



HAL
open science

Action progression and revision in multiagent belief structures

Andreas Herzig, Jérôme Lang, Pierre Marquis

► **To cite this version:**

Andreas Herzig, Jérôme Lang, Pierre Marquis. Action progression and revision in multiagent belief structures. 6th Workshop on Nonmonotonic Reasoning, Action, and Change (NRAC 2005) @ IJCAI 2005, Aug 2005, Edinburgh, Scotland, Royaume-Uni. hal-03519760

HAL Id: hal-03519760

<https://hal.science/hal-03519760>

Submitted on 12 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Action progression and revision in multiagent belief structures

Andreas Herzig

IRIT - Université Paul Sabatier
31062 Toulouse Cedex, France
herzig@irit.fr

Jérôme Lang

IRIT - Université Paul Sabatier
31062 Toulouse Cedex, France
lang@irit.fr

Pierre Marquis

CRIL - Université d'Artois
62307 Lens Cedex, France
marquis@cril.univ-artois.fr

Abstract

We present a model for progressing and revising doxastic belief states in a multiagent setting. The model is sophisticated enough to deal not only with both ontic and epistemic actions, but also to handle the case where it is not common belief that an action occurred or that an observation has been made. Our model includes a multiagent extension of AGM-style belief revision.

1 Introduction

Reasoning about action and change have received an enormous attention in the last fifteen years, and has resulted in several families of languages, such as propositional action languages or causal theories, the situation calculus, the fluent calculus, etc. All these languages have been extended so as to reason with incomplete knowledge and sensing.

The fact that the initial belief state may be incomplete and the actions may be nondeterministic leads to the necessity to distinguish formally between facts and beliefs. Epistemic or doxastic logic have been used in several places to reason about action and change., e.g. [Lakemeyer and Levesque, 1998; Baral and Son, 2001; Herzig *et al.*, 2003; Scherl and Levesque, 2003]. A common limitation of this series of works is that they consider only one agent (and nature, since the agent evolves in a nondeterministic world).

Now, many domains involve several agents, who are expected to interact (or more specifically negotiate, cooperate, communicate). Interacting in an efficient way requires to reason about other agents' beliefs and the evolution of these beliefs after the occurrence of some actions or events. What renders things complex when several agents are considered is the fact that agents' beliefs must include not only first-order beliefs, i.e., beliefs about the world, but also higher-order beliefs, i.e., beliefs about the agents' beliefs. Furthermore, agents must be able to make higher-order beliefs evolve in the light of new information, i.e., they must also possess update and revision procedures for such beliefs.

While some works exist so as to model the evolution of a multiagent belief model after some communication actions

are performed, there are only few works where usual physical and sensing actions – as considered in cognitive robotics – are considered in a multiagent setting. This paper contributes to fill the gap. For the sake of the exposition we suppose there are only two agents. It is not a loss of generality, in so far as all problems related to mutual belief, common belief, and communication are already present with two agents and are conceptually no more complex than with more than two agents. Moreover, all definitions and results of this paper can readily be extended to the N -agent case.

As it is often the case in the literature, we assume for the sake of simplicity that the set of actions available to the agents is partitioned into two subsets: *purely ontic* (or *physical*) actions may change the state of the world but do not bring any feedback, whereas *purely epistemic* (or *sensing*) actions leave the state of the world unchanged and may only bring some feedback about it. This does not induce a loss of generality, since more general actions, with both effects on the world and feedback, can be decomposed into two actions, one being purely ontic and the other one purely epistemic.

Belief models are taken to be Kripke structures of 2-agent doxastic logic $KD45_2^C$ on which Section 2 gives some background. Section 3 considers ontic actions (without feedback); we first define the progression of a belief state by an action whose occurrence is assumed to be common knowledge, and then we consider the general case where this assumption is relaxed. Section 4 focuses on observations and sensing actions. We define the progression of a belief model by an epistemic action model where agents may perform sensing actions (and thus gather observations) while others do not. Since initial beliefs of the agents are not required to be correct, agents have to perform a (mutual) belief revision process; accordingly, our progression by sensing actions involves a genuine extension of AGM revision. Section 5 discusses related work.

2 Belief models

We first give some basics about 2-agent propositional doxastic logic $KD45_2^C$ (see e.g. [Fagin *et al.*, 1995]).

Definition 1 (language) Let $AtProp = \{p_1, \dots, p_n\}$ be a finite set of propositional symbols (atoms). The language $\mathcal{L}_{AtProp}^{1,2}$ of $KD45_2^C$ is built inductively from $AtProp$, constant symbols \top and \perp , the connectives \neg , \wedge , \vee and modal operators \mathbf{B}_1 , \mathbf{B}_2 , $\mathbf{CB}_{1,2}$ in the standard way.

⁰A preliminary version of this paper has been presented at the 6th Workshop on Logic and the Foundations of Game and Decision Theory (LOFT-2004).

Formulas of $\mathcal{L}_{AtProp}^{1,2}$ are denoted by capital Greek letters Φ, Ψ , etc.; objective (i.e., modality-free) formulas are denoted by small Greek letters φ, ψ , etc. Objective formulas are interpreted in a classical way:

Definition 2 (states) $States = 2^{AtProp}$ is the set of propositional valuations, or states. $States(\varphi)$ denotes the set of states which classically satisfy the objective formula φ .

States are denoted by s, s' etc. Although states are formally sets of atoms, we prefer to denote them in the following form: if $AtProp = \{a, b, c\}$ then the state $\{a, c\}$ is denoted by $[a, -b, c]$ – it assigns a and c to true and b to false.

Definition 3 (pointed belief models) A pointed belief model (PBM) \mathcal{M} for $KD45_2^C$ is a tuple $\mathcal{M} = \langle W, val, R_1, R_2, w^* \rangle$ where

- $W = \{w, v, \dots\}$ is a nonempty set of possible worlds;
- $val : W \rightarrow States$ maps possible worlds to states;
- R_1, R_2 are binary relations on W satisfying **seriality** ($\forall i, w, R_i(w) \neq \emptyset$), **transitivity**, and **Euclideanity** (if $wR_i w'$ and $wR_i w''$ then $w'R_i w''$);
- $w^* \in W$ is a distinguished world (the actual world).

Note that R_i is not required to be reflexive (which implies that agents may hold wrong beliefs).

The subjective content of a pointed belief model $\mathcal{M} = \langle W, val, R_1, R_2, w^* \rangle$ is the tuple $M = \langle W, val, R_1, R_2 \rangle$ and is simply called a (nonpointed) belief model.

We are now ready to define the notion of satisfaction of a formula by a pointed belief model:

Definition 4 (satisfaction by a PBM) Satisfaction of formulas of $\mathcal{L}_{AtProp}^{1,2}$ in a (nonpointed) belief model $M = \langle W, val, R_1, R_2 \rangle$ at world $w \in W$ is defined inductively in the usual way:

- if Φ is objective then $M, w \models \Phi$ iff $val(w) \models \Phi$.
- $M, w \models \neg\Phi$ iff $M, w \not\models \Phi$.
- $M, w \models \Phi \wedge \Psi$ iff $M, w \models \Phi$ and $M, w \models \Psi$ (and similarly for the other connectives);
- for $i \in \{1, 2\}$, $M, w \models \mathbf{B}_i\Phi$ iff $M, w' \models \Phi$ for all $w' \in R_i(w)$
- $M, w \models \mathbf{CB}_{1,2}\Phi$ iff $M, w' \models \Phi$ for all $w' \in R_{CB}(w)$, where the accessibility relation for common belief R_{CB} is defined as usual by $R_{CB} = (R_1 \cup R_2)^*$ ($*$ denotes reflexive and transitive closure).

Satisfaction of a formula Φ of $\mathcal{L}_{AtProp}^{1,2}$ in the pointed belief model $\mathcal{M} = \langle W, val, R_1, R_2, w^* \rangle$ is finally given by $\mathcal{M} \models \Phi$ iff $M, w^* \models \Phi$.

3 Progression of belief models by ontic actions

3.1 Ontic actions

A (purely) *ontic* action is characterized by the fact that it does not bring any feedback. Examples are: tossing a coin without observing the outcome, sending an email to somebody without knowing whether it will be received by the addressee. In terms of belief states, there is no need to distinguish between the projection of agent i 's belief state by the action *before* the

action is performed, and i 's belief state *after* the action: what i foresees is what she gets. An ontic action α is a “state transformer”, i.e., a transition relation on states, alias valuations: to every α there is associated a function $(\cdot)^\alpha$ mapping valuations $s \in States$ to sets of valuations: $s^\alpha \subseteq States$ is the set of possible successor states of s after α .

Definition 5 (ontic actions) An ontic action α is a binary relation on States s.t. the direct image s^α of every state $s \in States$ by α is a nonempty set. $AtAct_O$ denotes the set of all ontic actions.

For the sake of simplicity, we assume that actions are *fully executable*, that is, $s^\alpha \neq \emptyset$.

In the rest of the paper we make use of the action β of switching p , defined on $AtProp = \{p\}$ by $[p]^\beta = \{\neg p\}$ and $[\neg p]^\beta = \{p\}$, and of the void action λ , which is defined by $s^\lambda = \{s\}$ for all $s \in States$.

3.2 Common knowledge of action occurrences

In this section we assume that the initial beliefs of the agents are expressed by a belief model \mathcal{M} , and that it is common knowledge to the agents that action α is being performed (for instance because one of them publicly announces that she is performing α). The *progression* of \mathcal{M} by α is the intended new belief model \mathcal{M}^α expressing the beliefs of the agents *after* action α has been performed. The intuition behind the progression of \mathcal{M} by α is the following: \mathcal{M}^α is obtained from \mathcal{M} by replacing in \mathcal{M} each world w by a set of worlds $\langle w, s \rangle$, such that s is a possible state resulting from the application of α in $val(w)$. Accessibility relations are then “transferred” from the old worlds to the new ones.

Definition 6 (progression of a PBM by an ontic action)

Let $\mathcal{M} = \langle W, val, R_1, R_2, w^* \rangle$ be a pointed belief model and $\alpha \in AtAct_O$. The progression of \mathcal{M} by α is the set of pointed belief models

$$\mathcal{M}^\alpha = \{ \langle W^\alpha, val^\alpha, R_1^\alpha, R_2^\alpha, \langle w^*, s \rangle \rangle \mid s \in val(w^*)^\alpha \}$$

where

- $W^\alpha = \{ \langle w, s \rangle \mid w \in W, s \in val(w)^\alpha \}$;
- $val^\alpha(\langle w, s \rangle) = s$;
- $\langle w, s \rangle R_i^\alpha \langle v, t \rangle$ if and only if $w R_i v$.

The reason why \mathcal{M}^α is generally not a single pointed belief model but a set is that actions may be nondeterministic, so that the actual state $val(w^*)$ may have several possible successor states by α . If α is deterministic, then obviously \mathcal{M}^α is a singleton.

3.3 The general case

In the previous section, we supposed that action occurrences are perceived completely and correctly by every agent, which is often unrealistic: some agents may be unaware that an action is being performed, or may just suspect that some action is being performed, or may know that an action from a given set is being performed, without knowing precisely which one (see [Baltag *et al.*, 1998] for an extended discussion on this issue, for the case of epistemic actions). In order to relax this constraint we extend Baltag *et al.*'s ideas and use *doxastic action structures*.

Definition 7 (pointed doxastic action structures)

A pointed doxastic action structure (PDAS) is a tuple

$$\mathcal{A} = \langle W_A, act, S_1, S_2, a^* \rangle$$

where

- $W_A = \{a, b, \dots\}$ is a set of possible “action worlds”;
- $act : W_A \rightarrow AtAct_O$ is a total function that maps possible action worlds to actions;
- S_1, S_2 are binary relations on W_A that are serial, transitive, and Euclidean.
- $a^* \in W_A$ is the actual action world – the occurrence of which some of the agents may be unaware.

Subsection 3.2 corresponds to the particular case where $W_A = \{a^*\}$ and $S_1(a^*) = S_2(a^*) = \{a^*\}$.

S_i relates an action a to agent i 's “subjective versions” of a : if $aS_i b$ and a occurs then in i 's view b is one of the actions that might have happened. In this way one can model incomplete and erroneous perception.

Example 1 Let $\beta = switch(p)$ and λ as defined in Section 3.1, and $\mathcal{A} = \langle W_A, act, S_1, S_2, a^* = a_0 \rangle$ where $W_A = \{a_0, a_1\}$, $act(a_0) = \beta$; $act(a_1) = \lambda$, $S_1(a_0) = \{a_1\}$; $S_2(a_0) = \{a_0\}$, and $S_1(a_1) = S_2(a_1) = \{a_1\}$. In the PDAS \mathcal{A} , agent 2 correctly believes that action β is performed (maybe because she is switching p); agent 1 is not aware of this action occurrence and thus believes that nothing happens, that is, the void action λ is performed, and that this is common belief; 2 is right about 1's beliefs.

Definition 8 (progression of a PBM by a PDAS) The progression of a PBM $\mathcal{M} = \langle W, val, R_1, R_2, w^* \rangle$ by a PDAS $\mathcal{A} = \langle W_A, act, S_1, S_2, a^* \rangle$ is a set of pointed belief models

$$\mathcal{M}^{\mathcal{A}} = \{ \langle W^{\mathcal{A}}, val^{\mathcal{A}}, R_1^{\mathcal{A}}, R_2^{\mathcal{A}}, \langle w^*, a^*, s \rangle \rangle \mid s \in val(w^*)^{act(a^*)} \}$$

where

- $W^{\mathcal{A}} = \{ \langle w, a, s \rangle \mid w \in W, a \in W_A, s \in val(w)^{act(a)} \}$
- $val^{\mathcal{A}}(\langle w, a, s \rangle) = s$;
- $\langle w, a, s \rangle R_i^{\mathcal{A}} \langle v, b, t \rangle$ iff $wR_i v$ and $aS_i b$.

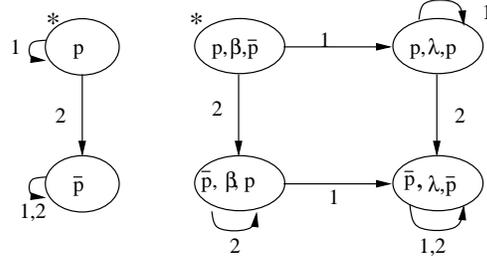
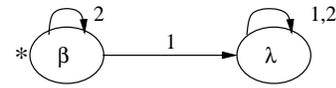
Intuitively, the world $\langle w, a, s \rangle$ encodes that the execution of action a in state $val(w)$ results in state s .

In order to make sure that this definition is well-founded we first have to check that each element of $\mathcal{M}^{\mathcal{A}}$ is actually a PBM, i.e., W_A is nonempty and each $R_i^{\mathcal{A}}$ is serial, transitive, and Euclidean.

Proposition 1 $\mathcal{M}^{\mathcal{A}}$ is a set of a $KD45_2^C$ pointed belief models.

Example 2 Let $AtProp = \{p\}$, \mathcal{A} as in Example 1, and $\mathcal{M} = \langle W, val, R_1, R_2, w^* = w_0 \rangle$, where $W = \{w_0, w_1\}$; $val(w_0) = p$; $val(w_1) = \neg p$; $R_1(w_0) = \{w_0\}$, $R_2(w_0) = \{w_1\}$; $R_1(w_1) = R_2(w_1) = \{w_1\}$.

The progression of \mathcal{M} by \mathcal{A} is the singleton $\mathcal{M}^{\mathcal{A}} = \{ \langle W^{\mathcal{A}}, val^{\mathcal{A}}, R_1^{\mathcal{A}}, R_2^{\mathcal{A}}, \langle w_0, a_0, [\neg p] \rangle \rangle \}$ depicted on the figure below; the left part of the figure is \mathcal{M} , the upper part \mathcal{A} , and the bottom-right part is $\mathcal{M}^{\mathcal{A}}$; actual worlds are labeled by $*$.



In $\mathcal{M}^{\mathcal{A}}$, both 1 and 2 (wrongly) believe p : 1 because he used to correctly believe p and believes nothing happened, and 2 because he used to believe $\neg p$ and knows that p has been switched. But this is not common belief since 1 believes that 2 believes $\neg p$ and 2 believes that 1 believes $\neg p$.

The progression of a belief model by a doxastic action structure recovers single-agent progression as a particular case, and slightly more generally, if in \mathcal{M} the agents have full common belief that the objective formula φ holds, and have common knowledge that action α occurs, then in $\mathcal{M}^{\mathcal{A}}$ the agents commonly believe $prog(\varphi, \alpha)$, where $prog$ refers here to classical propositional progression at the syntax level. (Semantically, we have $States(prog(\varphi, \alpha)) = \bigcup_{s \in States(\varphi)} s^{\alpha}$.)

4 Progression by epistemic actions

4.1 Observation actions

We now consider a set $AtAct_E$ of elementary epistemic actions, or observation actions, of the form $observe(\Phi)$, where Φ is a formula of $\mathcal{L}_{AtProp}^{1,2}$. $observe(\Phi)$ is the action of observing that Φ holds. For the sake of notation we simply write Φ instead of $observe(\Phi)$.

Observations can be made either “spontaneously” by the agent, or after performing a sensing action such as, typically, a test $sense(\Phi)$ sending back the truth value of Φ . (We shall see later that observation actions can express sensing actions).

We consider the possibilities of observing (and sensing) both objective formulas, as in cognitive robotics (such as “test whether this solution is an acid”), and subjective formulas (such as “ask agent 2 whether she believes the solution is an acid”, or “ask agent 2 whether she believes whether agent 3 knows whether the solution is acid or not”).

Importantly, the observation Φ is executable only in worlds where Φ is true (we’ll come back on this later.) This restriction merely expresses that observations are reliable.

We now introduce pointed observation structures (POS), which are similar to PDAS, except for one thing: while it is meaningful to talk about the “actual” ontic action being performed, this is no longer so for epistemic actions (observations), since they are intrinsically “subjective”, that is, it is meaningless to talk about an observation (resp. a sensing action) without referring to the agent who performs it. However, we may talk about the actual epistemic action performed by a given agent i . This leads us to consider the following pointed observations structures, which resemble our PDAS of Section

3, except for the valuation function, which now maps each world *and each agent* to an observation action.¹

Definition 9 (pointed observation structures)

A pointed observation structure (POS) is a tuple $\mathcal{O} = \langle W_{\mathcal{O}}, obs, T_1, T_2, o^* \rangle$, where

- $W_{\mathcal{O}}$ is a set of observation worlds, denoted o, o' etc.
- T_1 and T_2 are serial, transitive and Euclidean relations on $W_{\mathcal{O}}$;
- $obs : W_{\mathcal{O}} \times \{1, 2\} \rightarrow AtAct_E$ maps each observation world and each agent to an observation action;
- $o^* \in W_{\mathcal{O}}$.

\mathcal{O} must satisfy the following constraints, for all $o \in W_{\mathcal{O}}$:

1. $obs(o, 1) \wedge obs(o, 2)$ is consistent;
2. for all $i \in \{1, 2\}$, if $oT_i o'$ then $obs(o', i) = obs(o, i)$.

For $i \in \{1, 2\}$ and $o \in W_{\mathcal{O}}$, $obs(o, i)$ is the observation made by agent i at o . In particular, $obs(o^*, i)$ is the observation actually performed by i .

The constraint that $obs(o, 1) \wedge obs(o, 2)$ be consistent is required by our hypothesis that observations are truthful. It is not a sufficient condition: truthfulness is actually not expressible in a POS, since it refers as well to the states; it will only be manifest in the product of a POS and a belief model (see further). The second constraint ensures that every agent knows the observation she performs.

Although observations are truthful, the observation structure is not necessarily reflexive: while an agent is correct about her own observations, she might be wrong about other agents' observations. A nonreflexive epistemic observation structure may for instance represent the fact that agent 1 believes that agent 2 is sensing p while she is actually not – this is totally different from getting a wrong observation.

Pointed observation structures on observation actions can represent sensing actions. For instance, the POS $\mathcal{O} = \langle W_{\mathcal{O}}, obs, T_1, T_2, o^* \rangle$ where $W_{\mathcal{O}} = \{o_1, o_2\}$, $o^* = o_1$, $T_1(o_1) = T_1(o_2) = \{o_1, o_2\}$, $T_2(o_1) = \{o_1\}$, $T_2(o_2) = \{o_2\}$, $obs(o_1, 1) = \top$, $obs(o_1, 2) = p$, $obs(o_2, 1) = \top$, $obs(o_2, 2) = \neg p$ represents a situation where 1 does not observe anything and knows that 2 tests the truth value of p (and all this is common knowledge) – which means that she observes either p or $\neg p$, according to the actual truth value of p . The latter will be ensured by the fact that it is impossible to observe Φ when Φ is false.

We are now in position of defining the progression of a PBM by an POS. We first work out the case where agents have correct beliefs (i.e., \mathcal{M} is a $S5_2^C$ pointed model), because it is simpler and comes to a syntactical progression operator. The case of possibly incorrect beliefs needs a belief revision phase, and will be considered later.

¹This departs from observation structures in [Baltag et al., 1998] where observations are disconnected from the agent who performs them. Attaching observations to agents allows us (a) to talk about the actual observations performed by agents 1 and 2, and (b) to have a doxastic observation structure which remains $S5_2^C$ when agents do not hold incorrect beliefs about other agents' observations; both (a) and (b) are not possible with the structures in [Baltag et al., 1998].

4.2 Progression by epistemic actions in $S5_2^C$

We assume in this section that \mathcal{M} is a $S5_2^C$ PBM model (both R_1 and R_2 are reflexive), and \mathcal{O} is a $S5_2^C$ POS-model (both T_1 and T_2 are reflexive). We note M the (nonpointed) belief model associated to \mathcal{M} . Progression of \mathcal{M} by \mathcal{O} amounts to construct a restricted product of \mathcal{M} and \mathcal{O} . Unlike in the case for ontic actions, the result is a unique pointed belief model.

Definition 10 (progression of a $S5_2^C$ PBM by a POS)

Let $\mathcal{M} = \langle W, val, R_1, R_2, w^* \rangle$ be a pointed $S5_2^C$ -model and $\mathcal{O} = \langle W_{\mathcal{O}}, obs, T_1, T_2, o^* \rangle$ a pointed observation structure such that $\mathcal{M} \models obs(o^*, i)$ for $i = 1, 2$. The progression of \mathcal{M} by \mathcal{O} is

$$\mathcal{M}^{\mathcal{O}} = \langle W^{\mathcal{O}}, val^{\mathcal{O}}, R_1^{\mathcal{O}}, R_2^{\mathcal{O}}, w^{*\mathcal{O}} \rangle$$

where

- $W^{\mathcal{O}} = \{ \langle w, o \rangle \mid w \in W, o \in W_{\mathcal{O}}, \text{ and for } i = 1, 2, M, w \models obs(o, i) \}$;
- $val^{\mathcal{O}}(\langle w, o \rangle) = val(w)$;
- $\langle w, o \rangle R_i^{\mathcal{O}} \langle w', o' \rangle$ if and only if $w R_i w'$, $o T_i o'$, and $M, w' \models obs(o, i)$;
- $w^{*\mathcal{O}} = \langle w^*, o^* \rangle$.

Hence from all possible combinations of worlds and observations, only those from $W^{\mathcal{O}}$ which have truthful observations are retained. (As by hypothesis none of the agents can make erroneous observations this set is nonempty.) Moreover, for a given world $\langle w, o \rangle$ only those worlds $\langle w', o' \rangle$ are accessible for i where i 's observation $obs(o', i)$ is true in w . In this way the accessibility relations are constructed by filtering out accessible worlds (for agent i) where the observed formula (by agent i) is false: it makes that after observing Φ , agent i believes Φ (which is, again, justified by the assumption that observations are reliable).

For instance, consider a world where 1, after asking 2 whether she knows whether Mozart was left-handed or not (mlh), 2 answers “yes, I know” (i.e., the observation is $\mathbf{B}_2 \text{ mlh} \vee \mathbf{B}_2 \neg \text{mlh}$), then after progression, 1 knows that 2 knows whether Mozart was left-handed or not.

Proposition 2 *If both \mathcal{M} and \mathcal{O} are $S5_2^C$ belief models then $\mathcal{M}^{\mathcal{O}}$ is an $S5_2^C$ belief model.*

Example 3 *Suppose $AtProp = \{p\}$. Let $W = \{w_p, w_{\neg p}\}$, $val(w_p) = [p]$, $val(w_{\neg p}) = [\neg p]$, and for all $w \in W$, $R_1(w) = \{w\}$, and $R_2(w) = W$. Let $w^* = w_p$. Hence 1 knows that p , while 2 ignores whether p . The latter is common knowledge.*

Let $W_{\mathcal{O}} = \{o_p, o_{\neg p}\}$ with $obs(o_p, 1) = \top$ and $obs(o_p, 2) = \varphi$, for $\varphi = p, \neg p$; $T_1(o_p) = \{o_p\}$, and $T_2(o_p) = W_{\mathcal{O}}$. Suppose $o^ = o_p$. Hence 1 knows that 2 is testing whether p .*

We obtain $W^{\mathcal{O}} = \{ \langle w_p, o_p, [p] \rangle, \langle w_{\neg p}, o_{\neg p}, [\neg p] \rangle \}$, and $R_1^{\mathcal{O}}(\langle w, o, s \rangle) = R_2^{\mathcal{O}}(\langle w, o, s \rangle) = \{ \langle w, o, s \rangle \}$ for all $\langle w, o, s \rangle \in W^{\mathcal{O}}$. As expected, after the revision there is common belief that p .

4.3 Progression by epistemic actions in KD45₂^C

We now consider the case where initial beliefs may be incorrect. At that time, this requires to make a simplifying assumption on the nature of epistemic actions, namely, that what is observed are only facts (and not beliefs): an observation (an elements of $AtAct_E$) is an *objective formula* φ .

We consider an AGM preference structure, i.e., a collection of preference relations \preceq_X for every nonempty subset of states $X \subseteq S$ verifying the faithfulness condition: $\forall X \subseteq S, \forall x \in X, \forall y \in S, x \preceq_X y$. From that structure an AGM revision operation can be defined as follows.

Definition 11 (AGM revision on objective formulas)

$\varphi \star_{\preceq} \psi$ is the propositional formula – unique up to logical equivalence – such that

$$States(\varphi \star_{\preceq} \psi) = \min_{\preceq_{States(\varphi)}}(States(\psi)).$$

Such preference-based revision operators are characterized by the AGM postulates [Gärdenfors and Makinson, 1988; Katsuno and Mendelzon, 1991].

Similar to the case of action structures, the revision of a belief model by an observation structure is done by revising the belief states of each agent according to her view of the observation.

Definition 12 (revision of a PBM by a POS)

Let $\mathcal{M} = \langle W, val, R_1, R_2, w^* \rangle$ be a PBM and $\mathcal{O} = \langle W_O, obs, T_1, T_2, o^* \rangle$ a POS such that $\mathcal{M} \models obs(o^*, i)$ for $i = 1, 2$. Let \preceq be any preference relation. Then

$$\mathcal{M}^{\mathcal{O}} = \langle W^{\mathcal{O}}, val^{\mathcal{O}}, R_1^{\mathcal{O}}, R_2^{\mathcal{O}}, w^{*\mathcal{O}} \rangle$$

where

- $W^{\mathcal{O}} = \{\langle w, o, s \rangle \mid w \in W, o \in W_O, \text{ and for } i = 1, 2, s \in States(obs(o, i))\}$;
- $val^{\mathcal{O}}(\langle w, o, s \rangle) = s$;
- $R_i^{\mathcal{O}}(\langle w, o, s \rangle) = \{\langle w', o', s' \rangle \mid w' \in R_i(w), o' \in T_i(o), \text{ and } s' \in \min_{\preceq_{val(R_i(w))}}(States(obs(o, i)))\}$;
- $w^{*\mathcal{O}} = \langle w^*, o^* \rangle$.

In the definition of $W^{\mathcal{O}}$, the condition that $s \in States(obs(o, i))$ guarantees that observations are truthful. (Again, as by hypothesis none of the agents can make erroneous observations the set $W^{\mathcal{O}}$ is nonempty.) In the definition of $R_i^{\mathcal{O}}$, the minimization condition implements preference-based revision. As for S5₂^C, the result is a unique pointed belief model.

Our definition is well-founded since:

Proposition 3 $\mathcal{M}^{\mathcal{O}}$ is a KD45₂^C belief model.

The next proposition establishes that AGM belief revision is a particular case.

Proposition 4 Let φ be an objective formula, and let $\mathcal{M}(\varphi) = \langle W, val, R_1, R_2, w^* \rangle$ be a belief model such that $W = States(\varphi)$, $val(w) = w$, $R_i(w) = W$.

Let \preceq be a preference relation, and let $\mathcal{O} = \langle \{o\}, obs, T_1, T_2, o^* \rangle$ be a pointed observation structure such that for $i = 1, 2$, $obs(o, i) = \psi$, $T_i(o) = \{o\}$, and $o^* = o$.

Suppose o^* is reliable, i.e. $\mathcal{M} \models \psi$, and let $\mathcal{M}^{\mathcal{O}} = \langle W^{\mathcal{O}}, val^{\mathcal{O}}, R_1^{\mathcal{O}}, R_2^{\mathcal{O}}, w^{*\mathcal{O}} \rangle$ be the revised model. Then

$$States(\varphi \star_{\preceq} \psi) = R_1^{\mathcal{O}}(w^{*\mathcal{O}}) = R_2^{\mathcal{O}}(w^{*\mathcal{O}}).$$

Example 4 Suppose $AtProp = \{p\}$, and suppose \preceq is based on the Hamming distance $d_H(s, s')$, i.e. the number of symbols on which s and s' differ. For $S \subseteq X$, $d_H(s, X) = \min_{s'' \in X} d_H(s, s'')$, and \preceq is defined by:

$$s \preceq_X s' \text{ iff } d_H(s, X) \leq d_H(s', X).$$

Let $W = \{w_p, w_{\neg p}\}$, $val(w_p) = [p]$, $val(w_{\neg p}) = [\neg p]$, and for all $w \in W$, $R_1(w) = \{w\}$ and $R_2(w) = \{w_{\neg p}\}$. Let $w^* = w_p$. Then $\mathcal{M} \models \mathbf{B}_1(p \wedge \mathbf{B}_2 \mathbf{CB}_{1,2} p)$. Let \mathcal{O} such that $W_O = \{o_p\}$, $obs(o_p, i) = p$ and $T_i(o_p) = \{o_p\}$, $i = 1, 2$.

Then $W^{\mathcal{O}} = \{\langle w_p, o_p, [p] \rangle, \langle w_{\neg p}, o_p, [p] \rangle\}$, and for all possible worlds $\xi \in W^{\mathcal{O}}$, $R_1^{\mathcal{O}}(\xi) = \{\xi\}$ and $R_2^{\mathcal{O}}(\xi) = \{\langle w_{\neg p}, o_p, [p] \rangle\}$. Hence $\mathcal{M}^{\mathcal{O}} \models \mathbf{CB}_{1,2} p$.

5 Related work and conclusion

A number of papers have considered belief change operators on belief structures based on multiagent KD45 or S5.

A series of works considers *multiagent belief expansion*, starting with [Fagin *et al.*, 1995], who consider public announcements in S5, and express examples such as the Muddy Children Puzzle (cf. Section 4.2). The issue is then further studied in [Gerbrandy and Groeneveld, 1997] who develop a language for reasoning about information change relative to group announcements (public announcements within a given subset of agents.) [van Ditmarsch *et al.*, 2004] show that under the restriction to positive formulas public announcement coincides with expansion, and that they differ in general.

[Baltag *et al.*, 1998] go much more general and define complex epistemic actions as Kripke frames where worlds are valued by actions, to account for the case where agents have different information about which action is taking place. Our doxastic action and observation structures are taken from theirs (with a few differences explained further). See also [van Ditmarsch, 2002] who consider a more elaborate language of epistemic actions allowing for concurrent execution of epistemic actions by the agent.

[Baltag *et al.*, 1998] only partially allow for incorrect beliefs: more precisely, the initial belief model is S5, whereas complex epistemic actions are KD45: agents are assumed to hold initial correct beliefs, but may have misperceptions or wrong suspicions of communication actions between other agents. The fact that initial beliefs should be correct is important, since it allows for a simple belief expansion process, which filters out the worlds where the observation just made does not hold. However, expanding an S5 model by a KD45 complex epistemic action does generally not result in an S5 model: the final beliefs of the agents might be incorrect – which makes the process impossible to iterate. This lead Aucher [Aucher, 2004] to extend the latter approach so as to deal with possibly erroneous beliefs. His belief models are graded S5 models, consisting of an S5 model together with a function expressing the relative plausibility of the worlds. The revision process then works by retaining, for each set of indistinguishable worlds, the most plausible ones among those which satisfy the observation. The result is still a graded S5 model. The only problem with Aucher's semantics is that by enforcing that any conceivable world has some plausibility, models are extremely large, even for simple single-agent examples; furthermore, it may be difficult to assign plausibilities in the initial belief model. Our preference relations play

a similar role to Aucher's plausibilities, and remain more general and closer to the original AGM framework.

[Tallon *et al.*, 2004] propose an account of multi-agent revision, where agents may have initial erroneous beliefs, under the strong hypothesis that agents communicate publicly all their beliefs, i.e., their belief state as a whole.

The approach in [van der Meyden, 1994], older than and unrelated to the latter ones, allows for incorrect beliefs and revision, but remains very general and does not commit to a precise family of belief revision operators. Observations are common knowledge such as in group announcements.

So far, all reported approaches focus on deterministic episodic actions, namely *observation actions*; typical observation actions in these approaches are *communication actions* consisting in an agent telling another (or a group of others) something she claims to believe. With this class of actions, (a) the objective state of the world does not change and (b) actions are deterministic: they have *preconditions* (such as, in the case of sincere communication, an agent believing what she tells); when the precondition is satisfied, the action results in the same world, and when the precondition is not satisfied, the action is not executable. Therefore the progression problem (for ontic or sensing actions) is not considered.

Ontic actions, as well as sensing actions (which differ from observation actions in many aspects) are almost never considered in multiagent frameworks. An exception is [Martin *et al.*, 2004] who consider ontic, sensing and communication actions but without any account for higher-order beliefs (more precisely, there is a distinguished agent who holds beliefs about other agent's beliefs about facts, and nothing more). They assume that beliefs are correct and that actions are public; on the other hand, this work contains many results about the practical computation of the new beliefs. [Shapiro *et al.*, 1998] consider a multiagent extension of the situation calculus where communication actions are taken into account, but agents hold no beliefs about other agents. [Demolombe and Parra, 2002] consider a multiagent extension of the situation calculus with higher-order beliefs, building on the single-agent approach [Shapiro *et al.*, 2000]; they do not give a general model but study a few particular cases.

We mention three other approaches dealing with the dynamics of mutual beliefs; unlike the previous ones, they are not based on general Kripke structures but on simpler structures. [Kfir-Dahav and Tennenholtz, 1998] and [Su *et al.*, 2004] assume that each agent observes (at each instant) a specific subset of propositional variables, and who observed what is common knowledge, and [Liu and Williams, 2001], that each agent chooses to open some parts of its knowledge base to some other agents (*knowledge migration* consists then in revising one's knowledge base when accessing other agents' accessible knowledge.)

Further work will mainly focus on the practical computation of progression and regression, that ultimately would enable logic-based planning in multiagent domains.

References

[Aucher, 2004] G. Aucher. A combined system for update logic and belief revision. In *Working notes of the 7th Pacific Rim Int. Workshop on Multi-Agents (PRIMA'04)*, 2004.

- [Baltag *et al.*, 1998] A. Baltag, L. S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In *Proc. TARK'98*, pages 43–56, 1998.
- [Baral and Son, 2001] C. Baral and T. Son. Formalizing sensing actions – a transition function based approach. *Artificial Intelligence*, 125(1–2):19–91, 2001.
- [Demolombe and Parra, 2002] R. Demolombe and P. Pozos Parra. Formalisation de l'évolution des croyances dans le calcul de situations. *Information, Interaction, Intelligence*, pages 25–40, 2002.
- [Fagin *et al.*, 1995] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
- [Gärdenfors and Makinson, 1988] P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proc. TARK'88*, pages 83–95, Pacific Grove (CA), 1988.
- [Gerbrandy and Groeneveld, 1997] J. Gerbrandy and W. Groeneveld. Reasoning about information change. *J. of Logic, Language and Information*, 6(2), 1997.
- [Herzig *et al.*, 2003] A. Herzig, J. Lang, and P. Marquis. Action representation and partially observable planning using epistemic logic. In *Proc. IJCAI'03*, pages 1067–1072, 2003.
- [Katsuno and Mendelzon, 1991] H. Katsuno and A.O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proc. KR'91*, pages 387–394. Morgan Kaufmann, 1991.
- [Kfir-Dahav and Tennenholtz, 1998] N. Kfir-Dahav and M. Tennenholtz. Multi-agent belief revision. In *Proc. TARK'96*, pages 175–194. Morgan Kaufmann, 1998.
- [Lakemeyer and Levesque, 1998] G. Lakemeyer and H. Levesque. *AOL: a logic of acting, sensing, knowing, and only knowing*. In *Proceedings of KR'98*, pages 316–327, 1998.
- [Liu and Williams, 2001] W. Liu and M.-A. Williams. A framework for multi-agent belief revision. *Studia Logica*, 67(2):291–312, 2001.
- [Martin *et al.*, 2004] Y. Martin, I. Narasamdya, and M. Thielscher. Knowledge of other agents and communicative actions in the situation calculus. In *Proc. of KR2004*, 2004.
- [Scherl and Levesque, 2003] R. Scherl and H. Levesque. The frame problem and knowledge producing actions. *Artificial Intelligence*, 144(1–2), 2003.
- [Shapiro *et al.*, 1998] S. Shapiro, Y. Lespérance, and H. Levesque. Specifying communicative multi-agent systems. In *Agents and Multi-Agent Systems - Formalisms, Methodologies, and Applications*, volume 1441 of *LNAI*, pages 1–14. Springer-Verlag, 1998.
- [Shapiro *et al.*, 2000] S. Shapiro, M. Pagnucco, Y. Lespérance, and H. J. Levesque. Iterated belief change in the situation calculus. In *Proc. KR2000*, pages 527–538, 2000.
- [Su *et al.*, 2004] K. Su, G. Lv, and Y. Zhang. Reasoning about knowledge by variable forgetting. In *Proc. KR'04*, pages 573–586. Morgan Kaufmann, 2004.
- [Tallon *et al.*, 2004] J.M. Tallon, J.C. Vergnaud, and S. Zamir. Communication among agents: A way to revise beliefs in KD45 Kripke structures. *Journal of Applied Non-Classical Logics*, 14(4):477–500, 2004.
- [van der Meyden, 1994] R. van der Meyden. Mutual belief revision (preliminary report). In *Proc. KR'94*, pages 595–606, 1994.
- [van Ditmarsch *et al.*, 2004] H. van Ditmarsch, W. van der Hoek, and B. Kooi. Public announcements and belief expansion. In *Preproceedings of AiML-2004 (Advances in Modal Logic)*, 2004.
- [van Ditmarsch, 2002] H. van Ditmarsch. Descriptions of game actions. *J. of Logic, Language and Information*, 11:349–365, 2002.