

Belief Revision within Fragments of Propositional Logic*

Nadia Creignou and Odile Papini
 LIF UMR CNRS 7279, LSIS UMR CNRS 7296
 Université d'Aix-Marseille
 Campus de Luminy, 163 avenue de Luminy.
 13288 Marseille Cedex 9. FRANCE

Reinhard Pichler and Stefan Woltran
 Institute of Information Systems
 Vienna University of Technology
 Favoritenstraße 9–11
 A-1040 Vienna, Austria

Abstract

Belief revision has been extensively studied in the framework of propositional logic, but just recently revision within fragments of propositional logic has gained attention. Hereby it is not only the belief set and the revision formula which are given within a certain language fragment, but also the result of the revision has to be located in the same fragment. So far, research in this direction has been mainly devoted to the Horn fragment of classical logic. In this work, we present a general approach to define new revision operators derived from known operators (as for instance, Satoh's and Dalal's revision operators), such that the result of the revision remains in the fragment under consideration. Our approach is not limited to the Horn case but applicable to any fragment of propositional logic where the models of the formulas are closed under a Boolean function. Thus we are able to uniformly treat cases as dual-Horn, Krom and affine formulas, as well.

Introduction

Belief revision is a central topic in knowledge representation and reasoning. Belief revision consists in incorporating a new belief, changing as few as possible of the original beliefs while preserving consistency. Within the symbolic frameworks, where the beliefs are represented by logical formulas, the AGM paradigm (Alchourrón, Gärdenfors, and Makinson 1985) dedicated to the revision of theories, became a standard which provides rational postulates any reasonable revision operator should satisfy. Katsuno and Mendelzon (1991), when unifying semantic revision approaches, reformulated these postulates where a theory is represented by a propositional formula. Moreover they proposed a representation theorem that characterizes revision operations in terms of total pre-orders over interpretations.

Belief revision has been extensively studied within the framework of propositional logic and numerous concrete belief revision operators have been proposed according to either semantic or syntactic points of view, for example (Dalal 1988; Satoh 1988; Nebel 1998). Moreover, complexity results have been obtained (Eiter and Gottlob 1992;

Nebel 1998; Liberatore and Schaerf 2001). However, as far as we know, only few works have focused on belief revision within the framework of *fragments of propositional logic*, except for the Horn case (Delgrande and Peppas 2011; Wu, Zhang, and Zhang 2011; Van De Putte 2011).

The study of belief change within language fragments is motivated by two central observations:

- In many applications, the language is restricted a priori. For instance, a rule-based formalization of expert knowledge is much easier to handle for standard users. In case users want to revise some rules, they indeed expect that the outcome is still in the easy-to-read format they are used to.
- Many fragments of propositional logic allow for efficient reasoning methods. Suppose an agent who frequently has to answer queries about his beliefs. This should be done efficiently thus the beliefs are stored as a formula known to be in a tractable class. In case the beliefs of the agent are undergoing a revision, it is desired that the result of such an operation yields a formula in the same fragment. Hence, the agent still can use the dedicated solving method he is equipped with for this fragment. In case such changes are performed rarely, we do not bother whether the revision itself can be performed efficiently, but it is more important that the outcome can still be evaluated efficiently.

It seems thus natural to investigate how known operators can be refined such that they work properly within a language fragment. The main obstacle hereby is that for a language fragment \mathcal{L}' , given formulas¹ $\psi, \mu \in \mathcal{L}'$ there is no guarantee that the outcome of a revision, $\psi \circ \mu$, remains in \mathcal{L}' as well. Let, for example, $\psi = a \wedge b$ and $\mu = \neg a \vee \neg b$, be formulas expressed in conjunctive normal form (CNF) with Horn clauses (at most one positive literal), revising ψ by μ using Dalal's revision operator (Dalal 1988) does not remain in the Horn language fragment since $(a \vee b) \wedge (\neg a \vee \neg b)$ belongs to the result of the revision. The natural question arises whether there exist refinements \star of \circ such that $\psi \star \mu \in \mathcal{L}'$ always holds, but properties of \circ are retained whenever possible. For instance, for such a refined operator it seems reasonable that $\psi \star \mu$ is equivalent to $\psi \circ \mu$ whenever $\psi \circ \mu$

*This work was supported by the Austrian Science Fund (FWF) under grant P20704-N18 and ÖAD grant Amadée 17/2011/EGIDE PHC Amadeus project 24908NM.

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Here and throughout the paper, we will follow the Katsuno and Mendelzon's view of revision, cf. (Katsuno and Mendelzon 1991).

already yields a result from the desired fragment \mathcal{L}' . We introduce further natural criteria refined operators are expected to satisfy and we show general properties of these refined operators as well as their limits in satisfying postulates.

In fact, our **main contributions** are the following:

- We propose to adapt known belief revision operators to make them applicable in fragments of propositional logic. We provide natural criteria such operators should satisfy.
- Rather than restricting ourselves to the Horn fragment, we present a general framework which includes all fragments captured via closure properties on sets of models. In particular, (dual) Horn, Krom and affine formulas are thus covered.
- We characterize refined operators in a constructive way which allows us to study their properties in terms of the postulates by Katsuno and Mendelzon (1991). Most notably, we show that in case the initial operator satisfies certain postulates, then so does any of its refinements.
- We give a preliminary complexity analysis of selected refined operators.

Previous works dedicated to belief revision within fragments of propositional logic only focused on Horn fragments. The first mention of the Horn case for belief revision appears in (Eiter and Gottlob 1992), an analysis of belief revision complexity. In (Liberatore 2000) a compact representation for revision in the Horn case is proposed. In (Langlois et al. 2008) the study of belief revision in the Horn case provides a characterization of the existence of a complement of the Horn consequence which corresponds to a contraction operator. Horn contraction has been addressed in (Delgrande 2008; Booth, Meyer, and Varzinczak 2009; Delgrande and Wassermann 2010; Zhuang and Pagnucco 2011) however the results cannot help for defining revision operators since applying the Levi identity² produces a result which might not fit into the fragment of consideration. More recently, Delgrande and Peppas (2011) showed that classical AGM revision does not immediately generalize to the Horn case. They overcame this difficulty by restricting the rankings on interpretations, adding a closure under intersection condition on interpretations. Moreover, they added a new postulate to the set of AGM postulates. However they did not exhibit any concrete revision operator and they limited themselves to the Horn case.

Preliminaries

Propositional Logic. We consider \mathcal{L} as the language of propositional logic over some fixed alphabet \mathcal{U} of propositional atoms. We use standard connectives $\rightarrow, \oplus, \vee, \wedge, \neg$, and constants \top, \perp . A clause is a disjunction of literals. A clause is called (i) *Horn* if at most one of its literals is positive; (ii) *dual Horn* if at most one of its literals is negative; (iii) *Krom* if it consists of at most two literals. A \oplus -clause is defined like a clause but using exclusive- instead of standard-disjunction. We identify the following subsets

² $\psi \circ \mu = (\psi - \neg\mu) + \mu$, where $-$, resp. $+$, denotes the contraction, resp. the expansion, operator.

of \mathcal{L} : \mathcal{L}_{Horn} as the set of all formulas in \mathcal{L} being conjunctions of Horn clauses; \mathcal{L}_{DHorn} as the set of all formulas in \mathcal{L} being conjunctions of dual Horn clauses; \mathcal{L}_{Krom} as the set of all formulas in \mathcal{L} being conjunctions of Krom clauses; and \mathcal{L}_{Affine} as the set of all formulas in \mathcal{L} being conjunctions of \oplus -clauses. In what follows we sometimes just talk about arbitrary fragments $\mathcal{L}' \subseteq \mathcal{L}$. Hereby, we tacitly assume that any such fragment $\mathcal{L}' \subseteq \mathcal{L}$ contains at least the formula \top .

For any formula ϕ , let $\text{Var}(\phi)$ denote the set of variables occurring in ϕ . An interpretation is represented either by a set $I \subseteq \mathcal{U}$ of atoms (corresponding to the variables set to true) or by its corresponding characteristic bit-vector of length $|\mathcal{U}|$. For instance if we consider $\mathcal{U} = \{x_1, \dots, x_6\}$, the interpretation $x_1 = x_3 = x_6 = 1$ and $x_2 = x_4 = x_5 = 0$ will be represented either by $\{x_1, x_3, x_6\}$ or by $(1, 0, 1, 0, 0, 1)$. As usual, if an interpretation I satisfies a formula ϕ , we call I a model of ϕ . By $\text{Mod}(\phi)$ we denote the set of all models (over \mathcal{U}) of ϕ . Moreover, $\psi \models \phi$ if $\text{Mod}(\psi) \subseteq \text{Mod}(\phi)$ and $\psi \equiv \phi$ if $\text{Mod}(\psi) = \text{Mod}(\phi)$. For a set T of formulas, $Cn(T)$ denotes the closure of T under the consequence relation \models . A theory T is a deductively closed set of formulas such that $T = Cn(T)$. For fragments $\mathcal{L}' \subseteq \mathcal{L}$, we also use $T_{\mathcal{L}'}(\psi) = \{\phi \in \mathcal{L}' \mid \psi \models \phi\}$.

Revision. In the AGM paradigm (Alchourrón, Gärdenfors, and Makinson 1985), the underlying logic is assumed to be classical logic and the beliefs are modeled by a theory, called belief set. A revision operator $*$ is a function mapping a belief set T and a formula A to a new belief set $T * A$ which satisfies the following properties³:

- (K * 1) $T * A = Cn(T * A)$
- (K * 2) $A \in T * A$.
- (K * 3) $T * A \subseteq T + A$.
- (K * 4) If $\neg A \notin T$ then $T * A = T + A$.
- (K * 5) $T * A = \mathcal{L}$ only if A is unsatisfiable.
- (K * 6) If $A \equiv B$ then $T * A = T * B$.
- (K * 7) $T * (A \wedge B) \subseteq (T * A) + B$.
- (K * 8) If $\neg B \notin T * A$ then $(T * A) + B = T * (A \wedge B)$.

According to a semantic point of view, when a belief set is represented by a propositional formula ψ such that $T = \{\phi \in \mathcal{L} \mid \psi \models \phi\}$, revising ψ by μ amounts to finding the models of μ which are “closest” to the models of ψ . The closeness between models depends on the choice of the revision operator. In order to characterize different proposed semantic operators, Katsuno and Mendelzon (1991) reformulated the AGM postulates as follows:

- (R1) $\psi \circ \mu \models \mu$.
- (R2) If $\psi \wedge \mu$ is satisfiable, then $\psi \circ \mu \equiv \psi \wedge \mu$.
- (R3) If μ is satisfiable, then so is $\psi \circ \mu$.
- (R4) If $\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$, then $\psi_1 \circ \mu_1 \equiv \psi_2 \circ \mu_2$.
- (R5) $(\psi \circ \mu) \wedge \phi \models \psi \circ (\mu \wedge \phi)$.
- (R6) If $(\psi \circ \mu) \wedge \phi$ is satisfiable, then also $\psi \circ (\mu \wedge \phi) \models (\psi \circ \mu) \wedge \phi$.

³ $T + A$ is the smallest deductively closed set of formulas containing both T and A .

The (R1) postulate specifies that the added formula belongs to the revised belief set, (R2) gives the revised belief set when the added formula is consistent with the initial belief set, (R3) ensures that no inconsistency is introduced in the revised belief set, (R4) expresses the principle of irrelevance of the syntax, and (R5) and (R6) are the direct translation of both the ($K * 7$) and ($K * 8$) postulates and are the most controversial ones, as mentioned in (Dubois 2011).

Katsuno and Mendelzon showed that a revision satisfying the AGM postulates is equivalent to a total preorder on interpretations, which reflects a plausibility ordering on interpretations. More formally they provided the following representation theorem, stating that a revision operation satisfies the postulates (R1)–(R6) if and only if there exists a total pre-order \leq_ψ such that $\text{Mod}(\psi \circ \mu) = \text{Min}(\text{Mod}(\mu), \leq_\psi)$.

We now recall some well-known semantic revision operators for \mathcal{L} , the full version of propositional logic. Later we shall refine them towards revision operators for some fragments \mathcal{L}' . In the model-based revision operators recalled hereafter, the closeness between models relies on the symmetric difference between models, that is the set of propositional variables on which they differ.

Dalal (1988) measures minimal change by the cardinality of model change, i.e., let α and β be two propositional formulas and M and M' be two interpretations, $M\Delta M'$ denotes the symmetric difference between M and M' and $|\Delta|^{min}(\alpha, \beta)$ denotes the minimum number of propositional variables on which the models of α and β differ⁴ and is defined as $\min\{|M\Delta M'| : M \in \text{Mod}(\alpha), M' \in \text{Mod}(\beta)\}$. Dalal's operator is now defined as: $\text{Mod}(\psi \circ_D \mu) = \{M \in \text{Mod}(\mu) : \exists M' \in \text{Mod}(\psi) \text{ s.t. } |M\Delta M'| = |\Delta|^{min}(\psi, \mu)\}$. This operator satisfies (R1)–(R6).

Satoh (1988) interprets the minimal change in terms of set inclusion instead of cardinality on model difference. Thus let $\Delta^{min}(\alpha, \beta) = \min_{\subseteq} \{M\Delta M' : M \in \text{Mod}(\alpha), M' \in \text{Mod}(\beta)\}$ and define Satoh's operator as: $\text{Mod}(\psi \circ_S \mu) = \{M \in \text{Mod}(\mu) : \exists M' \in \text{Mod}(\psi) \text{ s.t. } M\Delta M' \in \Delta^{min}(\psi, \mu)\}$. Satoh's operator satisfies (R1)–(R5).

Another less known revision operation is due to Hegner. While Dalal's and Satoh's approaches deal with propositional variables possibly present in the models of ψ and μ , Hegner's operator focuses on variables occurring in μ and is defined as $\text{Mod}(\psi \circ_H \mu) = \{M \in \text{Mod}(\mu) : \exists M' \in \text{Mod}(\psi) \text{ s.t. } M\Delta M' \subseteq \text{Var}(\mu)\}$.

We are interested here in revision operators which are tailored for certain fragments. The following definition is very general. We shall later consider revision operators which satisfy several criteria and postulates.

Definition 1. A basic (revision) operator for $\mathcal{L}' \subseteq \mathcal{L}$ is any function $\circ : \mathcal{L}' \times \mathcal{L}' \rightarrow \mathcal{L}'$ satisfying $\top \circ \mu \equiv \mu$ for each $\mu \in \mathcal{L}'$. We say that \circ satisfies a KM postulate (Ri) ($i \in \{1, \dots, 6\}$) in \mathcal{L}' if the respective postulate holds when restricted to formulas from \mathcal{L}' .

⁴This is also expressed with Hamming distance when the interpretations are encoded as characteristic bit-vectors.

Refined Operators

The problem of known operators when applied in a fragment of propositional logic is illustrated by an example.

Example 1. Let $\psi, \mu \in \mathcal{L}_{Horn}$ (over $\mathcal{U} = \{a, b\}$) with $\psi = a \wedge b$ and $\mu = \neg a \vee \neg b$. We have $\text{Mod}(\psi) = \{\{a, b\}\}$ and $\text{Mod}(\mu) = \{\emptyset, \{a\}, \{b\}\}$. We obtain $\text{Mod}(\psi \circ_D \mu) = \text{Mod}(\psi \circ_S \mu) = \{\{a\}, \{b\}\}$. Thus, for instance, we can give $\phi = (a \vee b) \wedge (\neg a \vee \neg b)$ as a result of the revision. However, $\phi \notin \mathcal{L}_{Horn}$. In fact, there is no $\phi \in \mathcal{L}_{Horn}$ with $\text{Mod}(\phi) = \{\{a\}, \{b\}\}$, since each $\phi \in \mathcal{L}_{Horn}$ satisfies the following closure-property in terms of its models: for each $I, J \in \text{Mod}(\phi)$, also $I \cap J \in \text{Mod}(\phi)$.

In Example 1, to adapt \circ_D (or likewise, \circ_S) such that the outcome of the revision is from \mathcal{L}_{Horn} we have two options: (1) build the closure of the set of required models, in our case we have to add $\emptyset = \{a\} \cap \{b\}$; (2) remove models from the outcome. The disadvantage of the latter option is that there is no particular reason to prefer $\{a\}$ over $\{b\}$ or vice versa. However, removing both would yield the empty set and thus our revision would become inconsistent which is not desirable. The former approach looks also problematic since adding models reduces the number of formulas derivable from the revised formula, which might be in conflict with some KM postulates. In fact, one of the main goals of the paper is to understand the limits of such repairs in terms of the KM postulates. Note that in Example 1, $\psi, \mu \in \mathcal{L}_{Krom}$ holds, and the revision ϕ is also in \mathcal{L}_{Krom} .

The considerations of the above example can be generalized to the following problem statement. Given a known revision operator \circ and a fragment \mathcal{L}' of propositional logic, how can we adapt \circ to a new revision operator \star such that, for each $\psi, \mu \in \mathcal{L}'$, also $\psi \star \mu \in \mathcal{L}'$? Let us note here that for rule-based fragments like Horn, μ can be either a fact or a rule. This allows for two kinds of revision: either an agent obeys rules, discovers facts and revises the facts accordingly, or the belief base consists of facts and rules and the rules are revised by new rules.

Next, let us define a few natural desiderata for \star .

Definition 2. Let \mathcal{L}' be a fragment of classical logic and $\circ : \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{L}$ a revision operator. We call an operator $\star : \mathcal{L}' \times \mathcal{L}' \rightarrow \mathcal{L}'$ a \circ -refinement for \mathcal{L}' if it satisfies the following properties, for each $\psi, \psi', \mu, \mu' \in \mathcal{L}'$.

- *consistency:* $\psi \star \mu$ is satisfiable if and only if $\psi \circ \mu$ is satisfiable
- *equivalence:* if $\psi \circ \mu \equiv \psi' \circ \mu'$ then $\psi \star \mu \equiv \psi' \star \mu'$
- *containment:* $T_{\mathcal{L}'}(\psi \circ \mu) \subseteq T_{\mathcal{L}'}(\psi \star \mu)$
- *invariance:* If $\psi \circ \mu \in \mathcal{L}'$, then $T_{\mathcal{L}'}(\psi \star \mu) \subseteq T_{\mathcal{L}'}(\psi \circ \mu)$.

Containment and invariance jointly imply that for each $\psi, \mu \in \mathcal{L}'$ such that $\psi \circ \mu \in \mathcal{L}'$, $\psi \star \mu \equiv \psi \circ \mu$ holds.

Let us briefly discuss these properties. The first two conditions are rather independent from \mathcal{L}' , but relate the refined operator \star to the original revision \circ in certain ways. To be more precise, *consistency* states that the refined operator \star should yield a consistent revision exactly if the original operator \circ does so. *Equivalence* means that the definition of the \star -operator should not be syntax-dependent: revisions which

are equivalent w.r.t \circ are also equivalent w.r.t. \star . Note that this does not necessarily mean that $\psi \star \mu \equiv \psi \circ \mu$ holds for all formulas $\mu, \psi \in \mathcal{L}'$. The final two properties take more care of the fragment \mathcal{L}' . *Containment* ensures that \star can be seen as a form of approximation of \circ when applied in the \mathcal{L}' fragment, while *invariance* states that in case \circ behaves as expected (i.e. the revision is contained in \mathcal{L}') there is no need for \star to do something additional.⁵

Characterization of Refined Operators

In order to capture all \circ -refinements for a fragment \mathcal{L}' we need some formal machinery which we introduce next.

Formal Ingredients. We use k -ary Boolean functions $\beta: \{0, 1\}^k \rightarrow \{0, 1\}$ like

- the binary AND function denoted by \wedge ;
- the binary OR function denoted by \vee ;
- the ternary MAJORITY function, $\text{maj}_3(x, y, z) = 1$ if at least two of the variables x, y , and z are set to 1;
- the ternary XOR function $\oplus_3(x, y, z) = x \oplus y \oplus z$.

All of them satisfy the properties of *symmetry*, i.e., for all permutations σ , $\beta(x_1, \dots, x_k) = \beta(x_{\sigma(1)}, \dots, x_{\sigma(k)})$, and 0- and 1-*reproduction*, i.e., for every $x \in \{0, 1\}$, $\beta(x, \dots, x) = x$.

Recall that we consider interpretations also as bit-vectors. We thus extend Boolean functions to interpretations by applying coordinate-wise the original function. So, if $M_1, \dots, M_k \in \{0, 1\}^n$, then $\beta(M_1, \dots, M_k)$ is defined by $(\beta(M_1[1], \dots, M_k[1]), \dots, \beta(M_1[n], \dots, M_k[n]))$, where $M[i]$ is the i -th coordinate of the interpretation M .

Definition 3. Let \mathcal{B} denote the set of all Boolean functions over alphabet \mathcal{U} applied to interpretations over \mathcal{U} that satisfy symmetry as well as 0- and 1-reproduction.

Coming back to Example 1, recall that we mentioned that models of Horn formulas are closed under intersection. In terms of Boolean functions, this means that for any models I, J of a Horn-formula ϕ , also $I \wedge J$ is a model of ϕ . The next definition gives a general formal definition of closure.

Definition 4. Given a set $\mathcal{M} \subseteq 2^{\mathcal{U}}$ of interpretations and $\beta \in \mathcal{B}$, we define $Cl_\beta(\mathcal{M})$, the closure of \mathcal{M} under β , as the smallest set of interpretations that contains \mathcal{M} and that is closed under β , i.e., if $M_1, \dots, M_k \in Cl_\beta(\mathcal{M})$, then also $\beta(M_1, \dots, M_k) \in Cl_\beta(\mathcal{M})$.

Closures satisfy monotonicity: if $\mathcal{M} \subseteq \mathcal{N}$, then $Cl_\beta(\mathcal{M}) \subseteq Cl_\beta(\mathcal{N})$. Moreover, if $|\mathcal{M}| = 1$, then $Cl_\beta(\mathcal{M}) = \mathcal{M}$ (because by assumption β is 0- and 1-reproducing); finally, we always have $Cl_\beta(\emptyset) = \emptyset$.

We can now use these concepts to identify fragments of propositional logic. Additionally, we want fragments to fulfill some natural properties and for technical reasons we require closedness under conjunction.

⁵This prevents us from defining a Dalal-refinement which always selects a single interpretation. Thus, an operator as discussed in (Delgrande and Peppas 2011) does not fit into our concept.

Definition 5. Let $\beta \in \mathcal{B}$. A set $\mathcal{L}' \subseteq \mathcal{L}$ of propositional formulas is a β -fragment if:

1. for all $\psi \in \mathcal{L}'$, $\text{Mod}(\psi) = Cl_\beta(\text{Mod}(\psi))$
2. for all $\mathcal{M} \subseteq 2^{\mathcal{U}}$ with $\mathcal{M} = Cl_\beta(\mathcal{M})$ there exists a $\psi \in \mathcal{L}'$ with $\text{Mod}(\psi) = \mathcal{M}$
3. if $\phi, \psi \in \mathcal{L}'$ then $\phi \wedge \psi \in \mathcal{L}'$.

We call fragments $\mathcal{L}' \subseteq \mathcal{L}$ which are β -fragments for a $\beta \in \mathcal{B}$ also characterizable fragments (of propositional logic).

Well-known fragments of propositional logic can be captured now as follows (see e.g., (Horn 1951; Schaefer 1978)).

Proposition 1. \mathcal{L}_{Horn} is an \wedge -fragment, \mathcal{L}_{DHorn} is an \vee -fragment, \mathcal{L}_{Krom} is a maj_3 -fragment and \mathcal{L}_{Affine} is a \oplus_3 -fragment.

As suggested by their names the Horn fragment and the dual Horn fragment are dual in the following sense: a formula ϕ is Horn if and only if the formula $dual(\phi)$ obtained from ϕ in negating each literal is dual Horn. Moreover the set of models of ϕ is in one-to-one correspondence with the set of models of $dual(\phi)$. From now on we thus omit discussions about the dual Horn fragment. All the results stated below for the Horn fragment also hold for the dual Horn fragment in replacing the function \wedge by the function \vee .

First Examples of Refined Operators. First, let us consider Hegner's revision operator that has the interesting property to be well adapted to any characterizable fragment.

Proposition 2. Let \mathcal{L}' be a characterizable fragment of propositional logic. Then, Hegner's revision operator, \circ_H , restricted to formulas in \mathcal{L}' is a refinement of its own for \mathcal{L}' .

Proof. The properties of Definition 2 are obviously satisfied. We only have to prove that if ψ and μ are formulas in a β -fragment \mathcal{L}' , so is $\psi \circ_H \mu$. Suppose that β is of arity k . Let N_1, \dots, N_k be models of $\psi \circ_H \mu$. By definition of Hegner's revision operator there exist M_1, \dots, M_k models of ψ such that for every i , $N_i \Delta M_i \subseteq \text{Var}(\mu)$. Since β applies to interpretations coordinate-wise we have $\beta(N_1, \dots, N_k) \Delta \beta(M_1, \dots, M_k) \subseteq \text{Var}(\mu)$. Moreover, $\beta(N_1, \dots, N_k)$ is a model of μ (since $\mu \in \mathcal{L}'$ and \mathcal{L}' is a β -fragment), similarly $\beta(M_1, \dots, M_k)$ is a model of ψ . Therefore $\beta(N_1, \dots, N_k)$ is a model of $\psi \circ_H \mu$. Thus we have proved that the set $\text{Mod}(\psi \circ_H \mu)$ is closed under β for every $\beta \in \mathcal{B}$. Hence, by definition of a β -fragment, there exists a formula $\nu \in \mathcal{L}'$ such that $\text{Mod}(\nu) = \text{Mod}(\psi \circ_H \mu)$. \square

Even if we do not fix the revision operator, the ingredients defined above put us in a position to define for any operator \circ , a certain refinement in terms of Definition 2.

Definition 6. Let $\circ: \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{L}$ be a revision operator, $\mathcal{L}' \subseteq \mathcal{L}$ a fragment of classical logic, such that \mathcal{L}' is a β -fragment for some $\beta \in \mathcal{B}$. We define the closure-based \circ -refined operator \circ^{Cl_β} as

$$\text{Mod}(\psi \circ^{Cl_\beta} \mu) := Cl_\beta(\text{Mod}(\psi \circ \mu)). \quad (1)$$

Example 2. Recall Example 1, where we had $\psi, \mu \in \mathcal{L}_{Horn}$ with $\text{Mod}(\psi \circ \mu) = \{\{a\}, \{b\}\}$ ($\circ \in \{\circ_S, \circ_D\}$). Our refined operator \circ^{Cl_\wedge} is defined as $\text{Mod}(\psi \circ^{Cl_\wedge} \mu) = Cl_\wedge(\text{Mod}(\psi \circ \mu)) = \{\{a\}, \{b\}, \emptyset\}$ and thus yields a revision in \mathcal{L}_{Horn} .

Operators \circ^{Cl_β} are refined in the sense of Definition 2.

Proposition 3. For any revision operator, $\circ : \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{L}$ and any β -fragment fragment $\mathcal{L}' \subseteq \mathcal{L}$ of classical logic, \circ^{Cl_β} is a \circ -refinement for \mathcal{L}' .

Proof. For each \mathcal{M} such that $\mathcal{M} = Cl_\beta(\mathcal{M})$, there exists a $\phi \in \mathcal{L}'$ with $\text{Mod}(\phi) = \mathcal{M}$. Thus the above definition of \circ^{Cl_β} indeed yields a mapping $\mathcal{L}' \times \mathcal{L}' \rightarrow \mathcal{L}'$. It remains to show that \circ^{Cl_β} satisfies consistency, equivalence, containment and invariance.

Consistency for \circ^{Cl_β} holds by the fact that, for all $\beta \in \mathcal{B}$ we have $\mathcal{M} \subseteq Cl_\beta(\mathcal{M})$ and $Cl_\beta(\emptyset) = \emptyset$. Equivalence is clear by definition, since we operate on models. To show containment for \circ^{Cl_β} , let $\psi, \mu, \phi \in \mathcal{L}'$ such that $\psi \circ \mu \models \phi$. We have $\text{Mod}(\psi \circ^{Cl_\beta} \mu) = Cl_\beta(\text{Mod}(\psi \circ \mu)) \subseteq Cl_\beta(\text{Mod}(\phi)) = \text{Mod}(\phi)$, i.e., $\psi \circ^{Cl_\beta} \mu \models \phi$. The first equality is by definition, the containment is implied by the assumption $\text{Mod}(\psi \circ \mu) \subseteq \text{Mod}(\phi)$ and by monotonicity of Cl_β . The second equality holds since $\phi \in \mathcal{L}'$ and \mathcal{L}' is a β -fragment. Finally, invariance for \circ^{Cl_β} holds, since in case $\psi \circ \phi \in \mathcal{L}'$, we have $\text{Mod}(\psi \circ^{Cl_\beta} \mu) = Cl_\beta(\text{Mod}(\psi \circ \mu)) = \text{Mod}(\psi \circ \mu)$; the first equality is by definition; the second one since \mathcal{L}' is a β -fragment. \square

We will later show how closure-based refined operators behave in terms of the KM postulates. Before doing so, let us motivate the need for further refined operators.

Example 3. Consider the following example for $\circ \in \{\circ_D, \circ_S\}$ with formulas $\psi, \mu \in \mathcal{L}_{Horn}$, such that $\text{Mod}(\psi) = \{\{a, b, c, d\}, \{a, d\}\}$ and $\text{Mod}(\mu) = \{\{a, b\}, \{b, c\}, \{c, d\}, \{b\}, \{c\}, \emptyset\}$. We have $\text{Mod}(\psi \circ \mu) = \{\{a, b\}, \{b, c\}, \{c, d\}, \emptyset\} =: \mathcal{M}$. Note that $Cl_\wedge = \mathcal{M} \cup \{\{b\}, \{c\}\}$, thus we would have to add two further interpretations when applying the revision operator \circ^{Cl_\wedge} . On the other hand, we can do a smaller change in order to end up with a closed set of interpretations, since $Cl_\wedge(\mathcal{M} \setminus \{\{b, c\}\}) = \mathcal{M} \setminus \{\{b, c\}\}$. Thus, as a result of the revision, also $\mathcal{M} \setminus \{\{b, c\}\}$ should be a candidate.

Next, we show how to capture not only a specific refined operator but characterize the class of all refined operators.

Characterizing Refined Operators. Towards a more general approach to define revision operators we want to reduce the size of generated models, i.e., looking at (1) in Definition 6, we are interested in certain subsets of $Cl_\beta(\text{Mod}(\psi \circ \mu))$ instead of the whole set. For such a selection we formulate some basic properties in the next definition.

Definition 7. Given $\beta \in \mathcal{B}$, we define a β -mapping, f_β , as an application from sets of models into sets of models, $f_\beta : 2^{2^{\mathcal{U}}} \rightarrow 2^{2^{\mathcal{U}}}$, such that for every $\mathcal{M} \subseteq 2^{\mathcal{U}}$:

1. $Cl_\beta(f_\beta(\mathcal{M})) = f_\beta(\mathcal{M})$, i.e., $f_\beta(\mathcal{M})$ is closed under β
2. $f_\beta(\mathcal{M}) \subseteq Cl_\beta(\mathcal{M})$

3. if $\mathcal{M} = Cl_\beta(\mathcal{M})$, then $f_\beta(\mathcal{M}) = \mathcal{M}$
4. If $\mathcal{M} \neq \emptyset$, then $f_\beta(\mathcal{M}) \neq \emptyset$.

The underlying idea of functions f_β is to replace Cl_β by an arbitrary β -mapping f_β when defining refined operators as in (1). Note that Cl_β itself is a β -mapping for any $\beta \in \mathcal{B}$. Below we will provide three more β -mappings and the corresponding refined revision operators. In general, the concept of mappings allows us to define a family of refined operators for fragments of classical logic as follows.

Definition 8. Let $\circ : \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{L}$ be a revision operator and $\mathcal{L}' \subseteq \mathcal{L}$ be a β -fragment of classical logic with $\beta \in \mathcal{B}$. For a β -mapping f_β we denote with $\circ^{f_\beta} : \mathcal{L}' \times \mathcal{L}' \rightarrow \mathcal{L}'$ the operator for \mathcal{L}' defined as $\text{Mod}(\psi \circ^{f_\beta} \mu) := f_\beta(\text{Mod}(\psi \circ \mu))$. The class $[\circ, \mathcal{L}']$ contains all operators \circ^{f_β} where f_β is a β -mapping and $\beta \in \mathcal{B}$ such that \mathcal{L}' is a β -fragment.

The next proposition is central in reflecting that the above class captures all refined operators we had in mind.

Proposition 4. Let $\circ : \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{L}$ be a basic revision operator and $\mathcal{L}' \subseteq \mathcal{L}$ a characterizable fragment of classical logic. Then, $[\circ, \mathcal{L}']$ is the set of all \circ -refinements for \mathcal{L}' .

Proof. Since \mathcal{L}' is a characterizable fragment it is also a β -fragment for some $\beta \in \mathcal{B}$.

Let $\star \in [\circ, \mathcal{L}']$. We show that \star is a \circ -refinement for \mathcal{L}' . Since $\star \in [\circ, \mathcal{L}']$ there exists a $\beta \in \mathcal{B}$ and a β -mapping f_β , such that $\text{Mod}(\psi \star \mu) = f_\beta(\text{Mod}(\psi \circ \mu))$ for all $\psi, \mu \in \mathcal{L}'$. Since f_β satisfies property 1 in Definition 7 and \mathcal{L}' is a β -fragment, \star is indeed a mapping $\mathcal{L}' \times \mathcal{L}' \rightarrow \mathcal{L}'$.

Consistency for \star : Let $\psi, \mu \in \mathcal{L}'$. If $\text{Mod}(\psi \circ \mu) \neq \emptyset$ then $\text{Mod}(\psi \star \mu) = f_\beta(\text{Mod}(\psi \circ \mu)) \neq \emptyset$ by property 4 in Definition 7. In case, $\text{Mod}(\psi \circ \mu) = \emptyset$, we make use of the fact that $Cl_\beta(\emptyset) = \emptyset$ holds for all $\beta \in \mathcal{B}$. By property 2 in Definition 7, we get $\text{Mod}(\psi \star \mu) = f_\beta(\text{Mod}(\psi \circ \mu)) \subseteq Cl_\beta(\text{Mod}(\psi \circ \mu)) = \emptyset$. Equivalence for \star is clear by definition and since f_β is defined on sets of models. To show containment for \star , let $\phi \in T_{\mathcal{L}'}(\psi \circ \mu)$, i.e., $\phi \in \mathcal{L}'$ and $\text{Mod}(\psi \circ \mu) \subseteq \text{Mod}(\phi)$. We have $Cl_\beta(\text{Mod}(\psi \circ \mu)) \subseteq Cl_\beta(\text{Mod}(\phi))$ by monotonicity of Cl_β . By property 2 of Definition 7, $\text{Mod}(\psi \star \mu) \subseteq Cl_\beta(\text{Mod}(\psi \circ \mu))$. Since $\phi \in \mathcal{L}'$ we have $Cl_\beta(\text{Mod}(\phi)) = \text{Mod}(\phi)$. Thus, $\text{Mod}(\psi \star \mu) \subseteq \text{Mod}(\phi)$, i.e. $\phi \in T_{\mathcal{L}'}(\psi \star \mu)$. Finally, we require invariance for \star : In case $\psi \circ \mu \in \mathcal{L}'$, we have $Cl_\beta(\text{Mod}(\psi \circ \mu)) = \text{Mod}(\psi \circ \mu)$ since \mathcal{L}' is a β -fragment. By property 3 in Definition 7, we have $\text{Mod}(\psi \star \mu) = f_\beta(\text{Mod}(\psi \circ \mu)) = \text{Mod}(\psi \circ \mu)$. Thus $T_{\mathcal{L}'}(\psi \star \mu) \subseteq T_{\mathcal{L}'}(\psi \circ \mu)$ as required.

Let \star be a \circ -refinement for \mathcal{L}' . We show that $\star \in [\circ, \mathcal{L}']$. Let, for a set \mathcal{M} of interpretations such that $Cl_\beta(\mathcal{M}) = \mathcal{M}$, $(\psi_{\mathcal{M}}, \mu_{\mathcal{M}})$ be a pair of formulas from \mathcal{L}' such that $\text{Mod}(\psi_{\mathcal{M}} \circ \mu_{\mathcal{M}}) = \mathcal{M}$. Note that for any such $\mathcal{M} \subseteq 2^{\mathcal{U}}$ with $Cl_\beta(\mathcal{M}) = \mathcal{M}$ there exists such a pair. This is due to fact that \circ is a basic revision operator thus satisfying $\top \circ \mu \equiv \mu$ and since \mathcal{L}' is a β -fragment—thus for each such \mathcal{M} there exists a $\mu \in \mathcal{L}'$ with $\text{Mod}(\mu) = \mathcal{M}$. We show that a mapping f on such closed sets defined as $f(\mathcal{M}) := \text{Mod}(\psi_{\mathcal{M}} \star \mu_{\mathcal{M}})$ with $(\psi_{\mathcal{M}}, \mu_{\mathcal{M}})$ being a pair as discussed above is a β -mapping. Note that since \star is a

β -refinement, it satisfies the property of equivalence, thus the actual choice of the pair $(\psi_{\mathcal{M}}, \mu_{\mathcal{M}})$ is not relevant and, given \mathcal{M} , $\psi_{\mathcal{M}} \star \mu_{\mathcal{M}}$ is equivalent for all $(\psi_{\mathcal{M}}, \mu_{\mathcal{M}})$. Thus f is well-defined.

We continue to show that the four properties in Definition 7 hold for f . Property 1 is ensured since \star is defined as a mapping $\mathcal{L}' \times \mathcal{L}' \rightarrow \mathcal{L}'$ and \mathcal{L}' is a β -fragment. Property 3 is ensured jointly by containment and invariance of \star : Recall that in case $\psi \circ \mu \in \mathcal{L}'$, we have $\text{Mod}(\psi \circ \mu) = \text{Mod}(\psi \star \mu)$. Since $\psi \circ \mu \in \mathcal{L}'$ implies $\text{Mod}(\psi \circ \mu) = f_{\beta}(\text{Mod}(\psi \circ \mu))$ we get by definition of f that this property holds. Property 4 is ensured by consistency of \star . It remains to show that Property 2, i.e. $f(\mathcal{M}) \subseteq Cl_{\beta}(\mathcal{M})$ for any $\mathcal{M} \subseteq 2^{\mathcal{U}}$. Towards a contradiction suppose existence of an \mathcal{M} such that $f(\mathcal{M}) \not\subseteq Cl_{\beta}(\mathcal{M})$. Then there exists an $m \in \text{Mod}(\psi_{\mathcal{M}} \star \mu_{\mathcal{M}})$ such that $m \notin Cl_{\beta}(\text{Mod}(\psi_{\mathcal{M}} \circ \mu_{\mathcal{M}}))$. Let $\phi \in \mathcal{L}'$ such that $\text{Mod}(\phi) = Cl_{\beta}(\text{Mod}(\psi_{\mathcal{M}} \circ \mu_{\mathcal{M}}))$. We have $\psi_{\mathcal{M}} \circ \mu_{\mathcal{M}} \models \phi$, while $\psi_{\mathcal{M}} \star \mu_{\mathcal{M}} \not\models \phi$, which provides a contradiction to the assumption that \star satisfies containment. \square

Further Examples of Refined Operators. So far, we have considered the operator $\circ^{Cl_{\beta}}$ as one instantiation of a \circ -refined operator. Let us have a different operator next.

Definition 9. Let $\beta \in \mathcal{B}$ and suppose that \leq is a fixed linear order on the set $2^{\mathcal{U}}$ of interpretations. We define the function Min_{β} as $\text{Min}_{\beta}(\mathcal{M}) = \mathcal{M}$ if $Cl_{\beta}(\mathcal{M}) = \mathcal{M}$, and $\text{Min}_{\beta}(\mathcal{M}) = \text{Min}_{\leq}(\mathcal{M})$ otherwise.

For \mathcal{L}' a β -fragment and \circ a revision operator, the corresponding operators $\circ^{\text{Min}_{\beta}}$ are thus given as $\text{Mod}(\psi \circ^{\text{Min}_{\beta}} \mu) = \text{Min}_{\beta}(\text{Mod}(\psi \circ \mu))$. Clearly, Min_{β} is a β -mapping. Thus, by Proposition 4, $\circ^{\text{Min}_{\beta}}$ is a \circ -refined operator for \mathcal{L}' .

For the situation in Example 3, we so far have not found a satisfying instantiation of refined operators. In fact, we require a slightly more complicated concept here, which is based on the observation that, given a set \mathcal{M} of interpretations with $\mathcal{M} \neq Cl_{\beta}(\mathcal{M})$, there might be elements in \mathcal{M} which are “more responsible” for this inequation than others. To this end we define, for each element M in \mathcal{M} , its “repairset” as the interpretations missing in the closure of applying the operator β when M is involved. Then, the cost of M is the cardinality of its repairset.

Definition 10. For a k -ary Boolean function $\beta \in \mathcal{B}$ and $\mathcal{M} \subseteq 2^{\mathcal{U}}$, we define $\text{repairset}_{\beta}^{\mathcal{M}}(M) = \{\beta(M, M_1, \dots, M_{k-1}) \notin \mathcal{M} \mid M_i \in \mathcal{M}, 1 \leq i < k\}$, and define the cost of M (in \mathcal{M} in terms of β), $\text{cost}_{\beta}^{\mathcal{M}}(M)$, as the cardinality of $\text{repairset}_{\beta}^{\mathcal{M}}(M)$. Moreover, we define $\text{Max}_{\beta}^{\mathcal{M}} = \{M \in \mathcal{M} \mid \forall N \in \mathcal{M}, \text{cost}_{\beta}^{\mathcal{M}}(M) \geq \text{cost}_{\beta}^{\mathcal{M}}(N)\}$ as the set of elements in \mathcal{M} with the highest cost.

Example 4. Let $\mathcal{M} = \{\{a, b\}, \{b, c\}, \{c, d\}, \emptyset\}$ and consider the Boolean function \wedge . Then, $\text{repairset}_{\wedge}^{\mathcal{M}}(\{a, b\}) = \{\{b\}\}$, $\text{repairset}_{\wedge}^{\mathcal{M}}(\{b, c\}) = \{\{b\}, \{c\}\}$, $\text{repairset}_{\wedge}^{\mathcal{M}}(\{c, d\}) = \{\{c\}\}$, and $\text{repairset}_{\wedge}^{\mathcal{M}}(\emptyset) = \emptyset$. Thus, we have $\text{cost}_{\wedge}^{\mathcal{M}}(\{b, c\}) = 2$, $\text{cost}_{\wedge}^{\mathcal{M}}(\{a, b\}) = \text{cost}_{\wedge}^{\mathcal{M}}(\{c, d\}) = 1$, and $\text{cost}_{\wedge}^{\mathcal{M}}(\emptyset) = 0$.

The following definition defines a β -mapping based on the idea to get rid off the most costly interpretations.

Definition 11. Let $\beta \in \mathcal{B}$. We define the mapping t_{β}^1 as

$$t_{\beta}^1(\mathcal{M}) = \begin{cases} Cl_{\beta}(\mathcal{M}) & \text{if } \text{Max}_{\beta}^{\mathcal{M}} = \mathcal{M} \\ t_{\beta}^1(\mathcal{M} \setminus \text{Max}_{\beta}^{\mathcal{M}}) & \text{otherwise} \end{cases}$$

Informally, this operator functions along the lines of the following algorithm. We start with a set \mathcal{M} of interpretations. In case \mathcal{M} is already closed under β all elements have cost 0 and we return \mathcal{M} ; in case all $M \in \mathcal{M}$ have the same cost, we cannot objectively do better than build $Cl_{\beta}(\mathcal{M})$ and return that set; otherwise we remove the most costly elements from \mathcal{M} and restart with this reduced set. It can be shown that, for \mathcal{L}' a β -fragment and \circ a revision operator, the operator $\circ^{t_{\beta}^1}$, given as $\text{Mod}(\psi \circ^{t_{\beta}^1} \mu) = t_{\beta}^1(\text{Mod}(\psi \circ \mu))$, is a \circ -refined operator for \mathcal{L}' , basically since t_{β}^1 is a β -mapping in sense of Definition 7.

Example 5. Recall the setting from Example 3. There we had the situation that, for $\circ \in \{\circ_D, \circ_S\}$, $\text{Mod}(\psi \circ \mu) = \{\{a, b\}, \{b, c\}, \{c, d\}, \emptyset\} = \mathcal{M}$ with \mathcal{M} as analyzed in Example 4. Thus, $\text{Max}_{\wedge}^{\mathcal{M}} = \{\{b, c\}\}$. We obtain $t_{\wedge}^1(\mathcal{M}) = \mathcal{M} \setminus \{\{b, c\}\}$ (since $\mathcal{M} \setminus \{\{b, c\}\}$ is already closed under \wedge) and hence $\circ^{t_{\wedge}^1}$ behaves as suggested in Example 3.

Let us have a final example to motivate one further instantiation of refined operators.

Example 6. Consider $\psi, \mu \in \mathcal{L}_{\text{Horn}}$ with $\text{Mod}(\psi) = \{\{a, b, c\}, \{b, c, d\}, \{b, c\}\}$ and $\text{Mod}(\mu) = \{\{a, b, c, d\}, \{a, b\}, \{c, d\}, \emptyset\}$. For $\circ \in \{\circ_S, \circ_D\}$, we obtain $\text{Mod}(\psi \circ \mu) = \{\{a, b, c, d\}, \{a, b\}, \{c, d\}\}$, which is not closed under \wedge . The simplest “repair” would be to add interpretation \emptyset (as is done by $\circ^{Cl_{\wedge}}$), but $\circ^{t_{\wedge}^1}$ here removes both $\{a, b\}$ and $\{c, d\}$ (since these two have the highest cost) and we end up with $\text{Mod}(\psi \circ^{t_{\wedge}^1} \mu) = \{\{a, b, c, d\}\}$.

Definition 12. Let $\beta \in \mathcal{B}$, we define a mapping t_{β}^2 as

$$t_{\beta}^2(\mathcal{M}) = \begin{cases} t_{\beta}^2(\mathcal{M} \setminus \text{Max}_{\beta}^{\mathcal{M}}) & \text{if } |\text{Max}_{\beta}^{\mathcal{M}}| = 1 < |\mathcal{M}| \\ Cl_{\beta}(\mathcal{M}) & \text{otherwise} \end{cases}$$

where $\text{Max}_{\beta}^{\mathcal{M}}$ is as in Definition 11.

For \mathcal{L}' a β -fragment and \circ a revision operator, it holds that the operator $\circ^{t_{\beta}^2}$, given as $\text{Mod}(\psi \circ^{t_{\beta}^2} \mu) = t_{\beta}^2(\text{Mod}(\psi \circ \mu))$, is a \circ -refined operator for \mathcal{L}' .

It can be seen that $\circ^{t_{\wedge}^1}$ behaves as expected in all cases we have discussed in the previous examples. More precisely, in the case of Example 6, we now obtain for $\text{Mod}(\psi \circ^{t_{\wedge}^2} \mu)$ the simplest repair $Cl_{\wedge}(\{\{a, b, c, d\}, \{a, b\}, \{c, d\}\}) = \text{Mod}(\mu)$. For the other examples we had, it can be checked that t_{\wedge}^2 and t_{\wedge}^1 behave analogously. This, however, does not mean that t_{β}^2 is the “best” refinement we can get. In fact, it is up to the user to define a refined operator which is best suited for her purposes. Nonetheless, our generic results provide already basic properties for such operators. In the next section, we analyse the KM postulates for refined operators.

KM Postulates

In this section, we first show a positive result concerning the first four KM postulates. In fact, we prove that any operator

refined for a fragment \mathcal{L}' has good properties as long as the original operator has good properties. We show that our refinements satisfy the most basic postulates of belief revision (R1)–(R4).

We then show that particular refined operators even satisfy (R5). As a negative result, we show that for the four fragments we consider here, i.e. \mathcal{L}_{Horn} , \mathcal{L}_{DHorn} , \mathcal{L}_{Krom} and \mathcal{L}_{Affine} , there is no Dalal- or Satoh-refined operator that satisfies (R6). Finally, we also prove an impossibility result for (R5) when particular materializations of refined operators are considered.

Proposition 5. *Let \circ be a revision operator satisfying KM postulates (R1)–(R4), and $\mathcal{L}' \subseteq \mathcal{L}$ a characterizable fragment. Then each $\star \in [\circ, \mathcal{L}']$ satisfies (R1)–(R4) in \mathcal{L}' as well.*

Proof. Since \mathcal{L}' is characterizable there exists a $\beta \in \mathcal{B}$, such that \mathcal{L}' is a β -fragment. We thus can assume that $\star \in [\circ, \mathcal{L}']$ is an operator of form \circ^{f_β} where f_β is a suitable β -mapping. In what follows, note that we can restrict ourselves to $\psi, \mu \in \mathcal{L}'$, since we have to show that \circ^{f_β} satisfies (R1)–(R4) in \mathcal{L}' . (R1): Since \circ satisfies (R1), $\text{Mod}(\psi \circ \mu) \subseteq \text{Mod}(\mu)$. Thus, $Cl_\beta(\text{Mod}(\psi \circ \mu)) \subseteq Cl_\beta(\text{Mod}(\mu))$ by monotonicity of the closure. Hence, $Cl_\beta(\text{Mod}(\psi \circ \mu)) \subseteq \text{Mod}(\mu)$, since $\mu \in \mathcal{L}'$ and \mathcal{L}' is a β -fragment. According to property 2 in Definition 7 we have $f_\beta(\text{Mod}(\psi \circ \mu)) \subseteq Cl_\beta(\text{Mod}(\psi \circ \mu))$, and therefore by definition of \star , $\text{Mod}(\psi \star \mu) \subseteq \text{Mod}(\mu)$, which proves that $\psi \star \mu \models \mu$.

(R2): Suppose that $\psi \wedge \mu$ is satisfiable. We have $\text{Mod}(\psi \star \mu) = f_\beta(\text{Mod}(\psi \circ \mu)) = f_\beta(\text{Mod}(\psi \wedge \mu))$, since \circ satisfies (R2). Since $\psi \wedge \mu \in \mathcal{L}'$ (by definition of fragment) we have $\text{Mod}(\psi \star \mu) = \text{Mod}(\psi \wedge \mu)$ thanks to condition (3) in Definition 7.

(R3): Suppose μ satisfiable. Since \circ satisfies (R3), $(\psi \circ \mu)$ is satisfiable. Since \star is a \circ -refinement (Proposition 4), $(\psi \star \mu)$ is also satisfiable by the property of consistency (see Definition 2).

(R4): Let $\psi_1, \psi_2, \mu_1, \mu_2 \in \mathcal{L}'$ with $\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$. Since \circ satisfies (R4), $\psi_1 \circ \mu_1 \equiv \psi_2 \circ \mu_2$. Since \star is a \circ -refinement, $\psi_1 \star \mu_1 \equiv \psi_2 \star \mu_2$ by the property of equivalence (Definition 2). \square

A natural question is whether one can find refined operators for characterizable fragments that satisfy all postulates. Our next result answers negatively to this question in the sense that it shows that no matter which operator we choose from $[\circ, \mathcal{L}']$ in case of $\circ \in \{\circ_D, \circ_S\}$ and $\mathcal{L}' \in \{\mathcal{L}_{Horn}, \mathcal{L}_{DHorn}, \mathcal{L}_{Krom}, \mathcal{L}_{Affine}\}$, it will not satisfy (R6).

Proposition 6. *Let $\circ \in \{\circ_D, \circ_S\}$ and $\mathcal{L}' \in \{\mathcal{L}_{Horn}, \mathcal{L}_{DHorn}, \mathcal{L}_{Krom}, \mathcal{L}_{Affine}\}$. Then any refined operator $\star \in [\circ, \mathcal{L}']$ violates postulate (R6) in \mathcal{L}' .*

Proof. (R6) states that in case $(\psi \star \mu) \wedge \phi$ is satisfiable, then $\psi \star (\mu \wedge \phi) \models (\psi \star \mu) \wedge \phi$. We show in detail only the case $\mathcal{L}' = \mathcal{L}_{Horn}$. By definition, there is an \wedge -mapping f such that $\star = \circ^f$ and we have $f(\mathcal{M}) \subseteq Cl_\wedge(\mathcal{M})$ with $Cl_\wedge(f(\mathcal{M})) = f(\mathcal{M})$. Let $\text{Mod}(\psi) = \{\{a, b, c, d, e\}, \emptyset\}$ and $\text{Mod}(\mu) = \{\{a, b, c, d\}, \{a, b, c, e\}, \{a, b, c\}, \{a, b\}, \{a\}\}$. Note that

such $\psi, \mu \in \mathcal{L}_{Horn}$ exist. For $\circ \in \{\circ_D, \circ_S\}$, we have $\mathcal{M} = \text{Mod}(\psi \circ \mu) = \{\{a, b, c, d\}, \{a, b, c, e\}, \{a\}\}$. Let us consider the possibilities for $\text{Mod}(\psi \star \mu) = f(\mathcal{M})$. By the definition of refined operators, we know that $\{a, b\} \notin f(\mathcal{M})$ since $\{a, b\} \notin Cl_\wedge(\mathcal{M})$. We consider two cases:

First, assume $\{a, b, c\} \in f(\mathcal{M})$: let ϕ be such that $\text{Mod}(\phi) = \{\{a, b\}, \{a, b, c\}\} = \mathcal{N}$. Clearly, such a ϕ exists in \mathcal{L}_{Horn} . Also note that $\text{Mod}(\phi) \subseteq \text{Mod}(\mu)$. We get $\text{Mod}(\psi \star (\mu \wedge \phi)) = \text{Mod}(\psi \star \phi) = f(\text{Mod}(\psi \circ \phi)) = \mathcal{N}$ (\mathcal{N} is closed under \wedge , $f(\mathcal{N}) = \mathcal{N}$ holds by definition of refined operators) but $\text{Mod}((\psi \star \mu) \wedge \phi) = f(\mathcal{M}) \cap \mathcal{N} \subset \mathcal{N}$.

Otherwise, we have $\emptyset \subset f(\mathcal{M}) \subseteq \{\{a, b, c, d\}, \{a\}\}$ or $\emptyset \subset f(\mathcal{M}) \subseteq \{\{a, b, c, e\}, \{a\}\}$ (note that $\{\{a, b, c, d\}, \{a, b, c, e\}\} \subseteq f(\mathcal{M})$ would imply $\{a, b, c\} \in f(\mathcal{M})$). For the cases with $|f(\mathcal{M})| = 1$, we select $\phi \in \mathcal{L}_{Horn}$ where $\text{Mod}(\phi)$ is $\{\{a, b, c, d\}, \{a\}\}$ or $\{\{a, b, c, e\}, \{a\}\}$ such that $f(\mathcal{M}) \subseteq \text{Mod}(\phi)$ holds. Then, $\text{Mod}(\psi \star (\mu \wedge \phi)) = \text{Mod}(\psi \star \phi) = f(\text{Mod}(\psi \circ \phi)) = \text{Mod}(\phi)$ (again since $\text{Mod}(\phi)$ is closed under \wedge) but $\text{Mod}((\psi \star \mu) \wedge \phi) = f(\mathcal{M}) \cap \text{Mod}(\phi) \subset \text{Mod}(\phi)$. Two cases remain. Let us suppose $f(\mathcal{M}) = \{\{a, b, c, d\}, \{a\}\}$; the final case is then symmetric. We now use $\phi \in \mathcal{L}_{Horn}$ with $\text{Mod}(\phi) = \{\{a, b, c, e\}, \{a\}\}$. Again, $\text{Mod}(\psi \star (\mu \wedge \phi)) = \text{Mod}(\psi \star \phi) = f(\text{Mod}(\psi \circ \phi)) = \text{Mod}(\phi)$ (since $\text{Mod}(\phi)$ is closed under \wedge) but $\text{Mod}((\psi \star \mu) \wedge \phi) = \{\{a\}\}$. Thus, $\text{Mod}((\psi \star \mu) \wedge \phi) \neq \emptyset$ and $\text{Mod}(\psi \star (\mu \wedge \phi)) \not\subseteq \text{Mod}((\psi \star \mu) \wedge \phi)$. Hence, $\psi \star (\mu \wedge \phi) \not\models (\psi \star \mu) \wedge \phi$.

The case $\mathcal{L}' = \mathcal{L}_{DHorn}$ is dual. For $\mathcal{L}' = \mathcal{L}_{Krom}$, we can use formulas $\psi, \mu \in \mathcal{L}_{Krom}$ with the same set of models as for the case $\mathcal{L}' = \mathcal{L}_{Horn}$ and proceed similarly as above. Finally, for $\mathcal{L}' = \mathcal{L}_{Affine}$, formulas $\psi, \mu \in \mathcal{L}_{Affine}$ having as models $\text{Mod}(\psi) = \{\{a, b, c\}, \{b, c, d\}\}$ and $\text{Mod}(\mu) = \{\{a, b, c, d, e, f\}, \{a, b, c, d\}, \{a, b, e, f\}, \{c, d, e, f\}, \{a, b\}, \{c, d\}, \{e, f\}, \emptyset\}$ can be used to show the assertion. \square

The status of the 5th KM postulate, R5, is less clear. Indeed, on the one hand the next proposition shows that the β -mapping Min_β defined above allows to refine Dalal's operator for any characterizable fragment in satisfying the fifth postulate, whereas it is not the case for Satoh's operator. Moreover, we will show afterwards that the refinements of both Dalal's and Satoh's operators by any of the other mappings we have considered so far fail at satisfying R5.

Proposition 7. (1) *The refined operator $\circ_D^{\text{Min}_\beta}$ satisfies the KM postulate R5 in any β -fragment \mathcal{L}' .* (2) *The refined operator $\circ_S^{\text{Min}_\beta}$ violate postulate (R5) in any $\mathcal{L}' \in \{\mathcal{L}_{Horn}, \mathcal{L}_{DHorn}, \mathcal{L}_{Krom}, \mathcal{L}_{Affine}\}$.*

Proof. Let us first consider Dalal's operator. Let ψ, μ and ϕ be formulas in \mathcal{L}' . If $(\psi \circ_D^{\text{Min}_\beta} \mu) \wedge \phi$ is unsatisfiable, then obviously $(\psi \circ_D^{\text{Min}_\beta} \mu) \wedge \phi \models \psi \circ_D^{\text{Min}_\beta} (\mu \wedge \phi)$. Suppose now that $(\psi \circ_D^{\text{Min}_\beta} \mu) \wedge \phi$ is satisfiable. There are two cases to distinguish. First, if $Cl_\beta(\text{Mod}(\psi \circ_D \mu)) = \text{Mod}(\psi \circ_D \mu)$. Observe that in this case $\text{Mod}((\psi \circ_D \mu) \wedge \phi)$ is also closed under β . Thus, $\text{Mod}((\psi \circ_D^{\text{Min}_\beta} \mu) \wedge \phi) = \text{Min}_\beta(\text{Mod}((\psi \circ_D \mu) \wedge \phi)) = \text{Min}_\beta(\text{Mod}(\psi \circ_D (\mu \wedge \phi)))$ for Dalal's operator satisfies both R5 and R6. Hence,

$\text{Mod}((\psi \circ_D^{\text{Min}_\beta} \mu) \wedge \phi) = \text{Mod}((\psi \circ_D^{\text{Min}_\beta} (\mu \wedge \phi)))$. Second, suppose that $\text{Cl}_\beta(\text{Mod}(\psi \circ_D \mu)) \neq \text{Mod}(\psi \circ_D \mu)$. In this case $\text{Mod}((\psi \circ_D^{\text{Min}_\beta} \mu) \wedge \phi) = \text{Min}_\leq(\text{Mod}(\psi \circ_D \mu)) \cap \text{Mod}(\phi)$ (for $(\psi \circ_D^{\text{Min}_\beta} \mu) \wedge \phi$ is satisfiable), and we have $\text{Min}_\leq(\text{Mod}(\psi \circ_D \mu)) \cap \text{Mod}(\phi) \subseteq \text{Min}_\leq(\text{Mod}((\psi \circ_D \mu) \wedge \phi))$. Since $(\psi \circ_D \mu) \wedge \phi \equiv \psi \circ_D (\mu \wedge \phi)$ (for Dalal's operator satisfies both R5 and R6), we obtain $\text{Mod}((\psi \circ_D^{\text{Min}_\beta} \mu) \wedge \phi) \subseteq \text{Min}_\leq(\text{Mod}(\psi \circ_D (\mu \wedge \phi))) = \text{Mod}(\psi \circ_D^{\text{Min}_\beta} (\mu \wedge \phi))$, thus proving that $(\psi \circ_D^{\text{Min}_\beta} \mu) \wedge \phi \models \psi \circ_D^{\text{Min}_\beta} (\mu \wedge \phi)$.

Let us now consider Satoh's operator. Without loss of generality suppose that the linear order \leq on interpretations on which the operator Min_β is based verifies $\{a, b\} < \{d, e\} < \{c, d, e\} < \{a, b, c\}$.

We give a full proof only for $\mathcal{L}' = \mathcal{L}_{\text{Horn}}$. Let $\text{Mod}(\psi) = \{\{a, b, c, d, e\}\}$, $\text{Mod}(\mu) = \{\{a, b, c\}, \{a, b\}, \{d, e\}, \emptyset\}$, and $\text{Mod}(\phi) = \{\{a, b\}, \{d, e\}, \emptyset\}$. Indeed such $\psi, \mu, \phi \in \mathcal{L}_{\text{Horn}}$ exist. We have $\text{Mod}(\psi \circ_S \mu) = \{\{a, b, c\}, \{d, e\}\}$ -which is not closed under \wedge - thus $\text{Mod}(\psi \star \mu) = \text{Min}_\leq(\{\{a, b, c\}, \{d, e\}\}) = \{\{d, e\}\}$. (where \star denotes $\circ_D^{\text{Min}_\wedge}$). Hence, on the one hand, $\text{Mod}((\psi \star \mu) \wedge \phi) = \text{Min}_\leq(\{\{a, b, c\}, \{d, e\}\}) = \{\{d, e\}\}$. On the other hand, $\text{Mod}(\psi \star (\mu \wedge \phi)) = \text{Min}_\wedge(\{\{a, b\}, \{d, e\}\}) = \{\{a, b\}\}$.

The same proof works for the case $\mathcal{L}' = \mathcal{L}_{\text{Krom}}$. The proof for $\mathcal{L}' = \mathcal{L}_{\text{DHorn}}$ is dual. For $\mathcal{L}' = \mathcal{L}_{\text{Affine}}$, formulas ψ, μ, ϕ such that $\text{Mod}(\psi) = \{\{a, b, c, d, e\}\}$, $\text{Mod}(\mu) = \{\{a, b, c\}, \{a, b\}, \{d, e\}, \{c, d, e\}\}$, and $\text{Mod}(\phi) = \{\{a, b\}, \{c, d, e\}\}$ can be used. \square

Proposition 8. *Let $\circ \in \{\circ_D, \circ_S\}$ and $\mathcal{L}' \in \{\mathcal{L}_{\text{Horn}}, \mathcal{L}_{\text{DHorn}}, \mathcal{L}_{\text{Krom}}, \mathcal{L}_{\text{Affine}}\}$. Then the refined operators \circ^{Cl_β} , $\circ^{t_\beta^1}$ and $\circ^{t_\beta^2}$ violate postulate (R5) in \mathcal{L}' .*

Proof. We provide formulas $\psi, \mu, \phi \in \mathcal{L}'$ such that $\text{Mod}((\psi \star \mu) \wedge \phi) \not\subseteq \text{Mod}(\psi \star (\mu \wedge \phi))$, i.e., such that $(\psi \star \mu) \wedge \phi \not\models \psi \star (\mu \wedge \phi)$ for any $\star \in [\circ, \mathcal{L}']$.

In detail we only show that case $\mathcal{L}' = \mathcal{L}_{\text{Horn}}$: Let $\star = \circ^f$ for $f = \text{Cl}_\wedge, t_\wedge^1$ or t_\wedge^2 . Let $\text{Mod}(\psi) = \{\{a, b, c\}\}$, $\text{Mod}(\mu) = \{\{a, b\}, \{a, c\}, \{a\}\}$, and $\text{Mod}(\phi) = \{\{a, b\}, \{a\}\}$. Note that such $\psi, \mu, \phi \in \mathcal{L}_{\text{Horn}}$ exist. We have $\text{Mod}(\psi \star \mu) = f(\text{Mod}(\psi \circ \mu)) = f(\{\{a, b\}, \{a, c\}\})$ for $\circ \in \{\circ_D, \circ_S\}$. Hence, $\text{Mod}(\psi \star \mu) = \{\{a, b\}, \{a, c\}, \{a\}\}$ since $f(\{\{a, b\}, \{a, c\}\}) = \{\{a, b\}, \{a, c\}, \{a\}\}$ for all f under consideration. Therefore, $\text{Mod}((\psi \star \mu) \wedge \phi) = \{\{a, b\}, \{a\}\}$. On the other hand $\text{Mod}(\psi \star (\mu \wedge \phi)) = \text{Mod}(\psi \star \phi) = f(\text{Mod}(\psi \circ \phi)) = \{\{a, b\}\}$ (since $\text{Mod}(\psi \circ \phi) = \{\{a, b\}\}$ is already closed under \wedge and \star is a \circ -refinement for \mathcal{L}' , a \wedge -fragment).

The case $\mathcal{L}' = \mathcal{L}_{\text{DHorn}}$ is dual. For $\mathcal{L}' = \mathcal{L}_{\text{Krom}}$, we use $\psi, \mu, \phi \in \mathcal{L}_{\text{Krom}}$ with $\text{Mod}(\psi) = \{\{a, b, c\}\}$, $\text{Mod}(\mu) = \{\{a\}, \{b\}, \{c\}, \emptyset\}$, $\text{Mod}(\phi) = \{\{c\}, \emptyset\}$. For $\mathcal{L}' = \mathcal{L}_{\text{Affine}}$, formulas $\psi, \mu, \phi \in \mathcal{L}_{\text{Affine}}$ with $\text{Mod}(\psi) = \{\{a, b, c\}, \{b, c, d\}\}$, $\text{Mod}(\mu) = \{\{a, b, c, d\}, \{a, b\}, \{c, d\}, \emptyset\}$, and $\text{Mod}(\phi) = \{\{a, b, c, d\}, \emptyset\}$ can be employed. \square

Complexity Issues

Our goal in this section is to initiate a study of the computational complexity for refined operators tailored for charac-

terizable fragments of propositional logic. We focus on the complexity of model checking (see (Liberatore and Schaerf 2001)) which is the most basic computational problem in the belief revision context and which is defined as follows. Let \circ be a revision operator, \mathcal{L}' a β -fragment of classical logic and f_β a β -mapping. We consider the following problem:

Problem: MODEL-CHECKING($\circ, \mathcal{L}', f_\beta$)
Input: Two formulas $\psi, \mu \in \mathcal{L}'$, a model M
Question: $M \in \text{Mod}(\psi \circ^{f_\beta} \mu)$?

While the complexity of revision in the propositional case has been largely investigated (Eiter and Gottlob 1992; Nebel 1994; Eiter and Gottlob 1996; Nebel 1998; Liberatore 2000; Liberatore and Schaerf 2001) there are very few results on propositional sublanguages. As far as we know only the Horn fragment has been investigated. We first examine the complexity of model checking for Hegner's revision operator on any characterizable fragment. Then we focus on the Horn (and dual Horn) fragments to pinpoint the complexity of model checking for refined Dalal's and Satoh's operators.

Refined Hegner operator on characterizable fragments.

Recall that for any characterizable fragment \mathcal{L}' , if ψ and μ are two formulas in \mathcal{L}' , then so is $\psi \circ_H \mu$. As a consequence $[\circ_H, \mathcal{L}'] = \{\circ_H\}$. Therefore, in order to study the complexity of the model checking for Hegner's refined operators it is enough to consider MODEL-CHECKING($\circ_H, \mathcal{L}', Id$). Let $\text{SAT}(\mathcal{L}')$ denote the satisfiability problem for formulas in \mathcal{L}' .

Proposition 9. *Let \mathcal{L}' be a characterizable fragment, then MODEL-CHECKING($\circ_H, \mathcal{L}', Id$) \equiv SAT(\mathcal{L}') under log-space reductions.*

Proof. First, let X be the set of variables of ψ , Y the set of variables in μ , and $Z = X \cup Y$. As in (Liberatore and Schaerf 2001), we construct the new formula $\psi' = \psi \wedge \bigwedge \{z_i | z_i \in (M \setminus Y)\} \wedge \bigwedge \{\neg z_i | z_i \in (Z \setminus Y) \setminus M\}$. This proves MODEL-CHECKING($\circ_H, \mathcal{L}', Id$) \leq SAT(\mathcal{L}').

Conversely, let ψ be a formula in \mathcal{L}' . Construct $\mu \in \mathcal{L}'$ having the same set of variables as ψ and the empty set as unique model. Hence, ψ is satisfiable if and only if the empty set is a model of $\psi \circ_H \mu$. \square

By the above proposition we immediately have that MODEL-CHECKING($\circ_H, \mathcal{L}', Id$) is in P (respectively, NP-complete), whenever deciding the satisfiability of a formula in \mathcal{L}' is in P (respectively is NP-complete).

Refined Dalal operator for (dual) Horn. The complexity of model checking for the Dalal operator in the propositional case (resp. in the Horn case) is given by Liberatore and Schaerf (2001, Thm 7 and Thm 15). It is $P^{\text{NP}[O(\log n)]}$ -complete. We extend this hardness result to all refinements of Dalal's operator on $\mathcal{L}_{\text{Horn}}$ and, by duality, on $\mathcal{L}_{\text{DHorn}}$.

Proposition 10. MODEL-CHECKING($\circ_D, \mathcal{L}_{\text{Horn}}, f$) is $P^{\text{NP}[O(\log n)]}$ -hard for any \wedge -mapping f .

Proof. In (Liberatore and Schaerf 2001), $P^{\text{NP}[O(\log n)]}$ -hardness of Dalal-revision with arbitrary formulas is proved by reduction from the co-problem of UOCSAT (Kadin

$$\begin{aligned}
& (X \equiv X_1 \equiv \dots \equiv X_m) \wedge (X' \equiv X'_1 \equiv \dots \equiv X'_m) \wedge \\
& (\tilde{X} \equiv \tilde{X}_1 \equiv \dots \equiv \tilde{X}_m) \wedge (\tilde{X}' \equiv \tilde{X}'_1 \equiv \dots \equiv \tilde{X}'_m) \wedge \\
& (Y \equiv Y_1 \equiv \dots \equiv Y_m) \wedge (Y' \equiv Y'_1 \equiv \dots \equiv Y'_m) \wedge \\
& (\tilde{Y} \equiv \tilde{Y}_1 \equiv \dots \equiv \tilde{Y}_m) \wedge (\tilde{Y}' \equiv \tilde{Y}'_1 \equiv \dots \equiv \tilde{Y}'_m) \wedge \\
& (Z \equiv Z_1 \equiv \dots \equiv Z_m) \wedge (\tilde{Z} \equiv \tilde{Z}_1 \equiv \dots \equiv \tilde{Z}_m) \wedge \\
& (Y \equiv W) \wedge (Y' \equiv W') \wedge [(\bigwedge_{i=1}^p z_i) \rightarrow d] \wedge \\
& \bigwedge_{j=1}^n [(\neg x_j \vee \neg \tilde{x}_j) \wedge (\neg x'_j \vee \neg \tilde{x}'_j)] \wedge \\
& \bigwedge_{i=1}^p [(\neg y_i \vee \neg \tilde{y}_i) \wedge (\neg y'_i \vee \neg \tilde{y}'_i) \wedge (\neg z_i \vee \neg \tilde{z}_i)] \wedge \\
& \bigwedge_{i=1}^p [(y_i \wedge y'_i \rightarrow z_i) \wedge (\tilde{y}_i \wedge \tilde{y}'_i \rightarrow \tilde{z}_i)] \wedge \\
& \bigwedge_{i=1}^p [(y_i \wedge \tilde{y}'_i \rightarrow \tilde{z}_i) \wedge (\tilde{y}_i \wedge y'_i \rightarrow z_i)] \wedge \\
& \bigwedge_{i=1}^p [(\bigwedge A_i \wedge \bigwedge \tilde{B}_i \rightarrow y_i) \wedge (\bigwedge A'_i \wedge \bigwedge \tilde{B}'_i \rightarrow y'_i)]
\end{aligned}$$

Figure 1: Formula ψ as used in proof of Proposition 10.

1989), that is: given a set of clauses $\mathcal{C} = \{C_1, \dots, C_p\}$, decide whether its (cardinality) maximal consistent subset is unique. The $\text{P}^{\text{NP}[O(\log n)]}$ -hardness of Dalal-revision with Horn formulas is then proved by reduction from the model checking problem of \circ_D in the general case. Neither reduction is applicable in our case: the reduction in the general case clearly does not use Horn formulas and the reduction for Horn formulas yields a set of models with $\psi \circ_D^f \mu \neq \psi \circ_D \mu$ in general. Therefore, we present a new reduction from co-UOCSAT to $\text{MODEL-CHECKING}(\circ_D, \mathcal{L}_{\text{Horn}}, f)$.

Consider an arbitrary instance of co-UOCSAT, i.e., clause set $\mathcal{C} = \{C_1, \dots, C_p\}$ over alphabet $X = \{x_1, \dots, x_n\}$. Each clause C_i can be written as $C_i = (\bigwedge A_i \rightarrow \bigvee B_i)$, where A_i and B_i are subsets of X . We need a new variable d and new alphabets that are in one-to-one correspondence with X : $X_1, \dots, X_m, \tilde{X}, \tilde{X}_1, \dots, \tilde{X}_m, X', X'_1, \dots, X'_m, \tilde{X}', \tilde{X}'_1, \dots, \tilde{X}'_m$ where we set $m = 2p + 1$. Likewise, we need new alphabets $Y, Y_1, \dots, Y_m, \tilde{Y}, \tilde{Y}_1, \dots, \tilde{Y}_m, Y', Y'_1, \dots, Y'_m, \tilde{Y}', \tilde{Y}'_1, \dots, \tilde{Y}'_m, W, W', Z, Z_1, \dots, Z_m, \tilde{Z}, \tilde{Z}_1, \dots, \tilde{Z}_m$, which are in one-to-one correspondence with \mathcal{C} . Let us use U to denote the union of all these sets. Intuitively X' (likewise Y') will serve to rename the variables of X (resp. Y) while \tilde{X} (likewise \tilde{Y} and \tilde{Z}) is meant to represent the variables of X (resp. Y and Z) negated. Below we use e.g. $X \equiv X_1$ as a shorthand for $\bigwedge_{j=1}^n (x_j \equiv x_j^1)$ and we write $\neg W$ to denote $\bigwedge_{j=1}^p \neg w_i$. Consider the instance with

$$\begin{aligned}
& \psi \text{ as given in Figure 1,} \\
& \mu = \bigwedge_{u \in U \setminus (W \cup W' \cup \{d\})} u \wedge \neg W \wedge \neg W', \text{ and} \\
& M = U \setminus (W \cup W' \cup \{d\}).
\end{aligned}$$

All clauses are Horn. Moreover $\text{Mod}(\mu) = \{M, M \cup \{d\}\}$. Every subset of $\text{Mod}(\mu)$ is thus closed under any 0- and 1-reproducing Boolean function f . Hence, $\psi \circ_D^f \mu = \psi \circ_D \mu$. We claim that \mathcal{C} has a unique cardinality maximal consistent set if and only if $M \notin \text{Mod}(\psi \circ \mu)$. Indeed, let I and J be models of ψ and μ realizing the minimal distance. The copies of $X, X', \tilde{X}, \tilde{X}', Y, Y', \dots$ have been introduced, s.t. $I(\tilde{x}) = 1 - I(x)$, $I(\tilde{x}') = 1 - I(x')$, etc. holds. Hence, the distance between I and J becomes minimal if the number of clauses that are falsified in the interpretations represented by X and X' is minimized. If this minimum can only be

achieved in a single way then $I(d) = 1$ holds and only $M \cup \{d\}$ has minimal distance to the models of ψ . \square

Note that the upper complexity bounds (membership) are difficult to obtain for general fragments. The complexity of the problem “Does M belong to the closure of a set of models under β ?” is the key. The Horn fragment – which is an \wedge -fragment – has the following important property, which will allow us to prove the membership in NP.

Proposition 11. *Let $\mathcal{M} \subseteq 2^U$ and M an interpretation over variables $\{x_1, \dots, x_n\}$. Then $M \in \text{Cl}_\wedge(\mathcal{M})$ iff there exists M_1, \dots, M_k in \mathcal{M} ($k \leq n$) such that $M = M_1 \wedge \dots \wedge M_k$.*

Proof. This follows from the associativity of the \wedge function e.g., $(M_1 \wedge M_2) \wedge (M_3 \wedge M_4) = (M_1 \wedge M_2 \wedge M_3 \wedge M_4)$, and in observing that if $M_1 \not\subseteq M_2$ then $|M_1 \wedge M_2| < |M_1|$, thus justifying that $k \leq n$. \square

Proposition 12. $\text{MODEL-CHECKING}(\circ_D, \mathcal{L}_{\text{Horn}}, f)$ is $\text{P}^{\text{NP}[O(\log n)]}$ -complete, for $f \in \{\text{Cl}_\wedge, \text{Min}_\wedge\}$.

Proof. According to Proposition 10 only membership has to be proved. Let us sketch a polynomial time algorithm with a logarithmic number of calls to an NP-oracle.

We check whether ψ or μ is unsatisfiable with the oracle. If not, then we proceed in two steps. First we compute d the distance between ψ and μ by binary search with $O(\log n)$ calls to the NP-oracle “is $d(\psi, \mu) \leq k$?”.

In the case of $\text{MODEL-CHECKING}(\circ_D, \mathcal{L}_{\text{Horn}}, \text{Cl}_\wedge)$ we make then one call to the oracle “does M belong to $\text{Cl}_\wedge(\text{Mod}(\psi \circ \mu))$?”. This oracle is in NP. Indeed, according to Proposition 11 in order to decide whether M belongs to $\text{Cl}_\wedge(\text{Mod}(\psi \circ \mu))$ we have to guess k pairs of models $(M_1, N_1), \dots, (M_k, N_k)$, for every i check that M_i is a model of ψ , N_i is a model of μ and $d(M_i, N_i) = d$, and finally check that $M = M_1 \wedge \dots \wedge M_k$.

In the case of $\text{MODEL-CHECKING}(\circ_D, \mathcal{L}_{\text{Horn}}, \text{Min}_\wedge)$ we need more calls to oracles: “Is $\text{Mod}(\psi \circ \mu)$ closed under \wedge ?”. If yes, call to the oracle “does M belong to $\text{Mod}(\psi \circ \mu)$?”. Otherwise call to the oracle “ $M = \text{Min}_\leq(\text{Mod}(\psi \circ \mu))$?”. Since the distance between ψ and μ is known all these oracles are in NP (thanks to Proposition 11 for the first two ones). \square

Refined Satoh operator for (dual) Horn. We have similar results for Satoh’s operator in the Horn fragment.

Proposition 13. $\text{MODEL-CHECKING}(\circ_S, \mathcal{L}_{\text{Horn}}, f)$ is NP-hard for any \wedge -mapping f .

Proof. We can use the reduction in (Liberatore and Schaefer 2001, Thm 20). From a CNF-formula Π , the authors construct two Horn formulas ψ and μ , and a model M , s.t. Π is satisfiable iff $M \models \psi \circ_S \mu$. Formula μ has only two models M_1 and M_2 with $M_1 \subset M_2$. Any subset of $\{M_1, M_2\}$ is thus closed under any Boolean function f . Therefore, $\text{Mod}(\psi \circ_S \mu)$ is closed under any such f , and $\psi \circ_S^f \mu = \psi \circ_S \mu$. The same proof thus also shows the NP-hardness of $\text{MODEL-CHECKING}(\circ_S, \mathcal{L}_{\text{Horn}}, f)$. \square

Proposition 14. MODEL-CHECKING($\circ_S, \mathcal{L}_{Horn}, Cl_\wedge$) is NP-complete.

Proof. By Proposition 13 only membership has to be proved. According to Proposition 11 in order to decide whether M belongs to $Cl_\beta(\text{Mod}(\psi \circ \mu))$ we have to guess k pairs of models $(M_1, N_1), \dots, (M_k, N_k)$ such that for every i , N_i certifies that M_i is indeed a model of $\psi \circ_S \mu$ and finally check that $M = M_1 \wedge \dots \wedge M_k$. Verifying that N_i is a witness of the fact that $M_i \in \text{Mod}(\psi \circ_S \mu)$ comes down to verifying that M_i is indeed a model of μ , N_i is a model of ψ and there exist no pairs (M'_i, N'_i) with $N'_i \Delta M'_i \subset N_i \Delta M_i$. This last step can be performed in polynomial time if ψ and μ are Horn formulas. Indeed, this test is equivalent to verifying that the formula $\phi = \psi[X/Y] \wedge \mu \wedge \bigwedge_{x_j \notin N \Delta M} (x_j \equiv y_j)$ has no other solution than $N_Y \cup M$, where N_Y denotes the interpretation on the set Y defined by $N_Y(y_j) = I(x_j)$. This verification can be done in polynomial time. The correctness of checking the minimality of $N_i \Delta M_i$ via the formula ϕ crucially depends on the closure under intersection that is fulfilled by the set of models of a Horn formula. This property is needed to show that if ϕ has yet another model, then $N_i \Delta M_i$ is not minimal. \square

Conclusion

This paper contributes to the current line of research in belief change where particular fragments of propositional logic are considered as source and target language. In contrast to previous work which mainly was devoted to the case of Horn logic, we provided here a more general view which takes semantic properties of the language fragments into account. Our main goal was to understand to which extent established revision operators can be “refined” to work in particular fragments. As we have shown, this works well for the basic postulates while the more involved postulates (R5 and R6) are more problematic. We have illustrated that our generic framework captures many natural approaches of refinements of operators (we provided four concrete such operators) and thus can be used to analyze further proposals for concrete operator refinements. Finally, we have complemented our work with a preliminary complexity analysis.

Future work includes a more thorough investigation of the complexity of the refined revision operators. We thus want to extend our study both, to further revision operators and to further fragments of propositional logic, in particular to the Krom case. An important aspect, which has been neglected so far, is how to obtain the concrete result of a revision in terms of a concrete formula from the fragment under consideration. The work by Zanuttini and Hébrard (2002) provides valuable techniques which we want to further adapt for our purposes. Furthermore a full picture under which circumstances R5 can be satisfied is on our agenda. Another direction is to weaken the properties which we suggested for refined operators in Definition 2; indeed, giving up the property of invariance would allow us to define refinements which satisfy all postulates, but it is questionable whether those instances can still be understood as refinements of a given operator. Finally, we plan to apply the methodology presented here for revision to other major operations in the

area of belief change, in particular to contraction and merging.

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.
- Booth, R.; Meyer, T.; and Varzinczak, I. J. 2009. Next steps in propositional horn contraction. In *Proc. IJCAI*, 702–707.
- Dalal, M. 1988. Investigations into theory of knowledge base revision. In *Proc. AAAI*, 449–479.
- Delgrande, J. P., and Peppas, P. 2011. Revising horn theories. In *Proc. IJCAI*, 839–844.
- Delgrande, J. P., and Wassermann, R. 2010. Horn clause contraction functions: Belief set and belief base approaches. In *Proc. KR*, 143–152.
- Delgrande, J. P. 2008. Horn clause belief change: Contraction functions. In *Proc. KR*, 156–165.
- Dubois, D. 2011. Information fusion and revision in qualitative and quantitative settings - steps towards a unified framework. In *Proc. ECSQARU*, 1–18.
- Eiter, T., and Gottlob, G. 1992. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artif. Intell.* 57(2-3):227–270.
- Eiter, T., and Gottlob, G. 1996. The complexity of nested counterfactuals and iterated knowledge base revisions. *J. Comput. Syst. Sci.* 53(3):497–512.
- Horn, A. 1951. On sentences which are true of direct unions of algebras. *Journal of Symbolic Logic* 16:14–21.
- Kadin, J. 1989. $P^{NP[O(\log n)]}$ and sparse Turing-complete sets for NP. *J. Comput. Syst. Sci.* 39(3):282–298.
- Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52:263–294.
- Langlois, M.; Sloan, R. H.; Szörényi, B.; and Turán, G. 2008. Horn complements: Towards horn-to-horn belief revision. In *Proc. AAAI*, 466–471.
- Liberatore, P., and Schaefer, M. 2001. Belief revision and update: Complexity of model checking. *J. Comput. Syst. Sci.* 62(1):43–72.
- Liberatore, P. 2000. Compilability and compact representations of revision of horn knowledge bases. *ACM Trans. Comput. Log.* 1(1):131–161.
- Nebel, B. 1994. Base revision operations and schemes: Semantics, representation and complexity. In *Proc. ECAI*, 341–345.
- Nebel, B. 1998. How hard is it to revise a belief base? In Dubois, D.; Prade, H.; Gabbay, D. M.; and Smets, P., eds., *Belief Change*, volume 3 of *Handbook of Defeasible Reasoning and Uncertainty Management Systems*. Springer Netherlands. 77–145.
- Satoh, K. 1988. Nonmonotonic reasoning by minimal belief revision. In *Proc. FGCS*, 455–462.
- Schaefer, T. J. 1978. The complexity of satisfiability problems. In *Proc. STOC*, 216–226. ACM Press.

- Van De Putte, F. 2011. Prime implicates and relevant belief revision. *J. Log. Comput.*
- Wu, M.; Zhang, D.; and Zhang, M. 2011. Language splitting and relevance-based belief change in horn logic. In *Proc. AAAI*, 268–273.
- Zanuttini, B., and Hébrard, J.-J. 2002. A unified framework for structure identification. *Inf. Process. Lett.* 81(6):335–339.
- Zhuang, Z. Q., and Pagnucco, M. 2011. Transitively relational partial meet horn contraction. In *Proc. IJCAI*, 1132–1138.