

Multi-Agent Epistemic Planning with Inconsistent Beliefs, Trust and Lies

Supplementary Documentation

Author 1¹, Author 2²

Department 1, University 1, Address 1
 authour1@affiliation1.com
 Department 2, University 2, Address 2
 authour2@affiliation2.com

The following document provides supplementary information for the paper “Multi-Agent Epistemic Planning with Inconsistent Beliefs, Trust and Lies” submitted to 18th *The Pacific Rim International Conference on Artificial Intelligence* (PRICAI-2021).

We will start the presentation with some preliminary definitions in Section A. Then, in Section B we will demonstrate the properties of the transition function of the epistemic actions with attitudes, listed in the paper in Proposition 1. Finally, in Section C, we will present several examples of execution of the newly introduced transition function. These examples will show how the e-states structure is modified after the update.

A Preliminary Definitions

Before starting with the demonstrations we need to introduce some terminology that will help us avoid unnecessary clutter during the proofs. In particular, let a Domain \mathcal{D} , a $\mathbf{p} \in \mathcal{S}$ where \mathcal{S} is the set of all the possibilities reachable from $\mathcal{D}(\varphi_i)$ with a finite sequence of action instances and the set of agents $\alpha \subseteq \mathcal{D}(\mathcal{AG})$ be given. The operator $\mathcal{B}_\alpha^{\mathbf{p}}$ captures *all the reachable possibilities for α given a starting possibility \mathbf{p}* .

Let us describe now how this operator can be used to represent the notions of 1) agents’ belief; 2) common belief; and 3) nested beliefs.

A.1 Agents Beliefs Representation

To link the operator introduced above with the concept of belief let us start with the case where the set of agents α contains only one element i , *i.e.*, $\alpha = \{i\}$. We, therefore, use $\mathcal{B}_i^{\mathbf{p}}$ to identify the set of all the possibilities that i , starting from the possibility \mathbf{p} , cannot distinguish.

The construction of the set identified by $\mathcal{B}_i^{\mathbf{p}}$ is procedural and it is done by applying the operator $(\mathcal{B}_i^{\mathbf{p}})^k$, with $k \in \mathbb{N}$, until a *fixed point* is found. The operator $(\mathcal{B}_i^{\mathbf{p}})^k$ is defined as follows:

$$(\mathcal{B}_i^{\mathbf{p}})^k = \begin{cases} \mathbf{p}(i) & \text{if } k = 0 \\ \{\mathbf{q} \mid (\exists \mathbf{u} \in (\mathcal{B}_i^{\mathbf{p}})^{k-1})(\mathbf{q} \in \mathbf{u}(i))\} & \text{if } k \geq 1 \end{cases}$$

Finally, we can define $\mathcal{B}_i^{\mathbf{p}} = \bigcup_{k \geq 1} (\mathcal{B}_i^{\mathbf{p}})^k$. It is easy to see that this is equivalent to the set of possibilities reached by the operator \mathbf{B}_i starting from \mathbf{p} and, therefore, that it represents the beliefs of i in \mathbf{u} .

Let us note that fixed point of the succession $(\mathcal{B}_\alpha^{\mathbf{p}})^k$ is reached in finite iterations. This is because:

- $(\mathcal{B}_\alpha^S)^k$ is monotonic; namely that $(\mathcal{B}_\alpha^S)^k \subseteq (\mathcal{B}_\alpha^S)^{k+1}$ with $k \in \mathbb{N}$ (Lemma 1); and
- the set \mathcal{S} of all the possibilities reached by applying a finite action instances sequence Δ to a given possibility \mathfrak{p} has a finite number of elements (Proposition 1).

A.2 Common Belief Representation

Now, similarly to the single-agent case, we can define the set $\mathcal{B}_\alpha^{\mathfrak{p}}$. This represents the *common belief* of α (\mathbf{C}_α) starting from \mathfrak{p} . As before we introduce the operator $(\mathcal{B}_\alpha^{\mathfrak{p}})^k$ of which the fixed point will result in $\mathcal{B}_\alpha^{\mathfrak{p}}$.

$$(\mathcal{B}_\alpha^{\mathfrak{p}})^k = \begin{cases} \bigcup_{i \in \alpha} \mathfrak{p}(i) & \text{if } k = 0 \\ \{\mathfrak{q} \mid (\exists \mathfrak{u} \in (\mathcal{B}_\alpha^{\mathfrak{p}})^{k-1})(\mathfrak{q} \in \bigcup_{i \in \alpha} \mathfrak{u}(i))\} & \text{if } k \geq 1 \end{cases}$$

A.3 Nested Knowledge Representation

We can also express the concept of *nested knowledge* in a more compact way. Let two sets of agents $\alpha_1 \subseteq \mathcal{D}(\mathcal{AG}), \alpha_2 \subseteq \mathcal{D}(\mathcal{AG})$ be given; the set of possibilities reachable by applying $\mathbf{C}_{\alpha_1} \mathbf{C}_{\alpha_2}$ starting from \mathfrak{p} is:

$$\mathcal{B}_{\alpha_1, \alpha_2}^{\mathfrak{p}} = \{\mathfrak{q} \mid (\exists r \in \mathcal{B}_{\alpha_1}^{\mathfrak{p}})(\mathfrak{q} \in \mathcal{B}_{\alpha_2}^r)\}$$

Let us note that, when α_1 or α_2 contains only one agent i , \mathbf{C}_i and \mathbf{B}_i are equal.

Lemma 1 (Operator \mathcal{B}_α^S monotony). *The sequence (\mathcal{B}_α^S) is monotonic; meaning that, for every $k \in \mathbb{N}$, $(\mathcal{B}_\alpha^S)^k \subseteq (\mathcal{B}_\alpha^S)^{k+1}$.*

Proof. Without losing generality let a possibility \mathfrak{p} and an agent i be given. To demonstrate the monotonicity of $(\mathcal{B}_i^{\mathfrak{p}})$ we start by recalling that:

$$(\mathcal{B}_i^{\mathfrak{p}})^k = \{\mathfrak{q} \mid (\exists \mathfrak{u} \in (\mathcal{B}_i^{\mathfrak{p}})^{k-1})(\mathfrak{q} \in \mathfrak{u}(i))\}.$$

By construction each possibility respects the **KD45** logic (Table 1) and, therefore, some structural constraints. In particular, to comply with axioms **4** and **5**, if a possibility $\mathfrak{q} \in \mathfrak{p}(i)$ then $\mathfrak{q} \in \mathfrak{q}(i)$. In terms of our sequence, this translates into *if a possibility $\mathfrak{q} \in (\mathcal{B}_i^{\mathfrak{p}})^{k-1}$ then $\mathfrak{q} \in (\mathcal{B}_i^{\mathfrak{p}})^k$* . \square

It is easy to see that this property ensures that the agent’s reachability function respect introspection. That is; when an agent reaches \mathfrak{q} she has to ‘know’ that herself considers \mathfrak{q} possible. Thanks to this property we can now infer that each iteration of the sequence $(\mathcal{B}_i^{\mathfrak{p}})^k$ contains at least $(\mathcal{B}_i^{\mathfrak{p}})^{k-1}$ and, therefore, that the sequence (\mathcal{B}_α^S) is monotonic.

Proposition 1 (States Size Finiteness). *Given a finite action instances sequence Δ —namely a plan—and a starting point \mathfrak{p} —namely a plan—with a finite number of possible worlds, i.e., $|\bigcup_{\text{ag} \in \mathcal{D}(\mathcal{AG})} \mathfrak{p}(\text{ag})| = n$, the set \mathcal{S} of all the possibilities generated by applying Δ to \mathfrak{p} has a finite number of elements.*

Property of \mathcal{B}	Axiom
$(\mathbf{B}_i\varphi \wedge \mathbf{B}_i(\varphi \Rightarrow \psi)) \Rightarrow \mathbf{B}_i\psi$	K
$\neg\mathbf{B}_i\perp$	D
$\mathbf{B}_i\varphi \Rightarrow \mathbf{B}_i\mathbf{B}_i\varphi$	4
$\neg\mathbf{B}_i\varphi \Rightarrow \mathbf{B}_i\neg\mathbf{B}_i\varphi$	5

Table 1. KD45 axioms [2].

Proof. Following the definition of the transition function of $m\mathcal{A}^p$ (Definition 11 of [1] and Definition 6 of the presented paper) we can determine an upper bound for the number of new possibilities generated after the application of an action instance and, moreover, of an action instance sequence. In particular from a given possibility \mathbf{p} such that $|\mathcal{B}_\alpha^{\mathbf{p}}| = n$ (where α is the set of all the agents) the cardinality of the set $\mathcal{B}_\alpha^{\mathbf{p}'}$ will be, at most, equal to $3n$. That is because:

- when an *ontic* action is executed each possibility $\in |\mathcal{B}_\alpha^{\mathbf{p}}|$ can be either updated—if reached by a fully observant agent—or kept unchanged—if reached by an oblivious agent. This means that an upper bound to the size of $\mathcal{B}_\alpha^{\mathbf{p}'}$ in case of an ontic action execution is $2n$ where only the updated possibilities (n) are new elements of \mathcal{S} .
- The case with *sensing* and *announcement* actions is similar

This identifies $3n$ as upper bound for the growth of a state size and for the generation of new possibilities after an action execution. Therefore given the size n of the initial state and the length of the action sequence l we can conclude that $|\mathcal{S}| \leq (n \times 3^l)$ and it is indeed finite. \square

B Epistemic Actions Transition Function

In what follows we will demonstrate that the transition function for epistemic actions, introduced in Definition 6 of the paper, respects the properties listed in the manuscript (Proposition 1). Before starting the demonstrations, for the sake of readability, let us re-introduce this transition function.

Let a domain \mathcal{D} , its set of action instances $\mathcal{D}(\mathcal{AI})$, and the set \mathcal{S} of all the possibilities reachable from $\mathcal{D}(\varphi_i)$ with a finite sequence of action instances be given.

In the case of sensing actions we do not consider any specific executor j , as we assume that all the fully observant agents are executing the actions. Each agent trusts her senses, *i.e.*, $\mathbf{F}_a = \mathbf{T}_a$. Similarly, we assume partially observant agents to keep their beliefs about the physical world unchanged, *i.e.*, $\mathbf{P}_a = \mathbf{I}_a$. Hence, the refined frame of reference of sensing actions is $\rho_a = \langle \mathbf{T}_a, \mathbf{I}_a, \mathbf{O}_a \rangle$. In the case of announcement actions, specifying the executor $j \in \mathcal{D}(\mathcal{AG})$ and the attitudes is necessary to resolve inconsistent beliefs. Therefore, the frame of reference of announcement actions is $\rho_a = \langle (\{j\}, \mathbf{T}_a, \mathbf{M}_a, \mathbf{S}_a), (\mathbf{I}_a, \mathbf{D}_a), \mathbf{O}_a \rangle$. In the course of the e-state update the attitude of j depends on the point of view of the agent currently handled by the transition function. In particular, as mentioned before, the executor considers herself stubborn trustful and stubborn agents believe that the announcer is telling the truth, while mistrustful agents will consider the executor to be lying. Let us observe that the announcer is aware of the other agents' perspective on her attitude and so are the remaining agents w.r.t. each other.

Let ℓ be the (unique) fluent literal such that $[a \text{ senses/announces } \ell] \in \mathcal{D}$. With a slight abuse of notation, we define the *value* of ℓ in a possibility \mathbf{w} as $val(a, \mathbf{w}) = \mathbf{w}(\ell)$. The *effect* of action a is $e(a) = 1$ if ℓ is a positive fluent literal ($e(a) = 0$, otherwise). We use the following notations and

assumptions for the sake of readability: given a possibility p , 1) p' indicates the updated version of p ; 2) $p(\mathcal{F}) = \{f \mid f \in \mathcal{D}(\mathcal{F}) \wedge u \models f\} \cup \{\neg f \mid f \in \mathcal{D}(\mathcal{F}) \wedge u \not\models f\}$. 3) if not stated otherwise, we consider $p'(\mathcal{F}) = p(\mathcal{F})$.

Definition 7 (Transition function for epistemic actions). Let $i \in \mathcal{D}(\mathcal{AG})$. Applying an epistemic action instance a on the pointed possibility u results in the new pointed world $\Phi(a, u) = u'$ such that:

$$u'(i) = \begin{cases} u(i) & \text{if } i \in O_a \\ P(a, u) & \text{if } i \in P_a \\ F(a, u, 1) & \text{if } i \in T_a \\ F(a, u, 0) & \text{if } i \in M_a \\ S(a, u, e(a), 1) & \text{if } i \in S_a \\ S(a, u, e(a), 0) & \text{if } i = j \end{cases}$$

where P, F, S are defined below.

Let $w'_x = \chi(a, w, x)$ and $\bar{w}'_x = \bar{\chi}(a, w, \bar{x})$ where: i) w'_x and \bar{w}'_x represent the possibility w updated with χ and $\bar{\chi}$, respectively; and ii) x and \bar{x} represent opposite boolean values s.t. $x = \neg\bar{x}$. Moreover, let the boolean variable b be 1 and 0 when executing χ and $\bar{\chi}$, respectively. Then w'_x and \bar{w}'_x are defined as follows:

$$w'_x(\ell) = \begin{cases} x & \text{if } \ell = f \\ u(\ell) & \text{otherwise} \end{cases} \quad \bar{w}'_x(\ell) = \begin{cases} \bar{x} & \text{if } \ell = f \\ u(\ell) & \text{otherwise} \end{cases}$$

$$\left. \begin{array}{l} w'_x(i) \\ \bar{w}'_x(i) \end{array} \right\} = \begin{cases} w(i) & \text{if } i \in O_a \\ P(a, w) & \text{if } i \in P_a \\ \bigcup_{v \in w(i)} \chi(a, v, x) & \text{if } i \in T_a \vee (i = j \wedge b = 1) \\ \bigcup_{v \in w(i)} \bar{\chi}(a, v, \bar{x}) & \text{if } i \in M_a \vee (i = j \wedge b = 0) \\ S(a, w, x, 1) & \text{if } i \in S_a \end{cases}$$

1) Let $w'_p = P(a, w)$. Then:

$$w'_p(i) = \begin{cases} w(i) & \text{if } i \in O_a \\ \bigcup_{v \in w(i)} P(a, v) & \text{if } i \in I_a \\ \bigcup_{v \in w(i)} \chi(a, v, 0) \cup \chi(a, v, 1) & \text{if } i \in D_a \\ \bigcup_{v \in w(i)} \chi(a, v, val(a, v)) & \text{if } i \in T_a \cup M_a \cup \{j\} \\ \bigcup_{v \in w(i)} S(a, v, val(a, v), 1) & \text{if } i \in S_a \end{cases}$$

2) Let $w'_f = F(a, w, b)$. Then:

$$w'_f(i) = \begin{cases} w(i) & \text{if } i \in O_a \\ P(a, w) & \text{if } i \in P_a \\ \bigcup_{v \in w(i)} \chi(a, v, e(a)) & \text{if } i \in T_a \vee (i = j \wedge b = 1) \\ \bigcup_{v \in w(i)} \bar{\chi}(a, v, \neg e(a)) & \text{if } i \in M_a \vee (i = j \wedge b = 0) \\ \bigcup_{v \in w(i)} S(a, v, e(a), 1) & \text{if } i \in S_a \end{cases}$$

3) Let $w'_s = S(a, w, x, s)$. Then:

$$w'_s(i) = \begin{cases} w(i) & \text{if } i \in O_a \\ P(a, w) & \text{if } i \in P_a \\ \bigcup_{v \in w(i)} \chi(a, v, x) & \text{if } i \in T_a \vee (i = j \wedge s = 1) \\ \bigcup_{v \in w(i)} \bar{\chi}(a, v, \neg x) & \text{if } i \in M_a \\ \bigcup_{v \in w(i)} S(a, v, x, s) & \text{if } i \in S_a \vee (i = j \wedge s = 0) \end{cases}$$

B.1 Desired properties of the epistemic actions update

Proposition 2 (Epistemic Actions properties). Let $a(j)$ be an epistemic action instance where j **announces** ℓ (where ℓ can either be f or $\neg f$). Let u be an e -state and let u' be its updated version (that is, $\Phi(a, u) = u'$), then in our updated transition function it holds that:

1. $u' \models C_{F_a}(C_{T_a}(\ell \wedge B_j \ell))$;
2. $u' \models C_{F_a}(C_{M_a}(\neg \ell \wedge B_j \neg \ell))$;
3. $\forall i \in (S_a \cup \{j\})$, $u' \models \varphi$ if $u \models \varphi$ with $\varphi \in \{B_i \ell; B_i \neg \ell; (\neg B_i \ell \wedge \neg B_i \neg \ell)\}$;
4. $\forall i \in F_a$, $u' \models C_{P_a}(B_i \ell \vee B_i \neg \ell)$;
5. $\forall i \in D_a$, $u' \models C_{(F_a \cup P_a \cup \{j\})}(\neg B_i \ell \wedge \neg B_i \neg \ell)$;
6. for every pair of agents $i \in \mathcal{D}(\mathcal{AG})$ and $o \in O_a$, and a belief formula φ , $u' \models B_i(B_o \varphi)$ if $u \models B_i(B_o \varphi)$.

Let us demonstrate each point separately. Let us assume that a is “ j **announces** f ”. The case when “ j **announces** $\neg f$ ” is similar, and we will only highlight the differences when it is needed.

1. In the following we demonstrate Property 1.

- First of all we identify the set of all the possibilities reached by the *fully observant* agents in u as $\mathcal{B}_{F_a}^u$.

- We then re-apply the reachability function following the beliefs of the *trustful* agents. This means that the set of beliefs of the *trustful*, from the point of view of the *fully observant* ones, is represented by the set $\mathcal{B}_{\mathbf{F}_a, \mathbf{T}_a}^u = \{\mathbf{p} \mid \mathbf{p} \in \mathcal{B}_{\mathbf{T}_a}^q \wedge \mathbf{q} \in \mathcal{B}_{\mathbf{F}_a}^u\}$.
- Now to calculate $\mathcal{B}_{\mathbf{F}_a, \mathbf{T}_a}^{u'}$, following Definition 7, we apply “ $\chi(\mathbf{a}, \mathbf{p}, 1)$ ” to every element \mathbf{p} of $\mathcal{B}_{\mathbf{F}_a, \mathbf{T}_a}^u$. Let us note that if $e(\mathbf{a}) = 0$, that is if \mathbf{j} **announces** $\neg f$, we should apply “ $\chi(\mathbf{a}, \mathbf{p}, 0)$ ” instead.
- This means that the set of updated beliefs of *trustful* agents, from the point of view of the *fully observant* ones, is represented by the set $\mathcal{B}_{\mathbf{F}_a, \mathbf{T}_a}^{u'} = \{\mathbf{p}' \mid \mathbf{p}'(\mathcal{F}) = ((\mathbf{p}(\mathcal{F}) \setminus \{\neg f\}) \cup \{f\}) \wedge \mathbf{p} \in \mathcal{B}_{\mathbf{F}_a, \mathbf{T}_a}^u\}$. It is important to notice that the truth value of the fluent f in the set of possibilities $\mathcal{B}_{\mathbf{F}_a}^u$ is not important as the application of “ $\chi(\mathbf{a}, \mathbf{p}, 1)$ ” on all these possibilities forces their updated version to set the truth value of $f = 1$ (similarly, for the negated case, the fluent truth value is 0).
- It is then straightforward to see that the set $\mathcal{B}_{\mathbf{F}_a, \mathbf{T}_a}^{u'}$ entails f , as all the reached possibility have the truth value of f set to 1. Recalling that, as shown in Section A.3, the set $\mathcal{B}_{\alpha, \beta}^u$ corresponds to the possibilities reached by $\mathbf{C}_\alpha(\mathbf{C}_\beta)$ where $\alpha, \beta \subseteq \mathcal{D}(\mathcal{AG})$ it is clear that the updated e-state $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{T}_a}(f))$ (and similarly, in the negated case, $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{T}_a}(\neg f))$).
- Now, to demonstrate that $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{T}_a}(\mathbf{B}_j(f)))$ we need to recall that the *trustful* agents consider that the announcer \mathbf{j} to be *trustful* as well. This means that $\mathcal{B}_{\mathbf{F}_a, \mathbf{T}_a, \{j\}}^{u'}$ is equal to $\mathcal{B}_{\mathbf{F}_a, \mathbf{T}_a}^{u'}$ as the *trustful* agents believes that the announcer \mathbf{j} used “ $\chi(\mathbf{a}, \mathbf{p}, 1)$ ” to update her beliefs (being, from their perspective *trustful* agent). This means that all the possibility in $\mathcal{B}_{\mathbf{F}_a, \mathbf{T}_a, \{j\}}^{u'}$ have the truth value of f set to 1 (or 0 in the negated case).
- Following Section A.3 we know that $\mathcal{B}_{\mathbf{F}_a, \mathbf{T}_a, \{j\}}^{u'}$ is equal to the possibilities reached by applying $\mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{T}_a}(\mathbf{B}_j))$. Given that all these possibilities have the truth value of f set to 1 it is straightforward to see that $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{T}_a}(\mathbf{B}_j(f)))$.
- From the previous points now know that $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{T}_a}(f)) \wedge \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{T}_a}(\mathbf{B}_j(f)))$ and therefore that $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{T}_a}(f \wedge \mathbf{B}_j(f)))$ as stated in Property 1 (while for the negated case we can easily derive that $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{T}_a}(\neg f \wedge \mathbf{B}_j(\neg f)))$).

2. Let us proceed with Property 2.

- First we identify the set of all the possibilities reached by the *fully observant* agents in u as $\mathcal{B}_{\mathbf{F}_a}^u$.
- We then re-apply the reachability function following the beliefs of the *mistrustful* agents. This means that the set of beliefs of the *mistrustful*, from the point of view of the *fully observant* ones, is represented by the set $\mathcal{B}_{\mathbf{F}_a, \mathbf{M}_a}^u = \{\mathbf{p} \mid \mathbf{p} \in \mathcal{B}_{\mathbf{M}_a}^q \wedge \mathbf{q} \in \mathcal{B}_{\mathbf{F}_a}^u\}$.
- Now to calculate $\mathcal{B}_{\mathbf{F}_a, \mathbf{M}_a}^{u'}$, following Definition 7, we apply “ $\chi(\mathbf{a}, \mathbf{p}, 0)$ ” to every element \mathbf{p} of $\mathcal{B}_{\mathbf{F}_a, \mathbf{M}_a}^u$. Let us note that if $e(\mathbf{a}) = 0$, that is if \mathbf{j} **announces** $\neg f$, we should apply “ $\chi(\mathbf{a}, \mathbf{p}, 1)$ ” instead.

- This means that the set of updated beliefs of *mistrustful* agents, from the point of view of the *fully observant* ones, is represented by the set $\mathcal{B}_{\mathbf{F}_a, \mathbf{M}_a}^{u'} = \{\mathbf{p}' \mid \mathbf{p}'(\mathcal{F}) = ((\mathbf{p}(\mathcal{F}) \setminus \{\mathbf{f}\}) \cup \{\neg \mathbf{f}\}) \wedge \mathbf{p} \in \mathcal{B}_{\mathbf{F}_a, \mathbf{M}_a}^u\}$. It is important to notice that the truth value of the fluent \mathbf{f} in the set of possibilities $\mathcal{B}_{\mathbf{F}_a}^{u'}$ is not important as the application of “ $\chi(\mathbf{a}, \mathbf{p}, 0)$ ” on all these possibilities forces their updated version to set the truth value of $\mathbf{f} = 0$ (similarly, for the negated case, the fluent truth value is 1).
 - It is then straightforward to see that the set $\mathcal{B}_{\mathbf{F}_a, \mathbf{M}_a}^{u'}$ entails $\neg \mathbf{f}$, as all the reached possibility have the truth value of \mathbf{f} set to 0. Recalling that, as shown in Section A.3, the set $\mathcal{B}_{\alpha, \beta}^u$ corresponds to the possibilities reached by $\mathbf{C}_\alpha(\mathbf{C}_\beta)$ where $\alpha, \beta \subseteq \mathcal{D}(\mathcal{AG})$ it is clear that the updated e-state $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{M}_a}(\neg \mathbf{f}))$ (and similarly, in the negated case, $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{M}_a}(\mathbf{f}))$).
 - Now, to demonstrate that $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{M}_a}(\mathbf{B}_j(\neg \mathbf{f})))$ we need to recall that the *mistrustful* agents consider that the announcer j to be *mistrustful* as well. This means that $\mathcal{B}_{\mathbf{F}_a, \mathbf{M}_a, \{j\}}^{u'}$ is equal to $\mathcal{B}_{\mathbf{F}_a, \mathbf{M}_a}^{u'}$ as the *mistrustful* agents believes that the announcer j used “ $\chi(\mathbf{a}, \mathbf{p}, 0)$ ” to update her beliefs (being, from their perspective *mistrustful* agent). This means that all the possibility in $\mathcal{B}_{\mathbf{F}_a, \mathbf{M}_a, \{j\}}^{u'}$ have the truth value of \mathbf{f} set to 0 (or 1 in the negated case).
 - Following Section A.3 we know that $\mathcal{B}_{\mathbf{F}_a, \mathbf{M}_a, \{j\}}^{u'}$ is equal to the possibilities reached by applying $\mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{M}_a}(\mathbf{B}_j))$. Given that all these possibilities have the truth value of \mathbf{f} set to 0 it is straightforward to see that $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{M}_a}(\mathbf{B}_j(\neg \mathbf{f})))$.
 - From the previous points now know that $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{M}_a}(\neg \mathbf{f})) \wedge \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{M}_a}(\mathbf{B}_j(\neg \mathbf{f})))$ and therefore that $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{M}_a}(\neg \mathbf{f} \wedge \mathbf{B}_j(\neg \mathbf{f})))$ as stated in Property 1 (while for the negated case we can easily derive that $u' \models \mathbf{C}_{\mathbf{F}_a}(\mathbf{C}_{\mathbf{M}_a}(\mathbf{f} \wedge \mathbf{B}_j(\mathbf{f})))$).
3. To demonstrate Property 3 let us consider $i \in (\mathbf{S}_a \cup \{j\})$ and that u does entail φ (where $\varphi \in \{\mathbf{B}_i \ell; \mathbf{B}_i \neg \ell; (\neg \mathbf{B}_i \ell \wedge \neg \mathbf{B}_i \neg \ell)\}$). The case where $u \not\models \varphi$ is similar and, therefore, omitted.
- Let us start by recalling that the executor agent j consider herself as *stubborn*, given that announcing something should not affect her belief on what she has announced. This means that, to calculate the updated version of u' , agent j applies the sub-function \mathbf{S} as the *stubborn* agents do.
 - Now, being $i \in (\mathbf{S}_a \cup \{j\})$, we know from Definition 7 that the updated version of her reachable possibilities is represented by the set $\mathcal{B}_i^{u'} = \{\mathbf{p}' \mid \mathbf{p} \in \mathcal{B}_i^u \wedge \mathbf{p}' = \mathbf{S}(\mathbf{a}, \mathbf{u}, \ell, \mathbf{s})\}$ (The boolean value \mathbf{s} is either 1, if $i \in \mathbf{S}_a$, or 0, when $i = j$).
 - Following Definition 7 we know that each possibility in $\mathcal{B}_i^{u'}$ has the same fluent set of its previous version.
 - Moreover, we know that an *stubborn* agent preserves all the edges. In fact the unfolding of the execution of \mathbf{S} from u , when considered from an *stubborn* agent ag 's point of view, simply re-apply \mathbf{S} to all the possibilities in \mathcal{B}_{ag}^u . This means that if an agent reached a possibility \mathbf{q} from another possibility \mathbf{p} in u she will reach \mathbf{q}' from \mathbf{p}' in u' .

- From the last statement, and given that the updated version of each possibility maintains the same fluent set we can conclude that, if $u \models \varphi$ then $u' \models \varphi$ (similarly if $u \not\models \varphi$ then $u' \not\models \varphi$) with $\varphi \in \{\mathbf{B}_i \ell; \mathbf{B}_i \neg \ell; (\neg \mathbf{B}_i \ell \wedge \neg \mathbf{B}_i \neg \ell)\}$ and $i \in (\mathbf{S}_a \cup \{j\})$.

4. We identify the set of the possibilities reachable by *partial observants* agents with $\mathcal{B}_{\mathbf{P}_a}^u$. We also recall that this set is equal to $\mathbf{C}_{\mathbf{P}_a}$ in u .

- Now to calculate $\mathcal{B}_{\mathbf{P}_a}^{u'}$, following Definition 7, we apply “ $\mathbf{P}(a, \mathbf{p})$ ” to every element \mathbf{p} of $\mathcal{B}_{\mathbf{P}_a}^u$. This results in all the possibilities \mathbf{p}' of $\mathcal{B}_{\mathbf{P}_a}^{u'}$ to have the same fluent set of the corresponding possibility $\mathbf{p} \in \mathcal{B}_{\mathbf{P}_a}^u$.

- It is easy to identify two disjoint subsets $\mathcal{B}_{\mathbf{P}_a}^{u'_0}$ and $\mathcal{B}_{\mathbf{P}_a}^{u'_1}$ of $\mathcal{B}_{\mathbf{P}_a}^{u'}$ that contains only possibility such that:
 - $\mathcal{B}_{\mathbf{P}_a}^{u'_0} \not\models \ell$;
 - $\mathcal{B}_{\mathbf{P}_a}^{u'_1} \models \ell$;
 - $(\mathcal{B}_{\mathbf{P}_a}^{u'_0} \cup \mathcal{B}_{\mathbf{P}_a}^{u'_1}) = \mathcal{B}_{\mathbf{P}_a}^{u'}$; and
 - $(\mathcal{B}_{\mathbf{P}_a}^{u'_0} \cap \mathcal{B}_{\mathbf{P}_a}^{u'_1}) = \emptyset$.

- From these two sets, following Definition 7 we can now construct the sets $\mathcal{B}_{\mathbf{P}_a, i}^{u'_0}$ and $\mathcal{B}_{\mathbf{P}_a, i}^{u'_1}$, with $i \in (\mathbf{F}_a \cup \{j\})$, by applying the sub-functions $\chi(a, \mathbf{p}, 0) \forall \mathbf{p} \in \mathcal{B}_{\mathbf{P}_a}^{u'_0}$ and $\chi(a, \mathbf{p}, 1) \forall \mathbf{p} \in \mathcal{B}_{\mathbf{P}_a}^{u'_1}$ respectively. These two sets are simply the set of possibilities reachable from the *fully observant* agents (and the executor, considered *fully observant* by the *partially observants*) starting from $\mathcal{B}_{\mathbf{P}_a}^{u'_0}$ and $\mathcal{B}_{\mathbf{P}_a}^{u'_1}$ respectively.

- Let us note that *trustful*, *stubborn* and the executor are considered equally by the *partially observant* given that they do not identify a truth value but simply believe that the *fully observant* agents will know the truth value of the announced fluent.

- Given that the set $\mathcal{B}_{\mathbf{P}_a, i}^{u'_0}$ resulted from the application of the transition function from the point of view of *fully observant* agents, we know from Points 1 and 2 that for $\forall \mathbf{p} \in \mathcal{B}_{\mathbf{P}_a, i}^{u'_0}, \mathbf{p} \not\models \ell$.

- This implies that $\mathcal{B}_{\mathbf{P}_a, i}^{u'_0}$ reaches only possibilities where the interpretation of ℓ is false and, similarly, in $\mathcal{B}_{\mathbf{P}_a, i}^{u'_1}$ reaches only possibilities where the interpretation of ℓ is true.

- This means that $\mathcal{B}_{\mathbf{P}_a, i}^{u'_0} \models \neg \ell$ and $\mathcal{B}_{\mathbf{P}_a, i}^{u'_1} \models \ell$.

- It is easy to see, then, that $\mathcal{B}_{\mathbf{P}_a}^{u'_0} \models \mathbf{B}_i \neg \ell$ being $\mathcal{B}_{\mathbf{P}_a, i}^{u'_0} = \{\mathbf{p} \mid \mathbf{p} \in \bigcup_{q \in \mathcal{B}_{\mathbf{P}_a}^{u'_0}} q(i)\}$ (and similarly

$$\mathcal{B}_{\mathbf{P}_a}^{u'_1} \models \mathbf{B}_i \ell).$$

- Finally, being $\mathcal{B}_{\mathbf{P}_a}^{u'} = \mathcal{B}_{\mathbf{P}_a}^{u'_0} \cup \mathcal{B}_{\mathbf{P}_a}^{u'_1}$ we can conclude that $\mathcal{B}_{\mathbf{P}_a}^{u'} \models \mathbf{B}_i \neg \ell \vee \mathbf{B}_i \ell^1$ and therefore $u' \models \mathbf{C}_{\mathbf{P}_a}(\mathbf{B}_i \neg \ell \vee \mathbf{B}_i \ell)$.

5. Let us now illustrate the proof for Property 5.

- First, to avoid unnecessary clutter let us use i) \mathbf{V}_a to indicate the set of the *observant* agents, *i.e.*, $\mathbf{V}_a = (\mathbf{F}_a \cup \mathbf{P}_a \cup \{\mathbf{j}\})$ and; ii) \mathbf{i} to indicate a *doubtful* agent, *i.e.*, $\mathbf{i} \in \mathbf{D}_a$.
- We then identify the set of all the possibilities reached by the *observant* agents in \mathbf{u} as $\mathcal{B}_{\mathbf{V}_a}^u$.
- Next, we re-apply the reachability function following the beliefs of the *doubtful* agents. This means that the set of beliefs of a *doubtful* agent \mathbf{i} , from the point of view of the *observant* ones, is represented by the set $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^u = \{\mathbf{p} \mid \mathbf{p} \in \mathcal{B}_{\mathbf{i}}^q \wedge \mathbf{q} \in \mathcal{B}_{\mathbf{i}}^u\}$.
- Now to calculate $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'}$, following Definition 7, we apply both “ $\chi(\mathbf{a}, \mathbf{p}, 0)$ ” and “ $\chi(\mathbf{a}, \mathbf{p}, 1)$ ” to every element \mathbf{p} of $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^u$.
- This means that the set of updated beliefs of a *doubtful* agent, from the point of view of the *observant* ones, is represented by the union of the sets $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'_0} = \{\mathbf{p}' \mid \mathbf{p}'(\mathcal{F}) = ((\mathbf{p}(\mathcal{F}) \setminus \{\mathbf{f}\}) \cup \{\neg \mathbf{f}\}) \wedge \mathbf{p} \in \mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^u\}$ and $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'_1} = \{\mathbf{p}' \mid \mathbf{p}'(\mathcal{F}) = ((\mathbf{p}(\mathcal{F}) \setminus \{\neg \mathbf{f}\}) \cup \{\mathbf{f}\}) \wedge \mathbf{p} \in \mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^u\}$. It is important to notice that the truth value of the fluent \mathbf{f} in the set of possibilities $\mathcal{B}_{\mathbf{V}_a}^u$ is not important as the application of “ $\chi(\mathbf{a}, \mathbf{p}, 0/1)$ ” on all these possibilities forces their updated version to set the truth value of $\mathbf{f} = 0/1$.
- As all the reached possibility from $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'_0}$ and $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'_1}$ have the truth value of \mathbf{f} set to 0 and 1 respectively we can easily derive that the former entails $\neg \mathbf{f}$, while the latter entails \mathbf{f} .
- Moreover, being $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'} = (\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'_0} \cup \mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'_1})$, we know that $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'} \not\models \mathbf{f}$ and $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'} \not\models \neg \mathbf{f}$. This is true because the subset $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'_0} \not\models \mathbf{f}$ while $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'_1} \not\models \neg \mathbf{f}$.
- Recalling that, as shown in Section A.3, the set $\mathcal{B}_{\alpha, \mathbf{ag}}^u$ corresponds to the possibilities reached by $\mathbf{C}_\alpha(\mathbf{B}_{\mathbf{ag}})$, where $\alpha \subseteq \mathcal{D}(\mathcal{AG})$ and $\mathbf{ag} \in \mathcal{D}(\mathcal{AG})$, it is clear that the set $\mathcal{B}_{\mathbf{V}_a, \mathbf{i}}^{u'}$ corresponds to the possibilities reached by $\mathbf{C}_{\mathbf{V}_a}(\mathbf{B}_i)$ starting from u' .
- Since the set identified in the last point can derive both \mathbf{f} and $\neg \mathbf{f}$, following entailment rules *ii* and *iv* of Definition 2 of the paper, we can infer that both $\mathbf{C}_{\mathbf{V}_a}(\neg \mathbf{B}_i \neg \mathbf{f})$ and $\mathbf{C}_{\mathbf{V}_a}(\neg \mathbf{B}_i \mathbf{f})$ hold.

6. Finally, let us demonstrate for Property 6.

¹ The two sets are completely disjoint as one only contains possibilities that entails ℓ while the other only possibilities that do not. This means that that does not exist any fully-observant-edge between possibilities that belongs in two different sets.

- When an agent $o \in \mathbf{O}_a$, from Definition 7, we know that $p'(o) = p(o)$. This means that, independently from the how a possibility p' has been updated, the point of view any oblivious agent o from p' is equal to the one that the point of view of o from p .
- This implies that, $\forall p' \in \mathcal{B}_i^u$ with $i \in \mathcal{D}(\mathcal{AG})$, $p'(o) = p(o)$ where $o \in \mathbf{O}_a$.
- This means that for every element in \mathcal{B}_i^u we have an updated version in $\mathcal{B}_i^{u'}$ that has the same reachability function for each oblivious agent o .
- Then, it is easy to see, that $\mathcal{B}_{i,o}^u = \mathcal{B}_{i,o}^{u'}$ and, therefore, that these two sets contain the same possibilities.
- Given that the two sets of possibilities are the same, it means that the reachability functions that they represent are the same. Being the two functions the same it means that given a belief formula φ , $u \models \mathbf{B}_i(\mathbf{B}_o\varphi)$ iff $u' \models \mathbf{B}_i(\mathbf{B}_o\varphi)$.
- Finally, we can conclude that if $u \models \mathbf{B}_i(\mathbf{B}_o\varphi)$ then $u' \models \mathbf{B}_i(\mathbf{B}_o\varphi)$.

C Examples of Execution

In this section we will show some examples of execution to better illustrate how the newly introduced transition function works. When needed we will describe the agents' attitudes. We will present, through labeled graphs, the e-state before and after the update for each example. While, to capture all the combinations of attitudes we would need a far larger number of examples, we decided to provide only those that show the fundamental attitudes behavior and interaction. We invite the interested reader to further investigate the attitudes' relations by “playing” with our planner² that offers the possibility to print out the graphical representation of the e-states before and after an action execution. All the examples will be based on a simple variation of a domain used in the literature [1, 3, 4]: **Coin in the Box**. The sense of this Section is to show the e-states updated rather than providing a significant use case. Before presenting the examples of execution let us introduce the domain that we will use: **Rigged Coin in the Box**.

Domain 1 (Rigged Coin in the Box) *Five agents, l, m, r, s, c, are inside a room containing a box. Agents l, m stand to the left of the box; r, s are at the box's right and c is positioned in front of the box. Inside the box, agent c placed a rigged coin. Any agent might peek (sensing action) inside the box to learn the coin position. Since the coin is rigged, the actual position (either tails or heads up) is only visible when an agent is standing in front of the box (i.e., c) or at its right (i.e., r and s). On the other hand, any agent that stands at the box's left (i.e., l and m) will always see the coin laying heads up. All the agents can share information through the action shout(position) (announcement action).*

THE INITIAL CONFIGURATION: For the sake of readability, in all the following examples, we will use the same initial e-state while varying the agent's attitudes and the executed action. Let us now explain the initial configuration and then illustrate the corresponding e-state, in Figure 1, that from now on will be identified with u . An explanation on how to “read” an e-state graphical representation is presented right after the following initial e-state description.

² Available upon request.

In our initial configuration we assume that is common belief that agents l , m , r and s do not know the coin position, *i.e.*, $\mathbf{C}_{\mathcal{D}(\mathcal{AG})}(\neg\mathbf{B}_i(\text{tails}) \wedge \neg\mathbf{B}_i(\neg\text{tails}))$ with $i \in \{l, m, r, s\}$. On the other hands, agent c is aware of the coin position and the other agents know this, that is, in our initial e-state it holds $\mathbf{C}_{\mathcal{D}(\mathcal{AG})}(\mathbf{B}_c(\text{tails}) \vee \mathbf{B}_c(\neg\text{tails}))$. Assuming that the coin position is tails up the graphical representation of this e-state is as follows.




Figure 1. The initial e-state u .

Before exploring the examples, let us briefly illustrate how to interpret the graphical representation of e-states. Consider Figure 1. The *bold-lined* world represents the actual world. If a world u is connected by an edge labeled with agent i to a world v , this means that in the world u agent i believes v to be possible.

In the initial state, each agent except c admits both the worlds where tails holds and where $\neg\text{tails}$ holds. This means that such agents are uncertain about the coin position. On the other hand, agent c does know the actual configuration of the coin. We can understand this because in the actual world c admits only the world where tails hold.

Finally, observe that the remaining agents do not know what c knows. In fact, the formula $\mathbf{B}_m(\mathbf{B}_c\text{tails} \vee \mathbf{B}_c\neg\text{tails})$ is true. On the other hand, the formula $\mathbf{B}_m\mathbf{B}_c\text{tails}$ does not hold in the initial state.

Example 1 (Correct Sensing). This examples shows how u is updated after the execution of the action instance $\text{peek}\langle\{r, s\}\rangle$. As said in Domain 1 both the agents r and s are able to correctly determine whether the coin lies tails or heads up. Since we are executing a sensing action we are only interested in defining the oblivious, the fully and the partially observant agents. In particular, for this action instance we assume r and s to be fully observant, l and m to be partially observant and c to be oblivious. As we can see in the resulting e-state (Figure 2) r and s believe that the coin lies tails up. Moreover l and m still not know the coin position but believe that r and s know it. Finally, being c oblivious, she did not change her beliefs about anything.




Figure 2. The e-state u' obtained after the execution of a correct sensing on u .

Example 2 (Wrong Sensing). This examples shows how u is updated after the execution of the action instance $\text{peek}(\{l, m\})$. As said in Domain 1 both the agents l and m always see the coin lying heads up. Since we are executing a sensing action we are only interested in defining the oblivious, the fully and the partially observant agents. In particular, for this action instance we assume l and m to be fully observant, r and s to be partially observant and c to be oblivious. As we can see in the resulting e-state (Figure 3) l and m believe that the coin lies heads up. Moreover r and s still not know the coin position but believe that l and m know it. Finally, being c oblivious, she did not change her beliefs about anything.




Figure 3. The e-state u' obtained after the execution of a wrong sensing on u .

Example 3 (Trustful & Mistrustful listeners of a true announcement). This examples shows how u is updated after the execution of the action instance $\text{shout}(c)$ where c announces tails . In particular, for this action instance we assume:

- c to be the *executor*;
- l to be *trustful*;
- m to be *mistrustful*;
- r to be *impassive*; and
- s to be *doubtful*;

As we can see in the resulting e-state (Figure 4) l and m believe that the coin lies tails and heads up, respectively. Moreover, l and m believe that c shares their beliefs on the coin position. Finally, agents r and s , still not know the coin position but believe that c , l and m know it.




Figure 4. The e-state u' obtained after the announcement of tails in u with trustful & mistrustful listeners.

Example 4 ((Mis)Trustful & Stubborn listeners and true announcement). This examples shows how u is updated after the execution of the action instance $\text{shout}(c)$ where c announces **tails**. Differently from the previous example, the agents' attitudes are as follows:

- c to be the *executor*;
- l to be *trustful*;
- m to be *mistrustful*;
- r to be *doubtful*; and
- s to be *stubborn*;

As we can see in the resulting e-state (Figure 5) agents c and l believes that the coin lies tails up while m thinks that it lies heads up. Even if s did not change her beliefs on the coin position she knows what c, l believe that the coin is tails up while m think that it is heads up. Agent r , instead still not know the coin position but believe that c, l and m know it. We will use a dotted square to indicate that the edges that reach such square, transitively reach all the worlds contained.




Figure 5. The e-state u' obtained after the announcement of **tails** in u with (mis)trustful & stubborn listeners.

Example 5 (Lie). This examples shows how u is updated after the execution of the action instance $\text{shout}(c)$ where c announces $\neg\text{tails}$. Let us note that this announcement, since is performed by c that believes tails , is a lie. For this action instance we assume:

- c to be the *executor*;
- l to be *trustful*;
- m to be *mistrustful*;
- r to be *doubtful*; and
- s to be *oblivious*;

As we can see in the resulting e-state (Figure 6) agent l believed to the lie and now has a wrong belief about the coin position. On the other hand and m did not believe the announcer and, therefore, now correctly think that the coin lies tails up. As for the executor, c knew that she was laying and, therefore, still believe the that the coin is tails up. Moreover, l and m believe that c shares their beliefs on the coin position. Agents r , still not know the coin position but believe that c , l and m know it. Finally, being s oblivious, she did not change her beliefs about anything.




Figure 6. The e-state u' obtained after the execution of a lie (*i.e.*, c shout $\neg\text{tails}$) in u with trustful & mistrustful listeners.

References

1. Fabiano, F., Burigana, A., Dovier, A., Pontelli, E.: EFP 2.0: A multi-agent epistemic solver with multiple e-state representations. In: Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling, Nancy, France, October 26-30, 2020. pp. 101–109. AAAI Press (2020), <https://aaai.org/ojs/index.php/ICAPS/article/view/6650>
2. Fagin, R., Halpern, J.Y.: Reasoning about knowledge and probability. *Journal of the ACM (JACM)* **41**(2), 340–367 (1994). <https://doi.org/10.1145/174652.174658>
3. Le, T., Fabiano, F., Son, T.C., Pontelli, E.: 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. pp. 161–170. AAAI Press, Delft, The Netherlands (June 24–29 2018), <https://aaai.org/ocs/index.php/ICAPS/ICAPS18/paper/view/17733>
4. Muise, C.J., Belle, V., Felli, P., McIlraith, S.A., Miller, T., Pearce, A.R., Sonenberg, L.: Planning over multi-agent epistemic states: A classical planning approach. In: Proc. of AAAI. pp. 3327–3334 (2015)