



HAL
open science

Arbitrary Announcements in Propositional Belief Revision

Aaron Hunter, François Schwarzentruber

► **To cite this version:**

Aaron Hunter, François Schwarzentruber. Arbitrary Announcements in Propositional Belief Revision. DARE@IJCAI, 2015, Buenos Aires, Argentina. hal-02534088

HAL Id: hal-02534088

<https://hal.science/hal-02534088>

Submitted on 6 Apr 2020

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.

Arbitrary Announcements in Propositional Belief Revision

Aaron Hunter

British Columbia Institute of Technology
Burnaby, Canada
aaron_hunter@bcit.ca

Francois Schwarzenruber

ENS Rennes
Bruz, France
francois.schwarzenruber@ens-rennes.fr

Abstract

Public announcements cause each agent in a group to modify their beliefs to incorporate some new piece of information, while simultaneously being aware that all other agents are doing the same. Given some fixed goal formula, it is natural to ask if there exists an announcement that will make the formula true in a multi-agent context. This problem is known to be undecidable in a general modal setting, where the presence of nested beliefs can lead to complex dynamics. In this paper, we consider not necessarily truthful public announcements in the setting of propositional belief revision. We are given a goal formula for each agent, and we are interested in finding a single announcement that will make each agent believe the corresponding goal following AGM-style belief revision. If the goals are inconsistent, then this can be seen as a form of ampliative reasoning. We prove that determining if there is an arbitrary public announcement in this setting is not only decidable, but that it is simpler than the corresponding problem in the most simplified modal logics. Moreover, we argue that propositional announcements and beliefs are sufficient for modelling many practical problems, including simple robot controllers.

1 Introduction

We are concerned with the manner in which the beliefs of a group of agents can be manipulated by making group announcements. Following the highly influential work of [BMS98], this issue has been studied extensively in modal logics of public announcement. We are interested in the propositional case involving belief revision operators. As such, we assume that each agent i has some initial belief set K_i and as well as a goal formula ψ_i . We would like to find a single consistent formula ϕ that can be announced to all agents such that $K_i * \phi \models \psi_i$ for all i , where $*$ represents some suitable model of AGM-style belief revision¹. Note that the conjunction of the goal formulas $\bigwedge_i \psi_i$ need not to be

¹AGM stands for the highly influential approach to belief revision due to Alchourrón, Gärdenfors and Makinson [AGM85].

consistent, so announcing $\bigwedge_i \psi_i$ is not a solution in all cases. The problem that we address roughly corresponds to the notion of an *arbitrary public announcement* [BBvD⁺07] in the tradition of Dynamic Epistemic Logic (DEL) [vDvdHK08], except that we do not require announcements to be truthful. However, the problem is far simpler computationally in the setting of propositional revision. As such, we suggest that our approach is appropriate for practical applications, such as robot controllers, where the focus is on giving factual orders in the most efficient way possible.

This paper makes several contributions to existing work on the theory of belief change. First, we define arbitrary public announcements in an AGM setting. We demonstrate that the question is meaningful in this context, as there are natural applications where a reasonable treatment of announcements need not be concerned with nested beliefs. A non-trivial second contribution of the paper is a detailed analysis of the complexity of arbitrary announcements for a particular AGM revision operator, namely Dalal's revision operator [Dal88]. Finally, from a high-level perspective, this paper makes a contribution towards establishing the utility of AGM-style operators with respect to the (more expressive) DEL approach to belief change.

We do not define arbitrary public announcements for a full Dynamic epistemic logic for belief revision ([VB07], [BS06], [BS08]) because it would certainly lead to undecidability issues [FvD08]. We concentrate here on the propositional case: announcements are propositional and we are only interested in propositional beliefs. Many natural reasoning problems can be captured in this setting, including the motivating example below. We are interested in computational issues in this context.

Motivating Example Consider a domain involving a robot controller, and n robots that act independently. Each robot has beliefs about the current state of the world, and these beliefs are used to formulate goals. We assume that the goals are given as input to some form of planner, but we are not concerned with reasoning about actions here. Assume that the only way the controller can communicate with the robots is through broadcast messaging. If there are constraints on messaging in terms of cost or timing, then the controller may seek to minimize the number of messages sent to ensure all robots have the correct goals. Towards this end, it can be use-

ful to send messages that will impact the beliefs of each robot differently.

As a concrete example, suppose that there are two surveillance robots R_1 and R_2 that are tasked with patrolling a certain area. If there is a breach at the gate, we would like one robot to go check the gate while the other continues patrolling. In order to achieve the desired behaviour, the controller would like to broadcast a single alarm message “breach” that will lead each robot to simultaneously have the appropriate beliefs about how they should behave. In this paper, we formalize this kind of reasoning in an abstract setting.

2 Preliminaries

2.1 Belief Revision

A *vocabulary* is a countable set of propositional variables \mathbf{P} . An *interpretation* over \mathbf{P} is a function that maps each variable to a truth value, and assigns truth values to propositional formulas in the usual way. The study of *belief change* is concerned with the way that rational agents incorporate new information. One important form of belief change is *belief revision*, which is the change that occurs when the new information may be inconsistent with some of the original beliefs. One highly influential approach to belief revision has been the AGM approach, in which belief revision is captured by an operator that satisfies a particular set of rationality postulates [AGM85]. In this tradition, the beliefs of an agent are represented by a *belief set*, which is a logically closed set of formulas; if we have a finite vocabulary, this can equivalently be a single formula. A belief revision operator takes a belief set and a formula as input, and returns a new belief set.

It is well-known that every AGM revision operator $*$ has the property that, for any belief set K , there is an underlying total pre-order \prec_K such that $K * \phi = \min_{\prec_K}(\phi)$; so AGM revision involves finding minimal worlds consistent with a particular formula [KM92]. One problem with the AGM approach is that, in general, it does not handle the problem of *iterated belief revision*. The problem is that you have an ordering on states before the revision, but you only have a set of states after the revision.

Iterated revision requires a specification of exactly how the ordering changes when presented with a formula for revision [DP97; JT07]. As such, iterated belief revision operators are commonly defined in terms of *epistemic states*. An epistemic state is a total pre-order \prec over the set of states. A *belief change operator* $*$ takes a total pre-order \prec and a formula ϕ as input, and it returns a new epistemic state denoted $\prec * \phi$. The set $\min(\prec)$ is understood to represent the set of states that are currently believed by a given agent. A DP (Darwiche and Pearl) revision operator satisfies the postulates introduced in [DP97]. The key distinction here is that DP revision is concerned with returning a completely new ordering following revision, so we are able to revise again if necessary.

2.2 Dynamic Epistemic Logic

An alternative approach to reasoning about beliefs has been developed using variations of Dynamic epistemic logic [vDvdHK08]. The semantics of these logics is defined with respect to Kripke structures. A Kripke structure is a tuple

$M = \langle W, R_1, \dots, R_n, V \rangle$ where W is a non-empty set of *possible worlds*, each R_i is a binary *accessibility relation* on W , and V is a function that assigns an interpretation to each $w \in W$. Formulas in modal logic can take the form $\Box_i \phi$; and we have $M, w \models \Box_i \phi$ just in case $(M, v) \models \phi$ for every v with $(w, v) \in R_i$. Standard modalities include K (knowledge), B (belief), and the parametrized public announcement modality $[\phi]$ of dynamic epistemic logic. The semantics of $[\phi]$ is defined as a restriction to possible worlds where ϕ holds. We refer the reader to [vDvdHK08] for a survey of work in this area.

One important distinction between the AGM approach and the DEL approach is the fact that the DEL approach permits the representation of nested beliefs. As a result, it is much more expressive. Also, some of the properties we expect for belief change no longer hold. For example, consider the *Success* postulate for AGM revision, which states that $\phi \in K * \phi$. In a modal setting, this is essentially equivalent to the schema $[\phi]B\phi$. This is not valid in most multi-agent epistemic logics, because an announcement of ϕ often changes the world such that ϕ should not actually be believed.²

This leads to a natural question: how can we get another agent to believe some formula ϕ ? This question is addressed in DEL via the logic of arbitrary public announcements [BBvD⁺07]. Note that the problem is not identical, however, to the problem addressed in this paper. First of all, we do not have belief revision operators in public announcement logic; we have only hard updates. Furthermore, in DEL we require announcements to be truthful. In arbitrary public announcement logic, through a modality \diamond such that $\diamond\phi$ is true just in case there is some truthful formula ψ such that announcing ψ would make ϕ true. If ϕ has the form $B_1\phi_1 \wedge \dots \wedge B_n\phi_n$, then we have a DEL reformulation of the variant of the problem that we address in this paper, modulo hard updates and truth of announcements. We invite the reader to refer to [vDvdHK04] for a precise comparison between public announcement and belief revision.

2.3 Dynamic epistemic logic with plausibility models

A version of DEL has been defined to incorporate belief revision operators ([VB07], [BS06], [BS08]). In this logic, the epistemic relations for belief operators do not rely on fixed (traditionally KD45) relations on possible worlds. Instead, for each agent a , possible worlds are ordered in each equivalence class for a . The semantics of $[\phi]$ is defined in terms of model transformations that re-order the orderings over worlds. As far we know, arbitrary public announcements have never been considered in this context.

2.4 Complexity Results

In this section, we give some known complexity results related to public announcements. The first result is really the starting point for this work.

Theorem 1 ([FvD08]) *Satisfiability in the logic of arbitrary announcements is undecidable.*

²This is the case in so-called Moore sentences, for example.

Hence, if we have nested knowledge and a rich domain of Kripke structures, we can not hope to find public announcements to change the beliefs of agents in the desired manner. A simpler logic called *DLPA-APAL* (for dynamic logic of propositional assignments with arbitrary public announcement logic) is defined in [CS15]. In *DLPA-APAL*, Kripke structures are restricted to include just one world with each valuation, and epistemic relations are represented by programs over propositional assignments.

Theorem 2 ([CS15]) *Satisfiability in DLPA-APAL is NEXPTIME-complete for the single announcement case.*

This result shows that restricting to reasoning over beliefs on valuations is decidable. Our goal in this paper is to show that if we restrict further, by removing nested beliefs, we can get a better complexity even if we use belief revision operators and we do not force announcements to be true.

3 Propositional Effects of Announcements

3.1 The Basic Problem

Given n agents and an AGM revision operator $*$, we are looking for the existence of a consistent formula ϕ such that

$$\begin{aligned} K_1 * \phi &\models \psi_1 \\ K_2 * \phi &\models \psi_2 \\ &\vdots \\ K_n * \phi &\models \psi_n \end{aligned} \quad (1)$$

where K_i and ψ_i represent the belief set and the goal of the agent i , respectively. Obviously, if $\bigwedge_i \psi_i$ is consistent, then we can just revise by this. But this is not always the case.

Note that we do not require the announcement to be true as it is done in public announcement logic. We remark also that, in a propositional setting, there is no way to express an assertion of the form “if you are agent i , then revise by ϕ .” As such, there is no obvious way to specify a formula that explicitly specifies a different revision for each agent.

As formulated above, the announcement problem can be seen as a form of *distributed ampliative reasoning*. In general, if $\bigwedge_i \psi_i$ is not consistent, then it is not possible to find a consistent formula to announce that will result in a single agent believing each ψ_i . But in our setting, we are actually able to get some agent to believe each ψ_i . Hence, the announcement of ϕ is resulting in a situation where the *distributed beliefs* of the agents in the system are inconsistent; since the formula ϕ itself is consistent, this means that some of the revising agents are “jumping to conclusions” that are not strictly supported by the announcement.

Example Consider the robot controller example over the vocabulary $\{patrol, checkgate\}$. We think of this vocabulary as defining a state machine, where each interpretation represents a state; the robot can have actions that are triggered by transitions to given states. This is a standard control mechanism for simple agents in a video game setting, and it can function as a control mechanism for our simple robot agents as well. In our example, the *patrol* variable is true when the

robot should be patrolling their area and the *checkgate* variable is true when the robot is supposed to check the gate.

Suppose we have two robots R_1 and R_2 with initial belief states defined as follows:

$$\begin{aligned} Bel(R_1) &= \{\neg patrol \vee checkgate\} \\ Bel(R_2) &= \{\neg patrol \vee \neg checkgate\} \end{aligned}$$

The controller believes there is a problem at the gate. Is there a formula that can be broadcast to immediately get R_1 to check the gate while R_2 patrols the the grounds? In other words, is there a formula ϕ such that:

$$\begin{aligned} \{\neg patrol \vee checkgate\} * \phi &\models checkgate \\ \{\neg patrol \vee \neg checkgate\} * \phi &\models \neg checkgate \wedge patrol \end{aligned}$$

The answer here is clearly yes because we can set $\phi = patrol$.

The preceding example is framed in the context of the robot controller, but it also demonstrates an important case for propositional announcement. In particular, it shows that there are cases where the goals are inconsistent, yet a solution is possible. We are interested in determining how difficult it is to find such a solution.

3.2 Decision problem

In this section, we give an algorithm to solve the problem of interest. But, let us consider the corresponding decision problem.

The Propositional Announcement Problem ($PAP(*)$)

Input:

- An integer n
- A list K_1, \dots, K_n of formulas (initial beliefs).
- A list ψ_1, \dots, ψ_n of formulas (goals).

Output:

- Yes, if there exists ϕ satisfying (1)
- No, otherwise.

We refer to this problem as $PAP(*)$ to emphasize that it depends on some given operator $*$ on belief sets. In this paper, we will assume $*$ is an AGM revision operator, but that need not be the case in general.

We may reformulate $PAP(*)$ in a variant of the model checking problem of Dynamic epistemic logic with plausibility models ([VB07], [BS06]). We create a (potentially exponentially bigger than K_1, \dots, K_n) pointed Kripke model \mathcal{M}, w such that the valuations of the most plausible worlds in w are exactly models of K_i for all agents $i \in \{1, \dots, n\}$. Then we ask whether there exists a propositional formula ϕ such that $\mathcal{M}, w \models [\phi]_{BR}(B_1\psi_1, \dots, B_n\psi_n)$ where $[\phi]_{BR}$ is the announcement operator that semantically performs the belief revision for all agents and B_i is the belief operator for agent i .

For the moment, we are interested in analyzing the simplest case to obtain the most efficient algorithm possible. Hence, we consider Dalal’s revision operator based on the Hamming distance between interpretations [Dal88], which also has the advantage that it allows iterated revision. Let $*_d$ be the Dalal’s revision operator.

We present a non-deterministic algorithm that decides $PAP(*_d)$. In the algorithm, we use non-deterministic choice to select the minimum distance d_i between K_i and ψ_i . Also note that $d(K, v)$ denotes the minimum Hamming distance between the set of interpretations representing by formula K and interpretation v .

FIND_ANN ($K_1, \dots, K_n, \psi_1, \dots, \psi_n$)

0. Let m be the size of the underlying input vocabulary of $K_1, \dots, K_n, \psi_1, \dots, \psi_n$
 1. Guess $d_1, \dots, d_n \in \{0, 1, \dots, m\}$
 2. Guess valuations v_1, \dots, v_n
 3. For all $i, j \in \{1, \dots, n\}$
 - If $d(K_j, v_i) < d_j$, reject.
 4. For all $i \in \{1, \dots, n\}$
 - If $d(K_i, v_i) > d_i$, reject.
 5. For all $i, j \in \{1, \dots, n\}$
 - If $d(K_j, v_i) = d_j$ and $v_i \not\models \psi_j$, reject.
6. Accept.

We need to prove that **FIND_ANN**($\bar{K}, \bar{\psi}$) actually produces the desired result. In the proof, we write $\min_{\leq K_i}(\phi)$ for the minimal Hamming distance from a model of $\bar{\phi}$ to a model of K_i .

Theorem 3 *Let $\bar{K} = K_1, \dots, K_n$ and let $\bar{\psi} = \psi_1, \dots, \psi_n$ be sequences of formulas. Then **FIND_ANN**($\bar{K}, \bar{\psi}$) accepts if and only if there exists ϕ such that $K_i *_d \phi \models \psi_i$ for each i .*

Proof Suppose that **FIND_ANN**($\bar{K}, \bar{\psi}$) accepts. Let ϕ be a formula with $\|\phi\| = \bigcup_i \{v_i\}$, where the valuations v_i are those from an accepting run chosen on lines 1 and 2. By line 3,

$$\min_{\leq K_i}(\phi) \geq d_j$$

because if this is not the case, the valuations would have been rejected. But then, by line 4,

$$\min_{\leq K_i}(\phi) = d_j.$$

In other words, for each i , the models of ϕ that are models of K_i are distance d_i from K_i . Finally, since line 5 also is false for an accepting run, it follows that for every v such that $v \models \phi$, it must be the case that $d(K_j, v) = d_j$ implies $v \models \psi_j$. Hence, the minimal models of ϕ with respect to K_i are also models of ψ_i . So $K *_d \phi \models \psi_i$, which is what we wanted to show.

Now suppose that there exists ϕ such that $K_i *_d \phi \models \psi_i$ for each i . For each i , define $d_i = \min_{\leq K_i}(\phi)$. For each i , let v_i denote one of the valuations of ψ_i such that $d(K_i, v_i) = d_i$. We know that such a v_i exists, since $K_i *_d \phi \models \psi_i$. If we use $\bar{d} = d_1, \dots, d_n$ and $\bar{v} = v_1, \dots, v_n$ on lines 1 and 2 of the algorithm, then the choice of these values leads to an accepting run. \square

So **FIND_ANN**($\bar{K}, \bar{\psi}$) does what we want it to do: it determines if there is a public announcement ϕ such that revision by ϕ for each agent satisfies all ψ_i . We can also say something about the complexity. The initial guessing is a polynomial number of guesses with respect to the input length. Note that in the algorithm, tests $d(K_j, v_i) < d_j$ etc. can be answered by an NP oracle. Hence, the algorithm intuitively seems to lie

in the complexity class $\Sigma_2^P = NP^{NP}$. The following result makes this claim precise.

Theorem 4 *$PAP(*_d)$ is in Σ_2^P .*

Proof After the initial guesses, the algorithm clearly runs in polynomial time other than the checks of the form $d(K, v) < d$ at several stages. But recall that we are using the Hamming distance here, so this check can be performed as follows. Guess a set of atomic propositional variables of size less than d , and let v' be the interpretation obtained from v by switching the truth values of these variables. If $v' \models K$, then the minimum distance between K and v is less than d . This process is in NP .

Hence the entire algorithm runs in time NP with a polynomial number of calls to an NP oracle. This gives the complexity $NP^{NP} = \Sigma_2^P$, which was the desired result. \square

Thus, finding an announcement is far much efficient than in the modal case, even if we restrict possible worlds to be valuations. Note that hardness of $PAP(*_d)$ is left as an open issue.

So far, we have restricted attention to Dalal's revision operator. The primary motivation for this choice is the fact that Dalal's operator permits iterated revision, which we anticipate will be important in practice for a robot controller. Inspection of our proofs shows that the actual properties of the Hamming distance operator only occur in two places.

- When the minimum distance is guessed in the algorithm.
- In a number of checks of the form $d(K, v) < d$.

In these cases, we are able to have a polynomial number of operations with the Dalal's revision operator, because the distance only ranges from 0 to the size of the vocabulary. However, this is obviously not true for other revision operators; we return to this point in the conclusion.

4 Existence of Solutions

In this section, we give some basic results related to the existence of solutions for the propositional announcement problem. Note that the problem we have considered thus far assumes a single AGM revision operator that is used by all agents. This is, of course, not the most general setting. A more general version of our problem takes n agents and n belief revision operators $*_i$ as input; we are still looking for the existence of a consistent formula ϕ such that

$$\begin{aligned} K_1 *_1 \phi &\models \psi_1 \\ K_2 *_2 \phi &\models \psi_2 \\ &\vdots \\ K_n *_n \phi &\models \psi_n \end{aligned} \tag{2}$$

where K_i and ψ_i represent the belief set and the goal of the agent i , respectively. Note that this is a slightly more general problem than the problem we addressed previously; we will state our results in this section with respect to this formulation of the announcement problem.

The following result is immediate.

Theorem 5 For all $n > 1$, there are instances of (2) with no solution.

Proof Assume the vocabulary $\{p\}$, let $K_1 = \{p\}$, $K_2 = \{\neg p\}$, $\psi_1 = \neg p$ and $\psi_2 = p$. It follows immediately that any solution ϕ satisfies $\phi \models p \wedge \neg p$, so there is no consistent solution. \square

The following result may be less obvious, but it is similarly straightforward. The idea is simply that, as long as there are “enough” propositional variables, we can set up a set of inconsistent goals that can be guaranteed by an appropriate revision.

Theorem 6 Let n be a fixed natural number, and let \mathbf{P} be a propositional vocabulary. If $|\mathbf{P}| > \log_2 n$, then there is an instance of (2) over \mathbf{P} with inconsistent goals that has a solution.

Proof Assume the instance size n is given. Let P_0, \dots, P_n be a list of mutually inconsistent maximal conjunctions over literals of \mathbf{P} ; we know that such a list exists, because there are $2^{|\mathbf{P}|}$ distinct maximal conjunctions of literals, and we know that $2^{|\mathbf{P}|} > n$ by assumption. For each i , we set:

$$K_i = (\neg P_0 \vee P_i) \wedge \bigwedge_{j \neq i} \neg P_j$$

$$\psi_i = P_i \wedge \bigwedge_{j \neq i} \neg P_j$$

Note that the conjunction $\bigwedge_i \psi_i$ is clearly inconsistent. If we set $\phi = P_0$, then it is easy to verify that $K_i * \phi \models \psi_i$ for each i , which is what we wanted to show. \square

These results demonstrate that the general propositional announcement problem is non-trivial. There are instances with no solution and there are instances with solutions, and we can not easily identify which is the case without considering the possible revisions.


5 Prototype Development



As noted previously, we are interested in using propositional announcements in a practical implementation of a robot controller. Unfortunately, the algorithm provided in this paper is for a decision problem, so it does not actually return the announcement. However, it is easy to modify the algorithm to return the induced formula ϕ .

We have developed a simple implementation that calculates the formula suggested by our algorithm, and performs revision by the given announcement in the two agent case. In the demo, a number of variables are introduced to represent different weather conditions. Notably, there is a variable b indicating that an agent believes there is a rainbow and there is a variable c indicating that it is cloudy. A screenshot is provided in Figure 1. Suppose that initially agent 1 believes $c \rightarrow \neg b$ and agent 2 believes $c \rightarrow b$ (in the software, we may reach this situation by announcing privately $c \rightarrow \neg b$ to agent 1 and announcing privately $c \rightarrow b$ to agent 2). In the Figure, we would like agent 1 to believe c and $\neg b$ whereas we would like agent 2 to believe both c and b . The software

Arbitrary announcement in belief revision

We provide a pedagogical demonstration for belief revision. The following picture depicts the possible worlds for agent 1 and 2. You can make private announcements to agent 1 or agent 2 and public announcements. You can also ask for existence of a public announcement so agent 1 and agent 2 believe some properties.

Agent 1


Agent 2



Specified announcement

((((not b) and c) and ((not r) and (not s))) or ((b and c) and ((not r) and (not s))))

Announce privately to agent 1

Announce privately to agent 2

Announce publicly

Arbitrary announcement

We want:

- that agent 1 believes
- that agent 2 believes

Compute public announcement (if there is one)

Help on the syntax for formulas

- ((b or c) and (r or s)) ... that is propositions are b, c, r, s for respectively rainBow, Cloud, Rain and Sun. Formulas are fully parenthesized.
- (s <-> (r or b))
- (s -> (not r))

Figure 1: Demo Interface

has calculated an announcement that will achieve this goal. The announcement is computed with the algorithm presented in subsection 3.2 except that guesses are replaced by explicit loops.

We remark that, in most cases, the formula generated by our prototype implementation is very large and difficult to work with; it is really just a disjunction of maximal conjunctions describing all the suitable interpretations. In future work, we intend to improve the implementation to return a smaller formula. This could be done as a post-processing step following the calculation currently performed, or it could be done through the development of a new search algorithm.

6 Discussion

6.1 Future Work

In this paper, we have only addressed the complexity of arbitrary public announcement in the case of the Dalal’s revision operator, where the the number of levels in the plausibility ordering is at most $|\mathbf{P}|$ (the size of the vocabulary). In general, the worst case complexity to check all levels of a total pre-order over a vocabulary of size \mathbf{P} will be $2^{|\mathbf{P}|}$. If we are actually required to make this many checks, then the decision problem for other revision operators would lie at a higher level of the polynomial hierarchy. This would not necessarily be surprising, as it would follow the same pattern seen in Eiter and Gottlob’s classic paper on the complexity of various revision operators [EG92]. We do not know at this point if we actually need to check all possible levels in the underlying plausibility ordering for this particular application; it may be the case that we can perform fewer checks when the ordering is defined by an AGM revision operator. We leave the complexity of other AGM revision operators for future work.

It would also be useful to consider the announcement problem for iterated revision. In this context, the problem takes n agents and n DP revision operators $*_i$ as input. We are looking for the existence of a consistent formula ϕ such that

$$\begin{aligned}\psi_1 &\in \min(\prec_1 *_1 \phi) \\ \psi_1 &\in \min(\prec_2 *_2 \phi) \\ &\vdots \\ \psi_1 &\in \min(\prec_n *_n \phi)\end{aligned}$$

where \prec_i and ψ_i represent the epistemic state and the goal of the agent i , respectively. Following the results in [Lib97], we expect the complexity will be higher in this case.

6.2 Conclusion

We have considered arbitrary public announcements in the setting of propositional belief revision. We have proved that the decision problem for arbitrary announcements in this setting is computationally simpler than it is in a modal setting, even under very strong simplifying assumptions. We suggest that this result is important, because revision by announcements can be used in practical applications where the notion of fully nested beliefs is not required. We are currently developing a demonstration of this fact, in the context of simple robot controllers.

References

- [AGM85] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- [BBvD⁺07] P. Balbiani, A. Baltag, H. van Ditmarsch, A. Herzig, T. Hoshi, and T. De Lima. What can we achieve by arbitrary announcements? a dynamic take on fitch’s knowability. In *Proceedings of the Conference on Theoretical Aspects of Rationality and Knowledge (TARK-07)*, pages 42–51, 2007.
- [BMS98] A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. In *Proceedings of the Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98)*, pages 43–56, 1998.
- [BS06] Alexandru Baltag and Sonja Smets. Dynamic belief revision over multi-agent plausibility models. In *Proceedings of LOFT*, volume 6, pages 11–24, 2006.
- [BS08] Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. *Texts in logic and games*, 3:9–58, 2008.
- [CS15] T. Charrier and F. Schwarzentruber. Arbitrary public announcement logic with mental programs. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multi-Agent Systems*, 2015.
- [Dal88] M. Dalal. Investigations into a theory of knowledge base revision. In *Proceedings of the National Conference on Artificial Intelligence (AAAI88)*, pages 475–479, 1988.
- [DP97] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1-2):1–29, 1997.
- [EG92] T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence*, 57(2-3):227–270, 1992.
- [FvD08] T. French and H. van Ditmarsch. Undecidability for arbitrary public announcement logic. In *Advances in Modal Logic*, pages 23–42, 2008.
- [JT07] Y. Jin and M. Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171(1):1–18, 2007.
- [KM92] H. Katsuno and A.O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(2):263–294, 1992.
- [Lib97] Paolo Liberatore. The complexity of iterated belief revision. In *Proceedings of the 6th International Conference on Database Theory*, pages 276–290, 1997.
- [VB07] Johan Van Benthem. Dynamic logic for belief revision. *Journal of applied non-classical logics*, 17(2):129–155, 2007.
- [vDvdHK04] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. Public announcements and belief expansion. *Proceedings of AiML-2004 (Advances in Modal Logic)*, University of Manchester, pages 62–73, 2004.
- [vDvdHK08] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, 2008.