke structure using an update model. The paper shows that the new definition helps agents correct their false beliefs when they are full observers of a sensing action or a truthful announcement. Furthermore, the paper presents a sufficient condition for update models that guarantees that the resulting Kripke structure maintains the KD45_n property of the original Kripke structure if the update model is also KD45_n . In particular, the majority of update models recently described in the literature for reasoning about knowledge and beliefs of agents in multi-agent domains satisfy such sufficient condition. This implies that the KD45_n property will be maintained after the execution of an action sequence if the initial Kripke structure is KD45_n .

These results can help guide the design of update models for compound actions in applications dealing with knowledge and beliefs; they can also be used by epistemic planners that employ update models to correct agents' false beliefs.

KEYWORDS

Correcting False Belief; Maintaining KD45_n ; Update models; Update by update models; Reasoning about actions in multi-agent domains

ACM Reference Format:

Tran Cao Son, Loc Pham, and Enrico Pontelli. 2024. On Dealing with False Beliefs and Maintaining KD45_n Property. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

1 INTRODUCTION AND MOTIVATION

Update models¹, first proposed in [3], are useful tools to study actions' effects in multi-agent systems, such as in the development of dynamic epistemic logic [16] and in high-level action languages [6, 14]. The key idea behind an update model is that an action occurrence in a multi-agent setting can have many facets and can be perceived differently by different agents, which, ultimately, affects in different ways the beliefs of the agents. An update model encodes

¹Also known as *action models* or *event models*.



This work is licensed under a Creative Commons Attribution International 4.0 License.

different perspectives of agents given the state of the world and the state of knowledge and beliefs of agents and an action occurrence. Formally, an update model consists of a collection of *events*, each associated with a precondition and a substitution, and accessibility relations of agents between these events. Intuitively, each event encodes a possible view of an action occurrence by the agents; for example, an event can encode the fact that the agent believes that the action *does occur*, *does not occur*, or *might or might not occur*, etc.; the accessibility relation of an agent encodes its uncertainty about the events.

Fig. 1 graphically shows an update model (in the bottom left box), denoted by $\Sigma_{\text{look}(B)}$, that encodes the perspectives of agents *A*, *B*, and *C* when *B* looks at a coin and *A* and *C* are watching *B* but cannot see what *B* actually sees. This update model has two events, σ (corresponding to the event that *B* sees *head*) and τ (*B* sees *tail*), which are drawn as squares with their names and preconditions below the event. The accessibility relations of the agents are drawn as labeled links between events. Because *B*, who executes the action, can distinguish between σ and τ , the only links labeled *B* are the loops around σ and τ . On the other hand, *A* and *C* cannot distinguish between σ and τ . This fact is represented by the link labeled *A* and *C* between σ and τ and the loops around σ and τ .

Given an update model and a Kripke structure encoding the state

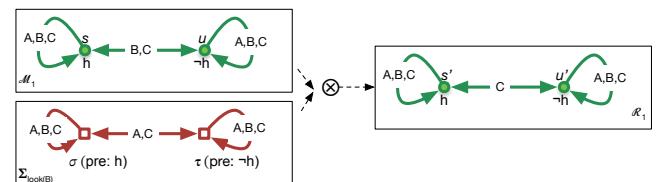


Figure 1: *B* looks at the coin while *A* and *C* are watching

of the world and the state of knowledge/beliefs of agents, the application of the update model to the Kripke structure results in a new Kripke structure. In the new Kripke structure, agents' beliefs are refined in accordance to their perspectives about the action occurrence. For example, the update of the Kripke structure in the top left corner of Fig. 1, denoted by \mathcal{M}_1 , using the update model $\Sigma_{\text{look}(B)}$ creates the Kripke structure on the right (\mathcal{R}_1). The Kripke structure \mathcal{M}_1 has two worlds (*s* and *u*) and *h* (representing head) is true (false) in *s* (*u*). The accessibility relations between *s* and *u* are drawn as labeled links between *s* and *u*. \mathcal{M}_1 encodes that it is common knowledge that *A* knows the status of the coin and that *B* and *C* do not. The resulting Kripke structure \mathcal{R}_1 indicates that it is common knowledge that *A* and *B* know the status of the coin but *C* does not.

A critical issue in the use of update models for reasoning about effects of actions in multi-agent domains is that agents can become ignorant after an update (or, the beliefs of the agents become incoherent). To see how this counter-intuitive result is possible, consider the application of the same update model $\Sigma_{\text{look}(B)}$ on the Kripke structure in the upper left corner (\mathcal{M}_2) in Fig. 2. Assume that the coin lies heads up, i.e., the true state of the world is represented by the world s on the left. The difference between \mathcal{M}_1 and \mathcal{M}_2 lies in that B has a false belief about the status of the coin (B believes that the coin lies tails up). The update of $\Sigma_{\text{look}(B)}$ results in the Kripke structure \mathcal{R}_2 in which B becomes ignorant at the true state of the world (s'). This is rather counter-intuitive since B , by virtue of looking at the coin, should have known that the coin lies heads up. We note that the problem arises because B has a false belief about the status of the coin to start with. Instead of being able to correct the false belief, B 's beliefs become incoherent as B would conclude that every formula is true! We refer to this problem as *inability to correct false beliefs*.

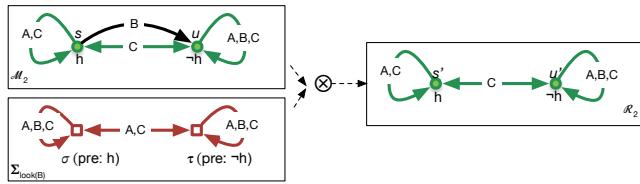


Figure 2: B looks at the coin and becomes ignorant

A consequence of losing all links labeled B in some world leads to the issue that the resulting Kripke structure (\mathcal{R}_2) no longer satisfies the KD45_n property (formal definition is provided in the next section) which is satisfied by both the original Kripke structure \mathcal{M}_2 and the update model $\Sigma_{\text{look}(B)}$. We will refer to this problem as *the problem of losing KD45_n property* (or *loss of KD45_n* , for short).

An important consequence of the loss of KD45_n is that if we would like to reason about knowledge and beliefs of agents then we will have to deal with two modalities, the knowledge modal operator K and the belief modal operator B or define knowledge formulae via belief formulae. This problem has been discussed in the literature, especially in the context of dynamic multi-agent systems (see, e.g., [1, 15]). We observe that, with the exception of the work by Buckingham et al. (2020) who explicitly used both B and K or Baltag and Smets (2016) who defined knowledge, belief, and conditional belief operators and work with them, the majority of proposals on formalizing actions in multi-agent domains [6–8, 14, 16, 17] only consider one modality and do not discuss how the other modality should be addressed, e.g., only B is used and no discussion on K is included. This is inadequate as, in general, an agent can have a false belief—which can easily arise when an agent is oblivious to the execution of an action—and *false beliefs are not equivalent to incorrect knowledge*.

In this paper, we take a different approach towards correcting false beliefs of agents and maintaining KD45_n after the execution of an action sequence. Instead of focusing on developing new definitions for update models, such as the *edge-conditioned* update models proposed in [8], or identifying conditions of update models

or the original Kripke structure, as done in [1, 4, 11, 15], we first propose a modification of how update models are applied to Kripke structures. We discuss the motivation of the new definition and prove that this definition helps correcting agents' false beliefs if the agents are full observers of a sensing action or a truthful announcement. We also propose a new syntactic characterization of update models to maintain the KD45_n property of a Kripke structure after the execution of an action sequence. In particular, we show that recently proposed update models for the study of action effects in multi-agent domains in the literature satisfy this condition. In summary, the contributions of the paper are:

- A novel definition of the application of an update model to a Kripke structure that allows full observer agents of a sensing or a truthful announcement action to correct their false beliefs;
- A sufficient condition to maintain the KD45_n property of a Kripke structure under the new definition of update;
- Propositions showing that several recently developed update models in the literature satisfy the proposed sufficient condition and therefore can be used in the development of epistemic planners that can plan for both knowledge and goals; and
- A discussion of potential applications of the proposed method.

2 BACKGROUND

2.1 Logics of Knowledge and Belief

We consider the standard logic of knowledge and belief with a set of modalities L_1, \dots, L_k and use the notation from [10]. Let \mathcal{P} be a set of *propositions*. We define $\mathcal{L}_{\mathcal{AG}}(L_1, \dots, L_k)$ to be the set of formulae defined as follows. Each $p \in \mathcal{P}$ is a formula. If φ and ψ are formulae, then so are $\neg\varphi$, $\varphi \rightarrow \psi$, and $L_i\varphi$ ($i = 1, \dots, k$). The connectives $\vee, \wedge, \leftrightarrow$ can be defined in terms of \neg and \rightarrow . An *atomic formula* is a formula that does not contain any modal operator L_i .

A logic Λ is a set of formulae in $\mathcal{L}_{\mathcal{AG}}(L_1, \dots, L_k)$ that: (i) contains all propositional tautologies; (ii) is closed under modus ponens; and (iii) is closed under substitution. A logic is *normal* if it contains the axioms $L_i(\varphi \rightarrow \psi) \rightarrow (L_i\varphi \rightarrow L_i\psi)$, referred to as (\mathbf{K}_L) for $i = 1, \dots, k$, and is closed under *generalization*, i.e., if it contains φ , then it will also contain $L_i\varphi$. A logic generated by a set A of formulae (*axioms*) is the smallest normal logic containing A . For two sets of axioms Λ_1 and Λ_2 , $\Lambda_1 + \Lambda_2$ is the smallest normal logic containing Λ_1 and Λ_2 .

Consider a set $\mathcal{AG} = \{1, 2, \dots, n\}$ of n agents, each agent i is associated with a belief operator B_i and a knowledge operator K_i . Our focus is the logic of belief, called KD45_n , over the language $\mathcal{L}_{\mathcal{AG}}(B_1, \dots, B_n)$ that is generated by the following axioms: (D) $B_i\varphi \rightarrow B_i\neg\varphi$, (4) $B_i\varphi \rightarrow B_iB_i\varphi$, and (5) $\neg B_i\varphi \rightarrow B_i\neg B_i\varphi$ where $i \in \mathcal{AG}$ and $\varphi \in \mathcal{L}_{\mathcal{AG}}(B_1, \dots, B_n)$. It is shown in [10] that knowledge can be reduced to true belief in KD45_n . More specifically, the knowledge modality K_i is reducible to B_i in KD45_n by the axiom $K_i\varphi \leftrightarrow (B_i\varphi \wedge \varphi)$. This also means that we can remove the knowledge modal operator from the language, yet still be able to derive conclusions about knowledge of agents.

A *Kripke frame* \mathcal{F} is a tuple $\langle S, \mathcal{B}_1, \dots, \mathcal{B}_n \rangle$ where S is a set of worlds (or *points*) and $\mathcal{B}_i \subseteq S \times S$ for $i \in \mathcal{AG}$, called the *accessibility*

relation for i . A Kripke structure² (or Kripke model) \mathcal{M} based on the frame \mathcal{F} is a pair (\mathcal{F}, π) , where $\pi : \mathcal{F}[S] \rightarrow 2^{\mathcal{P}}$ is a function that associates an interpretation of \mathcal{P} with each world in \mathcal{F} . For $\mathcal{M} = (\mathcal{F}, \pi)$, $\mathcal{M}[\pi]$ denotes π and $\mathcal{M}[S]$ and $\mathcal{M}[i]$ denote the set of worlds S and \mathcal{B}_i of \mathcal{F} , respectively. We say that an agent i is *ignorant* at s in \mathcal{M} if $\mathcal{M}[i](s) = \{u \mid (s, u) \in \mathcal{M}[i]\} = \emptyset$.

A pointed Kripke structure (or p-structure) is a pair (\mathcal{M}, s) , where \mathcal{M} is a Kripke structure and $s \in \mathcal{M}[S]$, called the *actual world*. We will often represent a Kripke model \mathcal{M} by a directed graph whose nodes are the worlds in $\mathcal{M}[S]$ and the labeled edges are the members of $\mathcal{M}[i]$. The interpretation associated to a world is often drawn below the corresponding node (e.g., as in Figures 1–2). Entailment of formulae in $\mathcal{L}_{\mathcal{AG}}(L_1, \dots, L_n)$ w.r.t. a p-structure is defined next.

DEFINITION 1. Given a formula φ and a p-structure (\mathcal{M}, s) :

- $(\mathcal{M}, s) \models \varphi$ if $\mathcal{M}[\pi](s) \models \varphi$ and φ is an atomic formula;
- $(\mathcal{M}, s) \models \neg\varphi$ if $(\mathcal{M}, s) \not\models \varphi$;
- $(\mathcal{M}, s) \models \varphi_1 \rightarrow \varphi_2$ if $(\mathcal{M}, s) \models \neg\varphi_1$ or $(\mathcal{M}, s) \models \varphi_2$;
- $(\mathcal{M}, s) \models L_i \varphi$ if, for each t s.t. $(s, t) \in \mathcal{B}_i$, $(\mathcal{M}, t) \models \varphi$.

Observe that if $\mathcal{M}[i](s) = \emptyset$ then $(\mathcal{M}, s) \models B_i \varphi$ for every atomic formula φ . As such, we also say that the belief of an ignorant agent at s in \mathcal{M} is *incoherent*.

A relation $R \subseteq S \times S$ is: *reflexive* if $(u, u) \in R$ for every $u \in S$; *serial* if for every $u \in S$ there exists some $v \in S$ such that $(u, v) \in R$; *transitive* if $(u, v) \in R$ and $(v, z) \in R$ imply that $(u, z) \in R$; *Euclidean* if $(u, v) \in R$ and $(u, z) \in R$ imply that $(v, z) \in R$.

Frames can be characterized by the properties of their accessibility relations. It is known that a frame $\mathcal{F} = (S, \mathcal{B}_1, \dots, \mathcal{B}_n)$ is a KD45_n frame iff for every $i = 1, \dots, n$, \mathcal{B}_i is serial, transitive, and Euclidean. A Kripke model $\mathcal{M} = (\mathcal{F}, \pi)$ is said to be a KD45_n model if its frame \mathcal{F} is a KD45_n frame.

2.2 Update Models

Update models describe transformations of (pointed) Kripke structures according to a predetermined pattern. An update model uses structures similar to pointed Kripke structures and they describe the effects of a transformation on p-structures using an *update operator* [2, 16].

A set $\{p \rightarrow \varphi \mid p \in \mathcal{P}, \varphi \in \mathcal{L}_{\mathcal{AG}}(\mathcal{B}_1, \dots, \mathcal{B}_n)\}$ is called an $\mathcal{L}_{\mathcal{AG}}(\mathcal{B}_1, \dots, \mathcal{B}_n)$ -substitution (or substitution, for short). For each substitution sub and each $p \in \mathcal{P}$, we assume that sub contains exactly one formula $p \rightarrow \varphi$. For simplicity of the presentation, we often omit $p \rightarrow p$ in a substitution. $SUB_{\mathcal{L}_{\mathcal{AG}}}$ denotes the set of all substitutions. A substitution is used to encode changes caused by an action occurrence. A formula $p \rightarrow \varphi$ in a substitution states the condition (φ) under which p will become true. For example, the action of flipping a coin can be represented by the substitution $\{h \rightarrow \neg h\}$ which says that h (the coin lies heads up) is true if $\neg h$ (the coin lies heads down) was true when the flip was performed.

DEFINITION 2 (UPDATE MODEL). An update model Σ is a tuple $\langle \Sigma, R_1, \dots, R_n, pre, sub \rangle$ where

- Σ is a non-empty set, whose elements are called events;
- each R_i is a binary relation over Σ ;

² We use Kripke structure interchangeably with Kripke model.

- $pre : \Sigma \rightarrow \mathcal{L}_{\mathcal{AG}}(\mathcal{B}_1, \dots, \mathcal{B}_n)$ is a function mapping each event $a \in \Sigma$ to a formula in $\mathcal{L}_{\mathcal{AG}}(\mathcal{B}_1, \dots, \mathcal{B}_n)$; and
- $sub : \Sigma \rightarrow SUB_{\mathcal{L}_{\mathcal{AG}}}$.

An update instance ω is a pair (Σ, e) where Σ is an update model and $e \in \Sigma$ (the designated event). An update template is a pair (Σ, Γ) where Σ is an update model with the set of events Σ and $\Gamma \subseteq \Sigma$.

The designated event in an update instance is the one that agents who are fully aware of the action occurrence will observe. Templates extend the notion of instance to capture non-deterministic actions and compound actions. The relation R_i describes agent i 's uncertainty about an action occurrence—i.e., if $(\sigma, \tau) \in R_i$ and event σ is performed, then agent i may believe that event τ is executed instead. pre defines the action precondition and sub specifies the changes of fluent values after the execution of an action. An update model is *serial* (resp., *reflexive*, *transitive*, *Euclidean*) if, for every $i \in \mathcal{AG}$, R_i is serial (resp., reflexive, transitive, Euclidean).

DEFINITION 3 (UPDATES USING AN UPDATE MODEL). Given an update model $\Sigma = \langle \Sigma, R_1, \dots, R_n, pre, sub \rangle$ and a Kripke structure $\mathcal{M} = (\mathcal{F}, \pi)$, the update operator induced by Σ defines a Kripke structure $\mathcal{M}' = \mathcal{M} \otimes \Sigma$, where:

- $\mathcal{M}'[S] = \{(s, \tau) \mid s \in \mathcal{M}[S], \tau \in \Sigma, (\mathcal{M}, s) \models pre(\tau)\}$;
- $((s, \tau), (s', \tau')) \in \mathcal{M}'[i]$ iff $\{(s, \tau), (s', \tau')\} \subseteq \mathcal{M}'[S], (s, s') \in \mathcal{M}'[i]$ and $(\tau, \tau') \in R_i$;
- $\forall f \in \mathcal{P}. [\mathcal{M}'[\pi](s, \tau)] \models f$ iff $f \rightarrow \varphi \in sub(\tau), (\mathcal{M}, s) \models \varphi$.

The update of a p-structure (\mathcal{M}, s) given an update template (Σ, Γ) is a set of p-structures, denoted by $(\mathcal{M}, s) \otimes (\Sigma, \Gamma)$, where $(\mathcal{M}', s') \in (\mathcal{M}, s) \otimes (\Sigma, \Gamma)$ iff it holds that $\mathcal{M}' = \mathcal{M} \otimes \Sigma$ and $s' = (s, \tau)$ where $\tau \in \Gamma$ and $s' \in \mathcal{M}'[S]$. Intuitively, the set $(\mathcal{M}, s) \otimes (\Sigma, \Gamma)$ is the set of p-structures encoding the result of the execution of the action, which is represented by the update template (Σ, Γ) , in the p-structure (\mathcal{M}, s) . It is easy to see that Kripke structure \mathcal{R}_1 in Fig. 1 is the result of the update of \mathcal{M}_1 by $\Sigma_{look(B)}$, i.e., $\mathcal{R}_1 = \mathcal{M}_1 \otimes \Sigma_{look(B)}$.

3 UPDATES USING AN UPDATE MODEL: PROPOSED CHANGES

Let us consider the Kripke structure and the update model in Fig. 2 and let us precisely describe \mathcal{M}_2 and $\Sigma_{look(B)}$:

- $\mathcal{M}_2 = (\langle \{s, u\}, \mathcal{B}_A, \mathcal{B}_B, \mathcal{B}_C \rangle, \pi)$,
 $\mathcal{B}_A = \{(s, s), (u, u)\}, \mathcal{B}_B = \{(s, u), (u, u)\},$
 $\mathcal{B}_C = \{(s, s), (u, u), (s, u), (u, s)\}$, and
 $\pi(s) = \{h\}, \pi(u) = \emptyset$.
- $\Sigma_{look(B)} = \langle \{\sigma, \tau\}, R_A, R_B, R_C, pre, sub \rangle$ where
 $pre(\sigma) = h, pre(\tau) = \neg h$,
 $R_A = R_C = \{(\sigma, \sigma), (\tau, \tau), (\sigma, \tau), (\tau, \sigma)\}$,
 $R_B = \{(\sigma, \sigma), (\tau, \tau)\}$, and
 $sub(\sigma) = sub(\tau) = \emptyset$ (or $sub(\sigma) = sub(\tau) = \{h \rightarrow h\}$ as noted above).

It is easy to see that both \mathcal{M}_2 and $\Sigma_{look(B)}$ are serial, transitive, and Euclidean. Furthermore, $\mathcal{R}_2 = \mathcal{M}_2 \otimes \Sigma_{look(B)}$ where $\mathcal{R}_2 = (R_2, \pi)$ with $R_2 = \langle \{s', u'\}, \mathcal{B}'_A, \mathcal{B}'_B, \mathcal{B}'_C \rangle$, $\mathcal{B}'_A = \{(s', s'), (u', u')\}$, $\mathcal{B}'_B = \{(u', u')\}$, $\mathcal{B}'_C = \{(s', s'), (u', u'), (s', u'), (u', s')\}$, and $\pi(s') = \{h\}, \pi(u') = \emptyset$.

In the definition of R_2 , we have that $s' = (s, \sigma)$ and $u' = (u, \tau)$.

Let us now examine the reason why B does not have any outgoing link from $s' = (s, \sigma)$. This is because, by Definition 3, the

outgoing link labeled B from s and the outgoing link labeled B from σ are *incompatible*, i.e., $(s, u) \in B_B$ and $(\sigma, \sigma) \in R_B$. On the other hand, the intuition behinds the definition of an update model says that $(\sigma, \sigma) \in R_B$ indicates that if σ occurs then B would recognize that this event happens. In addition, since s satisfies the precondition of σ , it is reasonable to expect that B would perceive that $s' = (s, \sigma)$ is one of the possible worlds in the Kripke structure that results from the execution of looking on the coin. Indeed, it is easy to verify that the following proposition holds.

PROPOSITION 1. *Let $\Sigma^* = \langle \{\sigma\}, R_1, \dots, R_n, \text{pre}, \text{sub} \rangle$ where $R_i = \{(\sigma, \sigma)\}$ for every i , $\text{pre}(\sigma) = \top$, and $\text{sub}(\sigma) = \emptyset$. Given a p -structure (\mathcal{M}, s) and an agent i who has a false belief about an atomic formula ψ , i.e., $(\mathcal{M}, s) \models \psi$ and $(\mathcal{M}, s) \models B_i \neg \psi$, then i becomes ignorant at (s, σ) in $\mathcal{M} \otimes \Sigma^*$.*

We note that Σ^* is a natural representation of an occurrence of a truthful public announcement (see, e.g., [2]). As such, the above proposition exposes a critical issue in using update models: agents can become ignorant after the execution of an action, even though, intuitively, agents must have been able to correct their false beliefs due to the action occurrence. A similar proposition can be proved for the occurrence of a sensing action.

The above discussion suggests that if (x, σ) is a new world (i.e., $(\mathcal{M}, x) \models \text{pre}(\sigma)$), i has false belief about the precondition of σ (i.e., $(\mathcal{M}, x) \models B_i \neg \text{pre}(\sigma)$), and $(\sigma, \sigma) \in R_i$ then we should have $((x, \sigma), (x, \sigma)) \in \mathcal{M}'[i]$. We observe that this might not be desirable in some situations. Recall that the addition of the loop is aimed at keeping the coherence of the agent's beliefs. Under this view, if there exists some world u and an event τ accessible from s and σ , respectively, and $(\mathcal{M}, u) \models \text{pre}(\tau)$ then we have $((x, \sigma), (u, \tau)) \in \mathcal{M}'[i]$ (see our discussion related to Figure 8 in Section 5). In other words, i 's belief is still coherent though it might be wrong. Taking this into consideration, we propose the following change to Definition 3:

Change #1: if (x, σ) is a new world (i.e., $(\mathcal{M}, x) \models \text{pre}(\sigma)$) and $C_i(x, \sigma)$ is true then $((x, \sigma), (x, \sigma)) \in \mathcal{M}'[i]$ if $(\sigma, \sigma) \in R_i$ where $C_i(x, \sigma)$ encodes the following statement:

for every u such that $(x, u) \in \mathcal{M}[i]$ there exists $\tau \in \Sigma$ such that $(\sigma, \tau) \in R_i$ and $(\mathcal{M}, u) \models \text{pre}(\tau)$.

From now on, we write $C_i(x, \sigma)$ and means

$$\forall u. [(x, u) \in \mathcal{M}[i] \Rightarrow \exists \tau \in \Sigma. [(\sigma, \tau) \in R_i \wedge (\mathcal{M}, u) \models \text{pre}(\tau)]]$$

Figure 3 shows the result of the update of \mathcal{M}_2 by $\Sigma_{\text{look}(B)}$ that takes into consideration **Change #1**. The link $(s', s') \in B_B$ is added due to **Change #1** that helps B to learn that the coin is heads up and realize that B 's initial belief that the coin lies tails up was wrong.

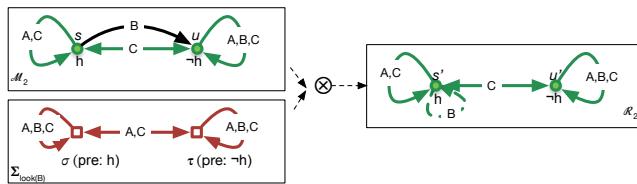


Figure 3: B looks at the coin and corrects its false belief

Change #1 considers the loop around σ . This might not be sufficient as there are links of the form (σ, τ) in R_i that express the similar intuition, i.e., if σ occurs then i could perceive that τ occurs as well. We motivate the second change using the following situation, taken from [15]. Let us consider a Kripke structure $\mathcal{M}_3 = (M_3, \pi)$ and an update model $\Sigma_{\text{ann}(o)}$ as shown in Figure 4. Here, $M_3[S] = \{s, u, t, w\}$, the accessibility relations, and the interpretations associated to the worlds are given in the figure.

The update model $\Sigma_{\text{ann}(o)}$ encodes a public and truthful announcement of o by A . Intuitively, everyone should believe (know) that o is true after the execution of this action.

It is easy to see that C has the false belief about o in s , i.e., $(M_3, s) \models o \wedge B_C \neg o$. Therefore, without **Change #1**, the update of \mathcal{M}_3 by $\Sigma_{\text{ann}(o)}$ will render that C 's belief is incoherent. As such, **Change #1** should be applied.

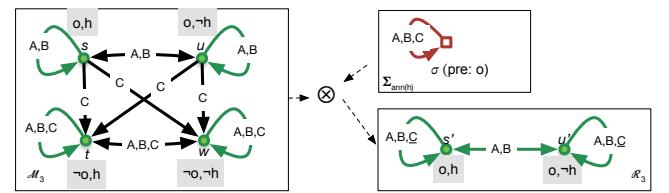


Figure 4: A announces o (with Change #1) - C miraculously knows h

Let us now consider the result of the update of \mathcal{M}_3 by $\Sigma_{\text{ann}(o)}$, with **Change #1**. It is shown in \mathcal{R}_3 in Figure 4. Because of **Change #1**, we have that $(s', s') \in \mathcal{R}_3[C]$ and $(u', u') \in \mathcal{R}_3[C]$ where, for $x \in \{s, u\}$, $x' = (x, \sigma)$. The question of interest is then should $(s', u') \in \mathcal{M}'[C]$? Looking at \mathcal{R}_3 , we can see that without $(s', u') \in \mathcal{R}_3[C]$ then $(\mathcal{R}_3, s') \models B_C h$. In other words, learning that o is true miraculously allows C to know the value of h . On the other hand, there is no reason for C to know h since before A announces o , C —like A and B —does not know h . After A announces o , both A and B still do not know h . Therefore, it is counterintuitive that C would know h , which would require (s', u') to belong to $\mathcal{R}_3[C]$. Similarly, we should have $(u', s') \in \mathcal{R}_3[C]$.

The above discussion stipulates that $((s, \sigma), (u, \tau))$ should be added to $\mathcal{M}'[i]$ if the conditions for adding $((s, \sigma), (s, \sigma))$ and $((u, \tau), (u, \tau))$ to $\mathcal{M}'[i]$ are satisfied and $(\sigma, \tau) \in R_i$. This is not always reasonable, however. Consider the Kripke structure \mathcal{M}'_3 (Figure 5) that is obtained from \mathcal{M}_3 by removing the set of links $\{(s, w), (u, t), (t, w), (w, t)\}$ from $\mathcal{M}_3[C]$.

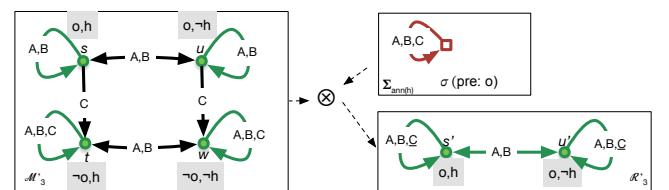


Figure 5: A announces o (with Change #1) - C maintains belief about h because it knows whether h before the announcement

In this situation, C 's belief is the world $\{\neg o, h\}$, i.e., C knows that h is true. Therefore, \mathcal{R}'_3 would be a reasonable result of updating \mathcal{M}'_3 by $\Sigma_{ann(o)}$. We observe that the difference between \mathcal{M}_3 and \mathcal{M}'_3 lies in that s and u are *not connected* by C . In addition, we observe that there is no link labeled C from s and u ; for otherwise, the link $((s, \sigma), (u, \tau))$ belongs to $\mathcal{R}_3[C]$ by definition.

DEFINITION 4. Let $\mathcal{M} = (M, \pi)$ be a Kripke structure over \mathcal{AG} and \mathcal{F} . We say that two worlds s, u in \mathcal{M} are connected by $i \in \mathcal{AG}$ if there exists a sequence of worlds $s = s_1, \dots, s_n = u$ such that for $j = 1, \dots, n-1$, $(s_j, s_{j+1}) \in \mathcal{M}[i]$ or $(s_{j+1}, s_j) \in \mathcal{M}[i]$.

The above discussion is summarized in the following change:

Change #2: if (x, σ) and (y, τ) are new worlds and $((x, \sigma), (x, \sigma))$ and $((y, \tau), (y, \tau))$ are added to $\mathcal{M}'[i]$ because of **Change #1**, $(x, y) \notin \mathcal{M}[i]$, $(y, x) \notin \mathcal{M}[i]$, and x and y are connected by i then $((x, \sigma), (y, \tau)) \in \mathcal{M}'[i]$.

Given the **Changes #1 & #2**, our new definition of the update using an update model is formalized as follows.

DEFINITION 5 (NEW: UPDATES USING AN UPDATE MODEL). Let \mathcal{M} be a Kripke structure, $\Sigma = \langle \Sigma, R_1, \dots, R_n, pre, sub \rangle$ be an update model as in Definition 3. The update induced by Σ defines a Kripke structure $\mathcal{M}' = \mathcal{M} \otimes \Sigma$, where:

- (i) $\mathcal{M}'[S] = \{(s, \tau) \mid s \in \mathcal{M}[S], \tau \in \Sigma, (\mathcal{M}, s) \models pre(\tau)\}$;
- (ii) For (s, τ) and (s', τ') in $\mathcal{M}'[S]$, $((s, \tau), (s', \tau')) \in \mathcal{M}'[i]$ iff
 - (a) $(s, s') \in B_i$ and $(\tau, \tau') \in R_i$; or
 - (b) $(s, \tau) = (s', \tau')$, $(\tau, \tau') \in R_i$, and $C_i(s, \tau)$ is true;
 - (c) $(s, \tau) \neq (s', \tau')$, $C_i(s, \tau)$ and $C_i(s', \tau')$ are true, $(\tau, \tau'), (\tau', \tau') \in R_i$, $(s, s'), (s', s) \notin \mathcal{M}[i]$, and s and s' are connected by i .
- (iii) For all $(s, \tau) \in \mathcal{M}'[S]$ and $f \in \mathcal{F}$, $\mathcal{M}'[\pi]((s, \tau)) \models f$ if $f \rightarrow \varphi \in sub(\tau)$ and $(\mathcal{M}, s) \models \varphi$.

We note that Item (ii.(a)) maintains the original update in Definition 3; Item (ii.(b)) considers **Change #1**, and Item (ii.(c)) records **Change #2**. It is easy to see that if a link $((s, \tau), (s', \tau'))$ belongs to R_i because of Item (ii.(a)) then $C_i(s, \tau)$ is false. As such, the three classes of links belonging to $\mathcal{M}'[i]$ because of Item (ii.(a)), (ii.(b)), and (ii.(c)), respectively, are mutually exclusive. We prove a property of an update by Definition 5 that is important for the maintenance of the Euclideanness of a Kripke structure.

LEMMA 1. Let Σ be an update model and \mathcal{M} be a Kripke structure. Consider an event σ in Σ , a world s in \mathcal{M} , and an agent i . Assume that $\mathcal{M}' = \mathcal{M} \otimes \Sigma$, $((s, \sigma), (u, \tau))$ and $((s, \sigma), (v, \delta))$ belong to $\mathcal{M}'[i]$ according to Definition 5. It holds that if $((s, \sigma), (u, \tau)) \in \mathcal{M}'[i]$ because of Item (ii.a) then $((s, \sigma), (v, \delta))$ cannot belong to $\mathcal{M}'[i]$ because of Item (ii.c).

Proof. $((s, \sigma), (u, \tau)) \in \mathcal{M}'[i]$ because of Item (ii.a) implies that $(s, u) \in \mathcal{M}[i]$, $(\sigma, \tau) \in R_i$, and $(\mathcal{M}, u) \models pre(\tau)$. Assume the contrary that $((s, \sigma), (v, \delta)) \in \mathcal{M}'[i]$ because of Item (ii.c). This means that $C_i(s, \sigma)$ is true, i.e., $(\mathcal{M}, u) \not\models pre(\tau)$. Contradiction! \square

We will next show that using the new definition of updates, agents can correct their false beliefs if they are full observers of a sensing action or a truthful announcement. Given an action occurrence, we will consider three groups of agents: *full observers*, *partial*

observers, and *oblivious agents*, which will be denoted by F , P , and O , respectively, as done in [6]. Intuitive, agents in F know the effects of the actions and should update their beliefs accordingly; agents in P know that the action occurs and that agents in F know the effects of the action occurrence but they do not; and agents in O are oblivious of the action occurrence. We will assume that F , P , and O are mutual exclusive and $F \cup P \cup O$ equals the set of agents.

3.1 Correcting False Beliefs by Sensing Actions

A *sensing action* is used by agent to learn certain properties of the world, by making direct observations. For example, when agents look at the coin, they will know exactly which face of the coin is up. Consider a sensing action a that helps agents to learn the truth value of a formula ψ whose precondition is φ . An occurrence of a affects an agent i in one of the three ways: i is a full observer ($i \in F$) who will learn ψ ; i is a partial observer ($i \in P$) who will know that full observers know the value of ψ , but i itself does not know the value of ψ ; or i is oblivious ($i \in O$). A reasonable update model for such an action occurrence (see, e.g., [6]), shown in Fig. 6, is

$$\Sigma^{Sensing}(a, \varphi, \psi) = \langle \{\sigma, \tau, \epsilon\}, R_1, \dots, R_n, pre, sub \rangle \text{ where}$$

- $R_i = \{(\sigma, \sigma), (\tau, \tau), (\epsilon, \epsilon)\}$ for $i \in F$;
- $R_i = \{(\sigma, \sigma), (\sigma, \tau), (\tau, \sigma), (\tau, \tau), (\epsilon, \epsilon)\}$ for $i \in P$;
- $R_i = \{(\sigma, \epsilon), (\tau, \epsilon), (\epsilon, \epsilon)\}$ for $i \in O$;
- $pre(\sigma) = \varphi \wedge \psi$, $pre(\tau) = \varphi \wedge \neg\psi$, $pre(\epsilon) = \top$; and
- $sub(\sigma) = sub(\tau) = sub(\epsilon) = \emptyset$.

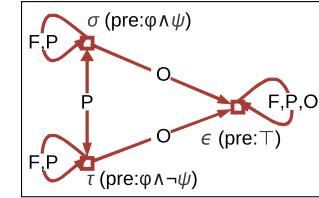


Figure 6: $\Sigma^{Sensing}(a, \varphi, \psi)$: update model for sensing action that sense ψ

PROPOSITION 2. Let $\Sigma^{Sensing}(a, \varphi, \psi)$ be an update model of a sensing action that sense the formula ψ , (\mathcal{M}, s) be a pointed Kripke structure such that $(\mathcal{M}, s) \models B_i \neg\psi$, $(\mathcal{M}, s) \models \psi$, and i a full observer of the action occurrence. It holds that $(\mathcal{M}', s') \models B_i \psi$ where $\mathcal{M}' = \mathcal{M} \otimes \Sigma^{Sensing}(a, \varphi, \psi)$ as in Definition 5 and $s' = (s, \sigma)$.

Proof. Because of $(\mathcal{M}, s) \models \psi$, we have $s' = (s, \sigma) \in \mathcal{M}'[S]$. Because σ is the only event in $\Sigma^{Sensing}$ satisfying that $(\sigma, \sigma) \in R_i$ and $(\mathcal{M}, s) \models B_i \neg\psi$, we have that $C_i(s, \sigma)$ is true. This implies that $\mathcal{M}'[i] = \{(s', s')\}$ where $s' = (s, \sigma)$. Since $sub(\sigma) = \emptyset$, we have that $(\mathcal{M}', s') \models \psi$, and therefore, $(\mathcal{M}', s') \models B_i \psi$. \square

3.2 Correcting False Beliefs by Truthful Announcements

In the literature, an announcement is classified into public, private, or semi-private announcement. We will prove that update model for semi-private announcement will help full observers to correct their false beliefs. The proof for public or private announcement is similar and omitted for brevity.

A *truthful semi-private announcement* of a formula ψ is an action that communicates to a group of agent $F \subseteq \mathcal{AG}$ that the formula is true, while another set of agent $P \subseteq \mathcal{AG}$ ($F \cap P = \emptyset$) are aware of its occurrence and the rest ($O = \mathcal{AG} \setminus (F \cup P)$) are unaware of the communication. Full observers should know that ψ is true after the announcement occurs, partial observers know that the full observers know ψ , and oblivious agents are unaffected.

An update model of a truthful semi-private announcement a of the formula ψ with the precondition φ (see, e.g., [6]) is identical to the update model of a sensing action. We denote it with $\Sigma^{Ann}(a, \varphi, \psi)$. We note that the key difference in the use of update models for sensing action and announcement action lies in the designated event. An update instance for a sensing action occurrence has two designated events, σ and τ , while only one is specified for an announcement action occurrence, σ . Similar to Proposition 2, we can show that the following holds.

PROPOSITION 3. *Let $\Sigma^{Ann}(a, \varphi, \psi)$ be an update model of an announcement action that announces ψ , (\mathcal{M}, s) be a pointed Kripke such that $(\mathcal{M}, s) \models B_i \neg \psi$ and $(\mathcal{M}, s) \models \varphi$, and i is a full observer of this action. It holds that $(\mathcal{M}', s') \models B_i \psi$ where $\mathcal{M}' = \mathcal{M} \otimes \Sigma^{Ann}(a, \varphi, \psi)$ as in Definition 5 and $s' = (s, \sigma)$.*

Proof. Similar to the proof of Prop. 2 and is omitted for brevity. \square

Propositions 2 and 3 show that full observers of a sensing action or a truthful announcement of ψ will actually correct their belief about ψ after the action occurrence. It is worth to point out that these propositions do not hold if the original definition of the update operation \otimes (Definition 3) is used.

4 A SUFFICIENT CONDITION FOR MAINTAINING $\mathbf{KD45}_n$ PROPERTY

In this section, we identify a condition for the update models that maintains the $\mathbf{KD45}_n$ property of Kripke structures. Specifically, we will introduce the notion of a $\mathbf{KD45}_n$ well-defined update model and prove that $\mathbf{KD45}_n$ well-defined update models maintain the $\mathbf{KD45}_n$ property of Kripke structures. In this paper, we will focus on update models whose preconditions are atomic formulae. Let us start by introducing some extra notations. For an atomic formula φ over \mathcal{F} , $Mod(\varphi)$ is the set of models of φ . For a set of formulae S , let $Mod(S) = \bigcup_{\varphi \in S} Mod(\varphi)$.

DEFINITION 6. *A collection S of atomic formulae over \mathcal{P} is complete if $Mod(S) = 2^{\mathcal{P}}$.*

Intuitively, if S is complete then for any possible world s and interpretation π over \mathcal{P} , there exists some $\varphi \in S$ such that $\pi[s] \models \varphi$.

DEFINITION 7. *Let $\Sigma = \langle \Sigma, R_1, \dots, R_n, pre, sub \rangle$ be an update model. Σ is said to be $\mathbf{KD45}_n$ well-defined with respect to i if for every $\sigma \in \Sigma$:*

- $(\sigma, \sigma) \in R_i$; or
- $S_\sigma = \{pre(\tau) \mid (\sigma, \tau) \in R_i\}$ is complete.

We say that Σ is well-defined if it is well-defined with respect to all agents $1, \dots, n$.

Given an update model Σ and a Kripke structure \mathcal{M} , the completeness of S_σ allows an agent i to maintain the coherence of its beliefs if σ is the true event that occurs (Proposition 4). We will next

prove that $\mathbf{KD45}_n$ well-defined update model maintains the $\mathbf{KD45}_n$ properties of Kripke structures. We need the following lemma.

LEMMA 2. *Let s and u be two worlds of a $\mathbf{KD45}_n$ Kripke structure. Assume that s and u are connected by i . Then, one of the following holds: (a) $(s, u) \in \mathcal{M}[i]$; (b) $(u, s) \in \mathcal{M}[i]$; or (c) there exist some world v such that $(s, v) \in \mathcal{M}[i]$ and $(u, v) \in \mathcal{M}[i]$.*

Proof. By definition of connectedness, there exists a sequence $s = s_1, \dots, s_n = u$ such that for $i = 1, \dots, n-1$, $(s_i, s_{i+1}) \in \mathcal{M}[i]$ or $(s_{i+1}, s_i) \in \mathcal{M}[i]$. We prove the lemma by induction over n .

Base: The Lemma is trivial for $n = 1$ or 2.

Step: Assume that the lemma is correct for $n > 2$. Since s and s_{n-1} is connected by i , by inductive hypothesis, we have the following cases:

- $(s, s_{n-1}) \in \mathcal{M}[i]$. If $(s_{n-1}, u) \in \mathcal{M}[i]$ then $(s, u) \in \mathcal{M}[i]$ because of the transitivity of $\mathcal{M}[i]$ (\mathcal{M} is $\mathbf{KD45}_n$). If $(u, s_{n-1}) \in \mathcal{M}[i]$ then s_{n-1} satisfies the conclusion (c) of the lemma.
- $(s_{n-1}, s) \in \mathcal{M}[i]$. If $(s_{n-1}, u) \in \mathcal{M}[i]$ then $(s, u) \in \mathcal{M}[i]$ because of the Euclideanness of $\mathcal{M}[i]$ (\mathcal{M} is $\mathbf{KD45}_n$). If $(u, s_{n-1}) \in \mathcal{M}[i]$ then $(u, s) \in \mathcal{M}[i]$ by transitivity of $\mathcal{M}[i]$.
- there exists v' such that $(s, v') \in \mathcal{M}[i]$ and $(s_{n-1}, v') \in \mathcal{M}[i]$. If $(s_{n-1}, u) \in \mathcal{M}[i]$ then $(v', u) \in \mathcal{M}[i]$ because of the Euclideanness of $\mathcal{M}[i]$ and thus $(s, u) \in \mathcal{M}[i]$ by transitivity of $\mathcal{M}[i]$. If $(u, s_{n-1}) \in \mathcal{M}[i]$ then $(u, v') \in \mathcal{M}[i]$ by transitivity of $\mathcal{M}[i]$ and thus v' satisfies the conclusion (c) of the lemma.

THEOREM 1. *Let \mathcal{M} be a $\mathbf{KD45}_n$ Kripke structure and Σ be a $\mathbf{KD45}_n$ well-defined update model. It holds that $\mathcal{M}' = \mathcal{M} \otimes \Sigma$ also satisfies the $\mathbf{KD45}_n$ property.*

Proof. We will prove that \mathcal{M}' satisfies the seriality, transitivity, and Euclidean property. This is done by three propositions 4–6. \square

PROPOSITION 4. *\mathcal{M}' is serial.*

Proof. Consider $i \in \mathcal{AG}$ and $(s, \tau) \in \mathcal{M}'[S]$, we have that $s \in \mathcal{M}[S]$, $\tau \in \Sigma$, and $(\mathcal{M}, s) \models pre(\tau)$. Let $S_\tau = \{\tau' \mid (\tau, \tau') \in R_i\}$. Since Σ satisfies $\mathbf{KD45}_n$, $S_\tau \neq \emptyset$. The proof is trivial if $(\tau, \tau) \in R_i$. Assume now that $(\tau, \tau) \notin R_i$. Because \mathcal{M} is serial, there exists some u such that $(s, u) \in \mathcal{M}[i]$. Since Σ is well-defined, there exists some $\tau' \in S_\tau$ such that $(\mathcal{M}, u) \models pre(\tau')$. This implies that $((s, \tau), (u, \tau')) \in \mathcal{M}'[i]$. Thus, we can conclude that $\forall s' \in \mathcal{M}'[S]$, $\exists u' \in \mathcal{M}'[S]$ such that $(s', u') \in \mathcal{M}'[i]$. This holds for arbitrary agent i . Thus, \mathcal{M}' is serial. \square

PROPOSITION 5. *\mathcal{M}' is transitive.*

Proof. Let $i \in \mathcal{AG}$ and $s_0, s_1, s_2 \in \mathcal{M}'[S]$ such that $(s_0, s_1) \in \mathcal{M}'[i]$ and $(s_1, s_2) \in \mathcal{M}'[i]$. We want to show that $(s_0, s_2) \in \mathcal{M}'[i]$. By definition, there exist $u_0, u_1, u_2 \in \mathcal{M}[S]$ and $\tau_0, \tau_1, \tau_2 \in \Sigma$ such that $s_0 = (u_0, \tau_0)$, $s_1 = (u_1, \tau_1)$, and $s_2 = (u_2, \tau_2)$. Furthermore, $(\tau_0, \tau_1) \in R_i$ and $(\tau_1, \tau_2) \in R_i$. This implies that $(\tau_0, \tau_2) \in R_i$. We consider the following cases:

- $(s_0, s_1) \in \mathcal{M}'[i]$ because of Item (ii.(a)) of Definition 5. It is easy to see that if $(s_1, s_2) \in \mathcal{M}'[i]$ because of Item (ii.(a)) or (ii.(b)) of Definition 5, then $(s_0, s_2) \in \mathcal{M}'[i]$ by Item (ii.(a)) because of the transitivity of \mathcal{M} and Σ or trivially. Consider now the case that $(s_1, s_2) \in \mathcal{M}'[i]$ because of Item (ii.(c)) of Definition 5, respectively. The first fact implies that $(u_0, u_1) \in \mathcal{M}[i]$. Connectedness between u_1 and u_2 and well-definiteness of

Σ , together with the second fact and Lemma 2, imply that there exists some $v \in \mathcal{M}[S]$ such that $(u_1, v) \in \mathcal{M}[i]$ and $(u_2, v) \in \mathcal{M}[i]$. Transitivity of \mathcal{M} and $(u_0, u_1) \in \mathcal{M}[i]$ and $(u_1, v) \in \mathcal{M}[i]$ imply $(u_0, v) \in \mathcal{M}[i]$. Euclidean property of \mathcal{M} and $(u_0, v) \in \mathcal{M}[i]$ and $(u_0, u_1) \in \mathcal{M}[i]$ imply $(v, u_1) \in \mathcal{M}[i]$. Transitivity of \mathcal{M} and $(u_2, v) \in \mathcal{M}[i]$ and $(v, u_1) \in \mathcal{M}[i]$ imply $(u_2, u_1) \in \mathcal{M}[i]$. This contradicts the assumption that (s_1, s_2) belongs to $\mathcal{M}'[i]$ because of case (c) (Item **ii**) of Definition 5 since this requires that $(u_2, u_1) \notin \mathcal{M}[i]$. In other words, this case cannot happen.

- $(s_0, s_1) \in \mathcal{M}'[i]$ because of Item **ii.(a)** of Definition 5. It means that $s_0 = s_1$ and therefore, $(s_0, s_2) \in \mathcal{M}'[i]$ because $(s_1, s_2) \in \mathcal{M}'[i]$.
- $(s_0, s_1) \in \mathcal{M}'[i]$ because of Item **ii.(c)** of Definition 5.
 - If $(s_1, s_2) \in \mathcal{M}'[i]$ also because of case Item **ii.(c)** of Definition 5. We have that $C_i(u_0, \tau_0)$, $C_i(u_1, \tau_1)$, and $C_i(u_2, \tau_2)$ are true. Furthermore, (τ_0, τ_0) , (τ_1, τ_1) , and (τ_2, τ_2) belong to R_i . $C_i(u_0, \tau_0)$ is true and $(\mathcal{M}, u_2) \models \text{pre}(\tau_2)$ implies that $(u_0, u_2) \notin \mathcal{M}[i]$. Similarly, $(u_2, u_0) \notin \mathcal{M}[i]$. Transitivity of connectedness allows us to conclude that u_0 and u_2 are connected by i . This allows us to conclude that $(s_0, s_2) \in \mathcal{M}'[i]$ because of Item **ii.(c)** of Definition 5.
 - Assume now that $(s_1, s_2) \in \mathcal{M}'[i]$ because of Item **ii.(a)** of Definition 5. It means that $(u_1, u_2) \in \mathcal{M}[i]$. Since (s_0, s_1) belongs to $\mathcal{M}'[i]$ because of Item **ii.(c)** of Definition 5, we have that u_0 and u_1 are connected by i but $(u_0, u_1) \notin \mathcal{M}[i]$ and $(u_1, u_0) \notin \mathcal{M}[i]$. Lemma 2 and the well-definiteness of Σ imply that there exists some v such that $(u_0, v) \in \mathcal{M}[i]$ and $(u_1, v) \in \mathcal{M}[i]$. Since $(u_1, u_2) \in \mathcal{M}[i]$, by Euclidean property of \mathcal{M} , $(v, u_2) \in \mathcal{M}[i]$ and hence, $(u_0, u_2) \in \mathcal{M}[i]$. This implies that $(s_0, s_2) \in \mathcal{M}'[i]$ because $(\tau_0, \tau_2) \in R_i$.

The above shows that transitivity holds in \mathcal{M}' . \square

PROPOSITION 6. \mathcal{M}' is Euclidean.

Proof. Let $i \in \mathcal{AG}$ and $s_0, s_1, s_2 \in \mathcal{M}'[S]$ such that $(s_0, s_1), (s_0, s_2) \in \mathcal{M}'[i]$. We want to show that $(s_1, s_2) \in \mathcal{M}'[i]$. Again, by definition of an update, there exist $u_0, u_1, u_2 \in \mathcal{M}[S]$ and $\tau_0, \tau_1, \tau_2 \in \Sigma$ such that $s_i = (u_i, \tau_i)$ and $(\mathcal{M}, s_i) \models \text{pre}(\tau_i)$ for $i = 0, 1, 2$. Furthermore, $(\tau_0, \tau_1) \in R_i$ and $(\tau_0, \tau_2) \in R_i$. This implies that $(\tau_0, \tau_2) \in R_i$ because Σ satisfies **KD45_n**. We consider the following cases:

- Both (s_0, s_1) and (s_0, s_2) belong to $\mathcal{M}'[i]$ because of Item **ii.(a)** of Definition 5 then $(s_1, s_2) \in \mathcal{M}'[i]$ because of the Euclidean property of \mathcal{M} and Σ and Definition 5, Item **ii.(a)**.
- Either $s_0 = s_1$ or $s_0 = s_2$ then, trivially, $(s_1, s_2) \in \mathcal{M}'[i]$.
- (s_0, s_1) belongs to $\mathcal{M}'[i]$ because of Item **ii.(c)** of Definition 5. Then, Lemma 1 shows that (s_0, s_2) also belongs to $\mathcal{M}'[i]$ because of Item **ii.(c)** of Definition 5. Similar to the argument made in the first Subcase of the third case in the proof of Proposition 5, we can conclude that $(s_1, s_2) \in \mathcal{M}'[i]$ because of Item **ii.(c)** of Definition 5.

Similar argument will allow us to conclude $(s_2, s_1) \in \mathcal{M}'[i]$.

The above shows that Euclidean property holds in \mathcal{M}' . \square

We will conclude this section with a discussion about the well-definiteness of update models that have been proposed recently in the literature. We start with the models discussed in [6].

- **Ontic Action:** An ontic action is executed to modify certain properties of the world. For example, when an agent opens a box, then the box will change its property from *closed* to *opened*. For simplicity

of the presentation, let us consider an ontic action a with a set C of effects of the form $p \rightarrow \varphi$, which states that p will be true if φ is true before the execution of a . Furthermore, assume that ψ is the precondition of a . The update model of an ontic action with the set of agent \mathcal{AG} is then defined by $\Sigma^{\text{Ontic}}(a, \psi, C) = \langle \{\sigma, \epsilon\}, R_1, \dots, R_n, \text{pre}, \text{sub} \rangle$ where $R_i = \{(\sigma, \sigma), (\epsilon, \epsilon)\}$ for $i \in F$ (F is the set of full observers), $R_i = \{(\sigma, \epsilon), (\epsilon, \epsilon)\}$ for $i \in O$ ($O = \mathcal{AG} \setminus F$), $\text{pre}(\sigma) = \psi$, $\text{pre}(\epsilon) = \top$, $\text{sub}(\tau) = \emptyset$, and $\text{sub}(\sigma) = C$ (Figure 7).

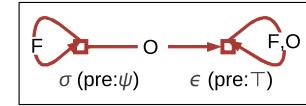


Figure 7: $\Sigma^{\text{Ontic}}(a, \psi, C)$: update model for ontic action

Consider an agent $i \in F$. It is easy to see that $\Sigma^{\text{Ontic}}(a, \psi, C)$ is well-defined with respect to i since $(\sigma, \sigma) \in R_i$ and $(\epsilon, \epsilon) \in R_i$. On the other hand, for $i \in O$, $\Sigma^{\text{Ontic}}(a, \psi, C)$ is well-defined with respect to i because for σ , $S_\sigma = \{\top\}$ is complete and $(\tau, \tau) \in R_i$. It is easy to verify that $\Sigma^{\text{Ontic}}(a, \psi, C)$ is **KD45_n** as well.

- **Sensing/Announcement Action:** Given $\Sigma^{\text{Sensing}}(a, \varphi, \psi)$ or $\Sigma^{\text{Ann}}(a, \varphi, \psi)$ (Section 3.1 or 3.2), we have that these are well-defined with respect to $i \in F \cup P$ since there are loops labeled i in all events; they are well-defined with respect to $i \in O$ at σ and τ since $S_\sigma = S_\tau = \{\top\}$ and at ϵ because $(\epsilon, \epsilon) \in R_i$. Again, we can check that these two update models are also **KD45_n** well-defined.

Another work that also employs update model in formalizing actions in multi-agent domains is [14], in the development of the language DER (Dynamic Epistemic Representation). It is easy to observe that the update models defined in this work also satisfy Definition 7. More specifically, for each update model Σ defined by Definition 4 in [14], we have that for every agent i and every event e in Σ , $(e, e) \in R_i$. As such, every update model Σ defined in [14] is **KD45_n** well-defined.

5 DISCUSSION

The main purpose of the present paper is to address the inability to correct false beliefs of agents. Propositions 2–3 show that the new update of a Kripke structure using an update model (Definition 5) achieves this goal. In the past, the action language $m\mathcal{A}^*$ includes a method for dealing with this problem [6]. Following their approach, the Kripke structure is revised *before* the update is applied. The revision aims at correcting the false beliefs of full and partial observers, thus preventing the update (via Definition 3) to create agents with incoherent beliefs. This approach does indeed help agents to correct their false beliefs. However, it sometimes allows agents to gain additional information that appears unreasonable. For example, it is easy to verify that the approach proposed in [6], when applied on \mathcal{M}_3 and $\Sigma_{\text{ann}(h)}$ shown in Figure 4, results in \mathcal{R}_3 , thus allowing C to learn h . This is because the approach in [6] does not deal with situations which require **Change #2**, as in Definition 5. Besides, we believe that there are situations in which the approach in [6] is too strong. An example is displayed in Figure 8.

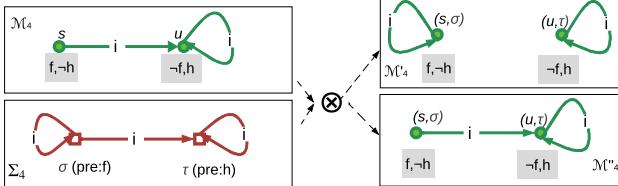


Figure 8: Difference between $m\mathcal{A}^*$ and our approach

In this example, we have a Kripke structure M_4 with two worlds s and u , $M_4[\pi](s) = \{f, \neg h\}$ and $M_4[\pi](u) = \{\neg f, h\}$. The update model Σ_4 has two events σ and τ with $pre(\sigma) = f$ and $pre(\tau) = h$. The accessibility relation of an agent i is given in the figure. In computing $M_4 \otimes \Sigma_4$, following the definition in [6], the link (s, u) is removed and thus, $M_4 \otimes \Sigma_4$ is the Kripke structure with two worlds (s, σ) and (u, τ) and the loops labeled i around these worlds (Figure 8, top right). It is easy to see that Definition 5 yields a different result for $M_4 \otimes \Sigma_4$, with the link labeled i connecting (s, σ) and (u, τ) and without the loop around (s, σ) since $C_i(s, \sigma)$ is false (Figure 8, bottom right).

The present paper is directly related to works that study properties of update models. In the past, [1, 11, 15] investigated the maintenance of the $KD45_n$ property using update models. The work [11] considers only ontic and sensing actions and assumes that actions are always executable. This is similar to the condition that $pre(\sigma) = \top$ for every event in the update model, which implies that the update models considered in [11] are well-defined. It is also easy to verify that the condition on update models proposed in [15], called *primitive*³, for maintaining the $KD45_n$ property is subsumed by the well-definiteness condition in Definition 7. Update models for ontic actions with oblivious agents, for example, are not primitive per [15]. In other words, the proposed sufficient condition for update models to maintain the $KD45_n$ property of Kripke structures in this paper is more general than those developed earlier in [11, 15]. In [1], the author identifies a semantical condition on the initial Kripke structure that guarantees that the result of its update by a serial update model is serial. In contrast, our condition is applied on the update model and only requires that the original Kripke structure is serial. Observe that [1] does not investigate transitivity and Euclideanness of the update result.

Propositions 2-3 showed that the majority of update models considered for formalizing actions in multi-agent domains (e.g., in [6, 14]) satisfy the well-definiteness condition. This implies that Definition 5 could be employed in the development of epistemic planners (e.g., [12, 13]) that work with one modality and plan for both knowledge and belief goals. To the best of our knowledge, there exists no such planning system with this capability yet. This gives rise to the question of how difficult it is to check for the well-definiteness of an update model and what the overhead will be for using the new update. It is easy to see that, in general, checking whether an update model is well-defined is a co-NP hard problem, since checking for the completeness of S_σ is equivalent to checking for the unsatisfiability of $\neg(\bigvee \varphi \in S_\sigma)$. Fortunately, this

³This requires that for every agent i and σ such that $(\sigma, \tau) \in R_i$, either $(\sigma, \sigma) \in R_i$ or $pre(\tau) = \top$ and $sub(\sigma) = sub(\tau) = \emptyset$.

is only theoretical, and checking this condition for update models in the literature is not difficult, as the majority of events are associated with \top (as their precondition) or satisfy the first condition (they have loops around them). With regard to overhead, additional computational tasks might be required, but they are fairly simple. First, for each agent i and $(\sigma, \sigma) \in B_i$, we need to check the condition in **Change #1**. However, this could be implemented as part of checking whether (s, σ) belongs to the set of worlds of the result of the update. Therefore, the overhead for this task is negligible. Second, for every pair of (s, σ) and (u, τ) such that $((s, \sigma), (s, \sigma))$ and $((u, \tau), (u, \tau))$ is added to B_i then the condition in **Change #2** needs to be checked. This will require a check for connectedness of s and u which is again, negligible, as it is linear in the size of the Kripke structure.

We note that in [8], a new type of update model called *edge-conditioned* update model has been proposed to deal with false beliefs of agents. In edge-conditioned update models, an accessibility relation between two events for an agent i is associated with a formula ϕ that represents the condition under which the link should exist. It is also noted that edge-conditioned update models can be converted to equivalent standard update models. Therefore, we believe that the well-definiteness of update models can also be formalized for this new type of update models as well.

Finally, it is worth mentioning that the language of serial Public Announcement Logic (sPAL) in [4] also maintains the $KD45_n$ of Kripke models after the execution of a truthful public announcement. The logic, however, does not employ update models and requires that no agent has false belief about the announced formula before the action is executed. This is different from what we proposed in this paper. In fact, Prop. 3 shows that the $KD45_n$ property can be maintained in arbitrary models if Def. 5 is employed.

6 CONCLUSION

We proposed a novel definition of updates of Kripke structures using update models. The new definition differs from the original one in that it introduces additional edges in the accessibility relations of the resulting Kripke structures. We showed that update models for actions proposed in the literature, when using the new definition, enable full observers of sensing actions or truthful announcements to correct existing false beliefs. This addresses a critical issue caused by the original definition, i.e., the issue of agents having incoherent beliefs. We also introduced the notion of a $KD45_n$ well-defined update model and proved that $KD45_n$ well-defined update models maintain the $KD45_n$ property of Kripke structures. This result is significant in that it allows us to reason about knowledge and beliefs in the $KD45_n$ logic, which requires only one modality. We believe that this result can be used by epistemic planners to generate plans with both knowledge and belief goals. To the best of our knowledge, the well-definiteness condition subsumes conditions proposed in earlier works, such as [11, 15], and is satisfied by the majority of update models that have recently been developed for the study of actions in multi-agent domains.

ACKNOWLEDGMENTS

The authors wish to acknowledge the partial support by NSF awards #1914635, #1812628, and by a contract with Wallaroo.

REFERENCES

[1] Guillaume Aucher. 2008. Consistency Preservation and Crazy Formulas in BMS. In *Logics in Artificial Intelligence, 11th European Conference, JELIA 2008, Dresden, Germany, September 28 - October 1, 2008. Proceedings (Lecture Notes in Computer Science, Vol. 5293)*, Steffen Hölldobler, Carsten Lutz, and Heinrich Wansing (Eds.). Springer, 21–33. https://doi.org/10.1007/978-3-540-87803-2_4

[2] A. Baltag and L. Moss. 2004. Logics for epistemic programs. *Synthese* (2004).

[3] A. Baltag, L. Moss, and S. Solecki. 1998. The logic of public announcements, common knowledge, and private suspicions. In *7th TARK*, 43–56.

[4] Alexandru Baltag and Bryan Renne. Winter 2016. Dynamic Epistemic Logic. In *The Stanford Encyclopedia of Philosophy (Winter 2016 Edition)*, Edward N. Zalta (Ed.).

[5] Alexandru Baltag and Sonja Smets. 2016. A Qualitative Theory of Dynamic Interactive Belief Revision. In *Readings in Formal Epistemology*, 813–858.

[6] Chitta Baral, Gregory Gelfond, Enrico Pontelli, and Tran Cao Son. 2022. An action language for multi-agent domains. *Artif. Intell.* 302 (2022), 103601. <https://doi.org/10.1016/j.artint.2021.103601>

[7] Adam Bjorndahl and Will Nalls. 2021. Endogenizing Epistemic Actions. *Stud Logica* 109, 5 (2021), 1049–1091. <https://doi.org/10.1007/S11225-020-09937-8>

[8] Thomas Bolander. 2018. *Seeing Is Believing: Formalising False-Belief Tasks in Dynamic Epistemic Logic*. Springer International Publishing, Cham, 207–236. https://doi.org/10.1007/978-3-319-62864-6_8

[9] David Buckingham, Daniel Kasenberg, and Matthias Scheutz. 2020. Simultaneous Representation of Knowledge and Belief for Epistemic Planning with Belief Revision. (9 2020), 172–181. <https://doi.org/10.24963/kr.2020/18>

[10] J.Y. Halpern, D. Samet, and E. Segev. 2009. Defining Knowledge in Terms of Belief: The Modal Logic Perspective. *Reviews of Symbolic Logic* 2, 3 (2009), 469–486.

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

[12] Tiep Le, Francesco Fabiano, Tran Cao Son, and Enrico Pontelli. 2018. EFP and PG-EFP: Epistemic Forward Search Planners in Multi-Agent Domains. In *International Conference on Automated Planning and Scheduling (ICAPS)*. AAAI Press.

[13] Loc Pham, Tran Cao Son, and Enrico Pontelli. 2023. Planning in Multi-Agent Domains with Untruthful Announcements. In *Proceedings of the Thirty-Third International Conference on Automated Planning and Scheduling, July 8–13, 2023, Prague, Czech Republic*, Sven Koenig, Roni Stern, and Mauro Vallati (Eds.). AAAI Press, 334–342. <https://doi.org/10.1609/icaps.v33i1.27211>

[14] David Rajaratnam and Michael Thielscher. 2021. Representing and Reasoning with Event Models for Epistemic Planning. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR)*, M. Bienvenu, E. Erdem, and G. Lakemeyer (Eds.). Hanoi, 519–528. <https://doi.org/10.24963/kr.2021/49>

[15] Tran Cao Son, Enrico Pontelli, Chitta Baral, and Gregory Gelfond. 2015. Exploring the KD45n Property of a Kripke Model after the Execution of an Action Sequence. In *AAAI*.

[16] Johan van Benthem, Jan van Eijck, and Barteld P. Kooi. 2006. Logics of communication and change. *Inf. Comput.* 204, 11 (2006), 1620–1662.

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