Simultaneous Representation of Knowledge and Belief for Epistemic Planning with Belief Revision

David Buckingham, **Daniel Kasenberg**, **Matthias Scheutz**
Tufts University
david.buckingham@tufts.edu, daniel.kasenberg@tufts.edu, matthias.scheutz@tufts.edu

Abstract

We propose a novel approach to the problem of false belief revision in epistemic planning. Our state representations are pointed Kripke models with two binary relations over possible worlds: one representing agents' necessarily true knowledge, and one representing agents' possibly false beliefs. State transition functions maintain $S5_n$ properties in the knowledge relation and $KD45_n$ properties in the belief relation. When new information contradicts an agent's beliefs, belief revision draws new possible worlds from the agent's knowledge relation. Our method also improves upon prior work by accommodating false announcements. We develop our system as an extension to the m$\mathcal{A}^*$ action language, presenting transition functions for ontic, sensing, and announcement actions.

1 Introduction

Epistemic planning is an approach to decision making for multi-agent task scenarios involving not only objective facts about the environment, but also agents' knowledge and beliefs, including iterated knowledge and belief. A popular and powerful EP paradigm is Dynamic Epistemic Logic (DEL) (Gerbrandy 1997; Bolander 2017), where states are represented as pointed Kripke models and actions are represented as structures called update-models. However, update models encode a specific execution instance of an action, including observability for each agent. The planning language m$\mathcal{A}^*$ (Baral et al. 2019) overcomes this limitation with compact action definitions that are dynamically instantiated to state-specific update models. Thus, it is particularly suited to use in epistemic planners. However, a limitation of m$\mathcal{A}^*$ is that it handles false belief revision (when an agent corrects a false belief) by granting the agent perfect correct beliefs, even about things not informed by the revision event. In contrast, our belief revision approach maintains an agent's uncertainty about facts that are not informed by a revision event.

Following Kraus and Lehmann (1988) and others, we employ Kripke structures that have two binary relations on worlds for each agent: one relation to represent knowledge and one relation to represent belief. Thus we maintain two tiers of information, a lower (belief) tier, and a strictly more general higher (knowledge) tier, for each agent. We assume that agents have access to a source of information about the environment that can reliably inform knowledge (e.g. non-noisy sensors), as well as less reliable sources that can only inform belief (e.g. possibly-false announcements). An agent accepts the most recent information received into the appropriate tier, except that lower tier information does not override higher tier information, e.g. an agent rejects an announcement that she knows to be untrue.

Belief revision draws on an agent's knowledge, adding to the set of worlds which the agent *believes* possible all worlds not previously *known* by the agent to be impossible. We refer to this process as belief *reset*. These representations can be used by multi-agent epistemic planning systems to explore problems involving false announcements, false beliefs, and recovery from false belief.

2 Motivation and Related Work

The 2017 Dagstuhl Seminar on Epistemic Planning (Baral et al. 2017) identified the need to explore the relationships between knowledge and belief, specifically because DEL does not allow agents to recover from false belief. The predecessor of m$\mathcal{A}^*$ (Baral et al. 2015) could not handle recovery from false belief. When an agent observes some fact φ to be true, the agent's beliefs are updated by removing from the set of worlds which the agent considers possible all worlds where φ is false. If the agent had wrongly believed $\neg\varphi$, then no possible worlds remain, and the agent is left "ignorant", believing all propositions to be trivially true. To address this, m$\mathcal{A}^*$ adds to the accessibility relation a reflexive connection on the designated world and removes all other connections from the designated world when an agent's false beliefs are corrected by a sensing or announcement action. At least one planner (Le et al. 2018) implements this method. The limitations of that approach with respect to this work are that it removes *all* of an agent's uncertainty, even about things unrelated to the belief being revised, it does not allow agents to make false announcements, and it only revises beliefs that are corrected directly by sensing or announcement actions, not indirectly as when an agent observes an action whose preconditions violate the agent's beliefs.

We refer the reader to van Benthem and Smets (2015) for a discussion of the development of logical systems of belief change. Baral et al. (2017) describes several belief revision approaches in EP. Most directly relevant to the present work are doxastic epistemic action models (Ditmarsch 2005), probabilistic action models (Baltag and Smets 2008a), syntactic revision (Huang et al. 2017), Proper Epistemic Knowledge

Base update (Miller and Muise 2016), and action plausibility models (Baltag and Smets 2008b). Some of these are more powerful and flexible than our approach in terms of the kinds of belief changes that can be represented. In particular, Andersen, Bolander, and Holm Jensen (2013) develop a framework and algorithms for single-agent Epistemic Planning with plausibility models. All of these approaches, however, lack the previously-mentioned qualities that make m$\mathcal{A}^*$ particularly useful. We therefore focus on extending m$\mathcal{A}^*$ and leave to future work the integration of features from these richer techniques.

3 Epistemic States

Given a set of Boolean fluents $\mathcal{F}$, a *fluent formula* is a formula built using $f \in \mathcal{F}$ and the propositional operators $\vee, \wedge, \rightarrow, \neg$. A *fluent atom* is a fluent formula containing a single element $f \in \mathcal{F}$. A *fluent literal* is either f or $\neg f$ where f is a fluent atom.

Consider a set of n agents $\mathcal{A} = \{1, \cdots, n\}$. We define a bimodal Kripke model as a Kripke model with two accessibility relations for each agent, a 4-tuple $\mathcal{M} = (W, V, \mathcal{B}_1, \ldots, \mathcal{B}_n, \mathcal{K}_1, \ldots, \mathcal{K}_n)$, where

- W is a non-empty finite set of worlds,
- $V : W \rightarrow 2^{\mathcal{F}}$ is a valuation function assigning a set of fluents to each world,
- $\mathcal{B}_i \subseteq W \times W$ is an accessibility relation representing agent i's belief, and
- $\mathcal{K}_i \subseteq W \times W$ is an accessibility relation representing agent i's knowledge.

Following (Kraus and Lehmann 1988), we will assume that $\mathcal{B}_i$ satisfies $KD45_n$, that $\mathcal{K}_i$ satisfies $S5_n$, and that $\mathcal{B}_i$ and $\mathcal{K}_i$ satisfy the $KB1$ ($\mathbf{K}_i\phi \rightarrow \mathbf{B}_i\phi$) and $KB2$ ($\mathbf{B}_i\phi \rightarrow \mathbf{K}_i\mathbf{B}_i\phi$) (Hintikka 1962) properties. In Section 7 we show that our state transitions preserve these properties. We will say that $\mathcal{B}_i$ (resp. $\mathcal{K}_i$) contains an *edge* from world u to world v if $(u, v) \in \mathcal{B}_i$ (resp. $\mathcal{K}_i$), or between worlds u and v if (u, v) or $(v, u) \in \mathcal{B}_i$ (resp. $\mathcal{K}_i$), or from a set of worlds X to a set of worlds Y if $\mathcal{B}_i$ (resp. $\mathcal{K}_i$) contains any $(u \in X, v \in Y)$, or between sets of worlds, similarly.

An epistemic-doxastic state, or *ed-state*, is a pair $(\mathcal{M}, d)$, where $\mathcal{M}$ is a bimodal Kripke model, and $d \in W$ is the *designated world*, representing the actual state of the environment. Given a set of agents $\mathcal{A}$, for $i \in \mathcal{A}$, a *belief formula* is either a fluent formula or a formula of one of the following forms: $\phi \vee \psi$, $\phi \wedge \psi$, $\neg\phi$, $\mathbf{B}_i\phi$, or $\mathbf{K}_i\phi$, where ϕ and ψ are belief formulae. Intuitively, $\mathbf{B}_i\varphi$ means that agent i believes φ and $\mathbf{K}_i\varphi$ means that agent i knows φ.

$$
\begin{array}{ll}
(\mathcal{M}, d) \models \varphi & \text{if } \varphi \text{ is a fluent formula and} \\
& \varphi \models \wedge_{f \in V(d)} \\
(\mathcal{M}, d) \models \neg\varphi & \text{iff } (\mathcal{M}, d) \not\models \varphi \\
(\mathcal{M}, d) \models \varphi \wedge \psi & \text{iff } (\mathcal{M}, d) \models \varphi \text{ and } (\mathcal{M}, d) \models \psi \\
(\mathcal{M}, d) \models \varphi \vee \psi & \text{iff } (\mathcal{M}, d) \models \varphi \text{ or } (\mathcal{M}, d) \models \psi \\
(\mathcal{M}, d) \models \mathbf{B}_i\varphi & \text{iff } \forall\, v \mid (d, v) \in \mathcal{B}_i : (\mathcal{M}, v) \models \varphi \\
(\mathcal{M}, d) \models \mathbf{K}_i\varphi & \text{iff } \forall\, v \mid (d, v) \in \mathcal{K}_i : (\mathcal{M}, v) \models \varphi
\end{array}
$$

4 State Transitions

Following Baral et al. (2015), we define state transition functions for three type of actions: ontic actions that cause changes in the real world, sensing actions, representing observations that change agents' knowledge, and announcement actions, representing communication between agents. Each action type is a 4-tuple, the first three elements of which are $a.\texttt{pre}$, $a.\texttt{obs}$, and $a.\texttt{aware}$. The fourth element of the tuple will be defined for each action type. Belief formula $a.\texttt{pre}$ is the action's precondition, $a.\texttt{obs}$ is a set of agent-fluent-formula pairs defining conditional full observability, and $a.\texttt{aware}$ is a set of agent-fluent-formula pairs defining conditional partial observability.

The execution of an action a in ed-state $(\mathcal{M}, d)$ assigns agents to one of three sets, F, P, and O for agents who are fully-aware of, partially-aware of, and oblivious to the action's occurrence:

$$
\begin{aligned}
F &= \{i \in \mathcal{A} \mid (i, \varphi) \in a.\texttt{obs}, (\mathcal{M}, d) \models \varphi\} \\
P &= \{i \in \mathcal{A} \mid i \notin F, (i, \varphi) \in a.\texttt{aware}, (\mathcal{M}, d) \models \varphi\} \\
O &= \{i \in \mathcal{A} \mid i \notin F \cup P\}
\end{aligned}
$$

Fully- and partially-observant agents know that the action has occurred and update their knowledge and belief to reflect that the action's preconditions hold. Furthermore, in the case of sensing and announcement actions, fully-observant agents become aware of the sensed or announced information.

Oblivious agents do not know or believe that the action has occurred, but do know that it *might* have. Thus, an oblivious agent's belief relation is not altered. However, an oblivious agent's knowledge relation *is* updated to express that the agent comes to know that the action might have occurred. This is in order to maintain the $S5_n$ property that agents do not have false knowledge, which is necessary for belief reset. On the contrary, an agent will not come to know that an action of which she is oblivious, and that did not in fact occur, might have occurred.

When an action a is applied in ed-state $(\mathcal{M}, d)$ where $\mathcal{M} = (W, V, \mathcal{B}_1, \ldots, \mathcal{B}_n, \mathcal{K}_1, \ldots, \mathcal{K}_n)$, the result is a new ed-state, $(\mathcal{M}', d')$, where $\mathcal{M}' = (W', V', \mathcal{B}'_1, \ldots, \mathcal{B}'_n, \mathcal{K}'_1, \ldots, \mathcal{B}'_n)$. The worlds W' of the new ed-state include the worlds of the previous ed-state W corresponding to the beliefs of oblivious agents, and a set of new worlds W^+ corresponding to the beliefs of full and partial observers. The new worlds W^+ only include worlds that satisfy $a.\texttt{pre}$ because full and partial observers of a learn that $a.\texttt{pre}$ must be true. Let $\mathrm{Pre}(\mathcal{M}, a) = \{w \in W \mid (\mathcal{M}, w) \models a.\texttt{pre}\}$. We construct a new world $w^+ \in W$ from each world $w \in \mathrm{Pre}(\mathcal{M}, a)$. Instead of defining additional worlds by writing "for any $u \in W^+$, there is some world w such that $u = w^+$ and $w \in \mathrm{Pre}(\mathcal{M}, a)$," when discussing worlds in W^+, we will immediately refer to them as, e.g., w^+ where $w \in \mathrm{Pre}(\mathcal{M}, a)$.

Let $W^+ = \{w^+ \mid w \in \mathrm{Pre}(\mathcal{M}, a)\}$. For all action types,

$$
\begin{aligned}
W' &= W \cup W^+ \\
V'(w) &= \begin{cases} V^+(w) & \text{if } w \in W^+ \\ V(w) & \text{otherwise} \end{cases} \\
d' &= d^+
\end{aligned}
$$

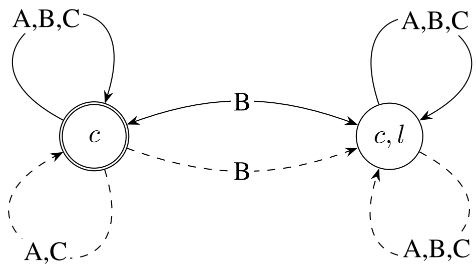

Figure 1: The ed-state before the example ontic action open_door is applied.

We will define V^+, ${\mathcal{B}_i}'$, and ${\mathcal{K}_i}'$ separately for each action type.

4.1 Ontic Actions

An ontic action is a 4-tuple, $(a.\texttt{pre}, a.\texttt{obs}, a.\texttt{aware}, a.\texttt{eff})$, where $a.\texttt{eff}$ is a set of fluent literals. We make the closed-world assumption that fluents not explicitly altered by the action remain unchanged. For an ontic action executed in ed-state $(\mathcal{M}, d)$, the valuation of the new worlds W^+ reflects the action's effects on the environment:

$$V^+(w^+) = (V(w) \setminus \{f \mid \neg f \in a.\texttt{eff}\}) \cup \{f \mid f \in a.\texttt{eff}\}$$

We define temporary relations $\mathcal{R}_i$ representing agents' beliefs after they have possibly been reset (updated with information from $\mathcal{K}_i$) due to observing an action whose precondition contradicts the agents' belief. For $i \in \mathcal{A}$:

$$\mathcal{R}_i = \{(u, v) \mid (u, v) \in \mathcal{B}_i \vee ((u, v) \in \mathcal{K}_i \wedge (\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre})\}$$

We update $\mathcal{B}_i$ to reflect changes to agents' beliefs after any possible belief reset and the influence of the action's preconditions and effects. The beliefs of fully- and partially-aware agents are determined according to $\mathcal{R}_i$. For $i \in F \cup P$:

$$\mathcal{B}'_i = \mathcal{B}_i \cup \{(u^+, v^+) \mid u, v \in \text{Pre}(\mathcal{M}, a), (u, v) \in \mathcal{R}_i\}$$

The beliefs of oblivious agents are not reset. For $i \in O$:

$$\mathcal{B}'_i = \mathcal{B}_i \cup \{(u^+, v) \mid u \in \text{Pre}(\mathcal{M}, a), (u, v) \in \mathcal{B}_i\}$$

We update $\mathcal{K}_i$ to reflect changes to agents' knowledge resulting from the action's preconditions and effects. Fully- and partially-aware agents know that the action occurred. For $i \in F \cup P$:

$$\mathcal{K}'_i = \mathcal{K}_i \cup \{(u^+, v^+) \mid u, v \in \text{Pre}(\mathcal{M}, a), (u, v) \in \mathcal{K}_i\}$$

Oblivious agents do not know whether or not the action occurred. For $i \in O$:

$$\mathcal{K}'_i = \mathcal{K}_i \cup \{(u^+, v^+) \mid u, v \in \text{Pre}(\mathcal{M}, a), (u, v) \in \mathcal{K}_i\} \cup \{(u^+, v), (v, u^+) \mid u \in \text{Pre}(\mathcal{M}, a), (u, v) \in \mathcal{K}_i\}$$

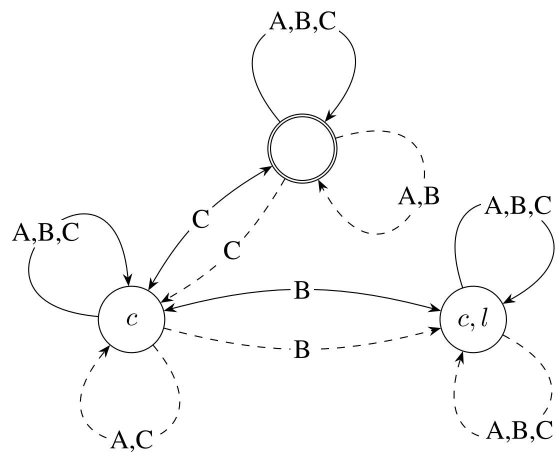

Figure 2: The ed-state after agent A performs the example ontic action open_door.

Example Ontic Action This example demonstrates three ways that an ontic action can affect an agent's beliefs: fully- and partially-observant agents are aware of changes to the world caused by the action, and also learn that the action's preconditions must be true, while the beliefs of ignorant agents may become false due to world changes caused by the action. Consider the ed-state illustrated in Figure 1. Graph nodes represent worlds. The designated world has a double border. Fluents true in a world are listed within its node. Solid arrows represent $\mathcal{K}_i$ and dashed arrows represent $\mathcal{B}_i$.

Three agents, A, B, and C, are in a room with a door. The door is closed (represented by the fluent c) but not locked ($\neg l$). Everyone believes (and knows) that the door is closed. Agents A and B are watching the door, but C is not. For each agent i let the fluent $watch(i)$, express that i is watching the door. The fluents $watch(A)$ and $watch(B)$ are true in all worlds, but are not shown in the figure due to space constraints. Agents A and C know that the door is not locked but agent B does not, and believes that it *is* locked (l). Agent A performs the ontic action *open_door*:

$$open\_door.\texttt{pre} = c \wedge \neg l$$
$$open\_door.\texttt{obs} = \{(A, \top), (B, watch(B)), (C, watch(C))\}$$
$$open\_door.\texttt{aware} = \emptyset$$
$$open\_door.\texttt{eff} = \{\neg c\}$$

Figure 2 shows the ed-state resulting from executing *open_door*. Agents A and B have beliefs updated to reflect that the door is now open. Furthermore, agent B now knows that the door is unlocked. Agent C, who was not looking, still believes (but does not know) that the door is closed.

4.2 Sensing Actions

A sensing action is a 4-tuple, $(a.\texttt{pre}, a.\texttt{obs}, a.\texttt{aware}, a.\texttt{det})$, where $a.\texttt{det}$ is a set of fluent atoms indicating fluents whose truth value is detected by an action. Since sensing actions do not change the actual environment, we have:

$$V^+(w^+) = V(w)$$

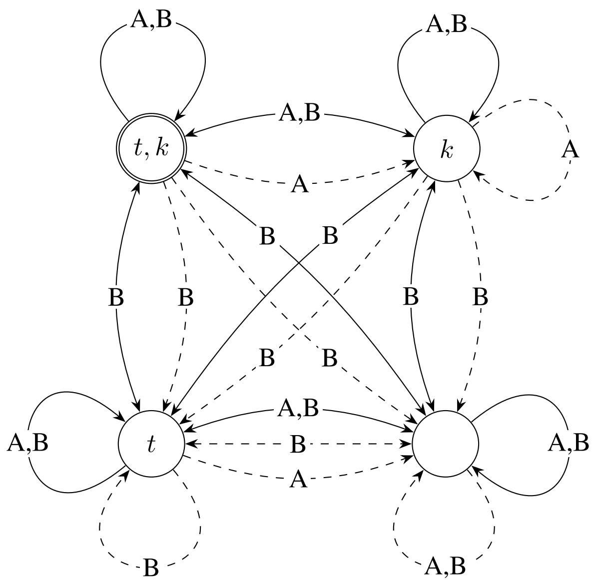

Figure 3: The ed-state before the sensing action look_in_box is applied.

Let $S(w)$ be a fluent formula expressing the values of the fluents in $a.\texttt{det}$ when a is executed in world w:

$$S(w) = \bigwedge_{f \in a.\texttt{det}} \begin{cases} f & \text{if } f \in V(w), \\ \neg f & \text{otherwise} \end{cases}$$

The temporary relation $\mathcal{R}_i$ is populated by $\mathcal{K}_i$ instead of $\mathcal{B}_i$ in worlds where the action preconditions (for full or partial observers) or the sensed facts (for full observers) violate an agent's beliefs. For $i \in F$:

$$\mathcal{R}_i = \{(u,v) \mid ((u,v) \in \mathcal{B}_i) \\ \vee ((u,v) \in \mathcal{K}_i \ \wedge (\mathcal{M},u) \models \mathbf{B}_i \neg(a.\texttt{pre} \wedge S(u)))\}$$

For $i \in P$:

$$\mathcal{R}_i = \{(u,v) \mid ((u,v) \in \mathcal{B}_i) \\ \vee ((u,v) \in \mathcal{K}_i \ \wedge (\mathcal{M},u) \models \mathbf{B}_i \neg a.\texttt{pre})\}$$

For full observers, the update to $\mathcal{B}_i$ removes edges between worlds that distinguish the sensed information. For $i \in F$:

$$\mathcal{B}'_i = \mathcal{B}_i \cup \{(u^+, v^+) \mid u, v \in \text{Pre}(\mathcal{M}, a), (u,v) \in \mathcal{R}_i, \\ S(u) = S(v)\}$$

Partial observers update their beliefs to reflect the preconditions. For $i \in P$:

$$\mathcal{B}'_i = \mathcal{B}_i \cup \{(u^+, v^+) \mid u, v \in \text{Pre}(\mathcal{M}, a), (u,v) \in \mathcal{R}_i\}$$

Oblivious agents do not reset their beliefs. For $i \in O$:

$$\mathcal{B}'_i = \mathcal{B}_i \cup \{(u^+, v) \mid u \in \text{Pre}(\mathcal{M}, a), (u,v) \in \mathcal{B}_i\}$$

The update to $\mathcal{K}_i$ mirrors that of $\mathcal{B}_i$, except that oblivious agents add edges expressing that they know the action *could* have occurred. For $i \in F$:

$$\mathcal{K}'_i = \mathcal{K}_i \cup \{(u^+, v^+) \mid u, v \in \text{Pre}(\mathcal{M}, a), (u,v) \in \mathcal{K}_i, \\ S(u) = S(v)\}$$

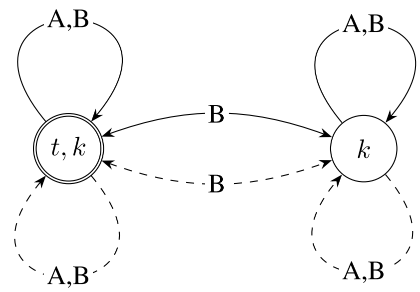

Figure 4: The ed-state after agent A performs the sensing action look_in_box.

For $i \in P$:

$$\mathcal{K}'_i = \mathcal{K}_i \cup \{(u^+, v^+) \mid u, v \in \text{Pre}(\mathcal{M}, a), (u,v) \in \mathcal{K}_i\}$$

For $i \in O$:

$$\mathcal{K}'_i = \mathcal{K}_i \cup \{(u^+, v^+) \mid u, v \in \text{Pre}(\mathcal{M}, a), (u,v) \in \mathcal{K}_i\} \\ \cup \{(u^+, v), (v, u^+) \mid u \in \text{Pre}(\mathcal{M}, a), (u,v) \in \mathcal{K}_i\}$$

Example Sensing Action We adapt the coin-in-the-box problem (Baral et al. 2015), to involve the correction of false belief. This example demonstrates two ways that a sensing action can correct a false belief: directly, when an agent senses the true value of a fluent about which she had had a false belief, and indirectly, when an agent observes or is aware of an action whose preconditions contradict a false belief.

Consider the ed-state illustrated in Figure 3. There are two agents, A and B, in a room with a box containing a coin, which lies tails-up (t). The box is locked and can only be opened with a key. Agent A has a key (k), and knows it, but believes wrongly that the coin lies heads-up ($\neg t$). Agent B does not believe that the coin is heads-up or that it is tails-up, but believes that Agent A does not have a key. Agent A performs the sensing action $look\_in\_box$:

$$look\_in\_box.\texttt{pre} = k$$
$$look\_in\_box.\texttt{obs} = \{(A, \top)\}$$
$$look\_in\_box.\texttt{aware} = \{(B, \top)\}$$
$$look\_in\_box.\texttt{det} = \{t\}$$

Figure 4 shows the state resulting from the $look\_in\_box$ action. Due to space constraints, only part of the state is illustrated: the new worlds created by the action (W^+) are shown, but none of the original worlds (or relations between them) are shown. This truncated state is equivalent with respect to formula entailment because there are no edges from W^+ to W in $\mathcal{B}'_i$ or $\mathcal{K}'_i$ for $i \in \{A, B\}$. Agent A knows that the coin is tails-up. As a partial observer, agent B is aware that the action has occurred, so now knows that agent A has the key and that agent A knows whether the coin lies tails-up. However, agent B still does not know (or have any belief about) whether the coin is tails-up.

4.3 Announcement Actions

An announcement action is a 4-tuple, $(a.\texttt{pre}, a.\texttt{obs}, a.\texttt{aware}, a.\texttt{ann})$, where $a.\texttt{ann}$ is a belief formula that is

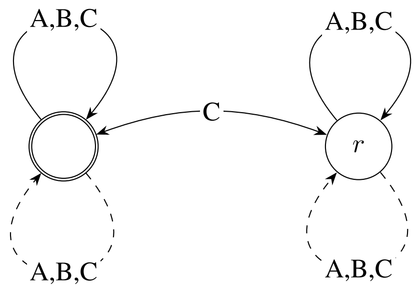

Figure 5: The ed-state before agent A performs the announcement action tell_raining.

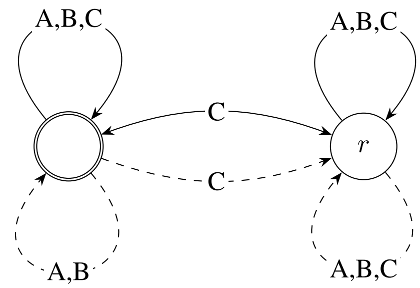

Figure 6: The ed-state after agent A performs the announcement action tell_raining.

asserted by the announcement. Like sensing actions, announcement actions do not affect the environment itself:

$$V^+(w^+) = V(w)$$

An agent will *reject* an announcement in world u if she knows it to be false, i.e. $(\mathcal{M}, u) \models \mathbf{K}_i \neg a.\mathtt{ann}$. In this case, the agent's belief will still be affected by the fact that the action has occurred (she will know its preconditions to be true), but not by what is announced. Full observers learn both the content of the announcement and that the action's preconditions hold. For $i \in F$:

$$\begin{aligned}\mathcal{R}_i = \{(u, v) \mid & ((u, v) \in \mathcal{B}_i) \\ & \vee ((u, v) \in \mathcal{K}_i \wedge (\mathcal{M}, u) \models \mathbf{B}_i \neg a.\mathtt{pre}) \\ & \vee ((u, v) \in \mathcal{K}_i \wedge (\mathcal{M}, u) \not\models \mathbf{K}_i \neg (a.\mathtt{pre} \wedge a.\mathtt{ann}) \\ & \quad \wedge (\mathcal{M}, u) \models \mathbf{B}_i \neg (a.\mathtt{pre} \wedge a.\mathtt{ann}))\}\end{aligned}$$

Partial observers learn that preconditions hold. For $i \in P$:

$$\begin{aligned}\mathcal{R}_i = \{(u, v) \mid & (u, v) \in \mathcal{B}_i \\ & \vee ((u, v) \in \mathcal{K}_i \ \wedge (\mathcal{M}, u) \models \mathbf{B}_i \neg a.\mathtt{pre})\}\end{aligned}$$

For full observers, the update to $\mathcal{B}_i$ removes edges to worlds where the announced information does not hold, unless the agent rejects the announcement. For $i \in F$:

$$\begin{aligned}\mathcal{B}'_i = \mathcal{B}_i \cup \{(u^+, v^+) \mid & u, v \in \mathrm{Pre}(\mathcal{M}, a), (u, v) \in \mathcal{R}_i, \\ & ((\mathcal{M}, u) \models \mathbf{K}_i \neg (a.\mathtt{pre} \wedge a.\mathtt{ann}) \vee (\mathcal{M}, v) \models a.\mathtt{ann})\}\end{aligned}$$

For $i \in P$:

$$\mathcal{B}'_i = \mathcal{B}_i \cup \{(u^+, v^+) \mid u, v \in \mathrm{Pre}(\mathcal{M}, a), (u, v) \in \mathcal{R}_i\}$$

For $i \in O$:

$$\mathcal{B}'_i = \mathcal{B}_i \cup \{(u^+, v) \mid u \in \mathrm{Pre}(\mathcal{M}, a), (u, v) \in \mathcal{B}_i\}.$$

Even though fully-observant agents know that the action's preconditions hold, they do not know the announcement holds, even if they do not reject it, because it could nevertheless be untrue. Thus, for $i \in F \cup P$:

$$\mathcal{K}'_i = \mathcal{K}_i \cup \{(u^+, v^+) \mid (u, v) \in \mathcal{K}_i, u, v \in \mathrm{Pre}(\mathcal{M}, a)\}$$

For $i \in O$:

$$\begin{aligned}\mathcal{K}'_i = \mathcal{K}_i & \cup \{(u^+, v^+) \mid (u, v) \in \mathcal{K}_i, u, v \in \mathrm{Pre}(\mathcal{M}, a)\} \\ & \cup \{(u^+, v), (v, u^+) \mid u \in \mathrm{Pre}(\mathcal{M}, a), (u, v) \in \mathcal{K}_i\}\end{aligned}$$

Example Announcement Action This example demonstrates how a lie can induce a false belief, or can be rejected by an agent who knows that it cannot be true. Consider the ed-state illustrated in Figure 5. There are three agents, A, B, and C. It is not raining ($\neg r$). Agents A and B know that it is not raining. Agent C believes that it is not raining but does not know this. Agent A performs the $tell\_raining$ action:

$$\begin{aligned}& tell\_raining.\mathtt{pre} = \top \\ & tell\_raining.\mathtt{obs} = \{(A, \top), (B, \top), (C, \top)\} \\ & tell\_raining.\mathtt{aware} = \emptyset \\ & tell\_raining.\mathtt{ann} = r\end{aligned}$$

Figure 6 shows the ed-state after agent A has performed the tell_raining action. As with the sensing action example, only the new worlds (W^+) and the edges between them are shown. Again, this truncated state is equivalent with respect to formula entailment because there are no edges from W^+ to W in $\mathcal{B}'_i$ or $\mathcal{K}'_i$ for $i \in \{A, B, C\}$. Agent C now has the false belief that it is raining. Agent B, however, who knew that it was not raining, rejected the announcement and still does not believe that it is raining.

5 Discussion

The state transitions defined in m$\mathcal{A}^*$ produce counterintuitive outcomes when applied to our examples. Figures 7 and 8 show the resulting epistemic states for our ontic and sensing action examples. We omit the announcement example because m$\mathcal{A}^*$ assumes that announcements are truthful. The relation $\mathcal{K}_i$ is ignored. In the ontic example, agent B is left in a state of ignorance. In this case, a minor adjustment to m$\mathcal{A}^*$ produces the same result as our technique: perform false belief revision when an agent observes an action whose preconditions violate the agent's beliefs, just as is done in the case of belief-correcting sensing and announcement actions. The second, sensing example is more problematic. Agent B's observation of the open_box action has left her in a state of ignorance, but if we try to correct this by again adding a reflexive connection to the designated world, agent B will come to believe that the coin lies tails-up, even though she was only a partial observer of the action.

One limitation of our approach is that an agent's knowledge is altered by an event even if the agent is oblivious to its occurrence. For that reason, applications of our technique

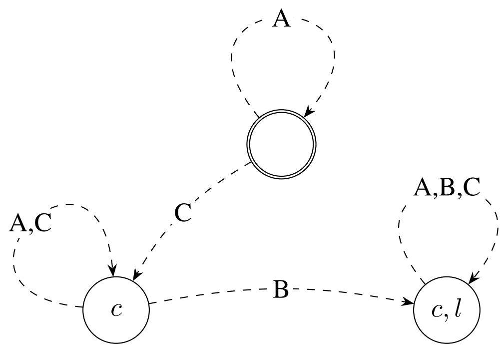

Figure 7: The result of applying the m$\mathcal{A}^*$ ontic action transition in the open_door example. Not shown are the fluents $w(A)$ and $w(B)$, which are true in all worlds.

may choose to rely primarily upon the belief relation to represent mental state, using the knowledge relation mainly for the purpose of belief reset. Another limitation is that belief reset can relax beliefs that are not informed by an event. If an agent believes ϕ and ψ, but doesn't know whether ϕ and doesn't know whether ψ, and learns that $\neg\psi$, then she will no longer believe ϕ, even though she ideally should only relax her belief that ψ.

As illustrated by the various examples, our proposed representations can express plans that involve deceit, recovery from false belief, and distinctions between agent belief and knowledge, e.g. in goals. A multi-agent epistemic planner using these representations could use the algorithm offered by Son et al. (2014) to construct an initial ed-state from a planning problem description. That algorithm builds a pointed Kripke model that satisfies $S5_n$ from a so-called finitary $S5_n$ theory. This can be extended to a bimodal pointed Kripke model by simply duplicating the accessibility relation to produce the relations $\mathcal{B}_i$ and $\mathcal{K}_i$, with both relations satisfying $S5_n$ and thus also $KD45_n$.

6 Conclusion

This work has defined state and transition representations for multi-agent epistemic planning that incorporate agents' knowledge as well as belief. Transitions include ontic, sensing, and announcement actions. We have illustrated these dynamics with examples of state transitions for each action type, and shown how our approach improves upon prior work. Our methods can be used by an epistemic planner to make plans that involve false beliefs, false announcements (including deceit), and recovery from false belief, with belief and knowledge interacting in complex multi-agent scenarios.

7 Proofs

In this section we prove that for agents $i \in \mathcal{A}$ actions preserve $S5_n$ properties in $\mathcal{K}_i$, $KD45_n$ properties in $\mathcal{B}$, and $KB1$ and $KB2$ properties in $\mathcal{K}_i$ and $\mathcal{B}_i$.

7.1 Preliminaries

Relations $\mathcal{R}_i$ for $i \in \mathcal{A}$ on worlds W satisfy $KD45_n$ if they are serial ($\forall\ u \in W\ : \exists\ v \in W, (u,v) \in \mathcal{R}_i$), transitive

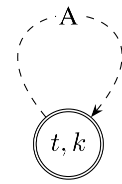

Figure 8: The result of applying the m$\mathcal{A}^*$ sensing action in the look_in_box example. Only the designated world is shown because the accessibility relation has no edges leaving the designated world.

($(u,v) \in \mathcal{R}_i \wedge (v,z) \in \mathcal{R}_i \rightarrow (u,z) \in \mathcal{R}_i$), and Euclidean ($(u,v) \in \mathcal{R}_i \wedge (u,z) \in \mathcal{R}_i \rightarrow (v,z) \in \mathcal{R}_i$). They satisfy $S5_n$ if they are reflexive ($\forall\ u \in W\ : (u,u) \in \mathcal{R}_i$), transitive, and symmetric ($\forall (u,v) \in \mathcal{R}_i : (v,u) \in \mathcal{R}_i$). Between a knowledge relation $\mathcal{K}_i$ and a belief relation $\mathcal{B}_i$ over a set W of worlds, the $KB1$ property specifies that $\mathcal{B}_i \subseteq \mathcal{K}_i$. The $KB2$ property specifies that for every $u, v, z \in W$, if $(u,v) \in \mathcal{K}_i$ and $(v,z) \in \mathcal{B}_i$, then $(u,z) \in \mathcal{B}_i$.

For each of these proofs we will assume that $S5_n$ and $KD45_n$ hold for $\mathcal{B}_i$ and $\mathcal{K}_i$ respectively, and that $KB1$ and $KB2$ hold between them. We will show inductively that the new relations $\mathcal{B}'_i$ and $\mathcal{K}'_i$ satisfy these properties over W'.

We will use $\mathcal{R}^+_i$ to refer to $\mathcal{R}_i \backslash \mathcal{B}_i$ (the portion of $\mathcal{R}_i$ corresponding to a belief reset). In many proofs, neither $\mathcal{B}'_i$ nor $\mathcal{K}'_i$ have any edges between W and W^+. In such cases, it will suffice to show the properties hold over $\mathcal{B}^+_i$ or $\mathcal{K}^+_i$, since the induction hypothesis shows they hold over $\mathcal{K}_i$ and $\mathcal{B}_i$. Only the proofs for oblivious agents ($i \in O$) do not make use of this principle.

7.2 Ontic Actions, $i \in F \cup P$

$\mathcal{K}_i$ Reflexivity Consider any $u^+ \in W^+$ where $u \in \text{Pre}(\mathcal{M}, a)$. By reflexivity of $\mathcal{K}_i$, $(u,u) \in \mathcal{K}_i$. Since $u \in \text{Pre}(\mathcal{M}, a)$, $(u^+, u^+) \in \mathcal{K}^+_i$.

$\mathcal{K}_i$ Symmetry Suppose $(u^+, v^+) \in \mathcal{K}^+_i$. Then $u, v \in \text{Pre}(\mathcal{M}, a)$ and $(u,v) \in \mathcal{K}_i$. By symmetry of $\mathcal{K}_i$, $(v,u) \in \mathcal{K}_i$. Since $v, u \in \text{Pre}(\mathcal{M}, a)$, $(v^+, u^+) \in \mathcal{K}^+_i$.

$\mathcal{K}_i$ Transitivity Suppose that $(u^+, v^+), (v^+, z^+) \in \mathcal{K}^+_i$. Then $u, v, z \in \text{Pre}(\mathcal{M}, a)$ and $(u,v), (v,z) \in \mathcal{K}_i$. By transitivity of $\mathcal{K}_i$, $(u,z) \in \mathcal{K}_i$. Since $u, z \in \text{Pre}(\mathcal{M}, a)$, it follows that $(u^+, z^+) \in \mathcal{K}^+_i$.

$\mathcal{B}_i$ Seriality Suppose $u^+ \in W^+$. Then $u \in \text{Pre}(\mathcal{M}, a)$. Either there is some $v \in \text{Pre}(\mathcal{M}, a)$ such that $(u,v) \in \mathcal{B}_i$, or there is not. If such a v exists, then $(u,v) \in \mathcal{R}_i$, so $(u^+, v^+) \in \mathcal{B}^+_i$. Otherwise, there is no such v, and thus $(\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}$ by definition. In this case, $\{(u,v) : v \in W, (u,v) \in \mathcal{K}_i\} \subseteq \mathcal{R}_i$, and by reflexivity of $\mathcal{K}_i$, we have that $(u,u) \in \mathcal{R}_i$. Thus $(u^+, v^+) \in \mathcal{B}^+_i$ with $u = v$.

$\mathcal{B}_i$ Transitivity Suppose that $(u^+, v^+), (v^+, z^+) \in \mathcal{B}^+_i$, with $u, v, z \in \text{Pre}(\mathcal{M}, a)$ and $(u,v), (v,z) \in \mathcal{R}_i$. There are four cases: (a) $(u,v), (v,z) \in \mathcal{B}_i$; (b) $(u,v) \in \mathcal{R}^+_i$, $(v,z) \in \mathcal{B}_i$; (c) $(u,v), (v,z) \in \mathcal{R}^+_i$; and (d) $(u,v) \in \mathcal{B}_i$,

$(v, z) \in \mathcal{R}_i^+$. In case (a), transitivity of $\mathcal{B}_i$ implies that $(u, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$. Since $u, z \in \text{Pre}(\mathcal{M}, a)$, it follows that $(u^+, z^+) \in \mathcal{B}_i^+$. In cases (b) and (c), $(u, v) \notin \mathcal{B}_i$, it must be that $(\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}$. Since we know by definition of $\mathcal{R}_i^+$ that $(u, v) \in \mathcal{K}_i$ and $(v, z) \in \mathcal{K}_i$, by transitivity of $\mathcal{K}_i$ we have that $(u, z) \in \mathcal{K}_i$. Because $(\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}$, we have that $(u, z) \in \mathcal{R}_i$. Thus, since $u, z \in \text{Pre}(\mathcal{M}, a)$, it follows that $(u^+, z^+) \in \mathcal{B}_i^+$. Case (d) cannot occur. Since $\mathcal{B}_i$ is serial and $(\mathcal{M}, v) \models \mathbf{B}_i \neg a.\texttt{pre}$, there must be some $w \in W \backslash \text{Pre}(\mathcal{M}, a)$ such that $(v, w) \in \mathcal{B}_i$. Because $\mathcal{B}_i$ is transitive and $(u, v), (v, w) \in \mathcal{B}_i$, we have $(u, w) \in \mathcal{B}_i$; because $\mathcal{B}_i$ is Euclidean and $(u, w), (u, v) \in \mathcal{B}_i$, we have $(w, v) \in \mathcal{B}_i$; and finally, because $\mathcal{B}_i$ is transitive and $(v, w), (w, v) \in \mathcal{B}_i$, we have $(v, v) \in \mathcal{B}_i$. But since $(\mathcal{M}, v) \models a.\texttt{pre}$, $(\mathcal{M}, v) \not\models \mathbf{B}_i \neg a.\texttt{pre}$, a contradiction.

$\mathcal{B}_i$ Euclideanness Suppose $(u^+, v^+), (u^+, z^+) \in \mathcal{B}_i^+$, so $u, v, z \in \text{Pre}(\mathcal{M}, a)$, $(u, v), (u, z) \in \mathcal{R}_i$. There are two cases: either (a) $(\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}$, or (b) $(\mathcal{M}, u) \not\models \mathbf{B}_i \neg a.\texttt{pre}$. In case (a), $KB2$ property for $\mathcal{K}_i, \mathcal{B}_i$ implies that $(\mathcal{M}, u) \models \mathcal{K}_i \mathcal{B}_i \neg a.\texttt{pre}$. Thus for any world $w \in W$ such that $(u, w) \in \mathcal{K}_i$, $(\mathcal{M}, w) \models \mathbf{B}_i \neg a.\texttt{pre}$. Since $(u, v) \in \mathcal{R}_i \subseteq \mathcal{K}_i$, it must be that $(\mathcal{M}, v) \models \mathbf{B}_i \neg a.\texttt{pre}$. Further, since $\mathcal{K}_i$ is Euclidean, we must have $(v, z) \in \mathcal{K}_i$. These two facts imply that $(v, z) \in \mathcal{R}_i$. In case (b), it must be that $(u, v), (u, z) \in \mathcal{B}_i$. By Euclideanness of $\mathcal{B}_i$, it follows that $(v, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$. In either case, because $(v, z) \in \mathcal{R}_i$ and $v, z \in \text{Pre}(\mathcal{M}, a)$, we have that $(v^+, z^+) \in \mathcal{B}_i^+$ as desired.

KB1 Suppose that $(u^+, v^+) \in \mathcal{B}_i^+$ where $u, v \in \text{Pre}(\mathcal{M}, a)$, $(u, v) \in \mathcal{R}_i$. By $KB1$, $\mathcal{R}_i \subseteq \mathcal{K}_i$, so $(u, v) \in \mathcal{K}_i$. Since $u, v \in \text{Pre}(\mathcal{M}, a)$, $(u^+, v^+) \in \mathcal{K}_i^+$ as desired.

KB2 Suppose that $(u^+, v^+) \in \mathcal{K}_i^+$ and $(v^+, z^+) \in \mathcal{B}_i^+$. Then $u, v, z \in \text{Pre}(\mathcal{M}, a)$, $(u, v) \in \mathcal{K}_i$, and $(v, z) \in \mathcal{R}_i$. Either (a) $(\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}$, or (b) $(\mathcal{M}, u) \not\models \mathbf{B}_i \neg a.\texttt{pre}$. In case (a), we know $(v, z) \in \mathcal{R}_i \subseteq \mathcal{K}_i$, so by transitivity of $\mathcal{K}_i$, $(u, z) \in \mathcal{K}_i$. Since the assumption was that $(\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}$, we thus have that $(u, z) \in \mathcal{R}_i$. Since $u, z \in \text{Pre}(\mathcal{M}, a)$, it follows that $(u^+, z^+) \in \mathcal{B}_i^+$ as desired. In case (b), it must be that $(v, z) \in \mathcal{B}_i$. Since $(u, v) \in \mathcal{K}_i$, by $KB2$ on $\mathcal{K}_i, \mathcal{B}_i$ it follows that $(u, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$. Since $u, z \in \text{Pre}(\mathcal{M}, a)$, it follows that $(u^+, z^+) \in \mathcal{B}_i^+$ as desired.

7.3 Ontic Actions, $i \in O$

For the knowledge relation on oblivious agents, we split $\mathcal{K}_i^+$ into three distinct sets:

$$\mathcal{K}_i^L := \{(u^+, v) : u \in \text{Pre}(\mathcal{M}, a), v \in W, (u, v) \in \mathcal{K}_i\}$$
$$\mathcal{K}_i^R := \{(v, u^+) : u \in \text{Pre}(\mathcal{M}, a), v \in W, (u, v) \in \mathcal{K}_i\}$$
$$\mathcal{K}_i^B := \{(u^+, v^+) : u \in \text{Pre}(\mathcal{M}, a), v \in W, (u, v) \in \mathcal{K}_i\}$$

$\mathcal{K}_i$ Reflexivity Consider any $u \in W'$. If $u \in W$, then $(u, u) \in \mathcal{K}_i \subseteq \mathcal{K}_i^+$ by reflexivity of $\mathcal{K}_i$. Otherwise, $u = w^+$ where $w \in \text{Pre}(\mathcal{M}, a)$, and by reflexivity of $\mathcal{K}_i$, $(w, w) \in \mathcal{K}_i$, and thus $(u, u) = (w^+, w^+) \in \mathcal{K}_i^B \subseteq \mathcal{K}_i'$.

$\mathcal{K}_i$ Symmetry Suppose $(u, v) \in \mathcal{K}_i'$. If $u, v \in W$ then $(u, v) \in \mathcal{K}_i$, and by symmetry of $\mathcal{K}_i$, $(v, u) \in \mathcal{K}_i \subseteq \mathcal{K}_i'$. If $u, v \in W^+$ then $(u, v) \in \mathcal{K}_i^B$, so $(u, v) = (w^+, z^+)$ for some $w, z \in \text{Pre}(\mathcal{M}, a)$ with $(w, z) \in \mathcal{K}_i$. By symmetry of $\mathcal{K}_i$, $(z, w) \in \mathcal{K}_i$ and thus $(z^+, w^+) \in \mathcal{K}_i^B$. If $u \in W, v \in W^+$ then $(u, v) \in \mathcal{K}_i^L$, and we can see that $(v, u) \in \mathcal{K}_i^R$. By the same logic, if $(u, v) \in \mathcal{K}_i^R$ then $(v, u) \in \mathcal{K}_i^L$.

$\mathcal{K}_i$ Transitivity Suppose $(u, v), (v, z) \in \mathcal{K}_i'$. We define $\hat{u}, \hat{v}, \hat{z}$ such that for $\hat{x} \in \{\hat{u}, \hat{v}, \hat{z}\}$, $\hat{x} = x$ if $x \in W, w \in \text{Pre}(\mathcal{M}, a) \mid w^+ = x$ otherwise. Thus $(\hat{u}, \hat{v}), (\hat{v}, \hat{z}) \in \mathcal{K}_i$ and it follows by transitivity of $\mathcal{K}_i$ that $(\hat{u}, \hat{z}) \in \mathcal{K}_i$, and by definition of $\mathcal{K}_i^+$ that $(u, z) \in \mathcal{K}_i^+$ as desired.

$\mathcal{B}_i$ Seriality For $u \in W$, seriality of $\mathcal{B}_i$ implies that there is some $v \in W \subseteq W'$ such that $(u, v) \in \mathcal{B}_i \subseteq \mathcal{B}_i'$ as desired. For $u \in W'$, there is some $w \in \text{Pre}(\mathcal{M}, a)$ such that $u = w^+$. Since $\mathcal{B}_i$ is serial, there is some $v \in W$ such that $(w, v) \in \mathcal{B}_i$. Thus by definition $(w^+, v) \in \mathcal{B}_i^+ \subseteq \mathcal{B}_i'$.

$\mathcal{B}_i$ Transitivity Suppose that $(u, v), (v, z) \in \mathcal{B}_i'$. Note that for every pair $(w, x) \in \mathcal{B}_i'$, $x \in W$. Thus $v, z \in W$ and by definition of $\mathcal{B}_i'$ it must be true that $(v, z) \in \mathcal{B}_i$. If also $u \in W$ then by definition of $\mathcal{B}_i^+$ $(u, v) \in \mathcal{B}_i$, and it follows by transitivity of $\mathcal{B}_i$ that $(u, z) \in \mathcal{B}_i \subseteq \mathcal{B}_i'$ as desired. Otherwise, $u = w^+$ for some $w \in \text{Pre}(\mathcal{M}, a)$, and $(w, v) \in \mathcal{B}_i$. By transitivity of $\mathcal{B}_i$, $(w, z) \in \mathcal{B}_i$, so $(u, z) = (w^+, z) \in \mathcal{B}_i^+ \subseteq \mathcal{B}_i'$ as desired.

$\mathcal{B}_i$ Euclideanness Suppose that $(u, v), (u, z) \in \mathcal{B}_i^+$. Certainly $v, z \in W$. Either $u \in W$ or $u \in W^+$. If $u \in W$, then $(u, v), (u, z) \in \mathcal{B}_i$, so since $\mathcal{B}_i$ is Euclidean we have that $(v, z) \in \mathcal{B}_i \subseteq \mathcal{B}_i'$ as desired. Otherwise, there is some $w \in \text{Pre}(\mathcal{M}, a)$ such that $u = w^+$, and $(w, v), (w, z) \in \mathcal{B}_i$. Again by Euclideanness of $\mathcal{B}_i$, $(v, z) \in \mathcal{B}_i \subseteq \mathcal{B}_i'$ as desired.

KB1 Suppose that $(u, v) \in \mathcal{B}_i'$. Either $(u, v) \in \mathcal{B}_i$, in which case $(u, v) \in \mathcal{K}_i \subseteq \mathcal{K}_i'$ as desired, by $KB1$ on $\mathcal{K}_i, \mathcal{B}_i$, or $u = w^+$ for some $w \in \text{Pre}(\mathcal{M}, a)$, and $(w, v) \in \mathcal{B}_i$. Thus $(w, v) \in \mathcal{K}_i$ by $KB1$ on $\mathcal{K}_i, \mathcal{B}_i$, and it follows that $(u, v) = (w^+, v) \in \mathcal{K}_i^L \subseteq \mathcal{K}_i$ as desired.

KB2 Suppose that $(u, v) \in \mathcal{K}_i'$ and $(v, z) \in \mathcal{B}_i'$. Since by definition of $\mathcal{B}_i'$, $z \in W$, there are four cases: (a) $u, v \in W$; (b) $u \in W, v \in W^+$; (c) $v \in W, u \in W^+$, or (d) $u, v \in W^+$. In case (a), $(u, v) \in \mathcal{K}_i$ and $(v, z) \in \mathcal{B}_i$, and by $KB2$ for $\mathcal{K}_i, \mathcal{B}_i$ we have that $(u, z) \in \mathcal{B}_i \subseteq \mathcal{B}_i'$ as desired. In case (b), $(u, v) \in \mathcal{K}_i^R$ and $(v, z) \in \mathcal{B}_i^+$. Here $v = w^+$ for some $w \in \text{Pre}(\mathcal{M}, a)$, where $(w, u) \in \mathcal{K}_i$, $(w, z) \in \mathcal{B}_i$. By $KB2$ for $\mathcal{K}_i, \mathcal{B}_i$ it follows that $(u, z) \in \mathcal{B}_i \subseteq \mathcal{B}_i'$ as desired. In case (c), $(u, v) \in \mathcal{K}_i^L$ and $(v, z) \in \mathcal{B}_i$. Here $u = w^+$ for some $w \in \text{Pre}(\mathcal{M}, a)$, where $(w, v) \in \mathcal{K}_i$, $(v, z) \in \mathcal{B}_i$. By $KB2$ for $\mathcal{K}_i, \mathcal{B}_i$ it follows that $(w, z) \in \mathcal{B}_i$, so by definition $(u, z) \in \mathcal{B}_i^+ \subseteq \mathcal{B}_i'$ as desired. In case (d), $(u, v) \in \mathcal{K}_i^B$ and $(v, z) \in \mathcal{B}_i^+$. Here $(u, v) = (w^+, x^+)$ for some $w, x \in \text{Pre}(\mathcal{M}, a)$, where $(w, x) \in \mathcal{K}_i$, $(x, z) \in \mathcal{B}_i$. By $KB2$ for $\mathcal{K}_i, \mathcal{B}_i$ it follows that $(w, z) \in \mathcal{B}_i$, so by definition $(u, z) \in \mathcal{B}_i^+ \subseteq \mathcal{B}_i'$ as desired.

7.4 Sensing Actions, $i \in F$

$\mathcal{K}_i$ Reflexivity Consider $u^+ \in W^+$, where $u \in \text{Pre}(\mathcal{M}, a)$. By reflexivity of $\mathcal{K}_i$, $(u, u) \in \mathcal{K}_i$. Since $u \in \text{Pre}(\mathcal{M}, a)$ and certainly $S(u) = S(u)$, it follows that $(u^+, u^+) \in \mathcal{K}_i^+$ as desired.

$\mathcal{K}_i$ Symmetry Suppose $(u^+, v^+) \in \mathcal{K}_i^+$. Then $u, v \in \text{Pre}(\mathcal{M}, a)$, $(u, v) \in \mathcal{K}_i$, and $S(u) = S(v)$. By symmetry of $\mathcal{K}_i$, we have that (v, u) and certainly $S(v) = S(u)$, so $(v^+, u^+) \in \mathcal{K}_i^+$ as desired.

$\mathcal{K}_i$ Transitivity Suppose $(u^+, v^+), (v^+, z^+) \in \mathcal{K}_i^+$. Then $u, v, z \in \text{Pre}(\mathcal{M}, a)$, $(u, v), (v, z) \in \mathcal{K}_i$, $S(u) = S(v)$ and $S(v) = S(z)$. By transitivity of $\mathcal{K}_i$, we have that $(u, z) \in \mathcal{K}_i$ and certainly $S(u) = S(z)$, so $(u^+, z^+) \in \mathcal{K}_i^+$ as desired.

$\mathcal{B}_i$ Seriality Suppose $u^+ \in W^+$ where $u \in \text{Pre}(\mathcal{M}, a)$. Either $(\mathcal{M}, u) \models \mathbf{B}_i\neg(a.\texttt{pre} \wedge S(u))$ or not. If $(\mathcal{M}, u) \models \mathbf{B}_i\neg(a.\texttt{pre} \wedge S(u))$, then it follows that for all $v \in W$ such that $(u, v) \in \mathcal{K}_i$, $(u, v) \in \mathcal{R}_i$. By reflexivity of $\mathcal{K}_i$, $(u, u) \in \mathcal{K}_i$. Thus $(u, u) \in \mathcal{R}_i$, so since $S(u) = S(u)$ trivially and $u \in \text{Pre}(\mathcal{M}, a)$ we have that $(u^+, v^+) \in \mathcal{B}_i^+$ as desired, with $v = u$. Otherwise, by seriality there must exist some $v \in W$ such that $(u, v) \in \mathcal{B}_i \subseteq \mathcal{R}_i$ and $(\mathcal{M}, v) \models a.\texttt{pre} \wedge S(u)$. The entailment of the conjunction implies that $(\mathcal{M}, v) \models a.\texttt{pre}$ (so $v \in \text{Pre}(\mathcal{M}, a)$) and $(\mathcal{M}, v) \models S(u)$ (so $S(u) = S(v)$). By definition it follows that $(u^+, v^+) \in \mathcal{B}_i^+$ as desired.

$\mathcal{B}_i$ Transitivity Suppose $(u^+, v^+), (v^+, z^+) \in \mathcal{B}_i^+$. Then $u, v, z \in \text{Pre}(\mathcal{M}, a)$ with $(u, v), (v, z) \in \mathcal{R}_i$, $S(u) = S(v) = S(z)$. Suppose $(u, v) \in \mathcal{B}_i$ (rather than $\mathcal{R}_i^+$). By seriality of $\mathcal{B}_i$, there exists some world w such that $(v, w) \in \mathcal{B}_i$. By transitivity of $\mathcal{B}_i$, we have that $(u, w) \in \mathcal{B}_i$, and by the Euclidean property of $\mathcal{B}_i$ we thus have that $(w, v) \in \mathcal{B}_i$. By the transitive property since $(v, w), (w, v) \in \mathcal{B}_i$, it follows that $(v, v) \in \mathcal{B}_i$. Since $v \in \text{Pre}(\mathcal{M}, a)$, it cannot be the case that $(\mathcal{M}, v) \models \neg(a.\texttt{pre} \wedge S(v))$. Thus, it must be that $(v, z) \in \mathcal{B}_i$. By transitivity of $\mathcal{B}_i$, we thus have $(u, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$, and since $S(u) = S(z)$ we then have that $(u^+, z^+) \in \mathcal{B}_i^+$ as desired. Otherwise, it must be the case that $(u, v) \in \mathcal{K}_i$ and $(\mathcal{M}, u) \models \mathbf{B}_i\neg(a.\texttt{pre} \wedge S(u))$. Since $(v, z) \in \mathcal{R}_i \subseteq \mathcal{K}_i$, we thus have $(u, z) \in \mathcal{K}_i$ by transitivity of $\mathcal{K}_i$, and because the entailment property still holds over u, it follows that $(u, z) \in \mathcal{R}_i$. Since we know $u, z \in \text{Pre}(\mathcal{M}, a)$ and $S(u) = S(z)$, it follows that $(u^+, z^+) \in \mathcal{B}_i^+$ as desired.

$\mathcal{B}_i$ Euclideanness Suppose $(u^+, v^+), (u^+, z^+) \in \mathcal{B}_i^+$. Then $S(u) = S(v) = S(z)$, and $u, v, z \in \text{Pre}(\mathcal{M}, a)$, and $(u, v), (u, z) \in \mathcal{R}_i$. Either $(\mathcal{M}, u) \models \mathbf{B}_i\neg(a.\texttt{pre} \wedge S(u))$, or not. If so, then since $(\mathcal{M}, u) \models \mathbf{B}_i\neg(a.\texttt{pre} \wedge S(u))$, we can deduce by $KB1$ on $\mathcal{K}_i, \mathcal{B}_i$ that $(\mathcal{M}, u) \models \mathbf{K}_i\mathbf{B}_i\neg(a.\texttt{pre} \wedge S(u))$. Thus, any $w \in W$ such that $(u, w) \in \mathcal{K}_i$ must be such that $(\mathcal{M}, w) \models \mathbf{B}_i\neg(a.\texttt{pre} \wedge S(u))$. In particular, since $(u, v) \in \mathcal{R}_i \subseteq \mathcal{K}_i$, we must have $(\mathcal{M}, v) \models \mathbf{B}_i\neg(a.\texttt{pre} \wedge S(v))$ since $S(v) = S(u)$. By Euclideanness of $\mathcal{K}_i$ (which follows from symmetry and transitivity of $\mathcal{K}_i$), we must have that $(v, z) \in \mathcal{K}_i$, and by the aforementioned entailment property it follows that $(v, z) \in \mathcal{R}_i$. Since $S(v) = S(z)$ and $v, z \in \text{Pre}(\mathcal{M}, a)$, we have that $(v^+, z^+) \in \mathcal{B}_i^+$ as desired. Otherwise, it must be that $(u, v), (u, z) \in \mathcal{B}_i$, so by Euclideanness of $\mathcal{B}_i$ it follows that $(v, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$. Since $v, z \in \text{Pre}(\mathcal{M}, a)$ and $S(v) = S(z)$, we have that $(v^+, z^+) \in \mathcal{B}_i^+$ as desired.

KB1 Suppose $(u^+, v^+) \in \mathcal{B}_i^+$. Then $u, v \in \text{Pre}(\mathcal{M}, a)$, $(u, v) \in \mathcal{R}_i$, and $S(u) = S(v)$. Since $\mathcal{R}_i \subseteq \mathcal{K}_i$, $(u, v) \in \mathcal{K}_i$ (by the $KB1$ assumption on $\mathcal{K}_i, \mathcal{B}_i$). Since $u, v \in \text{Pre}(\mathcal{M}, a)$ and $S(u) = S(v)$, we thus have that $(u^+, v^+) \in \mathcal{K}_i^+$ as desired.

KB2 Suppose $(u^+, v^+) \in \mathcal{K}_i^+, (v^+, z^+) \in \mathcal{B}_i^+$. Then $u, v, z \in \text{Pre}(\mathcal{M}, a)$, $S(u) = S(v)$, $S(v) = S(z)$, $(u, v) \in \mathcal{K}_i$ and $(v, z) \in \mathcal{B}_i$. Suppose $(v, z) \in \mathcal{R}_i$. Then by $KB2$ on $\mathcal{K}_i, \mathcal{B}_i$, we have that $(u, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$. Since $u, z \in \text{Pre}(\mathcal{M}, a)$ and $S(u) = S(z)$, it follows that $(u^+, z^+) \in \mathcal{B}_i^+$ as desired. Otherwise, it must be that $(v, z) \in \mathcal{K}_i$ and $(\mathcal{M}, v) \models \mathbf{B}_i\neg(a.\texttt{pre} \wedge S(v))$. By $KB2$ on $\mathcal{K}_i, \mathcal{B}_i$ it follows that $(\mathcal{M}, v) \models \mathbf{K}_i\mathbf{B}_i\neg(a.\texttt{pre} \wedge S(v))$. By symmetry of $\mathcal{K}_i$ and since $(u, v) \in \mathcal{K}_i$, we have that $(v, u) \in \mathcal{K}_i$. Thus, it must be that $(\mathcal{M}, u) \models \mathbf{B}_i\neg(a.\texttt{pre} \wedge S(u))$ since $S(u) = S(v)$. Also, by transitivity of $\mathcal{K}_i$, we have that $(u, z) \in \mathcal{K}_i$. These two facts imply that $(u, z) \in \mathcal{R}_i$. Since $S(u) = S(z)$ and $u, z \in \text{Pre}(\mathcal{M}, a)$, we thus have that $(u^+, z^+) \in \mathcal{B}_i^+$ as desired.

7.5 Sensing Actions, $i \in P \cup O$

Updates to $\mathcal{K}_i$ and $\mathcal{B}_i$ for sensing actions with $i \in P$ and $i \in O$ are identical to those for ontic actions with $i \in P$ and $i \in O$, respectively (Sections 7.2 and 7.3).

7.6 Announcement Actions, $i \in F$

$\mathcal{K}_i$ $S5_n$ Updates to $\mathcal{K}_i$ for announcement actions with $i \in F$ are identical to those for ontic actions with $i \in F$ (Section 7.2).

$\mathcal{B}_i$ Seriality Consider $u^+ = W^+$ where $u \in \text{Pre}(\mathcal{M}, a)$. Suppose $(\mathcal{M}, u) \not\models \mathbf{K}_i\neg(a.\texttt{ann} \wedge a.\texttt{pre})$, but $(\mathcal{M}, u) \models \mathbf{B}_i\neg(a.\texttt{ann} \wedge a.\texttt{pre})$. We will call these conditions respectively C1 and C2. By C1, there exists some $v \in W$ such that $(u, v) \in \mathcal{K}_i$ and $(\mathcal{M}, v) \models a.\texttt{ann} \wedge a.\texttt{pre}$. This implies that $v \in \text{Pre}(\mathcal{M}, a)$ and $(\mathcal{M}, v) \models a.\texttt{ann}$. Because $\{(u, w) : (u, w) \in \mathcal{K}_i \wedge \text{C1 holds} \wedge \text{ C2 holds}\} \subseteq \mathcal{R}_i$, we have that $(u, v) \in \mathcal{R}_i$. Thus, since $u, v \in \text{Pre}(\mathcal{M}, a)$ and $(\mathcal{M}, v) \models a.\texttt{ann}$, we have that $(u^+, v^+) \in \mathcal{B}_i^+$ as desired. Next, suppose that $(\mathcal{M}, u) \not\models \mathbf{B}_i\neg(a.\texttt{ann} \wedge a.\texttt{pre})$. This implies that there must be some $v \in W$ such that $(u, v) \in \mathcal{B}_i \subseteq \mathcal{R}_i$ and $(\mathcal{M}, v) \models a.\texttt{ann} \wedge a.\texttt{pre}$. Then $v \in \text{Pre}(\mathcal{M}, a)$ and $v \models a.\texttt{ann}$, so $(u^+, v^+) \in \mathcal{B}_i^+$ as desired. Next, suppose that $(\mathcal{M}, u) \models \mathbf{K}_i\neg(a.\texttt{ann} \wedge a.\texttt{pre})$ but $(\mathcal{M}, u) \not\models \mathbf{B}_i\neg a.\texttt{pre}$. Because $(\mathcal{M}, u) \not\models \mathbf{B}_i\neg a.\texttt{pre}$ there must be some $v \in W$ such that $(u, v) \in \mathcal{B}_i \subseteq \mathcal{R}_i$. Further, since $(\mathcal{M}, v) \models a.\texttt{pre}$, $v \in \text{Pre}(\mathcal{M}, a)$. Therefore, since $(\mathcal{M}, u) \models \mathbf{K}_i\neg(a.\texttt{ann} \wedge a.\texttt{pre})$, $(u^+, v^+) \in \mathcal{B}_i^+$ as desired. Finally, suppose that $(\mathcal{M}, u) \models \mathbf{K}_i\neg(a.\texttt{ann} \wedge a.\texttt{pre})$ and

$(\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}$. Because $\mathcal{K}_i$ is reflexive, we must have that $(u, u) \in \mathcal{K}_i$. Because $\{(u, v) : (u, v) \in \mathcal{K}_i \wedge (\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}\} \subseteq \mathcal{R}_i$, so $(u, u) \in \mathcal{R}_i$. Because $(u, u) \in \mathcal{R}_i$, $u \in \text{Pre}(\mathcal{M}, a)$, and $(\mathcal{M}, u) \models \mathbf{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$, it follows that $(u^+, u^+) \in \mathcal{B}_i^+$ as desired.

$\mathcal{B}_i$ Transitivity Suppose $(u^+, v^+), (v^+, z^+) \in \mathcal{B}_i^+$. Then $u, v, z \in \text{Pre}(\mathcal{M}, a)$, $(u, v), (v, z) \in \mathcal{R}_i$. Suppose that $(\mathcal{M}, u) \models \mathcal{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$ (call this C1) and $(\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}$ (call this C2). Certainly $(u, v), (v, z) \in \mathcal{K}_i$. By transitivity of $\mathcal{K}_i$, it follows that $(u, z) \in \mathcal{K}_i$. Because $\{(u, v) : (u, v) \in \mathcal{K}_i, \wedge \text{C2 holds}\} \subseteq \mathcal{R}_i$, we have that $(u, z) \in \mathcal{R}_i$. Since $u, z \in \text{Pre}(\mathcal{M}, a)$ and $(\mathcal{M}, u) \models \mathcal{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$, we thus have that $(u^+, z^+) \in \mathcal{B}_i^+$. Next, suppose C1 and ¬C2. C1 and ¬C2 imply that $(u, v) \in \mathcal{B}_i$ (since (u, v) cannot be in $\mathcal{R}_i^+$). ¬C2 implies that there is some $w \in W$ such that $(u, v) \in \mathcal{B}_i$ and $(\mathcal{M}, w) \models a.\texttt{pre}$. Since $\mathcal{B}_i$ is Euclidean, $(v, w) \in \mathcal{B}_i$, and thus $(\mathcal{M}, v) \not\models \mathbf{B}_i \neg a.\texttt{pre}$. If there was $x \in W$ such that $(v, x) \in \mathcal{K}_i$ and $(\mathcal{M}, x) \models a.\texttt{ann} \wedge a.\texttt{pre}$, then by the transitivity of $\mathcal{K}_i$ we have that $(u, x) \in \mathcal{K}_i$ and thus that C1 is false, a contradiction. Thus $(\mathcal{M}, v) \models \mathcal{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$, and (since C1 and ¬C2 hold at v) we must have $(v, z) \in \mathcal{B}_i$. By transitivity of $\mathcal{B}_i$, $(u, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$. Since C1 holds and $u, z \in \text{Pre}(\mathcal{M}, a)$, $(u^+, z^+) \in \mathcal{B}_i^+$ as desired. Otherwise, $(\mathcal{M}, u) \not\models \mathcal{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$, so by definition of $\mathcal{B}_i^+$ we must have that $(\mathcal{M}, v) \models a.\texttt{ann}$. Thus, since $v \in \text{Pre}(\mathcal{M}, a)$, we must have $(\mathcal{M}, v) \models a.\texttt{ann} \wedge a.\texttt{pre}$. Since $\mathcal{K}_i$ is reflexive, $(\mathcal{M}, v) \not\models \mathcal{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$. Thus $(\mathcal{M}, z) \models a.\texttt{ann}$, again by definition of $\mathcal{B}_i^+$. If $(\mathcal{M}, u) \models \mathbf{B}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$, note that $(u, z) \in \mathcal{K}_i$ by transitivity of $\mathcal{K}_i$ and thus $(u, z) \in \mathcal{R}_i$. Otherwise, $(u, v) \in \mathcal{B}_i$. By seriality of $\mathcal{B}_i$, there is some world w such that $(v, w) \in \mathcal{B}_i$. By transitivity of $\mathcal{B}_i$, we have that $(u, w) \in \mathcal{B}_i$, and by the Euclidean property of $\mathcal{B}_i$ we thus have that $(w, v) \in \mathcal{B}_i$. By the transitive property since $(v, w), (w, v) \in \mathcal{B}_i$, $(v, v) \in \mathcal{B}_i$. Since $v \in \text{Pre}(\mathcal{M}, a)$ and $(\mathcal{M}, v) \models a.\texttt{ann}$, it cannot be the case that $(\mathcal{M}, v) \models \neg(a.\texttt{ann} \wedge a.\texttt{pre})$. Thus, $(v, z) \in \mathcal{B}_i$. By transitivity of $\mathcal{B}_i$, we thus have $(u, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$. In either case, since $u, z \in \text{Pre}(\mathcal{M}, a)$ and $(\mathcal{M}, z) \models a.\texttt{ann}$ we have that $(u^+, z^+) \in \mathcal{B}_i^+$ as desired.

$\mathcal{B}_i$ Euclideanness Suppose $(u^+, v^+), (u^+, z^+) \in \mathcal{B}_i^+$, so $u, v, z \in \text{Pre}(\mathcal{M}, a)$, $(u, v), (u, z) \in \mathcal{R}_i$. Suppose that $(\mathcal{M}, u) \models \mathbf{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$. If there were $w \in W$ such that $(\mathcal{M}, w) \models (a.\texttt{ann} \wedge a.\texttt{pre})$ and $(v, w) \in \mathcal{B}_i^+$ then by transitivity of $\mathcal{B}_i$ $(u, w) \in \mathcal{B}_i^+$, which cannot occur because $(\mathcal{M}, u) \models \mathbf{B}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$. Thus it must be that $(\mathcal{M}, v) \models \mathbf{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$. If $(\mathcal{M}, u) \not\models \mathbf{B}_i \neg a.\texttt{pre}$, then $(u, v), (u, z) \in \mathcal{B}_i$ (since neither pair is in $\mathcal{R}_i^+$), so $(v, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$ as well. Otherwise, $(\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}$. In this case, suppose that for some $x \in W$ we had that $(\mathcal{M}, x) \models a.\texttt{pre}$ and $(v, x) \in \mathcal{B}_i$. Then by $KB2$ on $\mathcal{K}_i, \mathcal{B}_i$ we have that $(u, x) \in \mathcal{B}_i$, and thus $(\mathcal{M}, u) \not\models \mathbf{B}_i \neg a.\texttt{pre}$, a contradiction. Thus $(\mathcal{M}, v) \models \mathbf{B}_i \neg a.\texttt{pre}$, and thus $\{(v, v') : (v, v') \in \mathcal{K}_i\} \subseteq \mathcal{R}_i$. By Euclideanness of $\mathcal{K}_i$, $(v, z) \in \mathcal{K}_i$, so $(v, z) \in \mathcal{R}_i$. In either case, because $v, z \in \text{Pre}(\mathcal{M}, a)$ and $(\mathcal{M}, v) \models \mathbf{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$ we have that $(v^+, z^+) \in \mathcal{B}_i^+$ as desired. Now suppose that $(\mathcal{M}, u) \not\models \mathbf{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$ (call this C1) but $(\mathcal{M}, u) \models \mathbf{B}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$ (call this C2). C1 implies that there is some $w \in W$ such that $(u, v) \in \mathcal{K}_i$ and $(\mathcal{M}, w) \models a.\texttt{ann} \wedge a.\texttt{pre}$. Since $\mathcal{K}_i$ is Euclidean, $(v, w) \in \mathcal{K}_i$, and thus $(\mathcal{M}, v) \not\models \mathbf{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$. Suppose there was some $x \in W$ such that $(v, x) \in \mathcal{B}_i$ and $(\mathcal{M}, x) \models a.\texttt{pre}$. Then by transitivity of $\mathcal{B}_i$ we have that $(u, x) \in \mathcal{B}_i$ and thus that C2 is false, a contradiction. Thus $(\mathcal{M}, v) \models \mathcal{B}_i \neg a.\texttt{pre}$, and (since C1 and C2 hold at v) we must have $(v, z) \in \mathcal{R}_i$. Since C1 is true and $(u, z) \in \mathcal{B}_i^+$, it must be that $(\mathcal{M}, z) \models a.\texttt{ann}$. Thus we have that since $v, z \in \text{Pre}(\mathcal{M}, a)$, $(v^+, z^+) \in \mathcal{R}_i$. Otherwise, $(\mathcal{M}, u) \not\models \mathbf{B}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$. By $KB1$ on $\mathcal{K}_i, \mathcal{B}_i$ it follows that $(\mathcal{M}, u) \not\models \mathbf{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$, and thus (by definition of $\mathcal{B}_i^+$) that $(\mathcal{M}, z) \models a.\texttt{ann}$. Because $(\mathcal{M}, u) \not\models \mathbf{B}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$, by definition of $\mathcal{R}_i$ we must have that $(u, v), (u, z) \in \mathcal{B}_i$. By Euclideanness of $\mathcal{B}_i$, we have that $(v, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$ and since $v, z \in \text{Pre}(\mathcal{M}, a)$ and $(\mathcal{M}, z) \models a.\texttt{ann}$ it follows that $(v^+, z^+) \in \mathcal{B}_i^+$ as desired.

KB1 Suppose $(u^+, v^+) \in \mathcal{B}_i^+$. Then $u, v \in \text{Pre}(\mathcal{M}, a)$ and $(u, v) \in \mathcal{R}_i \subseteq \mathcal{K}_i$. Thus $(u^+, v^+) \in \mathcal{K}_i^+$ as desired.

KB2 Suppose that $(u^+, v^+) \in \mathcal{K}_i^+, (v^+, z^+) \in \mathcal{B}_i^+$. Then $u, v, z \in \text{Pre}(\mathcal{M}, a)$, $(u, v) \in \mathcal{K}_i$, and $(v, z) \in \mathcal{R}_i$. Suppose that $(\mathcal{M}, u) \models \mathbf{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$ (call this C1), and that $(\mathcal{M}, u) \models \mathbf{B}_i \neg a.\texttt{pre}$ (call this C2). By transitivity of $\mathcal{K}_i$, it follows that $(u, z) \in \mathcal{K}_i$. Since $\{(u, w) : (u, w) \in \mathcal{K}_i \wedge \text{C2 holds}\} \subseteq \mathcal{R}_i$, it follows that $(u, z) \in \mathcal{R}_i$. If C1 holds but C2 does not, then there is $w \in W$ such that $(u, v) \in \mathcal{B}_i$ and $(\mathcal{M}, w) \models a.\texttt{pre}$. Since $\mathcal{B}_i$ is Euclidean, $(v, w) \in \mathcal{B}_i$, and thus $(\mathcal{M}, v) \not\models \mathbf{B}_i \neg a.\texttt{pre}$. Suppose there was some $x \in W$ such that $(v, x) \in \mathcal{K}_i$ and $(\mathcal{M}, x) \models a.\texttt{ann} \wedge a.\texttt{pre}$. Then by the transitivity of $\mathcal{K}_i$ we have that $(u, x) \in \mathcal{K}_i$ and thus that C1 is false, a contradiction. Thus $(\mathcal{M}, v) \models \mathcal{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$, and (since C1 holds at v and C2 does not) we must have $(v, z) \in \mathcal{B}_i$. By $KB2$ on $\mathcal{K}_i, \mathcal{B}_i$, we must have that $(u, z) \in \mathcal{B}_i \subseteq \mathcal{R}_i$. In either case, since C1 holds and $u, z \in \text{Pre}(\mathcal{M}, a)$, it follows that $(u^+, z^+) \in \mathcal{B}_i^+$ as desired. Otherwise, $(\mathcal{M}, u) \not\models \mathbf{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$ (call this C3), and there exists some world w such that $(u, w) \in \mathcal{K}_i$ and $(\mathcal{M}, w) \models a.\texttt{ann} \wedge a.\texttt{pre}$. By Euclideanness of $\mathcal{K}_i$, $(v, w) \in \mathcal{K}_i$. Thus $(\mathcal{M}, v) \not\models \mathbf{K}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$. Thus, since $(v, z) \in \mathcal{B}_i^+$, it must be true that $(\mathcal{M}, z) \models a.\texttt{ann}$. If $(\mathcal{M}, u) \models \mathbf{B}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$ (call this C4) then since $(u, v) \in \mathcal{K}_i$ we have by C4 that $(u, v) \in \mathcal{R}_i^+$. By transitivity of $\mathcal{K}_i$, $(u, z) \in \mathcal{K}_i$ and since C3 and C4 continue to hold, $(u, z) \in \mathcal{R}_i$. Otherwise, we can apply the same logic as before to show that $(\mathcal{M}, v) \not\models \mathbf{B}_i \neg(a.\texttt{ann} \wedge a.\texttt{pre})$, and so it must be that $(v, z) \in \mathcal{B}_i$. By $KB2$ on $\mathcal{K}_i, \mathcal{B}_i$, it follows that $(u, z) \in \mathcal{B}_i \subset \mathcal{R}_i$. In either case, because $(u, z) \in \mathcal{R}_i$, $(\mathcal{M}, z) \models a.\texttt{ann}$, and $u, z \in \text{Pre}(\mathcal{M}, a)$, it follows that $(u^+, z^+) \in \mathcal{B}_i^+$ as desired.

7.7 Announcement Actions, $i \in P \cup O$

Updates to $\mathcal{K}_i$ and $\mathcal{B}_i$ for announcement actions with $i \in P$ and $i \in O$ are identical to those for ontic actions with $i \in P$ and $i \in O$, respectively (Sections 7.2 and 7.3).

Acknowledgements

This work was in part funded by AFOSR grant number FA9550-18-1-0465 and NASA grant number C17-2D00-TU.

References

Andersen, M.; Bolander, T.; and Holm Jensen, M. 2013. Don't plan for the unexpected: Planning based on plausibility models. *Logique et Analyse* 58.

Baltag, A., and Smets, S. 2008a. Probabilistic dynamic belief revision. *Synthese* 165(2):179.

Baltag, A., and Smets, S. 2008b. A qualitative theory of dynamic interactive belief revision. *Logic and the foundations of game and decision theory (LOFT 7)* 3:9–58.

Baral, C.; Gelfond, G.; Pontelli, E.; and Son, T. C. 2015. An action language for multi-agent domains: Foundations. *CoRR* abs/1511.01960.

Baral, C.; Bolander, T.; van Ditmarsch, H.; and McIlrath, S. 2017. Epistemic Planning (Dagstuhl Seminar 17231). *Dagstuhl Reports* 7(6):1–47.

Baral, C.; Gelfond, G.; Pontelli, E.; and Son, T. C. 2019. An action language for multi-agent domains. *CoRR*.

Bolander, T. 2017. A gentle introduction to epistemic planning: The DEL approach. *Electronic Proceedings in Theoretical Computer Science* 243:1–22.

Ditmarsch, H. P. V. 2005. Prolegomena to dynamic logic for belief revision. *Synthese* 147(2):229–275.

Gerbrandy, J. 1997. Dynamic epistemic logic.

Hintikka, J. 1962. Knowledge and belief: An introduction to the logic of the two notions. *Journal of Symbolic Logic* 29(3):132–134.

Huang, X.; Fang, B.; Wan, H.; and Liu, Y. 2017. A general multi-agent epistemic planner based on higher-order belief change. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, IJCAI'17, 1093–1101. AAAI Press.

Kraus, S., and Lehmann, D. 1988. Knowledge, belief and time. *Theoretical Computer Science* 58(1-3):155–174.

Le, T.; Fabiano, F.; Son, T.; and Pontelli, E. 2018. Efp and pg-efp: Epistemic forward search planners in multi-agent domains. In *Proceedings of the twenty-eighth international conference on automated planning and scheduling*.

Miller, T., and Muise, C. J. 2016. Belief update for proper epistemic knowledge bases. In *IJCAI*, 1209–1215.

Son, T. C.; Pontelli, E.; Baral, C.; and Gelfond, G. 2014. Finitary s5-theories. In Fermé, E., and Leite, J., eds., *Logics in Artificial Intelligence*, 239–252. Cham: Springer International Publishing.

van Benthem, J., and Smets, S. 2015. Dynamic logic of belief change. In van Ditmarsch, H.; Halpern, J.; van der Hoek, W.; and Kooi, B., eds., *Handbook of Logics for Knowledge and Belief*. College Publications. chapter 7, 299–368.