

This manuscript was published as:

Adrian Haret and Stefan Woltran. Deviation in Belief Change on Fragments of Propositional Logic. In *Proceedings of the 6th Workshop on Dynamics of Knowledge and Belief (DKB-2017) and the 5th Workshop KI and Kognition (KIK-2017)*, pages 64–76, Dortmund, Germany, September 26, 2017.

Deviation in Belief Change on Fragments of Propositional Logic^{*}

Adrian Haret and Stefan Woltran

DBAI group, TU Wien, A-1040, Vienna, Austria
{haret,woltran}@dbai.tuwien.ac.at

Abstract. It is known that prominent fragments of propositional logic are not closed under standard belief change operators. That is, applying such operators to knowledge bases in a fragment may produce results that have no equivalent in the same language. However, the potential range of such a deviation has not been investigated yet. In this paper, we give a systematic study of this problem by considering four prominent change operators (Dalal, Satoh, Winslett, and Forbus) and three important fragments (1CNF, Krom, and Horn). While all operators are shown to be closed under the 1CNF fragment, we observe that for the other two fragments the behavior of the operators significantly differs. We expect our considerations on deviation to play an important role in the design of change operators for concrete Knowledge Representation formalisms.

1 Introduction

Belief change [1, 11] occupies a central role in understanding the logic of modifying a knowledge base. The framework it provides allows formalization and comparison of various types of change, such as revision [15] and update [14]. Classical results in the field assume that the language in which knowledge is expressed subsumes propositional logic, but recent work has been increasingly focused on change in more specialized formalisms, suitable for use in concrete applications because of the kinds of things they can express, or their attractive computational properties [6, 19, 7, 20].

An obstacle in the application of standard belief change procedures to languages that do not subsume propositional logic is the fact that results are not guaranteed to be expressible in the same language. As an example consider the Horn fragment and the revision of a knowledge base $K = \{a \wedge b\}$ by formula $\mu = \neg a \vee \neg b$. Most belief change operators deliver the intuitive result that the outcome of this change should be equivalent to $a \oplus b$; the latter cannot be expressed in Horn although K and μ are from this fragment.

Existing research that takes this as its starting point has taken several directions. Papers in the spirit of [6] have studied how standard postulates have to be adapted such that representation theorems [1, 11] can be given for the Horn fragment. Another approach has looked into repairing inexpressible results like the one sketched above, in order to make them fit into the language of interest [2, 3]. Finally, the design of novel

^{*} This work was supported by the Austrian Science Fund (FWF) under grants P25521, P30168.

change operators, tailored specifically for fragments, has also been subject of recent research [8, 13].

In this paper we focus on certain fragments of propositional logic, namely the 1CNF, Krom and Horn fragment. Propositional fragments are of interest because: (i) they are closely related to the original framework for belief change, hence best suited to illustrate general principles, and (ii) some of these fragments serve as the basis for popular formalisms in Knowledge Representation (e.g., Horn logic is the backbone of DL-Lite). Our main aim is to investigate the interplay between these fragments and established belief change operators (Dalal, Satoh, Winslett, Forbus). These four operators are among the most well studied in the field, both with respect to their semantic representation and their complexity in fragments [9, 16, 4].

Thus, what is required is a closer look into the behaviour of established operators and the extent to which results can fall outside a given fragment. With the exception of Dalal’s operator [12], this has not yet been done. To this end, we introduce a measure for the degree to which a fragment is transformed through application of the studied operators. We call this measure the *deviation* of a revision operator from a particular fragment, and provide results on the deviation of established operators from the above-mentioned fragments. Our results indicate whether the range of possible results of an operator applied to a fragment is restricted to a subset of full propositional logic reasonably close to the original fragment, or if, by contrast, any propositional base can be obtained. Another issue is whether there exists any fragment which is closed under the operators and this, as we shall see, holds for the 1CNF fragment. We also show, through a comparative analysis, that different operators applied to the same fragment deviate in different ways. In other words, understanding deviation sheds light on the differences between major revision operators and on the expressiveness of the fragments considered. We expect considerations of this kind to play an important role in the design of change operators for concrete Knowledge Representation formalisms.

The rest of the paper is structured as follows. In Section 2 we present background notions. Section 3 presents our main results on the deviation of the 1CNF, 2CNF and Horn fragments. Section 4 provides a comparison between the ranges of the operators, and Section 5 offers conclusions and pointers to future work.

2 Background

We write \mathcal{L} for *the language of propositional logic* constructed from a finite alphabet \mathcal{U} of atoms using standard connectives \vee, \wedge, \neg , and constants \top, \perp . An *interpretation* is a set of atoms (the ones set to true), and the set of all interpretations is \mathcal{W} . The set of models of a formula φ is denoted by $[\varphi]$. If there is no danger of ambiguity, we write models as strings made up of their elements (e.g., abc instead of $\{a, b, c\}$). A knowledge base (knowledge base) is a finite set of formulas. We will usually identify a knowledge base K with $\bigwedge_{\varphi \in K} \varphi$. The set of models of a knowledge base K is $[K] = \bigcap_{\varphi \in K} [\varphi]$. We write $w \Delta u$ for the symmetric difference between interpretations w and u . If $\mathcal{M}, \mathcal{N} \subseteq \mathcal{W}$, then $\mathcal{M} \diamond \mathcal{N} = \{w \Delta u \mid w \in \mathcal{M}, u \in \mathcal{N}\}$. We use $w \diamond \mathcal{M}$ or $\mathcal{M} \diamond w$ to abbreviate $\{w\} \diamond \mathcal{M}$ or $\mathcal{M} \diamond \{w\}$, respectively. We write $\min_{\subseteq}(\mathcal{M}) = \{w \in \mathcal{M} \mid \exists w' \in \mathcal{M} \text{ s.t. } w' \subset w\}$, and $\min_{\text{card}}(\mathcal{M}) = \{w \in \mathcal{M} \mid \nexists w' \in \mathcal{M} \text{ s.t. } |w'| < |w|\}$.

We define a fragment of propositional logic as a set of propositional formulas characterized by a closure property on the set of models. Thus, a mapping $Cl: 2^{\mathcal{W}} \rightarrow 2^{\mathcal{W}}$ is called a *closure-operator* if, for any $\mathcal{M}, \mathcal{N} \subseteq \mathcal{W}$, it holds that (i) if $\mathcal{M} \subseteq \mathcal{N}$, then $Cl(\mathcal{M}) \subseteq Cl(\mathcal{N})$, (ii) if $|\mathcal{M}| = 1$, then $Cl(\mathcal{M}) = \mathcal{M}$ and (iii) $Cl(\emptyset) = \emptyset$. A *fragment* is a set $\mathcal{F} \subseteq \mathcal{L}$ closed under conjunction (i.e., $\varphi \wedge \psi \in \mathcal{F}$ for any $\varphi, \psi \in \mathcal{F}$) for which there exists an associated closure-operator Cl such that (i) for all $\varphi \in \mathcal{F}$, $[\varphi] = Cl([\varphi])$ and (ii) for all $\mathcal{M} \subseteq \mathcal{W}$ there is a $\varphi \in \mathcal{F}$ with $[\varphi] = Cl(\mathcal{M})$. We denote the closure-operator Cl associated to a fragment \mathcal{F} as $Cl_{\mathcal{F}}$. An \mathcal{F} -knowledge base is a finite set $K \subseteq \mathcal{F}$. A knowledge base $K \subseteq \mathcal{L}$ is \mathcal{F} -*expressible* if there exists an \mathcal{F} -knowledge base K' , such that $[K] = [K']$. A set of interpretations \mathcal{M} is \mathcal{F} -*expressible* if there exists an \mathcal{F} -knowledge base K such that $[K] = \mathcal{M}$.

Many well-known fragments of propositional logic are captured by this notion. To the Horn fragment (i.e., conjunctions of clauses with at most one positive literal) we associate the operator Cl_{Horn} , defined as the fixed point of the function $\bigcap(\mathcal{M}) = \{w_1 \cap w_2 \mid w_1, w_2 \in \mathcal{M}\}$. The Krom, or 2CNF, fragment (i.e., conjunctions of clauses of length at most 2) is linked to the operator $Cl_{2\text{CNF}}$, defined as the fixed point of the function $\text{maj}(\mathcal{M}) = \{\text{maj}_3(w_1, w_2, w_3) \mid w_1, w_2, w_3 \in \mathcal{M}\}$, where ternary majority $\text{maj}_3(w_1, w_2, w_3)$ yields an interpretation containing those atoms true in at least two out of w_1, w_2 and w_3 . Finally, the 1CNF fragment (i.e., conjunctions of literals) has as its operator $Cl_{1\text{CNF}}$, defined as the fixed point of the function $\text{Fill}(\mathcal{M}) = \{w_1 \cap w_2, w_1 \cup w_2 \mid w_1, w_2 \in \mathcal{M}\} \cup \{w_3 \mid w_1 \subseteq w_3 \subseteq w_2 \text{ and } w_1, w_2 \in \mathcal{M}\}$. Note that full classical logic is given via the identity closure operator $Cl_{\mathcal{L}}(\mathcal{M}) = \mathcal{M}$.

A *revision operator* \circ maps a knowledge base K and a formula μ to a knowledge base $K \circ \mu$. We consider here four standard operators: Winslett (\circ^W), Satoh (\circ^S), Forbus (\circ^F) and Dalal (\circ^D) [18, 17, 10, 5]. These operators are defined as follows:

$$\begin{aligned} [K \circ^W \mu] &= \{w \in [\mu] \mid \exists u \in [K] \text{ such that } w \Delta u \in \min_{\subseteq}([\mu] \diamond u)\}, \\ [K \circ^S \mu] &= \{w \in [\mu] \mid \exists u \in [K] \text{ such that } w \Delta u \in \min_{\subseteq}([\mu] \diamond [K])\}, \\ [K \circ^F \mu] &= \{w \in [\mu] \mid \exists u \in [K] \text{ such that } w \Delta u \in \min_{\text{card}}([\mu] \diamond u)\}, \\ [K \circ^D \mu] &= \{w \in [\mu] \mid \exists u \in [K] \text{ such that } w \Delta u \in \min_{\text{card}}([\mu] \diamond [K])\}. \end{aligned}$$

Example 1. Consider a knowledge base K with $[K] = \{abcd, a\}$ and a formula μ with $[\mu] = \{acd, bd, a\}$. For $K \circ \mu$, consult $[\mu] \diamond [K]$ depicted in Table 1. We have that $[K \circ^D \mu] = \{acd\}$, since $acd \Delta abcd$ is cardinality-minimal in the whole table; $[K \circ^S \mu] = \{acd, bd\}$, since $acd \Delta abcd$ and $bd \Delta abcd$ are \subseteq -minimal in the whole table; $[K \circ^F \mu] = \{acd, b\}$, since $acd \Delta abcd$ is cardinality-minimal on the $abcd$ -column, and $b \Delta a$ is cardinality-minimal on the a -column; $[K \circ^W \mu] = \{acd, bd, b\}$, since $acd \Delta abcd$ and $bd \Delta abcd$ are \subseteq -minimal on the $abcd$ -column, and $b \Delta a$ is \subseteq -minimal on the a -column.

This example illustrates a relationship between the operators which holds more generally, namely that for any knowledge base K and formula μ , we have $[K \circ^D \mu] \subseteq [K \circ^S \mu] \subseteq [K \circ^W \mu]$ and $[K \circ^D \mu] \subseteq [K \circ^F \mu] \subseteq [K \circ^W \mu]$, while $[K \circ^S \mu]$ and $[K \circ^F \mu]$ are not necessarily in a subset relationship to each other. We will be interested in all formulas that can be obtained by applying an operator to knowledge bases in a given fragment.

	$\overbrace{\quad\quad\quad}^{[K]}$		
	\diamond	$abcd$	a
$[μ]$	acd	b	cd
	bd	ac	abd
	b	acd	ab

Table 1. Revision: \subseteq -minimal elements in grey, cardinality-minimal elements in bold font

Definition 1. For a revision operator \circ and a fragment $\mathcal{F} \subseteq \mathcal{L}$, the image of \circ with respect to \mathcal{F} , denoted by $Im_{\mathcal{F}}(\circ)$, is defined as $Im_{\mathcal{F}}(\circ) = \{K \circ \mu \mid K, \mu \in \mathcal{F}\}$.

It is straightforward to see that for any operator $\circ \in \{\circ^D, \circ^S, \circ^F, \circ^W\}$ and any knowledge base K , it holds that $[K \circ K] = [K]$. It follows that $\mathcal{F} \subseteq Im_{\mathcal{F}}(\circ)$, for any fragment \mathcal{F} . However, as the following example illustrates, revision can produce results falling outside a given fragment.

Example 2. Consider a knowledge base $K = \{a \wedge b \wedge c\}$ with $[K] = \{abc\}$, and a formula $\mu = (\neg a \vee \neg b) \wedge (\neg a \vee \neg c) \wedge (\neg b \vee \neg c)$ with $[\mu] = \{\emptyset, a, b, c\}$. We have that K and μ are in both the 2CNF and the Horn fragment. However, for all $\circ \in \{\circ^D, \circ^S, \circ^F, \circ^W\}$ we get $[K \circ \mu] = \{a, b, c\}$. Since $Cl_{2CNF}(\{a, b, c\}) = Cl_{Horn}(\{a, b, c\}) = \{\emptyset, a, b, c\}$, we have that $K \circ \mu$ is neither 2CNF- nor Horn-expressible. Thus, $Horn \subset Im_{Horn}(\circ)$ and $2CNF \subset Im_{2CNF}(\circ)$.

Therefore the 2CNF and Horn fragments are not closed under any of the operators studied in this paper, though at this point we do not know anything more precise about what the image of these operators looks like. This motivates the following concept, which we study in the rest of the paper.

Definition 2. The deviation of \circ from \mathcal{F} is (i) total if $Im_{\mathcal{F}}(\circ) = \mathcal{L}$; (ii) partial if $\mathcal{F} \subset Im_{\mathcal{F}}(\circ) \subset \mathcal{L}$; and (iii) zero if $Im_{\mathcal{F}}(\circ) = \mathcal{F}$.

Deviation is closely related to the notion of *simplifiability* [12], which we briefly recall here. For a fragment \mathcal{F} and an operator \circ , a knowledge base $K^* \subseteq \mathcal{L}$ is \mathcal{F} -*simplifiable* w.r.t. \circ if there exists an \mathcal{F} -knowledge base K and $\mu \in \mathcal{F}$, such that $[K \circ \mu] = [K^*]$. It should be mentioned, however, that existing results on simplifiability apply only to Dalal’s operator. These results can be translated to our setting as follows: the deviation of \circ^D is (i) total for 2CNF, (ii) partial for Horn, and (iii) zero for 1CNF. Our results in Section 3 confirm these findings, though the reasoning proceeds along a different route.

3 Main results

Our results on deviation are summarized in Table 2, and the rest of the section is dedicated to justifying them. As mentioned in Section 2, $\mathcal{F} \subseteq Im_{\mathcal{F}}(\circ)$ for any fragment \mathcal{F} and all operators considered. In other words, it is non- \mathcal{F} -expressible knowledge bases

	1CNF	2CNF	Horn
\circ^S	zero	total	partial
\circ^D	zero	total	partial
\circ^F	zero	partial	partial
\circ^W	zero	partial	partial

Table 2. Precise results on deviation

K^* that are crucial in determining the extent of an operator’s deviation, and which will occupy our attention in the following. In general, $[K \circ \mu] \subseteq [\mu]$, for any K , μ and \circ . Thus, if K and μ are assumed to be in some fragment \mathcal{F} , we can only obtain K^* if $[\mu]$ contains at least $Cl_{\mathcal{F}}([K^*])$, and this is what we typically assume of μ . We present now results for the 1CNF, 2CNF and Horn fragments.

3.1 The 1CNF fragment

Let $[\mathcal{P}; \mathcal{I}] = \{\mathcal{P} \cup \mathcal{J} \mid \mathcal{J} \subseteq \mathcal{I}\}$, where $\mathcal{P} \cap \mathcal{I} = \emptyset$. Notice that if a knowledge base K is 1CNF-expressible, then $[K] = [\mathcal{P}; \mathcal{I}]$, where \mathcal{P} is the set of positive atoms in K and \mathcal{I} is the set of atoms that make no appearance in K (and on which K is, so to speak, indifferent). For example, if $U = \{a, b, c, d, e\}$ and $K = \{a \wedge b, \neg c\}$, then $[K] = [\{a, b\}; \{d, e\}]$ and, with our usual abbreviation, we may write $[K] = [ab; de]$. We now show that in revising a 1CNF knowledge base K by a 1CNF formula μ , the table of symmetric differences $[\mu] \Delta [K]$ always yields models of some 1CNF formula.

Lemma 1. *If K_1 and K_2 are 1CNF-expressible, then $[K_1] \diamond [K_2]$ is 1CNF-expressible.*

Proof. With the notation just introduced, let $[K_1] = [\mathcal{P}_1; \mathcal{I}_1]$ and $[K_2] = [\mathcal{P}_2; \mathcal{I}_2]$. We show, by double inclusion, that $[\mathcal{P}_1; \mathcal{I}_1] \diamond [\mathcal{P}_2; \mathcal{I}_2] = [(\mathcal{P}_1 \Delta \mathcal{P}_2) \setminus (\mathcal{I}_1 \cup \mathcal{I}_2); \mathcal{I}_1 \cup \mathcal{I}_2]$, which implies that $[K_1] \diamond [K_2]$ is 1CNF-expressible.

For one direction, take $(\mathcal{P}_1 \cup \mathcal{J}_1) \Delta (\mathcal{P}_2 \cup \mathcal{J}_2) \in [\mathcal{P}_1; \mathcal{I}_1] \diamond [\mathcal{P}_2; \mathcal{I}_2]$, with $\mathcal{J}_i \subseteq \mathcal{I}_i$, for $i \in \{1, 2\}$. We show that $(\mathcal{P}_1 \cup \mathcal{J}_1) \Delta (\mathcal{P}_2 \cup \mathcal{J}_2) = (\mathcal{P}_1 \Delta \mathcal{P}_2) \setminus (\mathcal{I}_1 \cup \mathcal{I}_2) \cup ((\mathcal{I}_1 \cup \mathcal{I}_2) \cap ((\mathcal{J}_1 \Delta \mathcal{J}_2) \setminus (\mathcal{P}_1 \Delta \mathcal{P}_2)))$.

“ \subseteq ” Take $w \in (\mathcal{P}_1 \cup \mathcal{J}_1) \Delta (\mathcal{P}_2 \cup \mathcal{J}_2)$ and without loss of generality assume that $w \in (\mathcal{P}_1 \cup \mathcal{J}_1) \setminus (\mathcal{P}_2 \cup \mathcal{J}_2)$. It follows that $w \in \mathcal{P}_1$ or $w \in \mathcal{J}_1$. *Case 1.* If $w \in \mathcal{P}_1$, then $w \in \mathcal{P}_1 \Delta \mathcal{P}_2$. Since $\mathcal{P}_1 \cap \mathcal{I}_1 = \emptyset$, it follows that $w \notin \mathcal{I}_1$ and $w \notin \mathcal{I}_1 \cup \mathcal{I}_2$. We conclude that $w \in (\mathcal{P}_1 \Delta \mathcal{P}_2) \setminus (\mathcal{I}_1 \cup \mathcal{I}_2)$. *Case 2.* If $w \in \mathcal{J}_1$, then $w \in \mathcal{J}_1 \Delta \mathcal{J}_2$. Second, $w \in \mathcal{I}_1$ and hence $w \in \mathcal{I}_1 \cup \mathcal{I}_2$ and $w \notin \mathcal{P}_1$. Since $w \notin \mathcal{P}_2$, we get that $w \notin \mathcal{P}_1 \Delta \mathcal{P}_2$ and thus $w \in (\mathcal{J}_1 \Delta \mathcal{J}_2) \setminus (\mathcal{P}_1 \Delta \mathcal{P}_2)$. We conclude that $w \in (\mathcal{I}_1 \cup \mathcal{I}_2) \cap ((\mathcal{J}_1 \Delta \mathcal{J}_2) \setminus (\mathcal{P}_1 \Delta \mathcal{P}_2))$.

“ \supseteq ” Take $w \in (\mathcal{P}_1 \Delta \mathcal{P}_2) \setminus (\mathcal{I}_1 \cup \mathcal{I}_2) \cup ((\mathcal{I}_1 \cup \mathcal{I}_2) \cap ((\mathcal{J}_1 \Delta \mathcal{J}_2) \setminus (\mathcal{P}_1 \Delta \mathcal{P}_2)))$. To get the conclusion, we perform a case analysis as to whether $w \in (\mathcal{P}_1 \Delta \mathcal{P}_2) \setminus (\mathcal{I}_1 \cup \mathcal{I}_2)$ or $w \in (\mathcal{I}_1 \cup \mathcal{I}_2) \cap ((\mathcal{J}_1 \Delta \mathcal{J}_2) \setminus (\mathcal{P}_1 \Delta \mathcal{P}_2))$.

For the other direction, take $(\mathcal{P}_1 \Delta \mathcal{P}_2) \setminus (\mathcal{I}_1 \cup \mathcal{I}_2) \cup ((\mathcal{I}_1 \cup \mathcal{I}_2) \cap \mathcal{J}) \in [(\mathcal{P}_1 \Delta \mathcal{P}_2) \setminus (\mathcal{I}_1 \cup \mathcal{I}_2); \mathcal{I}_1 \cup \mathcal{I}_2]$, where $\mathcal{J} \subseteq \mathcal{I}_1 \cup \mathcal{I}_2$. As above, we show by double inclusion that $(\mathcal{P}_1 \Delta \mathcal{P}_2) \setminus (\mathcal{I}_1 \cup \mathcal{I}_2) \cup ((\mathcal{I}_1 \cup \mathcal{I}_2) \cap \mathcal{J}) = (\mathcal{P}_1 \cup \mathcal{J}_1) \Delta (\mathcal{P}_2 \cup \mathcal{J}_2)$, where $\mathcal{J}_i = \mathcal{I}_i \cap (\mathcal{J} \Delta (\mathcal{P}_1 \Delta \mathcal{P}_2))$, for $i \in \{1, 2\}$.

If $[K_1] \diamond [K_2]$ is 1CNF-expressible, then it has a unique \subseteq -minimal element, which is also cardinality-minimal. We identify, then, $\min_{\subseteq}([K_1] \diamond [K_2]) = \min_{\text{card}}([K_1] \diamond [K_2])$ with this single element.

Example 3. If $[K_1] = \{a, ab, ac, abc\}$ and $[K_2] = \{\emptyset, b, c, bc\}$, then $[K_1] \diamond [K_2] = \{\emptyset, a, b, c, ab, ac, bc, abc\}$. We have that $\min_{\subseteq}([K_1] \diamond [K_2]) = \min_{\text{card}}([K_1] \diamond [K_2]) = \{\emptyset\}$, and write \emptyset instead of $\{\emptyset\}$.

In the following we use Lemma 1 to show that in the 1CNF fragment none of the operators can discriminate between $\{w_1, w_2\}$ and $\{w_1 \cap w_2, w_1 \cup w_2\}$, i.e., no operator can select one pair while leaving the other out.

Lemma 2. *If K and μ are 1CNF-expressible, then for any $\circ \in \{\circ^D, \circ^S, \circ^F, \circ^W\}$ and any $w_1, w_2 \in \mathcal{W}$, it holds that $\{w_1, w_2\} \subseteq [K \circ \mu]$ iff $\{w_1 \cap w_2, w_1 \cup w_2\} \subseteq [K \circ \mu]$.*

Proof. We show the result for Winslett's operator \circ^W .

(" \Rightarrow ") If $\{w_1, w_2\} \subseteq [K \circ^W \mu]$, take $u_i \in [K]$ such that $w_i \Delta u_i = \min_{\subseteq}([\mu] \diamond u_i)$, for $i \in \{1, 2\}$. Then consider $u_3, u_4 \in [K]$ such that $(w_1 \cap w_2) \Delta u_3 = \min_{\subseteq}((w_1 \cap w_2) \diamond [K])$ and $(w_1 \cup w_2) \Delta u_4 = \min_{\subseteq}((w_1 \cup w_2) \diamond [K])$. This is shown in Table 3, with arrows indicating whether an element is minimal on a row or on a column. We now argue that $(w_1 \cap w_2) \Delta u_3 = \min_{\subseteq}([\mu] \diamond u_3)$ and $(w_1 \cup w_2) \Delta u_4 = \min_{\subseteq}([\mu] \diamond u_4)$. For that, we first show that $(w_1 \cap w_2) \Delta u_3 \subseteq w \Delta u_3$, for some arbitrary $w \in [\mu]$. We take an atom $a \in (w_1 \cap w_2) \Delta u_3$ and do a case analysis.

Case 1. If $a \in (w_1 \cap w_2) \setminus u_3$, then by assumption $(w_1 \cap w_2) \Delta u_3 = \min_{\subseteq}((w_1 \cap w_2) \diamond [K])$, and thus $(w_1 \cap w_2) \Delta u_3 \subseteq (w_1 \cap w_2) \Delta u_1$. Since $a \in w_1 \cap w_2$, we get that $a \notin u_1$. With $a \in w_1$, this entails that $a \in w_1 \Delta u_1$. But $w_1 \Delta u_1 = \min_{\subseteq}([\mu] \diamond u_1)$, so $a \in w \Delta u_1$. With $a \notin u_1$, this implies that $a \in w$. Thus $a \in w \Delta (w_1 \cap w_2)$ and hence $(w_1 \cap w_2) \Delta u_3 \subseteq w \Delta u_3$. *Case 2.* If $a \in u_3 \setminus (w_1 \cap w_2)$, then $a \notin w_1$ or $a \notin w_2$. Without loss of generality, assume $a \notin w_1$. We again use the fact that $(w_1 \cap w_2) \Delta u_3 \subseteq (w_1 \cap w_2) \Delta u_1$ and conclude that $a \in u_1$, which further entails that $a \in w_1 \Delta u_1$. Then $a \in w \Delta u_1$, and thus $a \notin w$, which entails that $a \in w \Delta u_3$. With this we have shown that $(w_1 \cap w_2) \Delta u_3 = \min_{\subseteq}([\mu] \diamond u_3)$. The proof that $(w_1 \cup w_2) \Delta u_4 = \min_{\subseteq}([\mu] \diamond u_4)$ is similar.

(" \Leftarrow ") If $\{w_1 \cap w_2, w_1 \cup w_2\} \subseteq [K \circ^W \mu]$, there exist $u_3, u_4 \in [K]$ such that $(w_1 \cap w_2) \Delta u_3 = \min_{\subseteq}([\mu] \diamond u_3)$ and $(w_1 \cup w_2) \Delta u_4 = \min_{\subseteq}([\mu] \diamond u_4)$. Take $u_1, u_2 \in [K]$ such that $w_1 \Delta u_1 = \min_{\subseteq}(w_1 \diamond [K])$ and $w_2 \Delta u_2 = \min_{\subseteq}(w_2 \diamond [K])$. We then show that $w_i \Delta u_i = \min_{\subseteq}([\mu] \diamond u_i)$, for $i \in \{1, 2\}$. Indeed, take an arbitrary $w \in [\mu]$ and an $a \in w_1 \Delta u_1$ and consider the two possible cases. *Case 1.* If $a \in w_1 \setminus u_1$, then $a \in w_1 \Delta u_3$ and hence $a \notin u_3$. It follows that $a \in (w_1 \cup w_2) \Delta u_3$ and therefore $a \in w \Delta u_3$ and $a \in w$, which entails $a \in w \Delta u_1$. *Case 2.* If $a \in u_1 \setminus w_1$, then $a \in w_1 \Delta u_3$ and then $a \in (w_1 \cap w_2) \Delta u_3$ and hence $a \in w \Delta u_3$. This leads to $a \notin w$ and then to $a \in w \Delta u_1$. The proof for $w_2 \Delta u_2$ is entirely similar.

The reasoning for \circ^S, \circ^F and \circ^D is similar.

As we will show, what excludes non-1CNF-expressible knowledge bases from the range of the studied operators turns out to be the presence of so called *critical* interpretations [12], which we recapitulate here in an equivalent, though more concise formulation: $w_1, w_2 \in \mathcal{W}$ are *critical with respect to a knowledge base K^** if $w_1 \not\subseteq w_2$,




Table 3. The “ \Rightarrow ” direction for \circ^W in Lemma 2: arrows indicate area where shaded element is \subseteq -minimal.

$w_2 \not\subseteq w_1$, and $\{w_1, w_2\} \subseteq [K^*]$ or $\{w_1 \cap w_2, w_1 \cup w_2\} \subseteq [K^*]$, but $\{w_1, w_2, w_1 \cap w_2, w_1 \cup w_2\} \not\subseteq [K^*]$. The relevant result is that if a knowledge base K^* is non-1CNF-expressible, then there exist $w_1, w_2 \in Cl_{1CNF}([K^*])$ critical with respect to K^* [12].

Theorem 1. For any operator $\circ \in \{\circ^D, \circ^S, \circ^F, \circ^W\}$, the deviation of \circ from the 1CNF fragment is zero.

Proof. Let K^* be a knowledge base that is not 1CNF-expressible and assume existence of a 1CNF-expressible knowledge base K and a 1CNF formula μ such that $[K \circ \mu] = [K^*]$, and where $Cl_{1CNF}([K^*]) \subseteq [\mu]$. As mentioned, there are $w_1, w_2 \in \mathcal{W}$ critical with respect to K^* , thus $\{w_1, w_2\} \subseteq [K \circ \mu]$ or $\{w_1 \cap w_2, w_1 \cup w_2\} \subseteq [K \circ \mu]$. From Lemma 2 it follows that $\{w_1, w_2, w_1 \cap w_2, w_1 \cup w_2\} \subseteq [K \circ \mu]$, which is a contradiction.

An intuitive way to think about belief change in the 1CNF fragment is through what happens on the syntactic level, a view not often afforded for other fragments. For all of the operators considered, we have that if there are literals in K which are negated by μ , those literals are swapped with their negations in $K \circ \mu$. Thus, if $K = \{a, b, c\}$ and $\mu = \neg b \wedge \neg c \wedge \neg d$, then $K \circ \mu \equiv \{a \wedge \neg b \wedge \neg c \wedge \neg d\}$.

3.2 The 2CNF fragment

We obtain that Dalal and Satoh’s operators can produce any propositional knowledge base, but that this is not true for Forbus and Winslett.

Theorem 2. The deviation of \circ^D and \circ^S from the 2CNF fragment is total.

Proof. In [12] it is shown constructively that any propositional knowledge base is in $Im_{2CNF}(\circ^D)$. We show here that the same construction also works for \circ^S . Thus, take a non-2CNF-expressible knowledge base K^* such that $[K^*] = \{w_1, \dots, w_n\}$ and a formula μ such that $[\mu] = Cl_{2CNF}([K^*])$. Let $X = \{x_1, \dots, x_n\}$ be a set of n atoms that do not appear in K^* , and K be a knowledge base such that $[K] = Cl_{2CNF}(\{w_1 \cup \{x_2, \dots, x_n\}, \dots, w_n \cup \{x_1, \dots, x_{n-1}\}\})$. Then, for $i \in \{1, \dots, n\}$, we have that $\min_{\subseteq}([\mu] \circ [K]) = \{X \setminus \{x_i\} \mid i \in \{1, \dots, n\}\}$, where $X \setminus \{x_i\} =$

$w_i \triangle (w_i \cup (X \setminus \{x_i\}))$). To see why this holds, assume there exists $w_j \in [\mu]$ and $w_k \cup (X \setminus Y) \in [K]$ such that $w_j \triangle (w_k \cup (X \setminus Y)) \subset X \setminus \{x_i\}$, for some $i \in \{1, \dots, n\}$. Then it must be the case that $w_j \triangle w_k = \emptyset$, and $X \setminus Y \subset X \setminus \{x_i\}$. This last fact entails that $\{x_i\} \subset Y$, but such an element does not exist in $[K]$.

We illustrate this construction on a concrete example.

Example 4. Take a knowledge base K^* with $[K^*] = \{ab, ac, bc\}$. We have $Cl_{2\text{CNF}}([K^*]) = [K^*] \cup \{abc\}$, and to define K we introduce new variables x, y, z and take a 2CNF-expressible K such that $[K] = \{abyz, acxz, bcxy, abcxyz\}$. We have $[K \circ^S \mu] = [K \circ^D \mu] = [K^*]$ (see Table 4).

		[K]				
		\diamond	$abyz$	$acxz$	$bcxy$	$abcxyz$
{	[K*]	ab	yz	$bcxz$	$acxy$	$cxyz$
		ac	$bcyz$	xz	$abxy$	$bxyz$
		bc	$acyz$	$abxz$	xy	$axyz$
		\times	abc	cyz	bxz	axy

Table 4. Construction for the 2CNF fragment: \subseteq -minimal elements are highlighted in grey, elements that must not end up in the result are marked with \times .

Notice that in Example 4 we get $[K \circ^F \mu] = [K \circ^W \mu] = \{ab, ac, bc, abc\}$, so the construction in Theorem 2 does not work for \circ^F and \circ^W . Nonetheless, if we take a 2CNF-expressible knowledge base K' with $[K'] = \{\emptyset\}$ we get that $[K' \circ^F \mu] = [K' \circ^W \mu] = [K^*]$. This shows that the deviation of \circ^F and \circ^W from 2CNF is not zero. As we show now, their deviation is not total either.

Theorem 3. *The deviation of \circ^W and \circ^F from the 2CNF fragment is partial.*

Proof. We have already argued that the deviation of \circ^F and \circ^W from 2CNF is not zero. To show that it is not total, take a knowledge base K^* such that its set of models is $[K^*] = \{a, b, c, ab, ac, bc\}$, with $Cl_{2\text{CNF}}(K^*) = [K^*] \cup \{\emptyset, abc\}$, and assume, first, that there is a 2CNF-expressible knowledge base K and a 2CNF-expressible formula μ such that $Cl_{2\text{CNF}}(K^*) \subseteq [\mu]$ and $[K \circ^W \mu] = [K^*]$. Then there are $u_1, u_2, u_3 \in [K]$ such that $a \triangle u_1 = \min_{\subseteq}([\mu] \diamond u_1)$, $b \triangle u_2 = \min_{\subseteq}([\mu] \diamond u_3)$ and $c \triangle u_3 = \min_{\subseteq}([\mu] \diamond u_3)$. It must hold that $a \in u_1$, as otherwise $\emptyset \triangle u_1 \subset a \triangle u_1$. It must also hold that $b \notin u_1$, as otherwise $ab \triangle u_1 \subset a \triangle u_1$. In the same way it follows that $c \notin u_1$. We repeat this argument for u_2 and u_3 and obtain the configuration in Table 5, where we abbreviate $u'_1 \cup \{a\}$ as $a\langle u'_1 \rangle$, under the convention that if any of a, b or c does not make an appearance, that is because we know it cannot be there. Then neither of a, b or c is in $\text{maj}_3(u_1, u_2, u_3) = \langle u'_{123} \rangle$. It is straightforward to see, now, that we cannot have $[K \circ^W \mu] = [K^*]$. If there is some $w \in [\mu] \setminus Cl_{2\text{CNF}}(K^*)$ such that $w \triangle \langle u'_{123} \rangle \subset \emptyset \triangle \langle u'_{123} \rangle$, then none of the models of K^* gets selected in

		[K]					
		\diamond	$a\langle u'_1 \rangle$	$b\langle u'_2 \rangle$	$c\langle u'_3 \rangle$	$\langle u'_{123} \rangle$	
{	[K*]	a	$ab\langle u'_1 \rangle$	$ab\langle u'_2 \rangle$	$ac\langle u'_3 \rangle$	$a\langle u'_{123} \rangle$	
		b	$ab\langle u'_1 \rangle$	$\langle u'_2 \rangle$	$bc\langle u'_3 \rangle$	$b\langle u'_{123} \rangle$	
		c	$ac\langle u'_1 \rangle$	$bc\langle u'_2 \rangle$	$\langle u'_3 \rangle$	$c\langle u'_{123} \rangle$	
		ab	$b\langle u'_1 \rangle$	$a\langle u'_2 \rangle$	$abc\langle u'_3 \rangle$	$ab\langle u'_{123} \rangle$	
		ac	$c\langle u'_1 \rangle$	$abc\langle u'_2 \rangle$	$a\langle u'_3 \rangle$	$ac\langle u'_{123} \rangle$	
		bc	$abc\langle u'_1 \rangle$	$c\langle u'_2 \rangle$	$b\langle u'_3 \rangle$	$bc\langle u'_{123} \rangle$	
		\emptyset	$a\langle u'_1 \rangle$	$b\langle u'_2 \rangle$	$c\langle u'_3 \rangle$	$\langle u'_{123} \rangle$	
		\times	abc	$bc\langle u'_1 \rangle$	$ac\langle u'_2 \rangle$	$ab\langle u'_3 \rangle$	$abc\langle u'_{123} \rangle$
		\times	\dots	\dots	\dots	\dots	\dots
		\times	\dots	\dots	\dots	\dots	\dots

Table 5. $K^* \notin Im_{2CNF}(\circ^F)$ and $K^* \notin Im_{2CNF}(\circ^W)$

$[K \circ^W \mu]$. If this is not the case, then $\emptyset \triangle \langle u'_{123} \rangle \in \min_{\subseteq}([\mu] \diamond \langle u'_{123} \rangle)$, which entails that $\emptyset \in [K \circ^W \mu]$. This argument carries over to \circ^F by replacing \subseteq -minimality with cardinality-minimality.

3.3 The Horn fragment

Example 2 shows that the operators considered do not stay in the Horn fragment, so their deviation is not zero. Here we show that their deviation is partial, by finding knowledge bases which never show up as the result of revision in the Horn fragment, using the operators in question.

Theorem 4. *For any operator $\circ \in \{\circ^D, \circ^S, \circ^F, \circ^W\}$, the deviation of \circ from the Horn fragment is partial.*

Proof. Take a knowledge base K^* such that $[K^*] = \{ab, a, b\}$. We show that $K^* \notin Im_{Horn}(\circ^W)$. Assume there exists a Horn formula μ such that $Cl_{Horn}(K^*) \subseteq [\mu]$ and $[K \circ^W \mu] = [K^*]$. Then there are $u_1, u_2 \in [K^*]$ such that $a \triangle u_1 \in \min_{\subseteq}([\mu] \diamond u_1)$ and $b \triangle u_2 \in \min_{\subseteq}([\mu] \diamond u_2)$. We must have $a \in u_1$ (otherwise $\emptyset \triangle u_1 \subset a \triangle u_1$), and also $b \notin u_1$ (otherwise $ab \triangle u_1 \subset a \triangle u_1$). Similarly, we conclude that $b \in u_2$ and $a \notin u_2$. Since K is Horn-expressible, $u_1 \cap u_2 \in [K]$. But then neither a nor b are in $u_1 \cap u_2$. Table 6, with the notational convention used in the proof of Theorem 3, depicts this. It is now impossible to have $[K \circ^W \mu] = [K^*]$. Indeed, if there is some interpretation $w \in [\mu] \setminus Cl_{Horn}(K^*)$ such that $w \triangle (u_1 \cap u_2) \subset \emptyset \triangle (u_1 \cap u_2)$, then none of the models of K^* are in $[K \circ^W \mu]$. If no such w exists, then $\emptyset \in \min_{\subseteq}([\mu] \diamond (u_1 \cap u_2))$ and therefore $\emptyset \in [K \circ^W \mu]$. The argument carries over, with minor modifications, to the other operators.

4 Comparison over the Horn fragment

Operators whose deviations are total can produce any propositional knowledge base. Things become more complicated when deviation is partial, since a knowledge base

		[K]				
		◇	$a\langle u'_1 \rangle$	$b\langle u'_2 \rangle$	$\langle u'_1 \cap u'_2 \rangle$	
[μ]	{	[K*]	a	$\langle u'_1 \rangle$	$ab\langle u'_2 \rangle$	$a\langle u'_1 \cap u'_2 \rangle$
		b	$ab\langle u'_1 \rangle$	$\langle u'_2 \rangle$	$b\langle u'_1 \cap u'_2 \rangle$	$ab\langle u'_1 \cap u'_2 \rangle$
		ab	$b\langle u'_1 \rangle$	$a\langle u'_2 \rangle$	$ab\langle u'_1 \cap u'_2 \rangle$	$ab\langle u'_1 \cap u'_2 \rangle$
		×	∅	$a\langle u'_1 \rangle$	$b\langle u'_2 \rangle$	$\langle u'_1 \cap u'_2 \rangle$
×		

Table 6. $K^* \notin Im_{Horn}(\circ)$, for any of the operators

		[K]				
		◇	$abc\langle u'_1 \rangle$	$acd\langle u'_2 \rangle$	$ac\langle u'_1 \cap u'_2 \rangle$	
[μ]	{	[K*]	abcd	$d\langle u'_1 \rangle$	$b\langle u'_2 \rangle$	$bd\langle u'_1 \cap u'_2 \rangle$
		ab	$c\langle u'_1 \rangle$	$bcd\langle u'_2 \rangle$	$bc\langle u'_1 \cap u'_2 \rangle$	$bc\langle u'_1 \cap u'_2 \rangle$
		cd	$abd\langle u'_1 \rangle$	$a\langle u'_2 \rangle$	$ad\langle u'_1 \cap u'_2 \rangle$	$ad\langle u'_1 \cap u'_2 \rangle$
		×	∅	$abc\langle u'_1 \rangle$	$acd\langle u'_2 \rangle$	$ac\langle u'_1 \cap u'_2 \rangle$
×		

Table 7. $K^* \notin Im_{Horn}(\circ^F)$ and $K^* \notin Im_{Horn}(\circ^W)$

not obtainable with one operator may be obtainable with another. If this is the case, we would rightfully want to know it. In this section we present results on the differences between the images of our operators with respect to the Horn fragment.

Proposition 1. *If K^* is a knowledge base with $[K^*] = \{abcd, ab, cd\}$, then $K^* \notin Im_{Horn}(\circ^F)$ and $K^* \notin Im_{Horn}(\circ^W)$.*

Proof. Assume there exists a knowledge base K and formula μ , both Horn-expressible, such that $Cl_{Horn}([K^*]) \subseteq [\mu]$ and $[K \circ^W \mu] = [K^*]$. Then there exist $u_1, u_2 \in [K]$ such that $ab \Delta u_1 \in \min_{\subseteq}([\mu] \diamond u_1)$ and $bc \Delta u_2 \in \min_{\subseteq}([\mu] \diamond u_2)$. We infer that $\{a, b\} \subseteq u_1$, since otherwise it would hold that $\emptyset \Delta u_1 \in \min_{\subseteq}([\mu] \diamond u_1)$, which would imply that $\emptyset \in [K \circ^W \mu]$. We also infer that $\{c, d\} \not\subseteq u_1$, since otherwise it would hold that $abcd \Delta u_1 \subset ab \Delta u_1$, which contradicts the minimality of $ab \Delta u_1$. Thus, at most one of c and d can be in u_1 . Analogously, $\{c, d\} \subseteq u_2$ and at most one of a and b is in u_2 . A case analysis now reveals that this implies that $\emptyset \Delta (u_1 \cap u_2) \in \min_{\subseteq}([\mu] \diamond (u_1 \cap u_2))$. In Table 7 this is illustrated for the case when $c \in u_1$ and $a \in u_2$. If there is an element $w \in [\mu] \setminus Cl_{Horn}(K^*)$ such that $w \Delta (u_1 \cap u_2) \subset \emptyset \Delta (u_1 \cap u_2)$, then not all the models of K^* end up in $[K \circ^W \mu]$, which is a contradiction. If there is no such element, then $\emptyset \Delta (u_1 \cap u_2) \in \min_{\subseteq}([\mu] \diamond (u_1 \cap u_2))$ and thus $\emptyset \in [K \circ^W \mu]$, which is also a contradiction. The other cases are analogous. It follows that $K^* \notin Im_{Horn}(\circ^W)$. The same argument applies if we replace \subseteq -minimality with cardinality-minimality, so $K^* \notin Im_{Horn}(\circ^F)$.

This shows that there is a knowledge base which cannot be obtained with either Forbus' operator \circ^F or Winslett's operator \circ^W , though as we show in the proof of Theorem 5, it can be obtained with Dalal's operator \circ^D and Satoh's operator \circ^S .

Proposition 2. *If K^* is a knowledge base with $[K^*] = \{abcde, ab, cd, e\}$, then $K^* \notin \text{Im}_{\text{Horn}}(\circ^D)$.*

Proof. If there are Horn-expressible K and μ such that $\text{Cl}_{\text{Horn}}([K^*]) \subseteq [\mu]$ and $[K \circ^D \mu] = \{abcde, ab, cd, e\}$, then there must be $u_1, u_2, u_3, u_4 \in [K]$ such that $\{abcd \triangle u_1, ab \triangle u_2, cd \triangle u_3, e \triangle u_4\} \subseteq \min_{\text{card}}([\mu] \diamond [K])$. We conclude that at least three of $\{a, b, c, d, e\}$ must be in u_1 ; that $\{a, b\} \subseteq u_2$ and at most one of $\{c, d, e\}$ can be in u_2 ; that $\{c, d\} \subseteq u_3$ and at most one of $\{a, b, e\}$ can be in u_3 ; and $e \in u_4$. A case analysis of $\text{Cl}_{\text{Horn}}(K)$ leads to a contradiction.

The knowledge base K^* from Proposition 2 can be obtained with \circ^D , as we show below. But we need one more observation to derive our results on the relations between operators.

Proposition 3. *For any knowledge base K^* such that $\{a, b, ab\} \subseteq [K^*]$ and $\emptyset \notin [K^*]$, it holds that $K^* \notin \text{Im}_{\text{Horn}}(\circ^D) \cup \text{Im}_{\text{Horn}}(\circ^S)$.*

Proof. It is easy to see that the proof of Theorem 4 works for \circ^D and \circ^S and any knowledge base K^* such that $\{a, b, ab\} \subseteq [K^*]$ and $\emptyset \notin [K^*]$. However, as we show in the proof of Theorem 5, it does not carry over to \circ^F and \circ^W .

The following result gathers these results into a unified image of the differences between the operators over the Horn fragment.

Theorem 5. *For the operators we study, the following holds:*

- $\text{Im}_{\text{Horn}}(\circ^D) \not\subseteq \text{Im}_{\text{Horn}}(\circ^W) \cup \text{Im}_{\text{Horn}}(\circ^F)$,
- $\text{Im}_{\text{Horn}}(\circ^S) \not\subseteq \text{Im}_{\text{Horn}}(\circ^D) \cup \text{Im}_{\text{Horn}}(\circ^W) \cup \text{Im}_{\text{Horn}}(\circ^F)$,
- $\text{Im}_{\text{Horn}}(\circ^F) \not\subseteq \text{Im}_{\text{Horn}}(\circ^D) \cup \text{Im}_{\text{Horn}}(\circ^S)$,
- $\text{Im}_{\text{Horn}}(\circ^W) \not\subseteq \text{Im}_{\text{Horn}}(\circ^S) \cup \text{Im}_{\text{Horn}}(\circ^D)$.

Proof. From Proposition 1, we know that a knowledge base K_1^* with $[K_1^*] = \{abcd, ab, cd\}$ is in neither $\text{Im}_{\text{Horn}}(\circ^F)$ nor in $\text{Im}_{\text{Horn}}(\circ^W)$. However, a knowledge base K_1 with $[K_1] = \{abc, bcd, bc\}$ and a formula μ_1 with $[\mu_1] = \{abcd, ab, cd, \emptyset\}$ are both Horn-expressible and $[K_1 \circ^D \mu_1] = [K_1 \circ^S \mu_1] = [K_1^*]$.¹ Thus, $K_1^* \in \text{Im}_{\text{Horn}}(\circ^D) \cap \text{Im}_{\text{Horn}}(\circ^S)$. From Proposition 2 a knowledge base K_2^* with $[K_2^*] = \{abcde, ab, cd, e\}$ is not in $\text{Im}_{\text{Horn}}(\circ^S)$. However, if we take $[K_2] = \{ace\}$ and $[\mu_2] = \{abcde, ab, cd, \emptyset\}$, then $[K_2 \circ^D \mu_2] = [K_2^*]$ and thus $K_2^* \in \text{Im}_{\text{Horn}}(\circ^D)$. From Proposition 3 a knowledge base K_3^* with $[K_3^*] = \{ab, a, b, d\}$ is neither in $\text{Im}_{\text{Horn}}(\circ^D)$ nor in $\text{Im}_{\text{Horn}}(\circ^S)$. However, take $[K_3] = \{abcd, abcd, bcd, cd\}$ and $[\mu_3] = \{ab, a, b, d, \emptyset\}$. Then $[K_3 \circ^F \mu_3] = [K_3 \circ^W \mu_3] = [K_3^*]$ and thus $K_3^* \in \text{Im}_{\text{Horn}}(\circ^F) \cap \text{Im}_{\text{Horn}}(\circ^W)$.

The exact relationship between $\text{Im}_{\text{Horn}}(\circ^F)$ and $\text{Im}_{\text{Horn}}(\circ^W)$ is still open, as is the question whether $\text{Im}_{\text{Horn}}(\circ^D) \subset \text{Im}_{\text{Horn}}(\circ^S)$.

¹ This constitutes a counterexample to Proposition 11 in [12].

5 Conclusion and future work

We have presented results on the deviation of four major belief change operators from well-known fragments of propositional logic. These results are summarised in Table 2 of Section 3. We have found that for the 1CNF fragment, deviation is uniformly zero. The 1CNF fragment is ‘safe’, in that revision on 1CNF knowledge bases always produces 1CNF-expressible results. Future work would look for general conditions that make a fragment safe in this sense. For the 2CNF and Horn fragments deviation has been found to range from partial to total. If it is important that the result be expressible in the target language, some kind of repair (or refinement [2]) is required. Our results thus complement existing work on refinements, as they reveal when, and to which extent, such repairs are needed. On the other hand, if an application allows the result to be expressed in a richer language, then we would be interested in knowing whether results that cannot be obtained with one operator could be obtained with another. This applies to cases when deviation is partial, and is especially complicated for the Horn fragment. In Section 4 we have clarified some of the relationships between images with respect to the Horn fragment, but the question of whether there is any knowledge base that can be obtained with Dalal but not with Satoh, and the relationship between the images of Forbus and Winslett, on the other, are still open and left for future work. Also left for future work is the relationship between the images of Forbus and Winslett in the 2CNF fragment. To get a clearer idea of the images of these operators with respect to the more complicated fragments such as Horn, we would also consider the complexity of deciding whether a given knowledge base is in the image of an operator.

References

- [1] Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *J. Symb. Log.* 50(2), 510–530 (1985)
- [2] Creignou, N., Papini, O., Pichler, R., Woltran, S.: Belief revision within fragments of propositional logic. *J. Comput. Syst. Sci.* 80(2), 427–449 (2014)
- [3] Creignou, N., Papini, O., Rümmele, S., Woltran, S.: Belief Merging within Fragments of Propositional Logic. *ACM Trans. Comput. Log.* 17(3), 20:1–20:28 (2016)
- [4] Creignou, N., Pichler, R., Woltran, S.: Do Hard SAT-Related Reasoning Tasks Become Easier in the Krom Fragment? In: *Proc. of IJCAI 2013*. pp. 824–831 (2013)
- [5] Dalal, M.: Investigations into a Theory of Knowledge Base Revision. In: *Proc. of IJCAI 1988*. pp. 475–479 (1988)
- [6] Delgrande, J.P., Peppas, P.: Belief revision in Horn theories. *Artif. Intell.* 218, 1–22 (2015)
- [7] Delgrande, J.P., Schaub, T., Tompits, H., Woltran, S.: A Model-Theoretic Approach to Belief Change in Answer Set Programming. *ACM Trans. Comput. Log.* 14(2), 14:1–14:46 (2013)
- [8] Diller, M., Haret, A., Linsbichler, T., Rümmele, S., Woltran, S.: An Extension-Based Approach to Belief Revision in Abstract Argumentation. In: *Proc. of IJCAI 2015*. pp. 2926–2932 (2015)
- [9] Eiter, T., Gottlob, G.: On the Complexity of Propositional Knowledge Base Revision, Updates, and Counterfactuals. *Artif. Intell.* 57(2-3), 227–270 (1992)
- [10] Forbus, K.D.: Introducing Actions into Qualitative Simulation. In: *Proc. of IJCAI 1989*. pp. 1273–1278 (1989)

- [11] Gärdenfors, P.: Knowledge in Flux: Modelling the Dynamics of Epistemic States. The MIT Press, Cambridge, MA (1988)
- [12] Haret, A., Maily, J., Woltran, S.: Distributing Knowledge into Simple Bases. In: Proc. of IJCAI 2016. pp. 1109–1115 (2016)
- [13] Haret, A., Rümmele, S., Woltran, S.: Merging in the Horn Fragment. In: Proc. of IJCAI 15. pp. 3041–3047 (2015)
- [14] Katsuno, H., Mendelzon, A.O.: On the Difference between Updating a Knowledge Base and Revising It. In: Proc. of KR 1991). pp. 387–394 (1991)
- [15] Katsuno, H., Mendelzon, A.O.: Propositional Knowledge Base Revision and Minimal Change. *Artif. Intell.* 52(3), 263–294 (1992)
- [16] Liberatore, P., Schaerf, M.: Belief Revision and Update: Complexity of Model Checking. *J. Comput. Syst. Sci.* 62(1), 43–72 (2001)
- [17] Satoh, K.: Nonmonotonic Reasoning by Minimal Belief Revision. In: FGCS. pp. 455–462 (1988)
- [18] Winslett, M.: Updating Logical Databases. Cambridge University Press (1990)
- [19] Zhuang, Z.Q., Pagnucco, M., Zhang, Y.: Definability of Horn Revision from Horn Contraction. In: Rossi, F. (ed.) Proc. of IJCAI 2013). pp. 1205–1211. IJCAI/AAAI (2013)
- [20] Zhuang, Z., Wang, Z., Wang, K., Qi, G.: DL-Lite Contraction and Revision. *J. Artif. Intell. Res. (JAIR)* 56, 329–378 (2016)