A system for consistency preserving belief change

David Steiner*

Version of May 2006

Abstract

We present a dynamic doxastic logic which formalizes belief change for rational agents. It is based on the multi-agent modal logic $\mathsf{KD45}_n$, and therefore preserves the consistency of belief. The actions can be seen as announcements of arbitrary formulas to arbitrary groups of agents. We call it pure belief change because there is no change of propositional facts. We give a sound and complete axiomatization and argue that this logic works for communication where the source is not known to be secure. A possible application could be games where it can happen that the players lie.

1 Introduction

The discussion about belief revision started in the early eighties, and was fully launched with the publication of the AGM postulates in [1] by Alchourrón, Gärdenfors, and Makinson. At the beginning of the nineties, Katsuno and Mendelzon came up with different postulates (KM postulates) in [7]. Their intended application was updating a belief base instead of revising it. Herzig showed in [5] that the AGM postulates are incompatible with the KM postulates. Most researchers are in agreement with the fact, that one has to consider the application in order to decide, which postulates one wants to be satisfied.

In the AGM theory, the inconsistent belief state is a potential state, it can even happen that a contradiction in itself (e.g. the formula $\bot$) is accepted as

*Research partly supported by the Swiss National Science Foundation. Address: Institut für Informatik und angewandte Mathematik, Universität Bern, Neubrückstrasse 10, CH-3012 Bern, Switzerland. Email: `steiner@iam.unibe.ch` .

new belief. We think that the AGM postulates are rather stated for human belief than for the belief of rational agents. From our point of view, rational agents should never get inconsistent belief (cf. Halpern and Friedman in [3]). This is one of the reasons why our system for belief change does not fit into the AGM setting. On the other hand, we consider pure belief change where no factual change takes place. Therefore, our approach cannot satisfy the KM postulates either.

The question is, how we expect the agents to consistently expand their beliefs. Our answer is based on the idea that a rational agent is always aware of all the consequences of its beliefs: the agents accept incoming information, if they do not belief in its negation, and they reject it otherwise. This procedure ensures that the beliefs of the agents permanently remain consistent. But it can happen that the agents accept false statements, and after that, they will reject a true sentence. Nevertheless, this is what can happen in practice, if the agents do not know whether the source of the information is reliable or not. Our approach formalizes belief expansion for a huge set of formulas, and for propositional formulas in particular. But it is not belief expansion in the original sense, because there are beliefs that can be contracted as a consequence of accepted information. This phenomenon can be explained with the presence of the axioms of positive and negative introspection, which are always valid before and after an announcement.

There are many formalizations of belief change in modal logics with dynamic-style operators for incoming information (e.g. van Linder, van der Hoek, and Meyer [12], Gerbrandy and Groeneveld [4], Roorda, van der Hoek, and Meyer [8]). They all fit more or less into the AGM setting, but they get inconsistent theories by adding the D axiom $\neg B_i \bot$ (in order to avoid inconsistent belief). Other approaches compatible with $\mathsf{S5}_n$ (e.g. van Ditmarsch in [10]) are still not consistency preserving in our sense, although they satisfy the D axiom. The problem is, that every formula holds after the announcement of the formula $\bot$. A system for consistency preserving knowledge change extending the logic $\mathsf{S5}_n$ can be found in [9].

In section 2 we will define the language and semantics of our dynamic doxastic logic and we will show that performing an action always results in a serial, transitive, and Euclidean Kripke structure. We will give a sound and complete axiomatization in section 3 and prove that our logic has the same expressive strength as normal modal logic. In section 4 we discuss our results and give a short outlook to future work.

2 Language and semantics

It is the aim of this section to introduce the language $\mathcal{L}_n^{\mathsf{A}}$ of multi-agent modal logic with dynamic-style operators for group announcements. We will define the semantics of announcements via operations on Kripke structures and we will prove that we always operate on serial, transitive, and Euclidean Kripke structures.

Given a natural number $n \geq 1$, we fix the set $\mathcal{A} = \{1, \ldots, n\}$ of n rational agents. Further, we take a countable non-empty set $\mathcal{P}$ of atomic propositions denoted by $p, q, \ldots$, possibly with subscripts. The set of $\mathcal{L}_n^{\mathsf{A}}$ formulas is defined by the following grammar ($p \in \mathcal{P}$, $i \in \mathcal{A}$, $\emptyset \neq G \subseteq \mathcal{A}$),

$$\alpha \quad ::= \quad p \mid \neg\alpha \mid \alpha \wedge \alpha \mid B_i\alpha \mid [\alpha]_G\alpha \ .$$

The formula $B_i\alpha$ means *agent i believes α*, and $[\alpha]_G\beta$ expresses that *β holds after the announcement of α to the group G*. The connectives $\vee$, $\rightarrow$, and $\leftrightarrow$ are defined as usual. $\top := p_0 \vee \neg p_0$ and $\bot := p_0 \wedge \neg p_0$ for some fixed $p_0 \in \mathcal{P}$. The $\mathcal{L}_0$ formulas are the propositional formulas, the $\mathcal{L}_n$ formulas are the modal formulas without announcement operators. The length of a formula is inductively defined by

$$\begin{aligned}
|p| &:= 1 \ , \\
|\neg\alpha| &:= |\alpha| + 1 \ , \\
|\alpha \wedge \beta| &:= |\alpha| + |\beta| + 1 \ , \\
|B_i\alpha| &:= |\alpha| + 1 \ , \\
|[\alpha]_G\beta| &:= |\alpha| + |\beta| + 2 \ .
\end{aligned}$$

Iterated announcements $[\alpha]_G^k\beta$ are naturally defined by induction on k,

$$\begin{aligned}
[\alpha]_G^0\beta &:= \beta \ , \\
[\alpha]_G^{k+1}\beta &:= [\alpha]_G[\alpha]_G^k\beta \ .
\end{aligned}$$

Definition 1 *An n-Kripke structure $\mathsf{K} = (S, R_1, \ldots, R_n, V)$ is an $(n+2)$-tuple, where $S \neq \emptyset$ is a set of states, $R_i \subseteq S^2$ is a binary relation for all $i \in \mathcal{A}$, and $V : \mathcal{P} \mapsto Pow(S)$ is a valuation function.*

Since the n is given, we will always simply write Kripke structure. The set S is called the universe of K, denoted by $|\mathsf{K}|$. The relation R_i is the accessibility relation for agent i. In the sequel, we will write $\mathcal{K}_n^{stE}$ for the class of all Kripke structures with serial, transitive, and Euclidean accessibility relations,

serial: for all $u \in S$ there is a $v \in S$ s.t. uR_iv , transitive: for all $u, v, w \in S$, uR_iv and $vR_iw \;\Rightarrow\; uR_iw$,

Euclidean: for all $u, v, w \in S$, uR_iv and $uR_iw \;\Rightarrow\; vR_iw$.

We will now define the validity of an $\mathcal{L}_n^{\mathsf{A}}$ formula in an arbitrary Kripke-world, i.e. a pair K, s s.t. $s \in |\mathsf{K}|$. The crucial point of this definition is the case of $[\alpha]_G\beta$, where we simultaneously define an operation on the Kripke structure. The idea is, that we restrict the accessibility relations to the states where the announced formula holds, if and only if the agents do not belief in its negation.

Definition 2 *Let* $\mathsf{K} = (S, R_1, \ldots, R_n, V)$ *be an arbitrary Kripke structure and* $s \in S$ *be an arbitrary state. The validity of* $\mathcal{L}_n^{\mathsf{A}}$ *formulas in the Kripke-world* K, s *is inductively defined as follows.*

$$\begin{array}{rcl}
\mathsf{K}, s \models p & \textit{iff} & s \in V(p) \ , \\
\mathsf{K}, s \models \neg\alpha & \textit{iff} & \mathsf{K}, s \not\models \alpha \ , \\
\mathsf{K}, s \models \alpha \wedge \beta & \textit{iff} & \mathsf{K}, s \models \alpha \textit{ and } \mathsf{K}, s \models \beta \ , \\
\mathsf{K}, s \models B_i\alpha & \textit{iff} & \textit{for every } t \in S, sR_it \;\Rightarrow\; \mathsf{K}, t \models \alpha \ , \\
\mathsf{K}, s \models [\alpha]_G\beta & \textit{iff} & \mathsf{K}^{\alpha,G}, s_1 \models \beta \ .
\end{array}$$

For given α *and* G, *the Kripke structure* $\mathsf{K}^{\alpha,G} := (S', R_1^{\alpha,G}, \ldots, R_n^{\alpha,G}, V')$ *is simultaneously defined by*

$$\begin{array}{rcll}
S' & := & S \times \{0,1\} \ , & \\
V'(p) & := & V(p) \times \{0,1\} \ , & \\
R_i^{\alpha,G} & := & \{(s_0, t_0) \mid sR_it\} \cup & \\
 & & \{(s_1, t_1) \mid sR_it \textit{ and } (\mathsf{K}, s \models B_i\neg\alpha \textit{ or } \mathsf{K}, t \models \alpha)\} & (i \in G) \ , \\
R_i^{\alpha,G} & := & \{(s_0, t_0) \mid sR_it\} \cup \{(s_1, t_0) \mid sR_it\} & (i \notin G) \ .
\end{array}$$

We use s_0 *and* s_1 *as abbreviations for* $(s, 0)$ *and* $(s, 1)$, *respectively. The validity relation* $\models$ *is indeed inductively defined, because* $|[\alpha]_G\beta| = |\alpha| + |\beta| + 2$.

We say that an $\mathcal{L}_n^{\mathsf{A}}$ formula α is valid in the Kripke structure K ($\mathsf{K} \models \alpha$), if and only if for all $s \in |\mathsf{K}|$, $\mathsf{K}, s \models \alpha$. The formula α is valid with respect to $\mathcal{K}_n^{stE}$ ($\mathcal{K}_n^{stE} \models \alpha$), if and only if for all $\mathsf{K} \in \mathcal{K}_n^{stE}$, $\mathsf{K} \models \alpha$. Further, we say that α is satisfiable in $\mathcal{K}_n^{stE}$, if and only if there is a $\mathsf{K} \in \mathcal{K}_n^{stE}$ and an $s \in |\mathsf{K}|$, s.t. $\mathsf{K}, s \models \alpha$.

It is time now for an example, which illustrates how our approach guarantees the consistency of belief before and after an announcement, even if a lie has been told.

Example 3 *Imagine a game with the players* 1, 2, *and* 3, *as well as the cards Ace, Queen, and another Queen. One card is dealt to each player and the players can only see their own card. Then the situation is represented by the Kripke structure* K *as follows.*

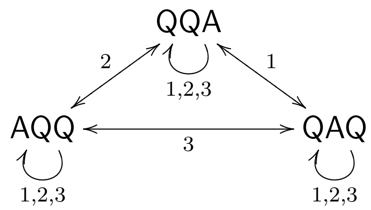

We need a proposition p_j *for each player* j *to express that this player has got the Ace. The valuation is given by*

$$\begin{aligned} V(p_1) &= \{\mathsf{AQQ}\} , \\ V(p_2) &= \{\mathsf{QAQ}\} , \\ V(p_3) &= \{\mathsf{QQA}\} . \end{aligned}$$

The only player who knows the deal is the player with the Ace.

Now player 1 *secretly tells player* 2, *that he does not have the Ace. No matter in which state we are, we can perform this announcement as described in definition 2. This means that* $\alpha = \neg p_1$, $G = \{1, 2\}$, *and hence,* $\mathsf{K}^{\alpha,G}$ *is the following Kripke structure.*

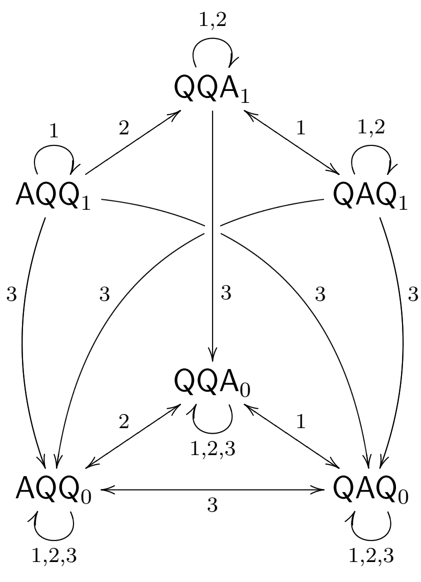

Observe that, no matter in which state we are before the announcement, player 2 *accepts the announcement and beliefs* $\neg p_1$ *afterwards. In the state*

QAQ*, player* 2 *learns nothing new, because she already believes that player* 1 *does not have the Ace. On the other hand, in the state* AQQ*, player* 2 *believes a lie, and player* 1 *rejects his own announcement, because he believes its negation* p_1*. This fact illustrates the essential difference of our procedure to the other approaches.*

Let us go back to the initial position and consider the case, that player 1 *secretly tells player* 2 *that he has got the Ace. Then player* 2 *will reject the announcement in the state* QAQ*, because she beliefs that player* 1 *does not have the Ace.*

The next lemma states, that $\mathsf{K}^{\alpha,G}$ is always in $\mathcal{K}_n^{stE}$, if the accessibility relations of K are serial, transitive, and Euclidean. This is what we mean by consistency preserving.

Lemma 4 *For all Kripke structures* K*, all* $\mathcal{L}_n^{\mathsf{A}}$ *formulas* α*, and all non-empty groups* $G \subseteq \mathcal{A}$ *of agents we have*

$$\mathsf{K} \in \mathcal{K}_n^{stE} \quad \Rightarrow \quad \mathsf{K}^{\alpha,G} \in \mathcal{K}_n^{stE} \ .$$

Proof Let K, α, and G be given as stated above and take an arbitrary $i \in \mathcal{A}$. We have to show that $R_i^{\alpha,G}$ is serial, transitive, and Euclidean. If $i \notin G$, it is easy to see that $R_i^{\alpha,G}$ has the desired properties. Now assume $i \in G$. The relation $R_i^{\alpha,G}$ is serial, because R_i is serial and definition 2 makes sure that agent i rejects the announcement, if he believes $\neg\alpha$. We concentrate on the proof of transitivity. Take $s, t, u \in |\mathsf{K}|$ and $k, l, m \in \{0, 1\}$ s.t. $s_k R_i^{\alpha,G} t_l$ and $t_l R_i^{\alpha,G} u_m$. Since $R_i^{\alpha,G}$ is always a subset of R_i, we know that sR_it, tR_iu, and, because of transitivity of R_i, we have sR_iu. If $k = 0$, it is immediate that $l = m = 0$, and we obviously get $s_0 R_i^{\alpha,G} u_0$. If $k = 1$, since $i \in G$, we have $l = m = 1$, and we distinct the following two cases. In the case $\mathsf{K}, s \models B_i\neg\alpha$, we immediately have $s_1 R_i^{\alpha,G} u_1$, because agent i rejects the announcement in the state s. In the case $\mathsf{K}, s \not\models B_i\neg\alpha$, we know that $\mathsf{K}, s \models B_i\neg B_i\neg\alpha$, because R_i is Euclidean. Since sR_it, we have $\mathsf{K}, t \not\models B_i\neg\alpha$. But now, we know that $\mathsf{K}, u \models \alpha$, because $t_1 R_i^{\alpha,G} u_1$. Therefore, we also have $s_1 R_i^{\alpha,G} u_1$ and we are done. The proof of Euclideanness is similar to the proof of transitivity. $\Box$

Finally, we want to mention that group announcements are in general not idempotent: there is no valid implication between $[\alpha]_G\beta$ and $[\alpha]_G^2\beta$ for arbitrary α, β, and G.

3 Axiomatization and properties

In this section we will give a sound and complete axiomatization $\mathsf{KD45}_n^{\mathsf{A}}$ for our announcement logic. We will define the notion of announcement-resistance for formulas with respect to a non-empty group of agents. These are the formulas which remain valid after arbitrary announcements to G. We will show that beliefs in announcement-resistant formulas can not be contracted by accepting a group announcement.

Definition 5 *The theory* $\mathsf{KD45}_n^{\mathsf{A}}$ *has the following axioms and rules.*

(PT)	*Every instance of a propositional tautology,*	
(K)	$B_i(\alpha \to \beta) \to (B_i\alpha \to B_i\beta)$,	
(D)	$B_i\neg\alpha \to \neg B_i\alpha$,	
(4)	$B_i\alpha \to B_iB_i\alpha$,	
(5)	$\neg B_i\alpha \to B_i\neg B_i\alpha$,	
(A1)	$[\alpha]_G p \leftrightarrow p$,	
(A2)	$[\alpha]_G(\beta \to \gamma) \to ([\alpha]_G\beta \to [\alpha]_G\gamma)$,	
(A3)	$[\alpha]_G\neg\beta \leftrightarrow \neg[\alpha]_G\beta$,	
(A4)	$\neg B_i\neg\alpha \to ([\alpha]_G B_i\beta \leftrightarrow B_i(\alpha \to [\alpha]_G\beta))$	$(i \in G)$,
(A5)	$B_i\neg\alpha \to ([\alpha]_G B_i\beta \leftrightarrow B_i[\alpha]_G\beta)$	$(i \in G)$,
(A6)	$[\alpha]_G B_i\beta \leftrightarrow B_i\beta$	$(i \notin G)$,

$$\text{(MP)}\ \frac{\alpha \qquad \alpha \to \beta}{\beta}\ , \qquad \text{(NEC.1)}\ \frac{\alpha}{B_i\alpha}\ , \qquad \text{(NEC.2)}\ \frac{\alpha}{[\beta]_G\alpha}\ .$$

The announcement axioms (A1) to (A6) are called reduction axioms, because they reduce the language $\mathcal{L}_n^{\mathsf{A}}$ to $\mathcal{L}_n$. We will see later, how the translation is established.

Lemma 6 *The system* $\mathsf{KD45}_n^{\mathsf{A}}$ *is sound with respect to* $\mathcal{K}_n^{stE}$*, i.e. for all* $\mathcal{L}_n^{\mathsf{A}}$ *formulas* α *we have*

$$\mathsf{KD45}_n^{\mathsf{A}} \vdash \alpha \quad \Rightarrow \quad \mathcal{K}_n^{stE} \models \alpha\ .$$

Proof The proof is by induction on the length of the derivation. In the base case, soundness of the axiom (A4) is proved as follows. Let $\mathsf{K} \in \mathcal{K}_n^{stE}$, $s \in |\mathsf{K}|$, $\emptyset \neq G \subseteq \mathcal{A}$, and $i \in G$ be given and assume that $\mathsf{K}, s \not\models B_i\neg\alpha$. Then we have

$$\begin{aligned}
\mathsf{K}, s \models [\alpha]_G B_i\beta \ \ &\text{iff} \ \ \mathsf{K}^{\alpha,G}, s_1 \models B_i\beta \\
&\text{iff} \ \ \text{for all } t \in S,\ s_1 R_i^{\alpha,G} t_1 \ \Rightarrow\ \mathsf{K}^{\alpha,G}, t_1 \models \beta \\
&\text{iff} \ \ \text{for all } t \in S,\ sR_it \text{ and } \mathsf{K}, t \models \alpha \ \Rightarrow\ \mathsf{K}, t \models [\alpha]_G\beta \\
&\text{iff} \ \ \mathsf{K}, s \models B_i(\alpha \to [\alpha]_G\beta)\ .
\end{aligned}$$

In the induction step, soundness of the rule (NEC.2) directly follows from lemma 4. □

As a preparatory step for the completeness proof, we will give some derivable formulas. The proof is left as an exercise.

Lemma 7 *The following axioms are derivable in* $\mathsf{KD45}_n^{\mathsf{A}}$.

1. $[\alpha]_G(\beta \wedge \gamma) \leftrightarrow ([\alpha]_G\beta \wedge [\alpha]_G\gamma)$,
2. $[\alpha]_G(\beta \vee \gamma) \leftrightarrow ([\alpha]_G\beta \vee [\alpha]_G\gamma)$,
3. $[\alpha]_G B_i\beta \leftrightarrow B_i[\alpha]_G\beta \vee (\neg B_i\neg\alpha \wedge B_i(\alpha \to [\alpha]_G\beta))$.

Our completeness proof can be done via a translation from $\mathcal{L}_n^{\mathsf{A}}$ to $\mathcal{L}_n$, because our announcement logic has the same expressive strength as normal modal logic.

Definition 8 *The translation t from $\mathcal{L}_n^{\mathsf{A}}$ formulas to $\mathcal{L}_n$ formulas is inductively defined as follows.*

$$\begin{aligned} t(p) &:= p \ , \\ t(\neg\alpha) &:= \neg t(\alpha) \ , \\ t(\alpha \wedge \beta) &:= t(\alpha) \wedge t(\beta) \ , \\ t(B_i\alpha) &:= B_i t(\alpha) \ , \\ t([\alpha]_G\beta) &:= h([t(\alpha)]_G t(\beta)) \ . \end{aligned}$$

In order to make t a function from $\mathcal{L}_n^{\mathsf{A}}$ to $\mathcal{L}_n$, we define the translation h from the set $\{[\alpha]_G\beta \mid \alpha, \beta \in \mathcal{L}_n\}$ to $\mathcal{L}_n$ by

$$\begin{aligned} h([\alpha]_G p) &:= p \ , \\ h([\alpha]_G\neg\beta) &:= \neg h([\alpha]_G\beta) \ , \\ h([\alpha]_G(\beta \wedge \gamma)) &:= h([\alpha]_G\beta) \wedge h([\alpha]_G\gamma) \ , \\ h([\alpha]_G B_i\beta) &:= B_i h([\alpha]_G\beta) \vee (\neg B_i\neg\alpha \wedge B_i(\alpha \to h([\alpha]_G\beta))) \quad (i \in G) \ , \\ h([\alpha]_G B_i\beta) &:= B_i\beta \quad (i \notin G) \ . \end{aligned}$$

It is obvious that for every $\mathcal{L}_n^{\mathsf{A}}$ formula α, its translation $t(\alpha)$ is a formula of $\mathcal{L}_n$. In addition, we can prove the equivalence of α and $t(\alpha)$ in $\mathsf{KD45}_n^{\mathsf{A}}$.

Lemma 9 *For every $\mathcal{L}_n^{\mathsf{A}}$ formula α we have*

$$\mathsf{KD45}_n^{\mathsf{A}} \ \vdash \ \alpha \leftrightarrow t(\alpha) \ .$$

Proof This lemma can be proved by induction on α. In the induction step, in the case α is of the form $[\beta]_G\gamma$, we need two more properties.

1. $\mathsf{KD45}_n^{\mathsf{A}} \vdash [\varphi]_G\delta \leftrightarrow h([\varphi]_G\delta)$ for every $\varphi, \delta \in \mathcal{L}_n$,
2. $\mathsf{KD45}_n^{\mathsf{A}} \vdash \varphi \leftrightarrow \psi \ \Rightarrow \ \mathsf{KD45}_n^{\mathsf{A}} \vdash [\varphi]_G\delta \leftrightarrow [\psi]_G\delta$ for every $\delta \in \mathcal{L}_n$.

Both properties can easily be proved by induction on δ. $\square$

The previous lemma is very helpful for proofs by induction on arbitrary $\mathcal{L}_n^{\mathsf{A}}$ formulas, because we need it in the last case of the induction step. For instance, the following property holds for arbitrary $\mathcal{L}_n^{\mathsf{A}}$ formulas γ,

$$\mathsf{KD45}_n^{\mathsf{A}} \vdash \alpha \leftrightarrow \beta \quad \Rightarrow \quad \mathsf{KD45}_n^{\mathsf{A}} \vdash [\alpha]_G\gamma \leftrightarrow [\beta]_G\gamma \ .$$

Lemma 9 is also helpful to prove completeness of $\mathsf{KD45}_n^{\mathsf{A}}$. We have indeed a completeness proof with maximal consistent sets, but it is much longer.

Lemma 10 *The system $\mathsf{KD45}_n^{\mathsf{A}}$ is complete with respect to $\mathcal{K}_n^{stE}$, i.e. for all $\mathcal{L}_n^{\mathsf{A}}$ formulas α we have*

$$\mathcal{K}_n^{stE} \models \alpha \quad \Rightarrow \quad \mathsf{KD45}_n^{\mathsf{A}} \vdash \alpha \ .$$

Proof From $\mathcal{K}_n^{stE} \models \alpha$ we get $\mathcal{K}_n^{stE} \models t(\alpha)$ by soundness and lemma 9. By completeness of $\mathsf{KD45}_n$, we have $\mathsf{KD45}_n \vdash t(\alpha)$ and $\mathsf{KD45}_n^{\mathsf{A}} \vdash t(\alpha)$, because $\mathsf{KD45}_n^{\mathsf{A}}$ is an extension of $\mathsf{KD45}_n$. By lemma 9 we immediately get $\mathsf{KD45}_n^{\mathsf{A}} \vdash \alpha$, and we are done. $\square$

In a next step, we will define the notion of announcement-resistance for $\mathcal{L}_n^{\mathsf{A}}$ formulas with respect to a non-empty group of agents. This definition differs from the notion of successful formulas in literature, cf. van Ditmarsch in [11]. The reason for the difference is, that in our approach, not even propositions are successful formulas in the original sense.

Definition 11 *Let $\emptyset \neq G \subseteq \mathcal{A}$. An $\mathcal{L}_n^{\mathsf{A}}$ formula α is called announcement-resistant for G, if for all $\mathcal{L}_n^{\mathsf{A}}$ formulas β we have*

$$\mathsf{KD45}_n^{\mathsf{A}} \ \vdash \ \alpha \rightarrow [\beta]_G\alpha \ .$$

Observe, that $\mathsf{KD45}_n^{\mathsf{A}}$ proves $\alpha \rightarrow [\beta]_G^k\alpha$ for all $k \geq 0$, if α is announcement-resistant for G. The following lemma shows, that for every G, a huge set of formulas is announcement-resistant for G.

Lemma 12 *Let $\emptyset \neq G \subseteq \mathcal{A}$. Then we have*

1. *for all $p \in \mathcal{P}$, the literals p and $\neg p$ are announcement-resistant for G.*
2. *for all $\alpha \in \mathcal{L}_n^{\mathsf{A}}$ and $i \in \mathcal{A} \setminus G$, the formulas $B_i\alpha$ and $\neg B_i\alpha$ are announcement-resistant for G,*
3. *all $\mathcal{L}_n^{\mathsf{A}}$ formulas provable in $\mathsf{KD45}_n^{\mathsf{A}}$ are announcement-resistant for G,*
4. *if α and β are announcement-resistant for G, the formulas $\alpha \wedge \beta$ and $\alpha \vee \beta$ are announcement-resistant for G,*

5. *if α is announcement-resistant for G and $i \in G$, the formula $B_i\alpha$ is announcement-resistant for G.*

PROOF The claims 1 to 3 directly follow from the axioms ($\mathsf{A1}$), ($\mathsf{A6}$), and the rule ($\mathsf{NEC.2}$). The Claims 4 and 5 can be established using lemma 7. $\Box$

As an immediate consequence of the previous lemma, we get the following fact. For all $\mathcal{L}_0$ formulas α, all non-empty $G \subseteq \mathcal{A}$, and all $i \in \mathcal{A}$, the formula $B_i\alpha$ is announcement-resistant for G. This means, that beliefs in propositional formulas can never be contracted by group announcements. We can therefore say, that our approach formalizes belief expansion for propositional belief.

The next lemma shows, that under certain conditions, the agents can really learn new sentences, as we have already seen in example 3. The proof is left to the reader.

Lemma 13 *Let $\emptyset \neq G \subseteq \mathcal{A}$ and α be announcement-resistant for G. Then for all $i \in G$ and $k, m \geq 1$ we have*

$$\mathsf{KD45}_n^{\mathsf{A}} \ \vdash \ \neg B_i \neg\alpha \to [\alpha]_G^k B_i^m \alpha \ .$$

We immediately get that $\mathsf{KD45}_n^{\mathsf{A}}$ proves $\neg B_i\alpha \wedge \neg B_i\neg\alpha \ \to \ [\alpha]_G^k B_i^m\alpha$, so the agents in G can learn the announcement-resistant formulas for G. This lemma shows, that with the right precondition ($\neg B_i\neg\alpha$), the announcement-resistant formulas for G are in a way the successful formulas for G.

4 Conclusions and future work

Since we present a system for modal belief change, we show that our approach can successfully formalize the muddy children puzzle.

Example 14 *It is easy to verify, that in the Kripke structure K for n children and in a state s where the children 1 to m are muddy, it holds that*

$$\mathsf{K}, s \ \models \ [\mathit{father}]_{\mathcal{A}}[\mathit{no}]_{\mathcal{A}}^{m-1} \bigwedge_{i=1}^{m} B_i \mathit{dirty}_i \ .$$

For the definition of the formulas father, no, and dirty_i, see [4].

For $n = 3$, the following Kripke structure illustrates the relevant connected *part of* $((\mathsf{K}^{father,\mathcal{A}})^{no,\mathcal{A}})^{no,\mathcal{A}}$*:*

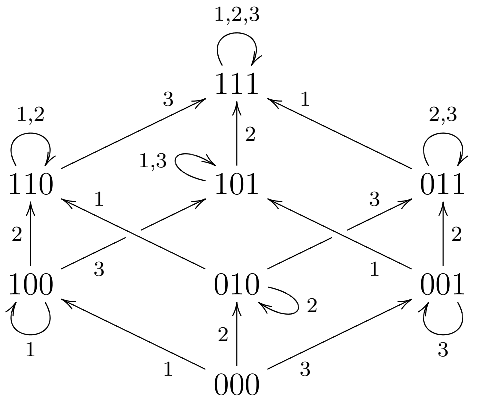

Observe, that no matter which children are muddy, every agent has consistent belief, even if some of the announcements were false.

As we have seen in the last section, propositional beliefs are always expanded. But arbitrary beliefs can be contracted: Since $\mathsf{KD45}_n^{\mathsf{A}}$ proves $\neg B_i p \wedge \neg B_i \neg p \rightarrow [p]_{\mathcal{A}} B_i p$, we can have $B_i \neg B_i p$ before an announcement and $\neg B_i \neg B_i p$ afterwards. Hence, on a given Kripke-world K, s, an announcement with the formula α to the group G defines a non-monotone operator on the set of agent i's beliefs. It will be interesting to investigate this operator in the future.

This paper can be seen as the first part of belief revision for $\mathsf{KD45}_n$. Until now, we have formalized expansion of propositional beliefs. The next step will be to define contraction $\dot{-}$ of propositional beliefs. It still has to be pure belief change without any factual change. Then, it will be straightforward to define revision $\dotplus$ for propositional beliefs using the Levi identity,

$$[\dotplus\alpha]_G\beta \quad := \quad [\dot{-}\neg\alpha]_G[\alpha]_G\beta \ .$$

Due to lemma 9 we know that the language $\mathcal{L}_n^{\mathsf{A}}$ has the same expressive strength like the language $\mathcal{L}_n$ of normal modal logic. In a worst case, the length of the translation of an $\mathcal{L}_n^{\mathsf{A}}$ formula is exponential. So we will do a complexity analysis of our announcement logic.

A next challenge will also be to extend our announcement logic with the notion of common belief. As in related approaches, we conjecture that adding announcement operators to common belief increases the expressive strength. We also think that in this context, group announcements are more expressive than announcements to single agents.

There is an approach of dynamic epistemic logic by Baltag, Moss, and Solecki in [2], where the actions are not only formulas, but action structures. It is possible to extend our approach to this more general definition, because the cartesian product of two $\mathcal{K}_n^{stE}$ structures is again a $\mathcal{K}_n^{stE}$ structure. For this purpose, the difficulty will be the complete axiomatization of this new logic expanded with common belief.

Last but not least, we want to mention that we will investigate, how our announcement logic can be applied to deal with security protocols. There is an approach of Hommersom, Meyer, and de Vink in [6], where announced formulas have to be true. We can also deal with false announcements, thus we belief our procedure could be some improvement.

References

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

[2] Baltag, A., Moss, L. S., and Solecki, S. The logic of public announcements, common knowledge, and private suspicions. In *TARK '98: Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge* (San Francisco, CA, USA, 1998), Morgan Kaufmann Publishers Inc., pp. 43–56.

[3] Friedman, N., and Halpern, J. Y. Belief revision: A critique. In *KR'96: Principles of Knowledge Representation and Reasoning*, L. C. Aiello, J. Doyle, and S. Shapiro, Eds. Morgan Kaufmann, San Francisco, California, 1996, pp. 421–431.

[4] Gerbrandy, J., and Groeneveld, W. Reasoning about information change. *Journal of Logic, Language and Information 6*, 2 (1997), 147–169.

[5] Herzig, A. Logics for belief base updating. In *Handbook of defeasible reasoning and uncertainty management* (1998), D. Dubois and H. Prade, Eds., vol. Belief Change, Kluwer Academic Publishers.

[6] Hommersom, A., Meyer, J.-J. C., and de Vink, E. P. Update semantics of security protocols. *Synthese 142* (2004), 229–267.

[7] Katsuno, H., and Mendelzon, A. O. On the difference between updating a knowledge base and revising it. In *KR'91: Principles of*

Knowledge Representation and Reasoning (1991), J. A. Allen, R. Fikes, and E. Sandewall, Eds., Morgan Kaufman, pp. 387–394.

[8] Roorda, J.-W., van der Hoek, W., and Meyer, J.-J. C. Iterated belief change in multi-agent systems. *Logic Journal of the IGPL 11*, 2 (2003), 223–246.

[9] Steiner, D., and Studer, T. Systems for public announcements with common knowledge (working title), 2006. In preparation.

[10] van Ditmarsch, H. *Knowledge games.* PhD thesis, University of Groningen, 2000. ILLC Dissertation Series 2000-06.

[11] van Ditmarsch, H. The Russian cards problem. *Studia logica 71* (2003), 1–32.

[12] van Linder, B., van der Hoek, W., and Meyer, J.-J. C. Actions that make you change your mind. In *Knowledge and Belief in Philosophy and Artificial Intelligence* (1995), A. Laux and H. Wansing, Eds., pp. 104–146.