

How Hard to Tell? Complexity of Belief Manipulation Through Propositional Announcements

Thomas Eiter¹, Aaron Hunter², François Schwarzentruber³

¹Vienna University of Technology (TU Wien)

²British Columbia Institute of Technology

³École Normale Supérieure de Rennes

eiter@kr.tuwien.ac.at, aaron_hunter@bcit.ca, francois.schwarzentruber@ens-rennes.fr

Abstract

Consider a set of agents with initial beliefs and a formal operator for incorporating new information. Now suppose that, for each agent, we have a formula that we would like them to believe. Does there exist a single announcement that will lead all agents to believe the corresponding formula? This paper studies the problem of the existence of such an announcement in the context of model-preference definable revision operators. First, we provide two characterisation theorems for the existence of announcements: one in the general case, the other for total preorders. Second, we exploit the characterisation theorems to provide upper complexity bounds. Finally, we also provide matching optimal lower bounds for the Dalal and Ginsberg operators.

1 Introduction

The beliefs of a group of agents can be manipulated through *announcements*. For instance, consider politicians that should announce the COVID-19 vaccine campaign organization. In order to prevent the call center to break down, they have to tailor a suitable public announcement to be sure that only people that need the vaccine will call. Such announcement can be difficult to conceive and has to take into account agents with different beliefs and that may have different inconsistent goals, depending on the population backgrounds.

Public announcements have previously been studied in modal logics [Plaza, 2007; Balbiani *et al.*, 2007; Ditmarsch *et al.*, 2008]. However, our focus here is strictly on propositional announcements in the context of formal *belief revision* operators. The problem is still significant in this context, as there are many natural applications involving announcements that do not involve nested beliefs.

We assume that each agent has an initial belief state as well as a formal mechanism for incorporating new information. When an announcement is made, each agent’s beliefs change accordingly. Now suppose that, for each agent i , we have an epistemic goal ψ_i ; that is, we would like the agent i to believe that ψ_i is true. The *propositional announcement problem* consists in finding a *single* formula ϕ that will ensure that each agent believes their epistemic goal following the announcement of ϕ . In this paper, we are concerned with

the computational complexity of the propositional announcement problem.

In this paper, we present a complete analysis of the computational complexity of the propositional announcement problem (PAP). Specifically, we make the following contributions:

1. We show that PAP as formulated in [Hunter *et al.*, 2017] is Σ_2^P -hard, thus closing a problem that was left open.
2. We characterize the existence of announcements, both in the general case and for all AGM-style operators.
3. We show that, for all AGM-style operators with complexity in P^{NP} , which include Dalal’s, Σ_2^P is an upper complexity bound.
4. For popular change operators like Ginsberg’s [1986] we show that PAP is EXPTIME-complete in general.

In addition to the complexity results, we demonstrate how to construct solutions to announcement finding problems.

Outline. The paper is organized as follows. In Section 2, we recall the basics about belief revision. In Section 3, we recall the propositional announcement problem and provide characterization results. In Section 4, we give the upper bound and lower bound complexity results. Finally, in Section 5 we first present related work and then we conclude.

2 Preliminaries

2.1 AGM Belief Revision

We assume a finite set P of propositional variables, called a *vocabulary*. A *belief set* over P is a deductively closed set K of propositional formulas, also called a *knowledge base*; we may view K as a single formula given by $\bigwedge\{\phi \in K\}$. The notions of satisfaction and entailment are as usual. Interpretations $\mu : P \rightarrow \{0, 1\}$ are also viewed as sets of variables assigned true, i.e., $\{p \in P \mid \mu(p) = 1\}$. We write $mod(\phi)$ resp. $mod(K)$ to denote the set of interpretations where ϕ resp. K is true (i.e., its models).

An AGM revision operator $*$ assigns each belief set K and formula ϕ a belief set $K * \phi$ that satisfies the AGM postulates [Alchourrón *et al.*, 1985]. Semantically, AGM revision operators have been characterized in terms of model selection as follows [Katsuno and Mendelzon, 1992]. For every K and ϕ ,

$$mod(K * \phi) = \min_{\preceq_K} (mod(\phi)), \quad (1)$$

i.e., the models of the revision are the minimal (closest) models of ϕ in some total preorder $\mu \preceq_K \mu'$ of the interpretations that fulfill further conditions, viz. (i) $\mu \models K$ implies for every μ' that $\mu \preceq \mu'$ and if $\mu' \preceq_K \mu$ then $\mu' \models K$ and (ii) $K \equiv K'$ implies $\preceq_K = \preceq_{K'}$. As usual, $\mu \prec_K \mu'$ denotes the strict version of $\mu \preceq_K \mu'$: it means $\mu \preceq_K \mu'$ and not $\mu' \preceq_K \mu$. On the other hand, each such total preorder \preceq_K induces an AGM revision operator.

A canonical example is the Dalal operator [Dalal, 1988], denoted by $*_d$. Here the ordering $\mu \preceq_K^D \mu'$ is based on the Hamming distance $d_H(\nu, \mu)$ between interpretations ν, μ , which is the cardinality of the set $\text{diff}(\nu, \mu) = \{p \in P \mid \nu \models p \Leftrightarrow \mu \not\models p\}$ of variables with different truth values in ν and μ . Then $\mu \preceq_K^D \mu'$ iff $d_H(\mu, K) \leq d_H(\mu', K)$, where $d_H(\mu, K) = \min_{\nu \models K} d_H(\mu, \nu)$ is this least Hamming distance between μ and any model ν of K . Then, $\text{mod}(K *_d \phi)$ are those models μ of ϕ such that $d_H(\mu, K)$ is minimal.

2.2 Beyond AGM Revision

Besides AGM revision, many further change operators have been proposed in the literature which do not satisfy the AGM postulates, even in a finite setting cf. [Eiter and Gottlob, 1992; Nebel, 1998]. A prominent such example is Ginsberg's [1986] operator $*_G$, which is defined for knowledge bases K and formulas ϕ by

$$K *_G \phi = \{K' \cup \{\phi\} \mid K' \in W(K, \phi)\}, \quad (2)$$

where $W(K, \phi)$ are the models comprising the maximal subsets $K' \subseteq K$ that are satisfiable with ϕ . Semantically, $K *_G \phi$ is regarded as the disjunction $K_1 \vee \dots \vee K_n$ of all K_i in $K *_G \phi = \{K_1, \dots, K_n\}$. Notably, this operator is known as the ‘‘flock’’ approach in databases [Fagin et al., 1986], and it underlies in varieties consistent query answering from data and knowledge bases [Bertossi, 2019; Bienvendu and Bourgaux, 2016].

It is possible to capture $*_G$ by model selection as in (1), if we view \preceq as a preference relation (and disregard (ii)). For each model μ and knowledge base K let $K_\mu = \{\alpha \in K \mid \mu \models \alpha\}$ be the formulas in K that are satisfied by μ .

Proposition 1. *For every knowledge base K and formula ϕ ,*

- (i) $W(K, \phi) = \max_{\subseteq} (\{K_\mu \mid \mu \models \phi\})$,
- (ii) $\text{mod}(K *_G \phi) = \{\mu \mid \mu = \arg \max_{\subseteq, \mu \models \phi} K_\mu\}$.

That is, we can compare models μ and μ' as follows: $\mu \preceq_K^G \mu'$ iff $K_\mu \supseteq K_{\mu'}$. This is in fact a partial ordering, and it leads to the following result:

Proposition 2. *For any K and formula ϕ , $\text{mod}(K *_G \phi) = \min_{\preceq_K^G} (\text{mod}(\phi))$.*

This is just one example. In a similar manner, we can consider generic preference-based operators, i.e., that can be defined by minimal model selection (1) in terms of a preorder \preceq_K that depends on K (possibly embracing conditions such as variable preference, formula protection, etc.).

3 The Propositional Announcement Problem

3.1 The Basic Problem

We define the announcement problem more formally. Given n agents and n AGM belief revision operators $*_i$ as input; we

are looking for a consistent formula ϕ such that

$$K_i *_i \phi \models \psi_i, \quad i = 1, \dots, n \quad (3)$$

where K_i and ψ_i represent the belief set and the goal of the agent i , respectively. If $\bigwedge_i \psi_i$ is consistent, simply announce it! So the problem is only of interest if $\bigwedge_i \psi_i$ is inconsistent.

Example 1. *Consider a pair of ‘‘herding’’ agents R_1 and R_2 that must collaborate to herd a flock of sheep. Each agent acts independently, influenced by a single controller making announcements. Assume the vocabulary is $\{\text{corner}, \text{retrieve}\}$. Informally, a herder will believe corner if the sheep must be cornered so they can not move. A herder will believe retrieve if there is a stray sheep that must be retrieved.*

Let the initial beliefs of R_1 and R_2 be as follows:

$$\begin{aligned} K_{R_1} &= \{\neg \text{corner} \vee \text{retrieve}\} \\ K_{R_2} &= \{\neg \text{corner} \vee \neg \text{retrieve}\} \end{aligned}$$

Suppose the controller sees one sheep has left the flock. Is there a formula that can be broadcast to immediately get R_1 to retrieve the sheep while R_2 corners the rest of the animals? In other words, is there a formula ϕ such that:

$$\begin{aligned} \{\neg \text{corner} \vee \text{retrieve}\} * \phi &\models \text{retrieve} \\ \{\neg \text{corner} \vee \neg \text{retrieve}\} * \phi &\models \neg \text{retrieve} \wedge \text{corner} \end{aligned}$$

If we assume both agents use the Dalal revision operator, the answer is yes; we can set $\phi = \text{corner}$. This instance of the announcement problem has a solution.

The announcement problem was previously considered in [Hunter et al., 2017], where the following basic results are shown.

Proposition 3. *For all $n > 1$, there are instances of (3) with no solution.*

Proposition 4. *Let P be a vocabulary and let the number of agents n satisfy $n < 2^{|P|}$. Then there is an instance of (3) over P with inconsistent goals that has a solution.*

Hence, there are cases with no solution and there are cases with a solution; the question is how hard it is to tell which is the case. We are therefore concerned with the complexity of the following decision problem.

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

- Input:** a list K_1, \dots, K_n of formulas (initial beliefs).
a list ψ_1, \dots, ψ_n of formulas (goals).

Output: yes, if some ϕ satisfying (3) exists, no, otherwise.

Note that an instance of the problem is written as PAP($*_i$) to emphasize that it depends on given revision operators $*_i$.

In [Hunter et al., 2017], it is proved that this decision problem is in $\Sigma_2^P = \text{NP}^{\text{NP}}$ if we restrict attention to the Dalal operator. The hardness of PAP is left open in that paper.

3.2 Characterisation

Recall that all AGM revision operators can be characterised by taking minimal models in some preference ordering. In fact, when we can obtain the revision models as the minimal models of a preorder \preceq (not necessarily a total one, i.e., models may be incomparable; in short by model preference), then we can characterize solutions ϕ to a propositional announcement problem (PAP) $P = (K_i, \psi_i)_{i=1}^n$ as follows.

Proposition 5. *Let ϕ be consistent and let P be a PAP with model-preference definable operators $*_i$, $1 \leq i \leq n$. Then a formula ϕ is a solution to P iff for all $i = 1, \dots, n$ and μ , if $\mu \models \phi \wedge \neg\psi_i$ then some $\mu' \prec_{K_i} \mu$ exists s.t. $\mu' \models \phi \wedge \psi_i$.*

The characterisation in Proposition 5 says that models μ violating the query ψ_i must be eliminated by models μ' which do satisfy the query and are closer to the knowledge base K_i .

In fact, without loss of generality, one can limit the solutions to fulfill $\text{mod}(\phi) \subseteq \bigcup_i \text{mod}(\psi_i)$ i.e., replace ϕ by $\phi \wedge \bigvee_i \psi_i$, because every model that does not satisfy any of the ψ_i must be eliminated under revision.

Proposition 6. *If ϕ is a solution to a model-preference definable PAP P , then $\phi \wedge \bigvee_i \psi_i$ is also a solution of P .*

In general, there can be multiple solutions ϕ , of course; however, there are interesting observations.

Proposition 7. *Every PAP P with model-preference definable operators $*_i$ has a unique weakest solution ϕ^* , i.e., for every solution ϕ , it holds that $\phi \models \phi^*$.*

The reason for this is simple: if ϕ_1 and ϕ_2 are solutions, then also $\phi_1 \vee \phi_2$ is a solution; in terms of models, the union of the set of models of solutions is also a solution.

Total Preorders

Smaller solutions with few models might exist, as in the case of Dalal's operator. There, at most n models are needed to find a solution. This can in fact be generalized to operators that are definable by total preorders \preceq_{K_i} , i.e., either $\mu \preceq_K \mu'$ or $\mu' \preceq_K \mu$ holds for all models μ and μ' .

Proposition 8. *Let P be a PAP with model-preference definable operators $*_i$ where all preorders \preceq_{K_i} are total. Then the following are equivalent:*

1. P has some solution,
2. There is (μ_1, \dots, μ_n) with $\mu_i \models \psi_i$ and for all j , either $\mu_j \models \psi_i$ or $\mu_i \prec_{K_i} \mu_j$.
3. $\{\mu_1, \dots, \mu_n\}$ as in point 2 is a solution of P .

Proof. (1 \Rightarrow 2) Let ϕ be a solution. Pick μ_i in $\min_{\preceq_{K_i}}(\text{mod}(\phi))$ which is non-empty because ϕ is consistent. As ϕ is a solution, $\mu_i \models \psi_i$. Suppose that $\mu_j \not\models \psi_i$. We also have $\mu_j \models \phi$. By Proposition 5, there exists μ' such that $\mu' \models \phi$, $\mu' \models \psi_i$ and $\mu' \prec_{K_i} \mu_j$. As μ_i is in $\min_{\preceq_{K_i}}(\text{mod}(\phi))$, and as μ_i and μ' are \preceq_{K_i} -comparable, we must have $\mu_i \preceq_{K_i} \mu'$. By transitivity, we have $\mu_i \prec_{K_i} \mu_j$.

(2 \Rightarrow 3). Let us show that $\min_{\preceq_{K_i}}(\{\mu_1, \dots, \mu_n\}) \models \psi_i$. Take a \preceq_{K_i} -minimal μ_j . As we do not have $\mu_i \prec_{K_i} \mu_j$, we must have $\mu_j \models \psi_i$. \square

4 Computational Complexity

It was shown in [Hunter *et al.*, 2017] that for Dalal's operator, the $\text{PAP}(*_d)$ problem is in Σ_2^P , by exhibiting a non-deterministic guess and check algorithm based on an ad-hoc instance of Proposition 8. However, the authors left the issue open whether this bound can be improved. We show that this bound is tight, proving Σ_2^P -hardness by a nontrivial construction. Furthermore, we establish that the PAP problem is Σ_2^P -complete in general for partial ordering definable operators $*_i$

with complexity in P^{NP} , while we show that it is EXPTIME -complete for model-preference operators $*_i$ in general. In particular, EXPTIME -hardness is shown for Ginsberg's operator. Hence, there is a considerable complexity gap between operators that are relevant in practice.

4.1 Upper Bounds

Total Preorders

For operators $*$ definable by total preorders, Proposition 8 implies that we can use a guess and check algorithm to decide the existence of a solution.

Theorem 1. *Let $*_i$ be operators that are definable by total preorders \preceq_{K_i} . If checking $\mu \preceq_{K_i} \mu'$ is feasible in complexity C , then $\text{PAP}(*_i)$ is decidable in NP^C .*

Proof. Consider the following non-deterministic algorithm:

```

for  $i := 1$  to  $n$  do
  |  $(\exists)$  choose  $\mu_i$  in  $\text{mod}(\psi_i)$ 
for  $i, j \in \{1, \dots, n\}$  do
  | if  $\mu_i \not\models \psi_j$  then
  |   | Check that  $\mu_j \preceq_{K_i} \mu_i \wedge \mu_i \not\preceq_{K_i} \mu_j$ 
    
```

The correctness follows by Proposition 8. \square

For all belief revision operators that satisfy the AGM postulates, it is by Proposition 8 a matter of the complexity of model comparison which may dominate the overall complexity. But in any case, we have a small result for the revision!

In particular for complexity $C = \text{P}^{\text{NP}}$, we obtain from Theorem 1 membership in $\text{NP}^{\text{P}^{\text{NP}}} = \text{NP}^{\text{NP}}$. Applied to Dalal's operator, we thus obtain the following result.

Corollary 1 (cf. [Hunter *et al.*, 2017]). *$\text{PAP}(*_d)$ is in Σ_2^P .*

Proof. Given an interpretation μ and an integer k , deciding $d_H(\mu, \nu) \leq k$ is in NP. Thus, computing $d_H(\mu, K)$ is feasible in polynomial time with an NP oracle, and so is deciding $\mu \preceq_K^D \mu'$; Theorem 1 implies membership in Σ_2^P . \square

For knowledge bases K in special form, the NP-oracle may be eliminated so that lower complexity results. For example, we have the following result.

Proposition 9. *$\text{PAP}(*_d)$ is in NP, if each K_i is rewritable into a DNF formula in polynomial time.*

Indeed, for a DNF $\alpha = D_1 \vee \dots \vee D_n$, determining the smallest distance $d_H(\mu, \mu')$ of a model μ to any model μ' of α is feasible in polynomial time.

General Case

Theorem 2. *Given model-preference definable operators $*_i$ such that $\mu' \preceq_{K_i} \mu$ is decidable in exponential time, $\text{PAP}(*_i)$ is in EXPTIME .*

Sketch of proof. Consider the following alternating algorithm in which the existential player builds ϕ on the fly, and that checks the assumption of Proposition 5:

(\exists) a model μ_0
 check($\mu_0, 0$)
procedure check($\mu, 0$)
 if $s < 2^{|P|}$ **then**
 (\forall) $j \in \{1, \dots, n\}$:
 if $\mu \not\models \psi_j$ **then**
 (\exists) μ' in $\text{mod}(\psi_j)$, $\mu' \prec_{K_j} \mu$:
 check($\mu', s+1$)

Note that at stage $s = 2^{|P|}$, a loop must have been entered, which allows for acceptance. As $\text{APSPACE} = \text{EXPTIME}$ [Chandra and Stockmeyer, 1976], the check $\mu' \preceq_{K_j} \mu$ (i.e., $\mu' \preceq_{K_j} \mu$ and $\mu \not\preceq_{K_j} \mu'$) is feasible with alternations in polynomial space. Overall, the algorithm runs with alternations in polynomial space, which shows EXPTIME -membership. \square

In the alternating PSPACE algorithm above the universal branching is only about polynomially many values. It may be tempting to believe that this branching can be eliminated and simulated deterministically. However, the recursion depth is not necessarily polynomially bounded, so this may end up in exponential space consumption. Notably, even if each relation \prec_{K_i} is acyclic, their interaction may create a cycle.

Example 2. Let $K_1 = \{\neg \text{evidence} \vee \text{guilty}\}$ be the prosecutor's and $K_2 = \{\neg \text{guilty} \vee \neg \text{evidence}\}$ be the accused's KB in a trial. The judge may announce proper account of evidence ($\phi = \text{evidence}$) to achieve $\psi_1 = \text{guilty}$ and $\psi_2 = \neg \text{guilty}$ under Dalal's operator. The models $\mu_1 = \{\text{evidence}, \text{guilty}\}$ and $\mu_2 = \{\text{evidence}\}$ of ϕ satisfy $\mu_1 \prec_{K_1} \mu_2$ and $\mu_2 \prec_{K_2} \mu_1$.

However, no acyclic paths of length $2^{|P|}$ are possible.

In case of just two agents and easier preference checking, we interestingly get a better complexity bound:

Theorem 3. For a PAP P with $n = 2$ and model-preference definable operators $*_i$ such that $\mu' \preceq_{K_i} \mu$ is decidable in polynomial space, deciding solution existence is in PSPACE .

Proof. In this case, no alternation in the algorithm above is needed: as $\mu_i \models \psi_i$ resp. $\mu' \models \psi_j$, for at most one $j \in \{1, 2\}$ a guess $\mu' \in \text{mod}(\psi_j)$ will be made; thus (\forall) can be eliminated deterministically. As checking $\mu' \preceq_{K_j} \mu$ is in PSPACE , membership in $\text{NPSPACE} = \text{PSPACE}$ follows. \square

4.2 Lower Bounds

PAP with Dalal's Operator

Theorem 4. The PAP($*_d$) problem (all $*_i$ are $*_d$) is Σ_2^P -hard.

Proof. We reduce the evaluation of QBFs of the form

$$\Phi = \exists X \forall Y E(X, Y), \quad (4)$$

where $X = X_1, \dots, X_m$ and $Y = Y_1, \dots, Y_l$ are lists (viewed as sets) of propositional variables in polynomial time to this problem. Without loss of generality, we assume that the following property (see Supplement):

(*) $|Y|$ is even, and if an assignment σ to $X \cup Y$ does not satisfy E , it sets exactly half of the variables in Y to 1, i.e., $|\{Y_i \in Y \mid \sigma(Y_i) = 1\}| = |\{Y_i \in Y \mid \sigma(Y_i) = 0\}|$.

We now construct knowledge bases K_i and formulas ψ_i , $i = 1, \dots, n$ such that deciding whether some formula ϕ exists such that $K_i *_d \phi \models \psi_i$ encodes the evaluation of Φ .

From Proposition 8, ϕ has w.l.o.g. the form $\mu_1 \vee \dots \vee \mu_n$, where each μ_i is a model of ψ_i and the following holds: no μ_j that violates ψ_i will be closer to K_i than μ_i , i.e., satisfy $d_H(\mu_j, K_i) < d_H(\mu_i, K_i)$.

In our encoding we need only $n = 2$. The idea is that intuitively, we use the pair (K_2, ψ_2) to generate a candidate assignment σ to X via a model μ_2 that will not be closer to K_1 than any model μ_1 for a solution $\phi = \mu_1 \vee \mu_2$ if $\neg E(\sigma(X), Y)$ is unsatisfiable, i.e., $\forall Y E(\sigma(X), Y)$ evaluates to true. Specifically, we define:

$$\begin{aligned} K_1 &:= [\neg E(X, Y) \wedge \bigwedge_{X_i \in X} EQ(X_i, X^i)] \vee (B \wedge Y), \\ \psi_1 &:= \overline{B} \wedge Y, \\ K_2 &:= B \wedge \overline{Y}, \\ \psi_2 &:= \overline{B} \wedge \overline{Y} \wedge \bigwedge_{X_i \in X} EQ(X_i, X^i). \end{aligned}$$

Here B and the X^i are disjoint sets of fresh variables, such that $2|B| = |Y|$ and $|X^i| = |B|$; we write B resp. Y for the conjunction of all variables in B resp. Y , and similarly \overline{B} resp. \overline{Y} for the conjunction of all negated variables in B resp. Y (e.g. for $B = \{B_1, B_2\}$ we have that B stands for $B_1 \wedge B_2$ and \overline{B} for $\neg B_1 \wedge \neg B_2$). Furthermore, $EQ(X_i, X^i)$ expresses that all variables in X^i have the same value as X_i :

$$EQ(X_i, X^i) = \bigwedge_{Z \in X^i} (X_i \leftrightarrow Z). \quad (5)$$

The purpose of B is to ensure that the distance of a candidate model μ_1 for ϕ to K_1 is at least $|B| = |Y|/2$; the purpose of the X^i is to ensure that the cost of flipping in a model μ_2 of ψ_2 a variable X_i to obtain a model of the left disjunct of K_1 is, due to the variables in X^i , at least $|B| + 1$.

We now argue that Φ evaluates to true if and only if the constructed PAP instance has a solution.

(\Rightarrow) Let σ be an assignment to X such that $\forall Y E(\sigma(X), Y)$ evaluate to true; hence, $\neg E(\sigma(X), Y)$ is unsatisfiable. We consider interpretations μ_1 and μ_2 as follows:

$$\mu_1 = Y \cup \mu_2 \text{ and } \mu_2 = \bigcup \{ \{X_i\} \cup X^i \mid X_i \in X : \sigma(X_i) = 1 \}.$$

Note that $\mu_i \models \psi_1 \wedge \neg \psi_2$ while $\mu_2 \models \psi_2 \wedge \neg \psi_1$. Consider now the distances $d_H(\mu_i, K_j)$ for $i, j = 1, 2$:

- $d_H(\mu_1, K_1) = |B|$, as flipping B is enough to arrive at a model of the right disjunct of K_1 ; there is no cheaper cost, as to satisfy the left disjunct of K_1 , by unsatisfiability of $\neg E(\sigma(X), Y)$ we would need to flip some X_i and the associated X^i , causing a distance $> |B|$;
- $d_H(\mu_2, K_1) > |B|$: to arrive at a model of the left disjunct of K_1 , since $\neg E(\sigma(X), Y)$ is unsatisfiable we need to flip some X_i and its associated variables X^i , which causes a distance $> |B|$; to arrive at a model of right disjunct of K_1 , flipping $B \cup Y$ has cost $|B| + |Y|$;
- $d_H(\mu_1, K_2) = |B| + |Y|$: just flip $B \cup Y$;
- $d_H(\mu_2, K_2) = |B|$: just flip B .

Hence, the disjunction $\phi = \mu_1 \vee \mu_2$ of the models μ_1 and μ_2 is a solution to the PAP from above.

(\Leftarrow) Suppose that without loss of generality some models μ_1 and μ_2 of ψ_1 and ψ_2 , respectively, exist such that their disjunction $\phi = \mu_1 \vee \mu_2$ yields a solution of the PAP instance.

Then μ_1 satisfies $\overline{B} \wedge Y$ and μ_2 satisfies $\overline{B} \wedge \overline{Y}$. Let us consider the distance of μ_1 to K_1 :

- $d(\mu_1, K_1) \leq |B|$, as it is sufficient to flip B to satisfy the right disjunct of K_1 ;
- $d(\mu_1, K_1) \geq |Y|/2$, as if we do not flip B , we must flip exactly half of Y and perhaps some X_i and X^i to satisfy $\neg E(X, Y)$ in the left disjunct of K_1 .

As $|B| = |Y|/2$ it follows $d(\mu_1, K_1) = |B| = |Y|/2$.

Now let σ be the assignment to X s.t. $\sigma(X_i) = 1$ iff $\mu_2 \models X_i$, $i = 1, \dots, n$. We argue that $\forall Y E(\sigma(X), Y)$ evaluates to true. As $\mu_2 \not\models \psi_1$, we have $d_H(\mu_1, K_1) < d_H(\mu_2, K_1)$; hence $d(\mu_2, K_1) > |Y|/2$. As K_1 is of the form $K_1 = K_{1,1} \vee K_{1,2}$, each $K_{1,j}$ must satisfy $d(\mu_2, K_{1,j}) > |Y|/2$, in particular $K_{1,1} = \neg E(X, Y) \wedge \bigwedge_{X_i \in X} EQ(X_i, X^i)$. As $\mu_2 \models \bigwedge_{X_i \in X} EQ(X_i, X^i)$, this means that whenever exactly half of Y is flipped, the resulting interpretation does not satisfy $\neg E(X, Y)$; hence, for the respective assignment μ to Y , $E(\sigma(X), \mu(Y))$ is satisfied. Hence $\forall Y E(\sigma(X), Y)$ evaluates to true, and so does Φ . This proves Σ_2^P -hardness. \square

Two remarks should be noticed:

- (1) We can rewrite the KBs K_i and queries ψ_i to CNFs in polynomial time, if we assume w.l.o.g. that $E(X, Y)$ is in DNF. Thus the Σ_2^P -hardness inherits to the CNF fragment.
- (2) If we assume that $E(X, Y)$ is in CNF and no Y variables occur in $E(X, Y)$, then we can rewrite K_1 and K_2 to DNF (note that we can rewrite in (5) $EQ(X_i, X^i)$ to a DNF $(X_i \wedge \bigwedge_{Z \in X^i} Z) \vee (\neg X_i \wedge \bigwedge_{Z \in X^i} \neg Z)$). This then shows NP-hardness as evaluating the QBF in (4) is NP-hard under these restrictions. Combined with Proposition 9, we thus obtain:

Theorem 5. *For PAP(*_d), deciding solution existence is NP-complete if all knowledge bases K_i are in DNF.*

It is possible to extend the construction in order to have only Horn KBs K_i and Horn ψ_i ; the idea is to simulate positive disjunction with choice of an atom for closer distance.

PAP with Ginsberg's Operator

We now provide matching lower bounds for the upper bounds for the PAP-problem in the general model preference case from above, specifically for Ginsberg's operator.

Theorem 6. *For PAP(*_G), deciding solution existence is (i) EXPTIME-hard in general and (ii) PSPACE-hard if $n = 2$.*

We sketch a proof of this result, which involves a reduction from the following problem.

A bipartite *and-or graph* is a tuple $G = (N_1, N_2, E_1, E_2)$ of disjoint sets N_1 and N_2 of nodes, and of edges E_1 and E_2 , respectively, where any edge in E_1 (resp., E_2) is of the form $(n_1, n_2, n'_2) \in N_1 \times N_2 \times N_2$ resp. $(n_2, n_1) \in N_2 \times N_1$. We assume that every node in N_1 occurs in at most one edge of E_1 . Intuitively, the nodes in N_1 are and-nodes and have either

no or two (not necessarily) distinct successors, while N_2 are or-nodes that can have multiple possible successors;

A nonempty set $N \subseteq N_1 \cup N_2$ of nodes is *serial* in G , if for every $n_1 \in N \cap N_1$, some edge $(n_1, n_2, n'_2) \in E_1$ exists and both n_2, n'_2 are in N , and for every $n_2 \in N \cap N_2$ some $(n_2, n_1) \in E_2$ exists such that $n_1 \in N$.

Lemma 1. *Let $\gamma_1(X, Y, Z)$ and $\gamma_2(Y, X)$, where $X = X_1, \dots, X_n$ and $Y = Y_1, \dots, Y_n$, $Z = Z_1, \dots, Z_n$ are Boolean formulas that describe the edges E_1 resp. E_2 of an and-or graph G with nodes $N_1 = \{0\} \times \{0, 1\}^{n-1}$ and $N_2 = \{1\} \times \{0, 1\}^{n-1}$. Deciding whether G has some serial node set N is EXPTIME-hard.*

Informally, this lemma holds as the and-or graph G may represent computations of an alternating Turing machine T , where nodes represent configurations c . In accepting c 's with no successors, T is modified to loop back to the initial configuration c_0 , while other loops are excluded. Then a serial set N exists iff the modified T has a looping run on all branches, i.e., the original T accepts on all branches. As T runs in polynomial space, deciding the latter is EXPTIME-hard.

We now construct a PAP instance as follows:

$$\begin{aligned} K_1 &= X \cup X' \cup \{\neg X_0\} \text{ and} \\ \psi_1 &= \neg X_0 \wedge Y_0 \wedge Z_0 \wedge \gamma_1(X, Y, Z) \wedge Op'(X; Y; Z); \\ K_2 &= Y \cup Y' \cup \{\neg Y_0\} \text{ and} \\ \psi_2 &= X_0 \wedge \neg Y_0 \wedge Z_0 \wedge \gamma_2(Y, X) \wedge Op'(X; Y; Z); \\ K_3 &= Z \cup Z' \cup \{\neg Z_0\} \text{ and} \\ \psi_3 &= X_0 \wedge Y_0 \wedge \neg Z_0 \wedge \gamma_2(Z, X) \wedge Op'(X; Y; Z); \end{aligned}$$

where

$$Op'(X; Y; Z) = (X \leftrightarrow \neg X') \wedge (Y \leftrightarrow \neg Y') \wedge (Z \leftrightarrow \neg Z').$$

Here X' , Y' and Z' are copies of X , Y , and Z , respectively, and X_0, Y_0, Z_0 are fresh variables.

Intuitively, ψ_1 describes edges from “universal” nodes in N_1 , while ψ_2 and ψ_3 describe edges from “existential” nodes in N_2 , where the variables X' , Y' and Z' serve to create incomparability among the encoding of nodes (to avoid clutter by possible world maximization revision $*_G$). The formulas ψ_1 , ψ_2 and ψ_3 are pairwise mutually inconsistent via the variables X_0, Y_0 , and Z_0 . Thus a move from an and-node n_0 to n_1 and n'_1 expressed by a model of ψ_1 must be followed by a move from n_1 (resp. n'_1) to some and-node n_2 (resp. n'_2) expressed by a model of ψ_2 . It can be shown that the and-or graph G has a nonempty serial set N iff the PAP(*_G) given by $P = (K_i, \psi_i)_{i=1}^3$ has a solution.

Case $n = 2$: with a simplified construction, we can obtain PSPACE-hardness: if each universal state has a single transition i.e., $(n_1, n_2, n'_2) \in E_1$ implies $n_2 = n'_2$, we have an ordinary graph (intuitively describing non-alternating computations). The Z variables can then be replaced by Y , and K_3 and ψ_3 can be dropped.

4.3 Constructing PAP Solutions

The algorithms above are constructive and can be used to obtain a solution to a PAP. In fact, Theorems 1 and 4 (respectively their proof) establish that computing some PAP solution for AGM-style revision operators is complete for the multi-valued function analog of Σ_2^P , i.e. the search problems

that can be solved by a nondeterministic Turing machine with an output tape and an NP oracle in polynomial time.

For model-preference based PAPs, the solutions obtained may not be of polynomial size. Indeed, as deciding $K *_G \phi \models \psi$ is Π_2^P -complete [Eiter and Gottlob, 1992], it follows that verifying whether a formula ϕ is a solution of a PAP($*_G$) P has the same complexity. Thus if always some polynomial size ϕ would exist, a guess and check algorithm would put solution existence in Σ_3^P and $\text{EXPTIME} \subseteq \Sigma_3^P$ would follow.

Polynomial-size solutions ϕ can be guessed and the conditions in Proposition 5 checked in exponential time. Note that the weakest solution ϕ^* of a model-preference based PAP P is an envelope, computable by an elimination algorithm:

1. Initially, set X_0 to the set of all interpretations.
2. Remove then iteratively from X_i a model μ , such that $\mu \not\models \psi_j$ for some j and no model $\mu' \in X_i$ with $\mu' \models \psi_j$ and $\mu' \prec_j \mu$ exists, i.e., set $X_{i+1} := X_i \setminus \{\mu\}$.

Any μ in Step 2 is a “bad model” and must be eliminated, as it would be a revision model but fails to satisfy the query. If $X(P)$ denotes the final X_i , then we have:

Proposition 10. *For a PAP P with model-preference operators \circ_i , the weakest solution ϕ^* fulfills $\text{mod}(\phi^*) = X(P)$, where no solution exists if $X(P) = \emptyset$.*

4.4 Reasoning About Solutions

It is of interest to know whether a formula α is a semantic core for the solutions, given by $\alpha \models \phi$ for every solution ϕ , and dually, whether a formula β is an envelope for the solutions, given by $\phi \models \beta$ for every solution ϕ . In a sense, a core describes possible models that must be compatible with every announcement for belief manipulation, while an envelope limits the models that can be used. Based on the results and techniques above, we can show:

Theorem 7. *Deciding whether α is a core (resp., β is an envelope) for a PAP P is (i) EXPTIME -complete for model-preference decidable in exponential time and (ii) Π_2^P -complete for operators \circ_i definable with total preorders \prec_{κ_i} decidable in P^{NP} .*

In fact, ϕ^* coincides with the logically strongest envelope β^* , i.e., $\beta^* \models \beta$ for each envelope β ; it is intuitively the vaguest statement that one can make to accomplish belief manipulation. A formula γ is thus implied by every solution ϕ iff it is implied by ϕ^* . Dually, a unique weakest core α^* exists, i.e., $\alpha \models \alpha^*$ for every core α , which is equivalent to the conjunction of all solutions. Intuitively, it is a lower bound of vagueness for the announcement, which can be void if e.g. solutions ϕ and ϕ' are incompatible, i.e., $\phi \wedge \phi'$ is unsatisfiable. For deciding whether $\beta^* \models \gamma$ resp. $\alpha^* \models \gamma$ holds for given PAP and γ , similar results as in Theorem 7 can be obtained, with the interesting exception that in (ii) deciding $\alpha^* \models \gamma$ is Σ_2^P -complete by exploiting Proposition 8.

5 Related Work

5.1 Dynamic Epistemic Logic

Dynamic Epistemic Logic (DEL) refers to a class of modal logics for reasoning about actions and belief [Ditmarsch *et*

al., 2008]. In DEL, the announcement problem is addressed in arbitrary public announcement logic (APAL) [Balbani *et al.*, 2007]. Contrary to us, APAL does not provide any revision mechanism, but allows for higher-order knowledge sentences. The model checking problem of APAL is in PSPACE [Ågotnes *et al.*, 2010], but contrary to our case, their input model is represented explicitly. Closer to our setting, the model checking of APAL in its succinct form is NEXPTIME-complete [Charrier and Schwarzentruher, 2015] (see also [Charrier *et al.*, 2019] for a more comprehensive study). The satisfiability problem of APAL is undecidable [French and van Ditmarsch, 2008], but gets decidable when announcements are propositional [van Ditmarsch and French, 2017].

5.2 Non-Modal Belief Revision

Schwind *et al.* [2019] considered a related problem called *inverse revision (IR)*: when is $K_i *_i \phi = \psi_i$ for all agents i possible, where ϕ and any AGM-style revision operators $*_i$ can be chosen; the set of all such ϕ is called the *frame*, which is complete if any ϕ and ϕ' are logically equivalent. Checking whether a formula ϕ is in the frame is co-NP-complete, and deciding frame-completeness is Π_2^P -complete.

Apparently, IR solutions are PAP solutions when the operators are fixed, but not the other way around. In fact, the aims of IR and PAP are different: the former wants to achieve an exact local belief state, while PAP aims to achieve goal fulfillment by an announcement. Our complexity results show that for AGM operators, the reasoning tasks can be intertranslated in polynomial time; studying this remains for future work.

6 Conclusion

We have considered the complexity of public announcement finding in a non-modal setting, where agents incorporate announcements using a revision operator. We have shown that for all AGM style operators, deciding existence of a solution is always in Σ_2^P and Σ_2^P -hard for Dalal’s operator, closing an open problem. Furthermore, we have shown that for all operators definable by preorders, the problem is in EXPTIME in general and EXPTIME -hard for Ginsberg’s operator. Moreover, we proved lower complexity under restrictions.

Our work can be extended in different directions. First, one may consider other revision operators, such as those in [Eiter and Gottlob, 1992] or iterated revision operators [Darwiche and Pearl, 1997; Jin and Thielscher, 2007]. In this context, extending our main results to relations \preceq_K that always allow for revision models seems possible. Second, besides AGM revision operators Katsuno and Mendelzon [1991] characterized update operators \diamond as in [Winslett, 1988; Forbus, 1989], which reflect change of the world rather than of the belief about the world, by model selection; specifically, $\text{mod}(K \diamond \phi) = \bigcup_{\mu \models K} \min_{\preceq_\mu}(\text{mod}(\phi))$, where \preceq_μ is a preorder. Using update operators for public announcements is natural; by our techniques, the respective solution existence problems can be solved in exponential time.

Finally, in addition to extending our approach to different revision operators, it would also be interesting to consider the application of the techniques used in this paper in other domains, such as preference aggregation.

References

- [Ågotnes *et al.*, 2010] Thomas Ågotnes, Philippe Balbiani, Hans van Ditmarsch, and Pablo Seban. Group announcement logic. *Journal of Applied Logic*, 8(1):62–81, 2010.
- [Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- [Balbiani *et al.*, 2007] Philippe Balbiani, Alexandru Baltag, Hans van Ditmarsch, Andreas Herzig, Tomohiro Hoshi, and Tiago 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.
- [Bertossi, 2019] Leopoldo E. Bertossi. Database repairs and consistent query answering: Origins and further developments. In Dan Suciu, Sebastian Skritek, and Christoph Koch, editors, *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2019, Amsterdam, The Netherlands, June 30 - July 5, 2019*, pages 48–58. ACM, 2019.
- [Bienvenu and Bourgaux, 2016] Meghyn Bienvenu and Camille Bourgaux. Inconsistency-tolerant querying of description logic knowledge bases. In Jeff Z. Pan, Diego Calvanese, Thomas Eiter, Ian Horrocks, Michael Kifer, Fangzhen Lin, and Yuting Zhao, editors, *Reasoning Web: Logical Foundation of Knowledge Graph Construction and Query Answering - 12th International Summer School 2016, Aberdeen, UK, September 5-9, 2016, Tutorial Lectures*, volume 9885 of *Lecture Notes in Computer Science*, pages 156–202. Springer, 2016.
- [Chandra and Stockmeyer, 1976] Ashok K. Chandra and Larry J. Stockmeyer. Alternation. In *Proceedings of FOCS’76*, pages 98–108, 1976.
- [Charrier and Schwarzenruber, 2015] Tristan Charrier and François Schwarzenruber. Arbitrary public announcement logic with mental programs. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multi-Agent Systems*, 2015.
- [Charrier *et al.*, 2019] Tristan Charrier, Sophie Pinchinat, and François Schwarzenruber. Symbolic model checking of public announcement protocols. *Journal of Logic and Computation*, 29(8):1211–1249, 2019.
- [Dalal, 1988] Mukesh Dalal. Investigations into a theory of knowledge base revision. In *Proceedings of AAAI 1988*, pages 475–479, 1988.
- [Darwiche and Pearl, 1997] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1-2):1–29, 1997.
- [Ditmarsch *et al.*, 2008] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer, 2008.
- [Eiter and Gottlob, 1992] Thomas Eiter and Georg Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence*, 57(2-3):227–270, 1992.
- [Fagin *et al.*, 1986] Ronald Fagin, Gabriel M. Kuper, Jeffrey Ullman, and Moshe Vardi. Updating logical databases. In P. Kanellakis and F. Preparata, editors, *Advances in Computing Research*, volume 3, pages 1–18, Greenwich, CT, 1986. JAI Press.
- [Forbus, 1989] Kenneth D. Forbus. Introducing actions into qualitative simulation. In *Proceedings IJCAI-89*, pages 1273–1278, 1989.
- [French and van Ditmarsch, 2008] Tim French and Hans van Ditmarsch. Undecidability for arbitrary public announcement logic. In *Advances in Modal Logic*, pages 23–42, 2008.
- [Ginsberg, 1986] Matthew L. Ginsberg. Counterfactuals. *Artificial Intelligence*, 30:35–79, 1986.
- [Hunter *et al.*, 2017] Aaron Hunter, François Schwarzenruber, and Eric Tsang. Belief manipulation through propositional announcements. In *Proceedings of IJCAI 2017*, pages 1109–1115, 2017.
- [Jin and Thielscher, 2007] Yi Jin and Michael Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171(1):1–18, 2007.
- [Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings KR-91*, pages 387 – 395, 1991.
- [Katsuno and Mendelzon, 1992] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(2):263–294, 1992.
- [Nebel, 1998] Bernhard Nebel. How hard is it to revise a belief base? In D. Gabbay and Ph.Smets, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, volume III: Belief Change, pages 77–145. Kluwer Academic, 1998.
- [Plaza, 2007] Jan Plaza. Logics of public communications. *Synthese*, 158(2):165–179, 2007.
- [Schwind *et al.*, 2019] Nicolas Schwind, Katsumi Inoue, Sébastien Konieczny, Jean-Marie Lagniez, and Pierre Marquis. What has been said? identifying the change formula in a belief revision scenario. In Sarit Kraus, editor, *Proceedings IJCAI 2019*, pages 1865–1871, 2019.
- [van Ditmarsch and French, 2017] Hans van Ditmarsch and Tim French. Quantifying over boolean announcements. *CoRR*, abs/1712.05310, 2017.
- [Winslett, 1988] Marianne Winslett. Reasoning about action using a possible models approach. In *Proceedings AAAI-88*, pages 89–93, 1988.