Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

On the Parameterized Complexity of Belief Revision*

Andreas Pfandler[1,2], **Stefan Rümmele**[1], **Johannes Peter Wallner**[1], and **Stefan Woltran**[1]
lastname@dbai.tuwien.ac.at
[1]Vienna University of Technology, Austria
[2]University of Siegen, Germany

Abstract

Parameterized complexity is a well recognized vehicle for understanding the multitude of complexity AI problems typically exhibit. However, the prominent problem of belief revision has not undergone a systematic investigation in this direction yet. This is somewhat surprising, since by its very nature of involving a knowledge base and a revision formula, this problem provides a perfect playground for investigating novel parameters. Among our results on the parameterized complexity of revision is thus a versatile fpt algorithm which is based on the parameter of the number of atoms shared by the knowledge base and the revision formula. Towards identifying the frontier between parameterized tractability and intractability, we also give hardness results for classes such as co-W[1], para-Θ_2^{P}, and $\mathsf{FPT}^{\mathsf{NP}}[f(k)]$.

1 Introduction

Belief revision [Alchourrón *et al.*, 1985] is a core formalism of Artificial Intelligence (see e.g. [Peppas, 2008]) aiming for a formal way of adapting one's beliefs in the light of new information. Following Katsuno and Mendelzon [1991], we consider revision in the setting where beliefs and new information are given as propositional formulas. Operators like the one by Satoh [1988] characterize the result of revision by a set of models representing the adapted belief base.

Understanding the computational complexity of reasoning is an indispensable step towards the design of practically efficient systems. Unfortunately, Eiter and Gottlob [1992] have shown that many revision operators suffer from a high computational complexity in the general case of propositional logic. For instance, deciding whether a formula follows from the result of a revision in terms of Satoh's operator has been shown to be Π_2^{P}-complete. Those results were later strengthened in terms of model checking by Liberatore and Schaerf [2001]. While these two papers studied easier cases in terms of the HORN fragment, other fragments that might lead to tractability have not been thoroughly investigated; a recent analysis [Creignou *et al.*, 2013] for KROM formulas being a notable exception.

Inspecting these results, it becomes clear that tractability requires rather strong restrictions on the involved formulas, i.e., the knowledge base and the revision formula. On the other hand, these components might be of vastly different size and shape. For exploiting this potential *heterogeneity* of the two components towards novel tractability results, parameterized complexity and parameterized algorithmics [Downey and Fellows, 1999] are the formal tools of choice. Here the complexity is not measured with respect to the input size alone but also with respect to one or multiple features of the instance that are expressed as an integer value, the so-called parameter. Parameterized complexity has been investigated already in several AI domains (see e.g. [Fellows *et al.*, 2012; Gottlob and Szeider, 2008; Gottlob *et al.*, 2010; Lackner and Pfandler, 2012]) and such results have been successfully put to practice, in particular in the area of argumentation [Dvořák *et al.*, 2012; 2014].

In this paper we utilize parameterized complexity theory for belief revision with the aim to keep the restrictions on the theory and new information as independent and relaxed as possible without destroying tractability. Our analysis not only gives new insights on the different sources of complexity of revision, but also provides the basis for novel implementations. We shall focus here on Satoh's operator and start our analysis by generalizing already established tractable cases from [Eiter and Gottlob, 1992] as well as the only fpt result we are aware of, viz. [Pichler *et al.*, 2009] which utilizes the combined treewidth parameter. The *main contributions* of our work are:

- We generalize existing tractability results due to Eiter and Gottlob [1992] via backdoors to Horn and show certain limits of the treewidth approach, in particular we provide para-coNP-hardness if treewidth is bound independently in the knowledge base T and the revision formula φ. Furthermore, also combining bounded treewidth of φ with the restriction of T to be Horn results in para-coNP-hardness.
- Towards a *generic* algorithm, we identify a novel parameter, namely the cardinality of the variables shared by knowledge base T and revision formula φ. Our algorithm is fpt when adding (possibly independent) parameters that make T's and φ's satisfiability fpt. We show that these additional bounds are necessary, since otherwise completeness for the class $\mathsf{FPT}^{\mathsf{NP}}[f(k)]$ is obtained. $\mathsf{FPT}^{\mathsf{NP}}[f(k)]$ contains prob-

*Supported by the Austrian Science Fund (FWF) under grants P25518, P25521, and Y698, and the German Research Foundation (DFG) under grant ER 738/2-1.

lems that can be computed by an fpt-algorithm that may make $f(k)$ many NP oracle calls. From a practical perspective, this result is interesting since it shows that the problem at hand can by implemented via an iterative SAT-procedure.

• Finally, we show that a relaxation of the bound on the shared vocabulary to a bound on the maximum symmetric distance between models of T and φ is less promising since this increases the complexity to para-Θ_2^{P}-hardness.

2 Background

Propositional Logic. We consider propositional logic over some fixed universe U of propositional atoms and standard connectives $\vee$, $\wedge$, and $\neg$. A literal is an atom a or a negated atom $\neg a$. A clause is a set of literals. A clause is called *Horn* if at most one of its literals is positive. Unless otherwise stated, formulas are in conjunctive normal form (CNF). For formula φ we denote the set of its clauses by $C(\varphi)$ and its set of atoms by $var(\varphi)$. A formula is called Horn if all of its clauses are Horn. We denote by HORN the class of Horn formulas. We represent an interpretation by a set $I \subseteq U$ containing all variables set to true. If interpretation I satisfies formula φ, we call I a model of φ, denoted by $I \models \varphi$. By $\mathrm{Mod}(\varphi)$ we denote the set of models of φ (over U). For formulas φ and ψ we denote by $\varphi \models \psi$ that $\mathrm{Mod}(\varphi) \subseteq \mathrm{Mod}(\psi)$. A truth assignment τ for a set $V \subseteq U$ is a function that assigns a truth value to each variable in V. We denote by $\varphi[\tau]$ the formula obtained from formula φ by deleting all clauses satisfied by τ and by deleting all literals that are set to false by τ from the remaining clauses.

Belief Revision. The problem of belief revision is specified as follows: Given a knowledge base (i.e., a formula) T and a formula φ, find a revised knowledge base $T \circ \varphi$, such that φ is true in all models of $T \circ \varphi$ and the change compared to the models of T is "minimal". Note that T, φ and ψ are throughout the paper assumed to represent theory, revision formula and query formula. The approach to belief revision we deal with in this paper relies on a so-called model-based operator. Such operators usually utilize a model distance $M \Delta M' = (M \setminus M') \cup (M' \setminus M)$ which yields the set of atoms differently assigned in interpretations M and M'. Assuming that T is consistent (we tacitly make this assumption throughout the paper), the operator due to Satoh [1988] can be defined as follows: $\mathrm{Mod}(T \circ \varphi) = \{J \in \mathrm{Mod}(\varphi) \mid \exists I \in \mathrm{Mod}(T) \text{ s.t. } I \Delta J \in \Delta^{min}(T, \varphi)\}$, $\Delta^{min}(T, \varphi) = min_{\subseteq}\{I \Delta J \mid I \in \mathrm{Mod}(T), J \in \mathrm{Mod}(\varphi)\}$, with $min_{\subseteq}$ selecting minimal elements with respect to set inclusion. The decision problem BR is Π_2^{P}-complete even in case ψ is a single atom [Eiter and Gottlob, 1992].

Cautious (or skeptical) reasoning. (BR)	
Instance:	CNF formulas T and φ, and arbitrary formula ψ with $var(\psi) \subseteq var(T) \cup var(\varphi)$.
Problem:	Decide if $T \circ \varphi \models \psi$ holds.

Parameterized Complexity. Recall that FPT denotes the class of parameterized problems for which there exists an algorithm that decides the problem in time $f(k) \cdot n^{\mathcal{O}(1)}$, where $f(\cdot)$ is an arbitrary computable function that only depends on the parameter k. A *parameterized reduction* of a parameterized problem Π to a parameterized problem Π' is an fpt algorithm that transforms an instance (I, k) of Π to an instance (I', k') of Π' such that: (i) (I, k) is a yes-instance of Π if and only if (I', k') is a yes-instance of Π', and (ii) $k' = g(k)$, where $g(\cdot)$ is an arbitrary computable function that only depends on k. Hardness and completeness with respect to parameterized complexity classes is defined analogously to the concepts from classical complexity theory, using parameterized reductions. The following parameterized classes will be needed in this paper: $\mathsf{FPT} \subseteq$ co-W[1] $\subseteq$ para-coNP $\subseteq \mathsf{FPT}^{\mathsf{NP}}[f(k)] \subseteq$ para-Θ_2^{P}. The class co-W[1] is the co-class of W[1], the first level of the so-called W-hierarchy. A prominent complete problem for W[1] is the INDEPENDENT SET problem parameterized by the size of the required independent set [Downey and Fellows, 1995]. Hence its co-problem is complete for co-W[1]. Parameterized problems in the W-hierarchy share the property that they can be solved in polynomial time if the parameter value is fixed to a constant. On the other hand para-coNP and para-Θ_2^{P} contain those parameterized problems that can be reduced via a parameterized reduction to a coNP-complete or Θ_2^{P}-complete problem respectively. The class $\mathsf{FPT}^{\mathsf{NP}}[f(k)]$ was recently introduced by de Haan and Szeider [2014b] and contains all problems that can be solved by an fpt-algorithm that can use $f(k)$ many calls to an NP oracle [de Haan and Szeider, 2014a]. Two important parameters that we use in this paper are treewidth (see e.g. [Bodlaender, 1993]) and backdoors (see e.g. [Gaspers and Szeider, 2012]).

Treewidth and Backdoors. A *tree decomposition* of a graph $G = (V, E)$ is a pair (H, χ), where H is a tree and χ maps each node t of H (we use $t \in H$ as a shorthand below) to a *bag* $\chi(t) \subseteq V$, such that (i) for each $v \in V$, there is an $t \in H$, s.t. $v \in \chi(t)$; (ii) for each $\{v, w\} \in E$, there is an $t \in H$, s.t. $\{v, w\} \subseteq \chi(t)$; and (iii) for each $t_1, t_2, t_3 \in H$, s.t. t_2 lies on the path from t_1 to t_3, $\chi(t_1) \cap \chi(t_3) \subseteq \chi(t_2)$ holds. The *width* of a tree decomposition is defined as the cardinality of its largest bag $\chi(t)$ minus one. The *treewidth* of a graph G, denoted as $\mathrm{tw}(G)$, is the minimum width over all tree decompositions of G. To build tree decompositions for propositional formulas, we use their incidence graphs. For formula χ such a graph consists of a vertex for each clause and each variable occurring in χ, and has as edges all unordered pairs $\{x, c\}$ where $x \in var(\chi)$ appears in clause $c \in C(\chi)$.

A *strong* HORN-*backdoor set* of a propositional formula χ is a set B of variables such that $\chi[\tau] \in$ HORN for each truth assignment τ over the variables B. This means, the cardinality of such a HORN-backdoor set measures the "distance" of χ to the class of HORN formulas. We denote with $\mathcal{B}(\chi)$ the smallest HORN-backdoor set of χ.

3 Horn, Treewidth and Beyond

We start with an fpt-result for the setting where $|var(\varphi)|$ is considered as the parameter and where T is HORN. Actually, this result is implicit in the proof of [Eiter and Gottlob, 1992, Theorem 8.3], where a polynomial time result was shown if the length of φ is bounded by a constant.

Proposition 1. *(implicit in [Eiter and Gottlob, 1992, Theorem 8.3]) BR parameterized by $|var(\varphi)|$ is fpt if $T \in$* HORN.

A successful technique to generalize tractability results is to move from the restrictions to HORN formulas (or other tractable fragments of SAT) to the backdoor approach and use, e.g., strong HORN-backdoor sets. Using this formalism it is possible to generalize the result of Proposition 1 from the restriction to HORN formulas to propositional formulas which are close to being HORN and hence have a small backdoor.

Theorem 2. *BR parameterized by* $|\mathcal{B}(T)| + |var(\varphi)|$ *is fpt.*

Proof. Let $\mathcal{B}$ be a strong HORN-backdoor set of T. The basic idea is to iterate over all ($2^{|\mathcal{B}|}$ many) assignments τ to $\mathcal{B}$ and to simplify T accordingly. Since $T[\tau] \in$ HORN, checking whether $(T[\tau] \circ \varphi) \models \psi$ holds is in FPT by Theorem 1. □

Notice that this result does not hold if BR is parameterized by $|\mathcal{B}(T)|$ only. In fact, BR remains coNP-complete if $T, \varphi \in$ HORN [Eiter and Gottlob, 1992, Theorem 7.2], which also immediately yields para-coNP-hardness. Another important structural parameter is the combined treewidth of T, φ, and ψ.

Proposition 3. *[Pichler* et al.*, 2009, Theorem 1] BR is fpt when parameterized by* $\mathrm{tw}(T \wedge \varphi \wedge \psi)$.

These initial fpt results possess quite strong conditions on T and φ. As we will see in the next section, the key to gain more versatile fpt results is to parameterize on joint properties of T and φ. Before that, we show that further *independent* relaxations on T and φ indeed increase complexity. Our first result in this direction is concerned with the situation when treewidth of T as well as of φ are bound individually. Then, we study the case of bound treewidth on φ where T is HORN.

Theorem 4. *BR is* para-coNP*-hard if parameterized by* $\mathrm{tw}(T) + \mathrm{tw}(\varphi)$.

Proof. We show coNP-hardness for the acyclic case, i.e., $\mathrm{tw}(T) = \mathrm{tw}(\varphi) = 1$ via reduction from the SAT problem. There we are given a formula χ in CNF and have to decide whether χ is satisfiable. Given an instance of SAT $\chi = c_1 \wedge c_2 \wedge \cdots \wedge c_m$ where $c_1, \ldots, c_m$ are clauses, we construct an instance (T, φ, ψ) of BR as follows. Let $var(\chi) = V$. We create $2m$ new copies of the set V, denoted by $V_1, \ldots, V_m$ and $V'_1, \ldots, V'_m$. For variable $x \in V$ we denote the i-th copy by x_i and the i-th copy of its primed version by x'_i. Additionally, let r be a new propositional variable. For clause c_i of χ we replace its variables with their i-th copy and add the literal $\neg r$. We call the resulting clause γ_i, e.g., for $c_2 = (x \vee \neg y \vee z)$, $\gamma_2 = (x_2 \vee \neg y_2 \vee z_2 \vee \neg r)$. Then $T = (\gamma_1 \wedge \gamma_2 \wedge \cdots \wedge \gamma_m) \wedge \bigwedge_{x \in V} \bigwedge_{1 \leq i \leq m} (x_i \rightarrow x'_i)$, $\varphi = \bigwedge_{x \in V} \bigwedge_{1 \leq i \leq m} (x'_i \rightarrow x) \wedge (x \rightarrow x_i)$, and $\psi = \neg r$. Note that the incidence graphs of T and φ are acyclic.

Next, we show the correctness of the reduction. Assume that χ is a yes-instance, i.e., there exists $M \in \mathrm{Mod}(\chi)$. We construct an interpretation M', which is a model of T and φ. For all $x \in V$ and all $1 \leq i \leq m$ we assign $x \in M'$ iff $x_i \in M'$ iff $x'_i \in M'$ iff $x \in M$. Additionally, $r \in M'$. Since M satisfies all clauses c_i, it holds that M' satisfies all clauses γ_i. The additional implication clauses in T and φ are satisfied by M' since all the different copies of the same variable agree on the truth value. Hence, M' is indeed a model of T and of φ. Therefore, $M' \in \mathrm{Mod}(T \circ \varphi)$. But since $r \in M'$, it follows that $\mathrm{Mod}(T \circ \varphi) \not\models \neg r$. Hence, (T, φ, ψ) is a no-instance.

For the other direction, assume that (T, φ, ψ) is a no-instance, i.e., there exists a model $M \in \mathrm{Mod}(T \circ \varphi)$ with $r \in M$. First we show that M is a model of $T \wedge \varphi$. To this end, we show that there is an interpretation M^*, which is a model of $T \wedge \varphi$. This immediately yields $\Delta^{min}(T, \varphi) = \{\emptyset\}$. We construct the interpretation M^* as follows: Let $r \notin M^*$ and for all $x \in V$ and all $1 \leq i \leq m$ we assign $x \in M^*$ iff $x_i \in M^*$ iff $x'_i \in M^*$. The latter part ensures that all implication clauses in T and φ are satisfied. Since $r \notin M^*$, it follows that M^* trivially satisfies all clauses γ_i, $1 \leq i \leq m$. Hence, M^* is indeed a model of $T \wedge \varphi$ and we have $\Delta^{min}(T, \varphi) = \{\emptyset\}$. Therefore, we know that M is a model of $T \wedge \varphi$, since $M \in \mathrm{Mod}(T \circ \varphi)$ holds by assumption. Next we construct the interpretation $M' = M \cap V$ and show that M' is a model of χ. Since $r \in M$ and each clause γ_i is satisfied by M, there is a literal from V_i in γ_i that is true in M. By construction, $x \in M$ iff $x_i \in M$ and hence x satisfies also clause c_i in M'. Since this holds for all clauses $c_1, \ldots, c_m$, interpretation M' is indeed a model of χ and χ is a yes-instance. □

Theorem 5. *BR is* para-coNP*-hard if parameterized by* $\mathrm{tw}(\varphi)$ *even if* $T \in$ HORN.

Proof. Actually, we show coNP-hardness in the case where φ is acyclic, i.e., $\mathrm{tw}(\varphi) = 1$. We give a reduction from the co-problem of the NP-complete INDEPENDENT SET problem.

Let an instance of INDEPENDENT SET be given by $G = (V, E)$ and $k > 0$. The question is, whether there exists a set of exactly k vertices $V' \subseteq V$ such that there exists no edge between vertices in V'. Let $V = \{v_1, \ldots, v_m\}$. By slight abuse of notation we use the vertices of G also as propositional variables. Thus, we fix the set of variables $var(T \cup \varphi) = \{r, v_1, \ldots, v_m\} \cup \{v_1^1, \ldots, v_m^1, \ldots, v_1^k, \ldots, v_m^k\}$, where for $1 \leq i \leq m$ the variables $v_i^1, \ldots, v_i^k$ are k new copies of the variable v_i and r is a new variable.

Similar to the proof of [Lackner and Pfandler, 2012, Theorem 14] we make use of the following subformulas to encode the INDEPENDENT SET problem. The crucial difference in this proof is the distribution of the subformulas to T and φ. We define $\chi_{\mathrm{IS}} = \bigwedge_{\{v_i, v_j\} \in E} (\neg v_i \vee \neg v_j)$, $\chi_1 = \bigwedge_{1 \leq l \leq k} (v_1^l \vee \cdots \vee v_m^l)$, $\chi_2 = \bigwedge_{1 \leq i \leq m} \bigwedge_{1 \leq l < l' \leq k} (\neg v_i^l \vee \neg v_i^{l'})$, and $\chi_3 = \bigwedge_{1 \leq i \leq m} \bigwedge_{1 \leq l \leq k} (v_i^l \rightarrow v_i)$.

The role of χ_{IS} is to permit only interpretations over $\{v_1, \ldots, v_m\}$ that correspond to an independent set in G. Formula χ_1 ensures that for each copy l with $1 \leq l \leq k$ there is an index i with $1 \leq i \leq m$ representing a vertex that is added to the independent set. In formula χ_2 it is ensured that at most one copy of each vertex can be set to true. Finally, in formula χ_3 the copies v_i^l are mapped back to the original vertices v_i in order to represent the independent set candidate.

The important step is now to distribute these subformulas as follows. In order to construct a BR instance (T, φ, ψ) we set $T = r \rightarrow (\chi_{\mathrm{IS}} \wedge \chi_2 \wedge \chi_3)$, $\varphi = \chi_1$, and $\psi = \neg r$. Notice that T can be transformed into CNF by adding the literal $\neg r$ to each clause of $\chi_{\mathrm{IS}} \wedge \chi_2 \wedge \chi_3$. Observe that $T \in$ HORN and $\mathrm{tw}(\varphi) = 1$, since each variable in $var(\varphi)$ occurs only once in φ.

For the correctness we show that there is a model $M \in \text{Mod}(T \circ \varphi)$ with $r \in M$ (i.e., (T, φ, ψ) is a no-instance) iff G has an independent set of size k (i.e., yes-instance of INDEPENDENT SET). To this end, consider the models in $\text{Mod}(T)$, of which there are two types. The first type t_1 contains the models M with $r \notin M$, whereas the second type t_2 contains models M' such that $r \in M'$ and $M' \models \chi_{\text{IS}} \wedge \chi_2 \wedge \chi_3$.

Assume there is w.l.o.g. an independent set $I \subseteq V$ with $I = \{v_1, \dots v_k\}$ (of size exactly k). Then, due to construction of the subformulas $\chi_{\text{IS}}, \chi_1, \chi_3, \chi_3$, we can construct a model M such that $v \in M$ for all $v \in \{r, v_1, \dots, v_k, v_1^1, \dots, v_k^k\}$, whereas all other variables are set to false. One can verify that M is also model of $T \wedge \varphi$ and that M is of type t_2. Notice that the first property yields that $\Delta^{min}(T, \varphi) = \{\emptyset\}$ and hence $M \in \text{Mod}(T \circ \varphi)$ whereas the second property yields $r \in M$.

For the reverse direction assume that there is a model $M^* \in \text{Mod}(T \circ \varphi)$ with $r \in M^*$. Observe that for all models $M \in \text{Mod}(\varphi)$ with $r \notin M$ it holds that $M \in \text{Mod}(T)$. Hence, we have $\Delta^{min}(T, \varphi) = \{\emptyset\}$. Combining this with the assumption that $M^* \in \text{Mod}(T \circ \varphi)$ it follows that $M^* \models T$. But since $r \in M^*$, i.e. it is of type t_2, we know that $M^* \models \chi_{\text{IS}} \wedge \chi_2 \wedge \chi_3$. From $M^* \models \varphi$ it follows that $M^* \models \chi_1$. Since the models of formula $\chi_{\text{IS}} \wedge \chi_1 \wedge \chi_2 \wedge \chi_3$ correspond to independent sets of G of size k [Lackner and Pfandler, 2012, Theorem 14], it follows that $V \cap M^*$ is an independent set of G of size k. □

With help of the construction used in the proof of Theorem 5, we obtain the following result where the number of clauses in φ, $|C(\varphi)|$, is used as a parameter.

Corollary 6. *BR is* co-W[1]*-hard if parameterized by* $|C(\varphi)|$ *even if* $T \in$ HORN.

Proof. We reduce from the co-problem of INDEPENDENT SET parameterized by the size of the independent set k, which is co-W[1]-complete. In the construction used in the proof of Theorem 5, $|C(\varphi)|$ is bounded by k. Therefore, the reduction is an fpt-reduction, which yields the result. □

4 Shared Variables

If we take a closer look at the constructions used in the proofs of Theorem 4 and 5, we realize that the number of shared variables, i.e., $|\mathit{var}(T) \cap \mathit{var}(\varphi)|$, is unbounded. As it turns out in our next result, using these shared variables as parameter yields a versatile fpt-result.

Theorem 7. *Let* $S = \mathit{var}(T) \cap \mathit{var}(\varphi)$ *and let* p_T *and* p_φ *be parameters such that for any assignment* τ *to* $S \cup \mathit{var}(\psi)$ *the* SAT *problem of* $T[\tau]$ *and* $\varphi[\tau]$*, respectively, can be decided by an fpt-algorithm. Then BR parameterized by* $|S| + p_T + p_\varphi$ *is fpt for queries* ψ *of constant size.*

Proof. The postulated algorithm is Algorithm 1. Intuitively, this algorithm computes first all partial models on the shared variables. The crucial observation is that the symmetric set difference of models of T and φ is restricted to shared variables. Hence, it suffices to restrict the expensive minimality check to those partial models.

We now show the correctness of this algorithm. We claim that M_T after line 5 contains all truth assignments τ to the variables in S and ψ such that τ can be extended to a model of T. Similarly, the set M_φ after line 5 contains all truth assignments τ to the variables in S and ψ such that τ can be extended to a model of φ. To this end, note that line 2 enumerates all subsets of the variables in S and ψ. Line 3 extends such a subset to a truth assignment. In lines 4-5 we test if this truth assignment can be extended to a model of T and φ respectively, and if this is true, we store the assignment in M_T and M_φ respectively. Next, we want to compute $\Delta^{min}(T, \varphi) = min_\subseteq(\{I \Delta J \mid I \in \text{Mod}(T), J \in \text{Mod}(\varphi)\})$. We claim that for the set D in line 6 it holds that $D = \{X \cup \neg X \mid X \in \Delta^{min}(T, \varphi)\}$, where $\neg X = \{\neg x \mid x \in X\}$. Note that the first difference between line 6 and the definition of $\Delta^{min}(T, \varphi)$ is that we consider the symmetric set difference between assignments (containing positive and negative literals) instead of models (containing positive literals only). But this does not change the correctness, since models I and J disagree on variable x iff the corresponding assignments I and J disagree on literals x and $\neg x$. The second difference between line 6 and the definition of $\Delta^{min}(T, \varphi)$ is that we consider only assignments over the variables of S and $\mathit{var}(\psi)$ instead of assignments over all variables. To see that this does not change the correctness, consider a arbitrary set $\delta \in \Delta^{min}(T, \varphi)$. We will show that $\mathit{var}(\delta) \subseteq S$. Assume towards a contradiction that this does not hold, i.e., there exists a variable $x \in \mathit{var}(\delta)$ with $x \in \mathit{var}(T)$ (the case $x \in \mathit{var}(\varphi)$ is symmetric) and $x \notin S$. By definition there exists $I \in \text{Mod}(T)$ and $J \in \text{Mod}(\varphi)$ such that $I \Delta J = \delta$. Since $x \in \delta$ we know that I and J do not agree on the truth value of x. But since $x \notin \varphi$ there exists $J' \in \text{Mod}(\varphi)$ which differs from J only in the truth value of x. Note that the resulting symmetric set difference $I \Delta J'$ is strictly smaller than δ, violating that $\delta \in \Delta^{min}(T, \varphi)$. Hence, indeed $\mathit{var}(\delta) \subseteq S$ and it is enough to compute the symmetric set differences over assignments over variables that contain at least all variables of S. Next, if we compare line 7 with the definition of $\text{Mod}(T \circ \varphi) = \{J \in \text{Mod}(\varphi) \mid \exists I \in \text{Mod}(T) \text{ s.t. } I \Delta J \in \Delta^{min}(T, \varphi)\}$, we can observe two differences. First, we get again a set of assignments instead of a set of models (containing positive literals only). Second, the resulting assignments are restricted to the variables in S and $\mathit{var}(\psi)$. Formally, we have the relation $R = \{(J \cap (S \cup \mathit{var}(\psi))) \cup \bigcup_{x \in J \cap (S \cup \mathit{var}(\psi))} \neg x \mid J \in \text{Mod}(T \circ \varphi)\}$. Finally, we have to check $T \circ \varphi \models \psi$. This means we go over each model of $T \circ \varphi$ and check if any of these models does not satisfy ψ (lines 8-9). Note that it is indeed enough to go over assignments $I \in R$ since the truth value of all variables in $\mathit{var}(\psi)$ is fixed in these assignments.

For the complexity, note that the loop in line 2 is executed $2^{|S \cup \mathit{var}(\psi)|}$ times. Since the two SAT calls can be executed by an fpt-algorithm with respect to parameters p_T and p_φ, we get an overall time for this loop of $\mathcal{O}^*(2^{|S \cup \mathit{var}(\psi)|}(f(p_T) + g(p_\varphi)))$, where f and g are computable functions only depending on p_T and p_φ, respectively. Furthermore, the size of M_T and M_φ after line 5 is bounded by $2^{|S \cup \mathit{var}(\psi)|}$ as well. Therefore, fpt for $|S|$, p_T and p_φ follows. □

Observe that the above result actually gives a family of fpt-results for BR with the parameter "number of shared variables" in common. In addition, since the revision formula φ usually does not contribute new knowledge to all parts of the theory,

Algorithm 1: Fpt-algorithm of Theorem 7

Input: Formulas T, φ and ψ.
Output: Decision whether $T \circ \varphi \models \psi$ holds.

$M_T = \emptyset$, $M_\varphi = \emptyset$
foreach $J \in 2^{S \cup var(\psi)}$ **do**
    $I = J \cup \bigcup_{x \in (S \cup var(\psi)) \setminus J} \neg x$
    if $T[I]$ *is satisfiable* **then** $M_T = M_T \cup \{I\}$
    if $\varphi[I]$ *is satisfiable* **then** $M_\varphi = M_\varphi \cup \{I\}$
$D = min_{\subseteq}(\{I \Delta J \mid I \in M_T, J \in M_\varphi\})$
$R = \{J \in M_\varphi \mid \exists I \in M_T \text{ s.t. } I \Delta J \in D\}$
foreach $I \in R$ **do**
    if $I \not\models \psi$ **then return** *no*
return *yes*

but rather introduces new information on a certain subject, one can hope that the number of shared variables is rather small even if theory T is large. The advantage of this result is that it becomes possible to combine *any* tractable fragments that make that satisfiability problem for T and φ fpt. Furthermore, notice that the choice of p_T and p_φ is completely independent. This yields a modular approach that can be instantiated with any reasonable combination of tractable fragments. For example, we obtain an fpt-result where we decouple the treewidth of T from the treewidth of φ if we additionally parameterize by $|var(T) \cap var(\varphi)|$. Compared to the hardness result in Theorem 4 we see that the number of shared variables is the essential parameter to obtain the fpt-result.

Although parameterizing by the number of shared variables provides an important bridge towards tractability, considering this parameter in isolation does not yield an fpt result.

Theorem 8. *BR is* $\mathsf{FPT}^{\mathsf{NP}}[f(k)]$*-complete if parameterized by* $|S|$*, where* $S = var(T) \cap var(\varphi)$*.*

Proof. We start by showing membership in $\mathsf{FPT}^{\mathsf{NP}}[f(k)]$. Let (T, φ, ψ) be an arbitrary instance of BR with $var(T) \cap var(\varphi) = S$ and $k = |S|$. As discussed in the proof of Theorem 7, $\Delta^{min}(T, \varphi)$ is restricted to variables in S. Hence, an upper bound for the number of minimal symmetric differences is given by $\left|\Delta^{min}(T, \varphi)\right| \leq 2^k$. We can test for each symmetric difference Z with a single NP oracle call if (i) it is minimal, or (ii) whether there exist models M and M' such that $M \models T$, $M' \models \varphi$ with $M \Delta M' = Z$ and $M' \not\models \psi$. Thus, we can decide whether $T \circ \varphi \models \psi$ holds by invoking $\mathcal{O}(2^k)$ calls to an NP oracle.

We prove hardness via reduction from the $\mathsf{FPT}^{\mathsf{NP}}[f(k)]$-complete problem BH(LEVEL)-SAT [de Haan and Szeider, 2014a] with parameter k, s.t. the constructed instance (T, φ, ψ) satisfies $|S| \leq k$. We recall first the definition of the unparameterized problem BH_n-SAT. An instance of this problem is given by $I = (\chi_1, \ldots, \chi_n)$, where each χ_i is a formula. If $n = 1$, then I is a yes-instance iff χ_1 is satisfiable. If $n \geq 2$ is odd then I is a yes-instance iff χ_n is satisfiable or $(\chi_1, \ldots, \chi_{n-1})$ is a yes-instance of BH_{n-1}-SAT. Otherwise, if $n \geq 2$ is even then I is a yes-instance iff χ_n is unsatisfiable and $(\chi_1, \ldots, \chi_{n-1})$ is a yes-instance of BH_{n-1}-SAT. A problem instance of the parameterized problem BH(LEVEL)-SAT is now such a sequence $I = (\chi_1, \ldots, \chi_k)$ with parameter k, and we ask whether I is a yes instance of BH_k-SAT.

Let now $I = (\chi_1, \ldots, \chi_k)$ be an arbitrary instance of BH(LEVEL)-SAT. W.l.o.g. we assume disjoint vocabularies for each pair of formulas. We define $T = \bigwedge_{1 \leq i \leq k} q_i$ and $\varphi = (r(I) \rightarrow q) \wedge \bigwedge_{1 \leq i \leq k} (q_i \rightarrow \chi_i)$ where q and the q_is are fresh variables, and $r(\cdot)$ is a recursive function. If $k = 1$, then $r(I) = q_1$. If $k \geq 2$ is odd, then $r(I) = (q_k \vee (r(\chi_1, \ldots, \chi_{k-1})))$. If $k \geq 2$ is even, then $r(I) = (\neg q_k \wedge (r(\chi_1, \ldots, \chi_{k-1})))$. We can transform φ straightforwardly to CNF. Finally let $\psi = q$.

We show that (T, φ, ψ) is a yes-instance of BR iff I is a yes-instance of BH(LEVEL)-SAT. First, consider the two cases for each χ_i: (i) χ_i is unsatisfiable, or (ii) χ_i is satisfiable. If χ_i is unsatisfiable, then $(q_i \rightarrow \chi_i) \equiv \neg q_i$. Thus, for all $M \in \mathrm{Mod}(T \circ \varphi)$ we have $q_i \notin M$. If χ_i is satisfiable, then there is a model $M_i \models \chi_i$. Now suppose that there is a model $M \in \mathrm{Mod}(T \circ \varphi)$, s.t. $q_i \notin M$. It follows that for an $M_T \models T$ we have $M_T \Delta M \in \Delta^{min}(T, \varphi)$. Let $M' = (M \setminus var(\chi_i)) \cup (M_i \cap var(\chi_i)) \cup \{q_i, q\}$. Clearly, $M' \models (q_i \rightarrow \chi_i)$ and $M' \models (r(I) \rightarrow q)$ and thus $M' \models \varphi$. Define the following model of T: $M'_T = M' \cup \{q_j \mid 1 \leq j \leq k\}$. It is straightforward to see that both $M' \Delta M'_T$ and $M \Delta M_T$ are subsets of the shared variables $\{q_j \mid \chi_j \in I\}$. We can conclude that $M' \Delta M'_T \subset M \Delta M_T$ (they differ only in q_i, since $q_i \in M_T$), contradicting that $M \in \mathrm{Mod}(T \circ \varphi)$.

It follows that $\Delta^{min}(T, \varphi) = \{\{q_i \mid \chi_i \text{ unsatisfiable }\}\}$. If I is a yes-instance of BH(LEVEL)-SAT, we can derive that for an $M \in \mathrm{Mod}(T \circ \varphi)$ it holds that $M \models r(I)$, and thus $q \in M$ and $M \models \psi$. If (T, φ, ψ) is a yes-instance of BR, then all models $M \in \mathrm{Mod}(T \circ \varphi)$ satisfy $\psi = q$, and thus $M \models r(I)$. This implies that I is a yes-instance of BH(LEVEL)-SAT. ☐

5 Maximum Hamming Distance

Finally, we investigate a further natural relaxation for capturing distance between T and φ. Define the *maximum Hamming distance* between T and φ as $|\Delta|^{max}(T, \varphi) = max(\{|(I \Delta J) \cap S| \mid I \in \mathrm{Mod}(T), J \in \mathrm{Mod}(\varphi)\})$, where $S = var(T) \cap var(\varphi)$. In words, $|\Delta|^{max}(T, \varphi)$ delivers the maximum number of shared variables that are assigned differently in a model of T and a model of φ. Bounding S implies a bound on $|\Delta|^{max}(T, \varphi)$. However, parameter $|\Delta|^{max}(T, \varphi)$ is not enough for membership in $\mathsf{FPT}^{\mathsf{NP}}[f(k)]$.

Theorem 9. *BR is* para-Θ_2^{P}*-hard if parameterized by maximum Hamming distance.*

Proof. We use a reduction from the Θ_2^{P}-complete problem UOCSAT [Kadin, 1989] such that the constructed instance has a maximum Hamming distance of 5. Let $\Gamma = c_1 \wedge \cdots \wedge c_n$ be an arbitrary instance of UOCSAT, where Γ is a formula in CNF with clauses $c_1, \ldots, c_n$. Γ is a yes-instance of UOCSAT if it is the case that for all interpretations satisfying a maximum number of clauses of Γ it holds that they satisfy the same clauses. For the reduction we utilize encodings for cardinality constraints. In particular we use so-called at-most constraints, $atMost(i, X)$, which evaluate to true under an interpretation M iff $|M \cap X| \leq i$. Furthermore, $exact(i, X)$ states that exactly i variables are assigned to true. One can construct

FPT	co-W[1] hard	para-coNP hard	$\mathsf{FPT}^{\mathsf{NP}}[f(k)]$ complete	para-Θ_2^{P} hard
T Horn, $\lvert var(\varphi)\rvert$ (Prop. 1) $\lvert\mathcal{B}(T)\rvert + \lvert var(\varphi)\rvert$ (Thm. 2) $\mathrm{tw}(T \wedge \varphi \wedge \psi)$ (Prop. 3) $\lvert S\rvert + p_T + p_\varphi$ (Thm. 7)	$\lvert C(\varphi)\rvert$ (Cor. 6)	$\mathrm{tw}(T) + \mathrm{tw}(\varphi)$ (Thm. 4) T Horn, $\mathrm{tw}(\varphi)$ (Thm. 5)	$\lvert S\rvert$ (Thm. 8)	max. Hamming dist. (Thm. 9)

Table 1: Parameterized complexity landscape for BR ($S = var(T) \cap var(\varphi)$).

formulas for such constraints which are polynomial in size (see [Roussel and Manquinho, 2009, Section 22.2.3.]).

For our reduction, we introduce three sets of variables: $L = \{l_i \mid 1 \leq i \leq n\}$, $R = \{r_i \mid 0 \leq i \leq n\}$, and $A = \{a_i \mid 0 \leq i \leq n\}$. We translate the given formula to $\Gamma^* = (l_1 \vee c_1) \wedge \cdots \wedge (l_n \vee c_n)$. One of the main subformulas for the reduction is $\chi = \Gamma^* \wedge \bigwedge_{r_i \in R}(r_i \leftrightarrow exact(i, L))$. That is, it contains the instrumented original formula with the fresh literals l_j and variables r_i such that any model of χ assigns r_i to true iff i variables from L are assigned to true. A basic observation is that, if $M \models \chi$ then there is an i with $0 \leq i \leq n$ such that $M \models r_i$ and for $j \neq i$ it holds that $M \not\models r_j$. In addition, M satisfies at least $n - i$ clauses of Γ. The formula χ can be transformed to an equivalent formula in CNF. The problematic conjuncts are $(r_i \leftrightarrow exact(i, L))$. We can change these to $(r_i \rightarrow exact(i, L))$, which can be transformed to CNF in polynomial time, and add further clauses to ensure that exactly one r_i is true in a model of χ. All the formulas in the sequel of this proof can be easily transformed to CNF, but for the sake of readability we will use non-CNF formulas.

Construct $\varphi = \chi \wedge \chi' \wedge ((\alpha \rightarrow \beta) \rightarrow q) \wedge \bigwedge_{a_i \in A}(r_i \leftrightarrow a_i)$ with $\alpha = (\bigwedge_{r_i \in R} r_i \leftrightarrow r'_i)$, $\beta = (\bigwedge_{l_i \in L} l_i \leftrightarrow l'_i)$, and q as a fresh variable. Within this proof if G is a formula, then we denote by G' the formula obtained by uniformly renaming all atoms with a prime. Define $T = (\bigwedge_{a_i \in A} \neg a_i) \wedge atMost(1, R') \wedge R' < R \wedge exact(1, R)$, where $R' < R = \bigwedge_{r'_j \in R'}(r'_j \rightarrow (\bigvee_{j<i} r_i))$. To finish our instance let $\psi = q$.

We now state some properties of the instance (T, φ, ψ). The instance can be computed in polynomial time w.r.t. Γ. All three formulas are satisfiable. It holds that $|\Delta|^{max}(T, \varphi) \leq 5$, since the shared vocabulary is $A \cup R \cup R'$ and all models of φ may assign only one variable per set to true, while all models of T may assign at most one variable of R' to true and exactly one variable of R to true (and none of A). Further, it cannot be the case that both r_i and r'_i are in a model of T. It cannot be the case that $\emptyset \in \Delta^{min}(T, \varphi)$. For each $D \in \Delta^{min}(T, \varphi)$ it holds that $a_i \in D$ for some $0 \leq i \leq n$. For two models M_φ, M_T s.t. $M_\varphi \models \varphi$ and $M_T \models T$ we have $\{a_i\} = M_\varphi \Delta M_T$ only if $\{r_i, r'_j\} \subseteq M_\varphi$ for some $j < i$. A last crucial observation is that if $M_\varphi \models \varphi$ and $\{r_i, r'_i\} \subseteq M_\varphi$, then there is a model $M_T \models T$ such that $M_\varphi \Delta M_T = \{a_i, r'_i\}$, and there is no M'_T such that $M'_T \models T$ and $M_\varphi \Delta M'_T \subset M_\varphi \Delta M_T$ (the set $\{a_i, r'_i\}$ is a minimal symmetric difference we can achieve).

Let $n - d$ be the maximum number of clauses of Γ that can be simultaneously satisfied. Further let M_1, M_2 be two interpretations such that $M_1 \models \varphi$ with $M_1 \models r_d \wedge r'_d$ and $M_2 \models \varphi$ with $M_2 \models r_{d_1} \wedge r'_{d_1}$ such that $d < d_1$. We show that (i) $M_1 \in \mathrm{Mod}(T \circ \varphi)$, and (ii) $M_2 \notin \mathrm{Mod}(T \circ \varphi)$. For (i), we first show that there is no M_φ such that $M_\varphi \models \varphi$ and $\{r_d, r'_j, a_d\} \subseteq M_\varphi$ for $j < d$. Suppose to the contrary that such an M_φ exists. Then M_φ satisfies $n - j$ clauses of Γ' where $n - j > n - d$. This is a contradiction to $n - d$ being the maximum number of clauses simultaneously satisfiable for Γ. By the observations above, we know that $\{a_d, r'_d\}$ is a minimal symmetric difference we can find for M_1 when comparing it to all models of T. We now show that there are no $M_\varphi \models \varphi$ and $M_T \models T$ such that $M_\varphi \Delta M_T \subset \{a_d, r'_d\}$. Due to the observations above, we cannot find such models such that the symmetric difference is a proper subset and does not contain a_d. To achieve the symmetric difference of $\{a_d\}$ we would need a model of φ that assigns r_d and r'_j to true for some $j < d$. This is a contradiction. Thus $M_1 \in \mathrm{Mod}(T \circ \varphi)$. The second case (ii) follows analogously, just note that we can find in this case a model of φ and model of T such that the symmetric difference contains only a_{d_1}. It follows that $\{M \in \mathrm{Mod}(T \circ \varphi) \mid M \models \alpha\} \cap var(\Gamma) = \{M' \mid M'$ satisfies a maximum number of clauses of $\Gamma\} \cap var(\Gamma)$.

We now show that (T, φ, ψ) is a yes-instance of BR iff Γ is a yes-instance of UOCSAT. Assume Γ is a yes-instance of UOCSAT. This implies that if an $M \in \mathrm{Mod}(T \circ \varphi)$ satisfies α ($M \models \alpha$) it holds that $M \models \beta$ (otherwise M would encode two different interpretations of Γ satisfying a maximum number of clauses, but satisfying different clause sets). Thus $M \models \psi$. The other direction can be shown analogously: if two models satisfying a maximum number of clauses for Γ exist, s.t. they satisfy different clauses, then there is an $M \in \mathrm{Mod}(T \circ \varphi)$, s.t. $M \models \alpha$, $M \not\models \beta$, and $q \notin M$. □

We complement this result by an upper bound on the required number of calls to a SAT-solver in order to solve the problem. The number of possible symmetric differences projected to the shared variables S is bounded by $\sum_{i=0}^{d} \binom{|S|}{i}$, with d the maximum Hamming distance. We can bound this by $\sum_{i=0}^{d} \binom{|S|}{i} \leq (1 + |S|)^d$. Thus, we can bound the number of required NP oracle calls by $\mathcal{O}((1 + |S|)^d)$.

6 Discussion

In this work we have studied the parameterized complexity of belief revision for parameters beyond treewidth. We summarize our results in Table 1. In particular, we have explored to what extend it is possible to decouple the theory from the revision formula. The parameter "shared variables" turned out to be crucial to obtain fpt-results in the decoupled setting. Ultimately, this parameter led to a versatile, modular fpt-result where tractable fragments of the theory and the revision formula can be combined arbitrarily. Moreover, if we do not rely on tractable fragments the problem becomes complete for the recently introduced class $\mathsf{FPT}^{\mathsf{NP}}[f(k)]$. This result is interesting as it adds a natural problem to this rather unexplored class, and, furthermore, it shows the applicability of a SAT-based procedure that can solve the problem with a restricted number of calls if the parameter is bounded. Our result also complements the work of Delgrande *et al.* [2007], who presented algorithms and implementations based on SAT-solvers for belief change problems, and provides further evidence that the usage of such solvers is a promising and interesting direction for devising systems in the area of belief revision.

We envisage the study of further belief revision operators, such as Dalal's operator [Dalal, 1988], in a parameterized setting. For other operators, the results might be very different as the computational complexity is heavily influenced by the choice of the operator (cf. [Creignou *et al.*, 2013]). Furthermore, we want to investigate to what extend the shared variables are key towards tractability for other belief change formalisms, like belief merging, contraction, or update.

References

[Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.*, 50:510–530, 1985.

[Bodlaender, 1993] Hans L. Bodlaender. A tourist guide through treewidth. *Acta Cybern.*, 11(1-2):1–21, 1993.

[Creignou *et al.*, 2013] Nadia Creignou, Reinhard Pichler, and Stefan Woltran. Do hard SAT-related reasoning tasks become easier in the Krom fragment? In *Proc. IJCAI 2013*, pages 824–831. AAAI Press, 2013.

[Dalal, 1988] Mukesh Dalal. Investigations into a theory of knowledge base revision. In *Proc. AAAI 1988*, pages 475–479, 1988.

[de Haan and Szeider, 2014a] Ronald de Haan and Stefan Szeider. Compendium of parameterized problems at higher levels of the polynomial hierarchy. *Electronic Colloquium on Computational Complexity (ECCC)*, 21:143, 2014.

[de Haan and Szeider, 2014b] Ronald de Haan and Stefan Szeider. Fixed-parameter tractable reductions to SAT. In *Proc. SAT 2014*, volume 8561 of *LNCS*, pages 85–102. Springer, 2014.

[Delgrande *et al.*, 2007] James P. Delgrande, Daphne H. Liu, Torsten Schaub, and Sven Thiele. COBA 2.0: A consistency-based belief change system. In *Proc. ECSQARU 2007*, volume 4724 of *LNCS*, pages 78–90. Springer, 2007.

[Downey and Fellows, 1995] Rodney G. Downey and Michael R. Fellows. Fixed-parameter tractability and completeness II: on completeness for W[1]. *Theor. Comput. Sci.*, 141(1&2):109–131, 1995.

[Downey and Fellows, 1999] Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*. Springer, 1999.

[Dvořák *et al.*, 2012] Wolfgang Dvořák, Reinhard Pichler, and Stefan Woltran. Towards fixed-parameter tractable algorithms for abstract argumentation. *Artif. Intell.*, 186:1–37, 2012.

[Dvořák *et al.*, 2014] Wolfgang Dvořák, Matti Järvisalo, Johannes P. Wallner, and Stefan Woltran. Complexity-sensitive decision procedures for abstract argumentation. *Artif. Intell.*, 206:53–78, 2014.

[Eiter and Gottlob, 1992] Thomas Eiter and Georg Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artif. Intell.*, 57(2-3):227–270, 1992.

[Fellows *et al.*, 2012] Michael R. Fellows, Andreas Pfandler, Frances A. Rosamond, and Stefan Rümmele. The parameterized complexity of abduction. In *Proc. AAAI 2012*, pages 743–749. AAAI Press, 2012.

[Gaspers and Szeider, 2012] Serge Gaspers and Stefan Szeider. Backdoors to satisfaction. In *The Multivariate Algorithmic Revolution and Beyond - Essays Dedicated to Michael R. Fellows on the Occasion of His 60th Birthday*, volume 7370 of *LNCS*, pages 287–317. Springer, 2012.

[Gottlob and Szeider, 2008] Georg Gottlob and Stefan Szeider. Fixed-parameter algorithms for artificial intelligence, constraint satisfaction and database problems. *Comput. J.*, 51(3):303–325, 2008.

[Gottlob *et al.*, 2010] Georg Gottlob, Reinhard Pichler, and Fang Wei. Bounded treewidth as a key to tractability of knowledge representation and reasoning. *Artif. Intell.*, 174(1):105–132, 2010.

[Kadin, 1989] Jim Kadin. $P^{(NP[O(logn)])}$ and sparse turing-complete sets for NP. *J. Comput. Syst. Sci.*, 39(3):282–298, 1989.

[Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artif. Intell.*, 52:263–294, 1991.

[Lackner and Pfandler, 2012] Martin Lackner and Andreas Pfandler. Fixed-parameter algorithms for finding minimal models. In *Proc. KR 2012*, pages 85–95. AAAI Press, 2012.

[Liberatore and Schaerf, 2001] Paolo Liberatore and Marco Schaerf. Belief revision and update: Complexity of model checking. *J. Comput. Syst. Sci.*, 62(1):43–72, 2001.

[Peppas, 2008] Pavlos Peppas. Belief revision. In *Handbook of Knowledge Representation*, pages 317–359. Elsevier, 2008.

[Pichler *et al.*, 2009] Reinhard Pichler, Stefan Rümmele, and Stefan Woltran. Belief revision with bounded treewidth. In *Proc. LPNMR 2009*, volume 5753 of *LNCS*, pages 250–263. Springer, 2009.

[Roussel and Manquinho, 2009] Olivier Roussel and Vasco M. Manquinho. Pseudo-Boolean and cardinality constraints. In *Handbook of Satisfiability*, volume 185 of *FAIA*, pages 695–733. IOS Press, 2009.

[Satoh, 1988] Ken Satoh. Nonmonotonic reasoning by minimal belief revision. In *Proc. FGCS 1988*, pages 455–462, 1988.