

Representation Results for Belief Update in Closed Fragments of Propositional Logic

NADIA CREIGNOU, Aix Marseille Université, CNRS, LIS, France

ADRIAN HARET*, Royal Holloway, University of London, United Kingdom

ODILE PAPINI, Aix Marseille Université, CNRS, LIS, France

STEFAN WOLTRAN, TU Wien, Austria

Fragments of propositional logic, i.e., tailored sub-languages designed for neatly structured data, are relevant in many practical settings. This paper studies *belief update* in fragments (e.g., Horn, Krom, affine) that obey a desirable *semantic closure* condition. We assume update is guided by the well-known Katsuno-Mendelzon (KM) postulates, which in full propositional logic characterize update operators as choice functions guided by total or partial preorders over possible worlds. Because many useful fragments cannot express every connective (e.g., they often lack closure under disjunction), the KM axioms must be rephrased and supplemented to keep updates rational in these less expressive environments. Our main result is a set of representation theorems: once the KM postulates are adjusted, they capture exactly the update operators generated by suitably constrained total or partial preorders within the fragment. In addition, we clarify how *revision* works in fragments when partial preorders are allowed and also present concrete, fragment-friendly update operators.

JAIR Associate Editor: Martin Gebser

JAIR Reference Format:

Nadia Creignou, Adrian Haret, Odile Papini, and Stefan Woltran. 2026. Representation Results for Belief Update in Closed Fragments of Propositional Logic. *Journal of Artificial Intelligence Research* 85, Article 47 (April 2026), 29 pages. DOI: [10.1613/jair.1.20040](https://doi.org/10.1613/jair.1.20040)

1 Introduction

Our paper focuses on *belief update*, a type of belief change that keeps an agent's beliefs in step with a changing world. First proposed in the context of deductive databases (Fagin et al. 1983), belief update was later absorbed into the larger theory of logical belief change, which lays out formal rules for revising beliefs in light of new information. This framework unifies various belief change operators under a common logical structure, using both syntactic constraints and semantic models to define what counts as a rational change operator (Katsuno and Mendelzon 1991a,b).

Belief update is closely related to *belief revision*, but is motivated by a different purpose. Revision aims to incorporate new information in a static world, adjusting beliefs to better reflect the true state while preserving as much of the prior belief as possible. By contrast, update assumes the world itself is changing, sometimes forcing the agent to discard more of its prior beliefs. Formally, if an agent's beliefs are captured by a propositional formula ψ , revision selects models of the new information μ that are closest, based on some distance measure,

*Corresponding Author.

Authors' Contact Information: Nadia Creignou, ORCID: [0009-0005-6522-6363](https://orcid.org/0009-0005-6522-6363), nadia.creignou@univ-amu.fr, Aix Marseille Université, CNRS, LIS, Marseille, France; Adrian Haret, ORCID: [0000-0002-0247-5590](https://orcid.org/0000-0002-0247-5590), adrian.haret@rhul.ac.uk, Royal Holloway, University of London, London, United Kingdom; Odile Papini, ORCID: [0000-0002-7537-5822](https://orcid.org/0000-0002-7537-5822), odile.papini@univ-amu.fr, Aix Marseille Université, CNRS, LIS, Marseille, France; Stefan Woltran, ORCID: [0000-0003-1594-8972](https://orcid.org/0000-0003-1594-8972), woltran@dbai.tuwien.ac.at, TU Wien, Vienna, Austria.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2026 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.20040](https://doi.org/10.1613/jair.1.20040)

to the models of ψ . Update, on the other hand, transforms ψ bit by bit, aligning its models with those of μ in response to changes in the environment.

Example 1. Consider a smarthome assistant that not only follows the user’s preferred settings but also adapts to their behavior. The assistant tracks three variables: whether the temperature is above 15° (a), whether the Wi-Fi is on after 21:00 (b), and whether the user’s family is online after 21:00 (c). The user’s initial instructions are simple: keep the temperature above 15° and switch off the Wi-Fi after 21:00, captured by $\psi = a \wedge \neg b$. Over time, the assistant notices a pattern: the user sometimes turns the Wi-Fi back on after 21:00, typically when the family is online. This behavioral rule is represented as $\mu = (b \leftrightarrow c)$. To adapt, the assistant updates its instructions: intuitively, the desirable outcome is a change from ψ to $\psi' = a \wedge (b \leftrightarrow c)$, meaning the assistant should maintain the temperature and enable Wi-Fi exactly when the family is online. Note that revision would fail here: since μ is consistent with ψ , a standard revision operator would produce $\psi \wedge \mu \equiv a \wedge \neg b \wedge \neg c$, which absurdly concludes that if the Wi-Fi is off, the family members must be offline.

The scenario in Example 1 shows that clinging to prior beliefs is not always appropriate: sometimes, update is needed instead of revision. Formally, the distinction between update and revision is captured by logical postulates grounded in choice mechanisms based on plausibility rankings over possible world states (Katsuno and Mendelzon 1991a). We follow this approach here and treat the classic representation theorems as the benchmark for evaluating any new proposal.

Example 1 also brings out a second key topic: belief change depends on the formalism used to express beliefs. The smarthome assistant reasons in terms of simple rules (if this, then that), reflecting the fact that agents often operate under limited expressive power and computational resources. In many cases, the belief language is intentionally restricted, either to reduce cognitive load for the user or to keep reasoning computationally feasible. As a result, the agent’s epistemic states, i.e., prior beliefs, new information, and their updated form, must conform to a predefined format.

Such concern for the practical aspects of belief change has led to an interest in belief change in languages, or *fragments*, that are less expressive than full propositional logic. In this paper, we focus on fragments such as Horn, Krom, affine, and 1CNF, each defined by its own semantic closure property. These fragments are of particular interest for two main reasons: (i) they often allow for efficient reasoning, such as tractable satisfiability checking, and (ii) they correspond closely to formalisms widely used in knowledge representation, including logic programming, databases, and description logics. As a result, understanding belief change in these fragments provides valuable insights for designing semantic approaches to change in practical applications (Delgrande, Schaub, et al. 2013; Kharlamov et al. 2013; Zhuang, Z. Wang, K. Wang, and Qi 2016).

However, moving from full propositional logic to a less expressive fragment introduces two main challenges. First, the standard change postulates become weaker in fragments, which allows for undesirable behavior by update operators. As a result, classic representation theorems may no longer hold. Second, many standard update operators from propositional logic are not suitable for fragments, because their outputs cannot always be expressed within the restricted language of the fragment. Thus, special care is required to find the right postulates for each fragment of interest, and new types of operators are needed to provide outputs of the required format. Such problems occur for all types of belief change operators in fragments, and we encounter them in the case of update as well.

Contributions. We lay out how representation results and standard update operators break down in fragments less expressive than propositional logic, a conclusion that mirrors previous research in fragment-based belief change (Creignou, Ktari, et al. 2018; Delgrande and Peppas 2015; Delgrande, Peppas, and Woltran 2018; Haret, Rümmele, et al. 2017). We address this issue by tailoring the core set of update postulates to the fragments we are working with, weakening them as needed and adding additional postulates to make up for the loss of information

incurred by switching to fragments. The additional postulates revolve around the semantic property of Suzumura consistency (Suzumura 1976). Familiar in the social choice literature, this property proves to be crucial in ensuring transitivity of the corresponding relations on truth-value assignments. Thus, rather than handling each fragment in a piecemeal fashion, we find a formulation that works for any fragment characterized by a *single semantic closure* property, and obtain two representation results for update in closed fragments: one using total preorders, the other using partial preorders.

Building on these results, we provide concrete update operators that fit into the fragments. The key property is *compliance*, which ensures that the rankings are suitable for producing an output representable in the fragment. Our contribution here lies in finding ways to make the rankings compliant.

Finally, we use the insights from update to provide two representation results for *revision* in fragments: one using total preorders, the other using partial preorders. The result for total preorders rehearses a known result (Delgrande and Peppas 2015; Delgrande, Peppas, and Woltran 2018), but we reformulate it here with explicit focus on the role of Suzumura consistency. The result for partial preorders is new.

Our work extends and corrects errors in previous work (Creignou, Haret, et al. 2018). The most significant change is the generalization of our results to general closed fragments, rather than just the Horn fragment. We also clarify the role of Suzumura consistency in the results, and fix the restriction on fragments without which the representation result for partial preorders does not go through.

Related work. Historically, the Horn fragment is the propositional fragment that has received most attention in belief change, with results on all major operators. For contraction, efforts to find suitable Horn operators tend to focus on defining AGM-style operators whose results remain Horn, with the main approach being to adapt standard AGM constructions (Booth, Meyer, Varzinczak, and Wassermann 2011; Booth, Meyer, and Varzinczak 2009; Delgrande 2008; Delgrande and Wassermann 2013, 2010; Zhuang and Pagnucco 2012; Zhuang and Pagnucco 2014; Zhuang, Pagnucco, and Zhang 2017; Zhuang, Z. Wang, K. Wang, and Delgrande 2015). One general lesson emerging from this line of work is that Horn contraction operators must sometimes introduce non-Horn formulas, or violate some of the AGM postulates. The typical workaround is to find operators that satisfy modified postulates, obtained by dropping or weakening some of the more problematic AGM postulates.

Revision is typically thought to be definable from contraction via the Levi identity, but in languages that lack some of the usual connectives this equivalence is harder to get off the ground (Zhuang, Pagnucco, and Zhang 2013), hence revision requires attention on its own. Notably, two additions to the AGM framework (an acyclicity postulate and a constraint on the total preorders used in the representation) ensure that outputs are Horn-expressible, and deliver a representation result (Delgrande and Peppas 2015). A similar approach works for Horn belief merging (Haret, Rümmele, et al. 2017). Remarkably, this strategy can be extended to a broad class of languages with limited expressive means (Delgrande, Peppas, and Woltran 2018; Zhuang, Z. Wang, K. Wang, and Delgrande 2019), and provides a template for our previous work on Horn update (Creignou, Haret, et al. 2018), which is the immediate precursor to the results here.

Belief change has also been studied for fragments of propositional logic other than Horn (e.g., dual Horn, Krom and affine), often motivated by Schaefer’s classification of satisfiability subclasses (Schaefer 1978). A notable approach uses a post-processing step to translate the (possibly non-fragment) output of a classical revision operator back into the target syntactic class, achieving closure at the cost of (occasionally) violating standard change postulates (Creignou, Papini, Pichler, et al. 2014). As such, it provides essential insight into the trade-offs involved in balancing expressibility and the satisfaction of postulates, and has been successfully extended to merging (Creignou, Papini, Rümmele, et al. 2016), contraction and erasure (Creignou, Ktari, et al. 2022) and update (Creignou, Ktari, et al. 2018; Ktari 2016).

A conspicuous application of fragment-based belief change has been to model the dynamics of argumentation frameworks (Delobelle et al. 2016; Diller et al. 2018; Haret, Wallner, and Woltran 2018), whose semantics can

be thought of as a fragment of propositional logic, logic programs under the answer set semantics (Delgrande, Schaub, et al. 2013; Zhuang, Delgrande, et al. 2016), and preference orders (Haret and Wallner 2022). Beyond propositional fragments, a major line of work studies belief change in description logics (DLs), which are decidable fragments of first-order logic used to formalize ontologies (Aiguier et al. 2018; Kharlamov et al. 2013; Qi and Du 2009; Qi, Liu, et al. 2006; Zhuang, Z. Wang, K. Wang, and Delgrande 2015; Zhuang, Z. Wang, K. Wang, and Qi 2016), and for general logics with minimal requirements on their expressive means (Falakh et al. 2025).

The choice-theoretic perspective used implicitly by most semantic approaches to belief change (and highlighted explicitly by us here) sees change operators as choice functions on possible worlds, subject to postulates that mirror classical rationality axioms (Bossert and Suzumura 2010; Sen 1971). While this connection has been remarked upon before (Bonanno 2009; Haret 2020; Rott 2001), there is likely ample space left to explore its full consequences. Most relevant to our current results is *Suzumura consistency* (Suzumura 1976), which we use as a guarantee that a (partial) relation extends to a transitive order when the fragment’s limited connectives make total preorders infeasible. Partial preorders have some precedent in belief change, showing up in the original KM proposal (Katsuno and Mendelzon 1991b), as well as in subsequent work (Benferhat et al. 2005; Ma et al. 2012; Sérayet et al. 2009). However, apart from our own previous work on Horn update (Creignou, Haret, et al. 2018), they have yet to make their way into work on fragments.

Outline. In Section 2 we outline the belief update model, together with the main existing results we want to emulate in fragments, as well as the fragments of interest to us in the rest of the paper. In Section 3 we lay down update postulates adapted to the fragments and work out representation results. In Section 4 we put forward concrete operators for the fragments we work with. In Section 5 we adapt the results in Section 3 to revision in fragments. We conclude with general comments and thoughts for future work in Section 6.

2 Preliminaries

We start by introducing the main notions used throughout the rest of the paper.

Propositional logic. We assume a finite set \mathcal{A} of propositional atoms. The language of *propositional logic* is the set \mathcal{L} of formulas constructed with the atoms in \mathcal{A} using the usual propositional connectives (\wedge , \vee , \neg , \rightarrow and \leftrightarrow) and the constants \perp and \top . An *interpretation*, or *truth-value assignment*, w is a function mapping every atom in \mathcal{A} to either *true* or *false*. In the interest of readability, we typically write w as a word consisting of the atoms set to true. The *universe* \mathcal{U} is the set of all interpretations for formulas in \mathcal{L} . The *symmetric difference* $w_1 \Delta w_2$ is defined as $w_1 \Delta w_2 = (w_1 \setminus w_2) \cup (w_2 \setminus w_1)$, i.e., as the set of atoms on which w_1 and w_2 differ. The models of a propositional formula φ are the interpretations that satisfy it, and we write $[\varphi]$ for the set of models of φ . If φ_1 and φ_2 are propositional formulas, we say that φ_1 *entails* φ_2 , written $\varphi_1 \models \varphi_2$, if $[\varphi_1] \subseteq [\varphi_2]$, and that they are *equivalent*, written $\varphi_1 \equiv \varphi_2$, if $[\varphi_1] = [\varphi_2]$. A propositional formula φ is *consistent* if $[\varphi] \neq \emptyset$. The models of \perp and \top are $[\perp] = \emptyset$ and $[\top] = \mathcal{U}$, i.e., \perp has no model and \top is satisfied by every interpretation in the universe. A formula φ is *complete* if for any formula ψ , either $\varphi \models \psi$ or $\varphi \models \neg\psi$. Equivalently, a formula φ is complete if it has exactly one model. If M is a set of interpretations and φ is a formula, we say that φ *represents* M and, alternatively, that M is *represented* by φ , if $[\varphi] = M$.

Closed Fragments of Propositional Logic. A *literal* ℓ is either an atom p , in which case ℓ is a *positive literal*, or its negation $\neg p$, in which case ℓ is a *negative literal*. The basic building block for the fragments we are interested in is a *clause*, which is a disjunction of literals. A *Horn clause* is a clause that has at most one positive literal. A *Krom clause* consists of at most two literals. An \oplus -*clause* is an exclusive disjunction of literals. A *unit clause* is a clause that consists of only one literal. The Horn, Krom, affine and 1CNF fragments, denoted $\mathcal{L}_{\text{Horn}}$, $\mathcal{L}_{\text{Krom}}$, $\mathcal{L}_{\text{affine}}$ and $\mathcal{L}_{\text{1CNF}}$, respectively, are the sets of all formulas in \mathcal{L} that are conjunctions of Horn, Krom, \oplus -clauses and unit clauses, respectively.

It is well known that the Horn fragment consists of formulas whose sets of models are closed under intersection. It turns out that the other fragments afford similar semantic characterizations, using appropriate closure operators. We define a *closure operator* as a mapping $cl: 2^{\mathcal{U}} \rightarrow 2^{\mathcal{U}}$ such that, for any sets of interpretations $M, N \subseteq \mathcal{U}$ it holds that: (i) if $M \subseteq N$, then $cl(M) \subseteq cl(N)$; (ii) if $|M| = 1$, then $cl(M) = M$; and (iii) $cl(\emptyset) = \emptyset$. If there is no danger of ambiguity we omit the brackets around sets, i.e., $cl(\{w_1, w_2\})$ is written as $cl(w_1, w_2)$.

We are interested in a particular set of closure operators, defined using the following functions. The *ternary majority function* $maj_3(w_1, w_2, w_3)$ yields an interpretation containing those atoms true in at least two of w_1, w_2 and w_3 . The *ternary XOR function* $\oplus_3(w_1, w_2, w_3)$ yields an interpretation containing the atoms true in one or three out of w_1, w_2 and w_3 . We now extend these functions to a series of operations on sets of interpretations:

$$\begin{aligned} int(M) &= \{w_1 \cap w_2 \mid w_1, w_2 \in M\}, \\ maj(M) &= \{maj_3(w_1, w_2, w_3) \mid w_1, w_2, w_3 \in M\}, \\ \oplus_3(M) &= \{\oplus_3(w_1, w_2, w_3) \mid w_1, w_2, w_3 \in M\}, \\ fill(M) &= \{w_1 \cap w_2, w_1 \cup w_2 \mid w_1, w_2 \in M\} \cup \{w_3 \mid w_1 \subseteq w_3 \subseteq w_2 \text{ and } w_1, w_2 \in M\}, \end{aligned}$$

for any set of interpretations $M \subseteq \mathcal{U}$. The closure operators we focus on in this paper are the *fixed points* of these functions: it is straightforward to check that the fixed point of each of the int , maj , \oplus_3 and $fill$ functions, as defined above, fit the definition of a closure operator.

If cl_* is a closure operator, a set \mathcal{L}_* of propositional formulas is a cl_* -fragment if it holds that:

- (i) $[\varphi] = cl_*([\varphi])$, for all $\varphi \in \mathcal{L}_*$;
- (ii) if $M \subseteq \mathcal{U}$ and $cl_*(M) = M$, then there exists $\varphi \in \mathcal{L}_*$ with $[\varphi] = M$;
- (iii) if $\varphi, \psi \in \mathcal{L}_*$, then $\varphi \wedge \psi \in \mathcal{L}_*$.

In other words, \mathcal{L}_* is a cl_* -fragment if the set of interpretations closed under cl_* are exactly those sets represented by some formula in \mathcal{L}_* and, in addition, \mathcal{L}_* is closed under intersection.

Notably, the propositional fragments of interest to us can be thought of as closed fragments under the closure operators introduced earlier. Specifically, the Horn fragment \mathcal{L}_{Horn} is closed with respect to the cl_{Horn} operator, defined as the fixed point of the int function; the Krom fragment \mathcal{L}_{Krom} is closed with respect to cl_{Krom} , the fixed point of maj ; the affine \mathcal{L}_{affine} fragment is closed with respect to cl_{affine} , the fixed point of \oplus_3 ; and the 1CNF fragment \mathcal{L}_{1CNF} is closed with respect to cl_{1CNF} , the fixed point of $fill$. Note that the full set of propositional formulas \mathcal{L} is also a closed fragment, under the identity closure operator $id(M) = M$.

None of the closed fragments defined above, apart from full propositional logic, is fully expressive: for any such closed fragment \mathcal{L}_* there are sets of interpretations M not representable by any formula in \mathcal{L}_* . However, $cl_*(M)$, where cl_* is a closure operator corresponding to \mathcal{L}_* , is the set of models of some formula in \mathcal{L}_* . Borrowing some terminology (Borchmann et al. 2020; Dechter and Pearl 1992; Kautz et al. 1995), we refer to such a formula as the \mathcal{L}_* -approximation, or \mathcal{L}_* -envelope of M , and denote it by $\varepsilon_*(M) \in \mathcal{L}_*$. Thus, $\varepsilon_*(M)$ is an \mathcal{L}_* -formula such that $[\varepsilon_*(M)] = cl_*(M)$. Intuitively, $\varepsilon_*(M)$ allows us to approximate M in the fragment \mathcal{L}_* through the set $cl_*(M)$, which is ‘close’ to M while still being representable in \mathcal{L}_* . In the interest of readability, and since the fragment is usually clear from the context, we write simply ε_M .

Example 2. The set $M = \{ab, ac, bc\}$ of interpretations is not closed under any of the operators considered (cl_{Horn} , cl_{Krom} , cl_{affine} , or cl_{1CNF}), and is thus not representable in any of the corresponding fragments. For instance, $a = ab \cap ac$ is not an element of M , i.e., M is not closed under intersection and thus is not representable by a Horn formula. However, the closure of M under intersection, given as $cl_{Horn}(M) = M \cup \{\emptyset, a, b, c\}$, can be represented in the Horn fragment by the \mathcal{L}_{Horn} -approximation $\varepsilon_{Horn}(M) = \neg a \vee \neg b \vee \neg c$.

Similarly, we have that $cl_{Krom}(M) = M \cup \{abc\}$, $cl_{affine}(M) = M \cup \{\emptyset\}$ and $cl_{ICNF}(M) = M \cup \{\emptyset, a, b, c, abc\}$. The approximations of M for the target fragments are $\varepsilon_{Krom}(M) = (a \vee b) \wedge (b \vee c) \wedge (a \vee c)$, $\varepsilon_{affine}(M) = (\neg a \oplus b \oplus c)$, and $\varepsilon_{ICNF}(M) = \top$.

By definition, any singleton set of interpretations is representable in all fragments presently considered. Leveraging this fact, we will use ψ_w , or, depending on the context, ε_w , to denote an \mathcal{L}_* -formula such that $[\psi_w] = [\varepsilon_w] = \{w\}$. Recall that, according to our terminology, such a formula is said to be complete. Similarly, for a pair of interpretations w_i, w_j , we write $\varepsilon_{i,j}$ for an \mathcal{L}_* -formula formula such that $[\varepsilon_{i,j}] = cl_*(w_i, w_j)$.

Propositional update. A *propositional update operator* \diamond is a function $\diamond: \mathcal{L} \times \mathcal{L} \rightarrow \mathcal{L}$, mapping a *prior belief* $\psi \in \mathcal{L}$ and *new information* $\mu \in \mathcal{L}$ to a new formula $\psi \diamond \mu \in \mathcal{L}$, referred to as the *updated*, or *posterior belief*. The following core postulates, formulated for any propositional formulas, are typically put forward (Katsuno and Mendelzon 1991a) as constraints on a rational update operator:

- (U₁) $\psi \diamond \mu \models \mu$.
- (U₂) If $\psi \models \mu$, then $\psi \diamond \mu \equiv \psi$.
- (U₃) If ψ and μ are satisfiable, then $\psi \diamond \mu$ is satisfiable.
- (U₄) If $\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$, then $\psi_1 \diamond \mu_1 \equiv \psi_2 \diamond \mu_2$.
- (U₅) $(\psi \diamond \mu_1) \wedge \mu_2 \models \psi \diamond (\mu_1 \wedge \mu_2)$.
- (U₆) If $\psi \diamond \mu_1 \models \mu_2$ and $\psi \diamond \mu_2 \models \mu_1$, then $\psi \diamond \mu_1 \equiv \psi \diamond \mu_2$.
- (U₇) If ψ is complete, then $(\psi \diamond \mu_1) \wedge (\psi \diamond \mu_2) \models \psi \diamond (\mu_1 \vee \mu_2)$.
- (U₈) $(\psi_1 \vee \psi_2) \diamond \mu \equiv (\psi_1 \diamond \mu) \vee (\psi_2 \diamond \mu)$.
- (U₉) If ψ is complete and $(\psi \diamond \mu_1) \wedge \mu_2$ is satisfiable, then $\psi \diamond (\mu_1 \wedge \mu_2) \models (\psi \diamond \mu_1) \wedge \mu_2$.

Intuitively, postulate U₁ states that models of the updated belief have to imply the new information; U₂ states that if μ was a consequence of ψ before update, then the agent's beliefs do not change after update, i.e., inertia is preferred to spontaneous evolution; U₃ states that if the original beliefs and new information are consistent, then update can always be performed; U₄ enforces irrelevance of syntax; U₅ expresses minimality of change. Postulate U₆ says that if updating ψ by μ_1 guarantees μ_2 and updating ψ by μ_2 guarantees μ_1 then the two updates have the same effect. Postulate U₇ says that when ψ is complete, a model of ψ updated by μ_1 and of ψ updated by μ_2 must be a model of ψ updated by $\mu_1 \vee \mu_2$. Postulate U₈ states that an update operator gives each model of the initial beliefs equal consideration. Postulates U₇ and U₈ are especially relevant here, as they need to be adapted in closed in fragments. Finally, postulate U₉ is the converse of U₅, restricted to complete formulas ψ . Note, finally, that postulates U₁₋₅ and U₉ imply postulates U₆ and U₇, and thus the full set U₁₋₉ is stronger than U₁₋₈.

Preorders. It is useful to think of update operators as encoding an implicit choice function over possible worlds ranked in terms of their plausibility, where plausibility is encoded as a binary relation \leq on the set \mathcal{U} of interpretations, with $w \leq w'$ read as saying that w is *at least as plausible* as w' . We write $w < w'$ if $w \leq w'$ and $w' \not\leq w$, and $w \approx w'$ if $w \leq w'$ and $w' \leq w$. A *preorder* \leq on \mathcal{U} is a reflexive and transitive binary relation on \mathcal{U} . A preorder \leq on \mathcal{U} is *total* if $w \leq w'$ or $w' \leq w$ for any $w, w' \in \mathcal{U}$. If $M \subseteq \mathcal{U}$ is a set of interpretations, the set $min_{\leq} M$ of *minimal elements of M with respect to \leq* is defined as:

$$min_{\leq} M = \{w \in M \mid \text{there is no } w' \in M \text{ such that } w' < w\}.$$

Note that the definition of $min_{\leq} M$ works for both total and partial preorders. If \leq is total, then $min_{\leq} M$ is equivalent to the set of interpretations $w \in M$ such that $w \leq w'$, for all $w' \in M$; this equivalence breaks down if \leq is partial and contains incomparable elements.

A *pointwise assignment* maps every interpretation w to a (possibly partial) preorder \leq_w , with $w_1 \leq_w w_2$ read as saying that w_1 is *at least as plausible as w_2 from the point of view of w* . A pointwise assignment *represents*

operator \diamond and, alternatively, \diamond is *represented* by the assignment if, for any formulas ψ and μ , it holds that:

$$[\psi \diamond \mu] = \bigcup_{w \in [\psi]} \min_{\leq_w} [\mu].$$

A preorder \leq_w is *faithful* if, for any interpretation w' , it holds that if $w \neq w'$, then $w <_w w'$, and a pointwise assignment is faithful if every preorder \leq_w is faithful.

The landmark results in propositional belief update state that: (i) any propositional update operator satisfying postulates U_{1-9} can be represented by a pointwise faithful assignment mapping interpretations to total preorders, and (ii) an update operator satisfying the weaker set of postulates U_{1-8} can be represented by a pointwise faithful assignments using *partial* preorders (Katsuno and Mendelzon 1991b). In full propositional logic, postulate U_9 ensures totality of the preorders; in closed fragments, where U_7 and U_8 need to be adapted, we will see that totality can fail even if U_9 is assumed.

Concrete update operators. An important device for generating preorders on interpretations is a function $d: \mathcal{U} \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$, which in a slight abuse of language we call a *distance function*, used to quantify the disagreement between two interpretations. A distance function d is expected to guarantee, at the very least, that $d(w, w) = 0$ and $d(w, w') > 0$ for any distinct interpretations w and w' . Given an interpretation w and a distance d , the ranking \leq_w^d is defined as:

$$w_1 \leq_w^d w_2 \text{ if } d(w, w_1) \leq d(w, w_2),$$

i.e., w_1 is considered at least as plausible as w_2 from the point of view of w if w_1 is at least as close to w as w_2 . For a distance function d , the update operator \diamond^d is defined as a propositional formula $\psi \diamond^d \mu$ such that:

$$[\psi \diamond^d \mu] = \bigcup_{w \in [\psi]} \min_{\leq_w} [\mu].$$

The standard examples for distances between interpretations are the *Hamming distance* d_H and the *drastic distance* d_{dr} , defined, for any interpretations w_1 and w_2 , as:

$$d_H(w_1, w_2) = |w_1 \Delta w_2|, \quad d_{dr}(w_1, w_2) = \begin{cases} 0, & \text{if } w_1 = w_2, \\ 1, & \text{otherwise.} \end{cases}$$

The Hamming distance d_H counts the number of atoms that w_1 and w_2 differ on. The corresponding operator \diamond^H was first introduced in the context of qualitative physics (Forbus 1989; Herzig and Rifi 1999), and is one of the go-to options for illustrating the mechanics of model-based belief change. The drastic distance d_{dr} is coarser, keeping track only of whether w_1 and w_2 are distinct. Note that both \leq_w^H and \leq_w^{dr} are total, pointwise faithful preorders (indeed, any preorder generated with a distance function satisfies these properties), and thus the corresponding operators satisfy postulates U_{1-9} .

Operators based on partial preorders can be generated using non-distance based measures of similarity on interpretations. Winslett's operator (Winslett 1990) is induced by the partial preorders \leq_w^{Win} , defined as:

$$w_1 \leq_w^{Win} w_2 \text{ if } w \Delta w_1 \subseteq w \Delta w_2,$$

i.e., from the point of view of w , w_1 is at least as plausible as w_2 if w and w_2 differ on at least the same atoms as the atoms that w and w_1 differ on. The resulting operator \diamond^{Win} satisfies postulates U_{1-8} , but not U_9 (Ktari 2016).

Example 3. Consider propositional formulas $\psi = a \wedge \neg b$ and $\mu = b \leftrightarrow c$ from Example 1, with $[\psi] = \{a, ac\}$ and $[\mu] = \{\emptyset, a, bc, abc\}$. To get the result of updating ψ by μ with \diamond^H and \diamond^{Win} we start by computing the symmetric differences between each model of ψ and each model of μ . These symmetric differences are depicted in Figure 1a. We then order the models of μ according to these symmetric differences, one order \leq_w for each model w of μ . For \diamond^H , the individual rankings are based on the size of the symmetric differences, with the resulting total preorders

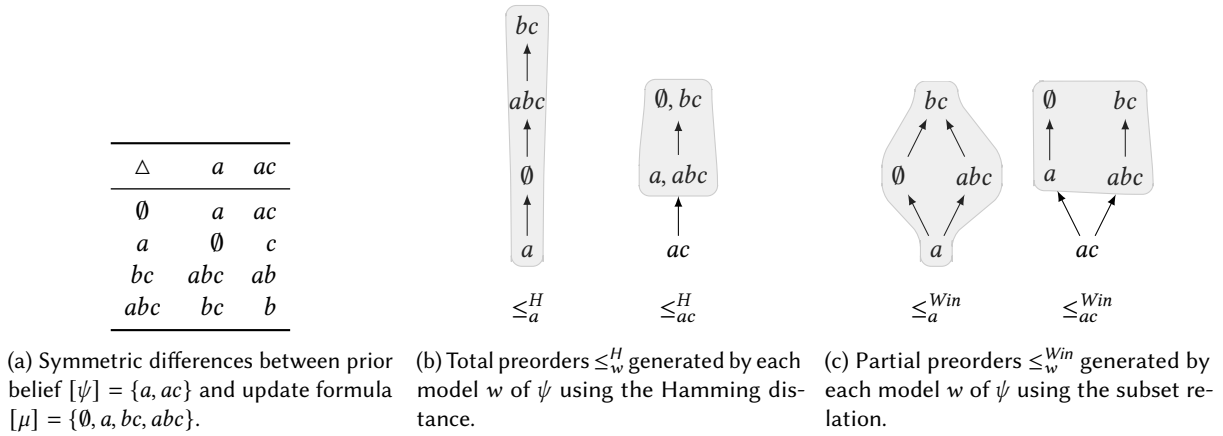


Fig. 1. The standard update operators \diamond^H and \diamond^{Win} work by combining results from the total and partial orders, respectively, generated by each model of ψ . An arrow from interpretation w_1 to w_2 indicates that $w_1 <_w w_2$, i.e., w_1 is considered strictly more plausible than w_2 from the point of view of w . If w_1 and w_2 are on the same level and separated by a comma they are considered equally plausible, i.e., $w_1 \approx_w w_2$. If w_1 and w_2 are depicted apart from each other, then they are incomparable with respect to \leq_w . Models of μ are depicted in the shaded region.

depicted in Figure 1b. The output of \diamond^H consists of the minimal models of μ across preorders \leq_a and \leq_{ac} . We obtain:

$$\begin{aligned}
 [\psi \diamond^H \mu] &= \min_{\leq_a^H} [\mu] \cup \min_{\leq_{ac}^H} [\mu] \\
 &= \{a\} \cup \{a, abc\} \\
 &= \{a, abc\},
 \end{aligned}$$

which can be represented by the formula $a \wedge (b \leftrightarrow c)$. With the drastic distance every interpretation in \leq_w other than w sits at a distance of 1, so the output is $[\psi \diamond^{dr} \mu] = \{a\} \cup \{\emptyset, a, bc, abc\} = [\mu]$.

For \diamond^{Win} the rankings are based on the subset relation between the symmetric differences, resulting in the partial preorders \leq_w^{Win} in Figure 1c. We obtain $[\psi \diamond^{Win} \mu] = \{a, abc\}$.

3 Update in Closed Fragments

If $\mathcal{L}_* \subseteq \mathcal{L}$ is a closed fragment of propositional logic, an \mathcal{L}_* -update operator \diamond is a function $\diamond: \mathcal{L}_* \times \mathcal{L}_* \rightarrow \mathcal{L}_*$. Our aim in this section is to understand and characterize the class of rational \mathcal{L}_* -update operators in terms of pointwise faithful assignments that represent them. Two challenges arise straightaway, one pertaining to the formulation of the postulates, the other pertaining to the nature of the preorders used by the assignments.

Postulates for update in closed fragments. The natural place to start in order to axiomatize \mathcal{L}_* -update operators is to use the set of postulates U_{1-9} , with the obvious caveat that they should be restricted to \mathcal{L}_* -formulas. This approach is standard in other areas of fragment-based belief change, e.g., Horn revision (Delgrande and Peppas 2015), and we embrace it here as well. Thus, we use U_i^* to denote an update postulate applied exclusively to \mathcal{L}_* -formulas. Next, we propose the following postulates, for any \mathcal{L}_* -formulas $\psi, \psi_1, \psi_2, \mu$ and $\mu_i, i \in \{1, \dots, n\}$, such that $\mu \equiv \bigvee_{i=1}^n \mu_i$, i.e., μ is an \mathcal{L}_* -formula equivalent to the disjunction of any number of \mathcal{L}_* -formulas:

$$(U_1^*) \quad \psi \diamond \mu \models \mu.$$

$$(U_2^*) \quad \text{If } \psi \models \mu, \text{ then } \psi \diamond \mu \equiv \psi.$$

- (U₃^{*}) If ψ and μ are satisfiable, then $\psi \diamond \mu$ is satisfiable.
- (U₄^{*}) If $\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$, then $\psi_1 \diamond \mu_1 \equiv \psi_2 \diamond \mu_2$.
- (U₅^{*}) $(\psi \diamond \mu_1) \wedge \mu_2 \models \psi \diamond (\mu_1 \wedge \mu_2)$.
- (U₆^{*}) If $\psi \diamond \mu_1 \models \mu_2$ and $\psi \diamond \mu_2 \models \mu_1$, then $\psi \diamond \mu_1 \equiv \psi \diamond \mu_2$.
- (U₇^{*}) If ψ is complete and $\mu \equiv \bigvee_{i=1}^n \mu_i$, then $\bigwedge_{i=1}^n (\psi \diamond \mu_i) \models \psi \diamond \mu$.
- (U₈^{*}) $\psi \diamond \mu \equiv \bigvee_{w \in [\psi]} (\psi_w \diamond \mu)$.
- (U₉^{*}) If ψ is complete and $(\psi \diamond \mu_1) \wedge \mu_2$ is satisfiable, then $\psi \diamond (\mu_1 \wedge \mu_2) \models (\psi \diamond \mu_1) \wedge \mu_2$.

Some comments are immediately in order. Other than being restricted to \mathcal{L}_* -formulas, postulates U₁^{*} – U₆^{*} and U₉^{*} are formulated exactly like their propositional counterparts. Postulates U₇^{*} and U₈^{*} are based on the same intuitions as the standard, propositional updates, but make some concessions to the reality of working in weaker fragments.

Postulate U₇^{*} is formulated to apply *only* to \mathcal{L}_* -formulas μ that are equivalent to the disjunction of (any number of) \mathcal{L}_* -formulas. The motivation for this rewriting comes from the desire to preserve the original intention of postulate U₇, which is the idea that $\psi \diamond \mu$ preserves the outcome of decomposing μ into distinct sub-components μ_i that, together, cover μ . A straightforward way to do this in a general manner would be to decompose μ into the ‘smallest’ sub-components that can be expressed in \mathcal{L}_* . The following example illustrates this.

Example 4. Consider a formula μ with $[\mu] = \{\emptyset, a, b, c\}$, and take a complete formula ψ . In propositional logic we can decompose μ into the disjunction of $[\mu_1] = \{\emptyset, a, b\}$ and $[\mu_2] = \{a, c\}$, and apply postulate U₇ to $\psi \diamond \mu_1$, $\psi \diamond \mu_2$ and $\psi \diamond \mu$: if, for instance, a is a model of $\psi \diamond \mu_1$ and $\psi \diamond \mu_2$, then postulate U₇ guarantees that a is also a model of $\psi \diamond \mu$. This does not work in the Horn fragment since, even though μ and μ_1 are Horn formulas, μ_2 is not. What we can do in the Horn fragment is to decompose μ into the disjunction of $[\mu_1] = \{\emptyset, a, b\}$, and $[\mu_3] = \{\emptyset, a, c\}$, the latter now being a Horn formula.

Example 4 shows that we may need to decompose a formula μ into sub-formulas μ_i that, together, cover μ . The following result shows that there is a general way to do this in any closed fragment.

Lemma 1. If \mathcal{L}_* is a closed fragment of propositional logic, then, for any $\mu \in \mathcal{L}_*$ such that $[\mu] = \{w_1, \dots, w_n\}$, it holds that $\mu \equiv \varepsilon_{1,2} \vee \dots \vee \varepsilon_{1,n}$.

PROOF. Recall that, for any two interpretations w_i and w_j , the \mathcal{L}_* -envelope $\varepsilon_{i,j}$ is an \mathcal{L}_* -formula such that $[\varepsilon_{i,j}] = cl_*(\{w_i, w_j\})$, guaranteed to exist. For one direction, note that, since we have that $\{w_i, w_j\} \subseteq cl_*(\{w_i, w_j\})$, for any two interpretations w_i and w_j , it follows immediately that $[\mu] \subseteq [\varepsilon_{1,2} \vee \dots \vee \varepsilon_{1,n}]$. For the other direction, take $w \in [\varepsilon_{1,j}]$, for some $j \in \{2, \dots, n\}$, which is equivalent to saying that $w \in cl_*(\{w_1, w_j\})$. Since $\{w_1, w_j\} \subseteq [\mu]$, then, by monotonicity of cl_* , we have that $cl_*(\{w_1, w_j\}) \subseteq cl_*([\mu])$, and hence $w \in [\mu]$. \square

By Lemma 1, any \mathcal{L}_* -formula μ is equivalent to the disjunction of \mathcal{L}_* -envelopes of pairs of models of μ , and our tailored postulate U₇^{*} is intended to make use of this option.

Example 5. Consider a Horn formula μ with $[\mu] = \{\emptyset, a, b, c\}$ from Example 4. By Lemma 1, we can decompose μ into the disjunction of $\varepsilon_{a,\emptyset}$, $\varepsilon_{a,b}$ and $\varepsilon_{a,c}$, where $[\varepsilon_{a,\emptyset}] = \{a, \emptyset\}$, $[\varepsilon_{a,b}] = \{a, b, \emptyset\}$ and $[\varepsilon_{a,c}] = \{a, c, \emptyset\}$. We can now use postulate U₇^{*} to conclude that if a is a model of $\psi \diamond \varepsilon_{a,\emptyset}$, $\psi \diamond \varepsilon_{a,b}$ and $\psi \diamond \varepsilon_{a,c}$, then a is also a model of $\psi \diamond \mu$.

A similar motivation underlies our formulation of U₈^{*}, which sees an \mathcal{L}_* -formula ψ decomposed into complete formulas, one for each of ψ 's models. The closed fragments we look at make this possible, since they can all represent individual interpretations.

Before moving on, we show, as a sanity check, that postulates U₇^{*} and U₈^{*} are equivalent to the standard postulates U₇ and U₈, respectively, when the fragment is full propositional logic \mathcal{L} . Recall that propositional logic is a closed fragment under the operator *id*.

Proposition 1. If a propositional update operator \diamond satisfies postulate U_7^{id} (respectively, U_8^{id}), \diamond also satisfies postulate U_7 (respectively, U_8). Conversely, if \diamond satisfies U_7 and U_4 (respectively, U_8 and U_4), then \diamond also satisfies U_7^{id} (respectively, U_8^{id}).

PROOF. Consider, first, postulate U_7^{id} . If \diamond satisfies postulate U_7^{id} , then applying it to any propositional formulas ψ , μ_1 and μ_2 immediately yields U_7 . Conversely, suppose that \diamond satisfies postulates U_7 and U_4 , and take $\mu \equiv \bigvee_{i=1}^n \mu_i$. We use induction on n to show that $\bigwedge_{i=1}^n (\psi \diamond \mu_i) \models \psi \diamond \mu$. For $n = 2$, this follows immediately from U_7 and U_4 . For $n > 2$, we use the fact that $\mu \equiv \mu' \vee \mu_n$ where $\mu' \equiv \bigvee_{i=1}^{n-1} \mu_i$. As per postulate U_7 , we have that $(\psi \diamond \mu') \wedge (\psi \diamond \mu_n) \models \psi \diamond (\mu' \vee \mu_n)$, which, together with postulate U_4 , implies that $(\psi \diamond \mu') \wedge (\psi \diamond \mu_n) \models \psi \diamond \mu$. By the induction hypothesis we have that $\bigwedge_{i=1}^{n-1} (\psi \diamond \mu_i) \models \psi \diamond \mu'$, which, coupled with the previous result, delivers the conclusion.

Consider, next, postulate U_8^{id} and recall that, for any interpretation w , ψ_w is a formula such that $[\psi_w] = \{w\}$. If \diamond satisfies U_8^{id} , we obtain satisfaction of U_8 as follows:

$$\begin{aligned} (\psi_1 \vee \psi_2) \diamond \mu &\equiv \bigvee_{w \in [\psi_1 \vee \psi_2]} (\psi_w \diamond \mu) \\ &\equiv \left(\bigvee_{w \in [\psi_1]} (\psi_w \diamond \mu) \right) \vee \left(\bigvee_{w \in [\psi_2]} (\psi_w \diamond \mu) \right) \\ &\equiv (\psi_1 \diamond \mu) \vee (\psi_2 \diamond \mu). \end{aligned} \quad // \text{ by } U_8^{id}$$

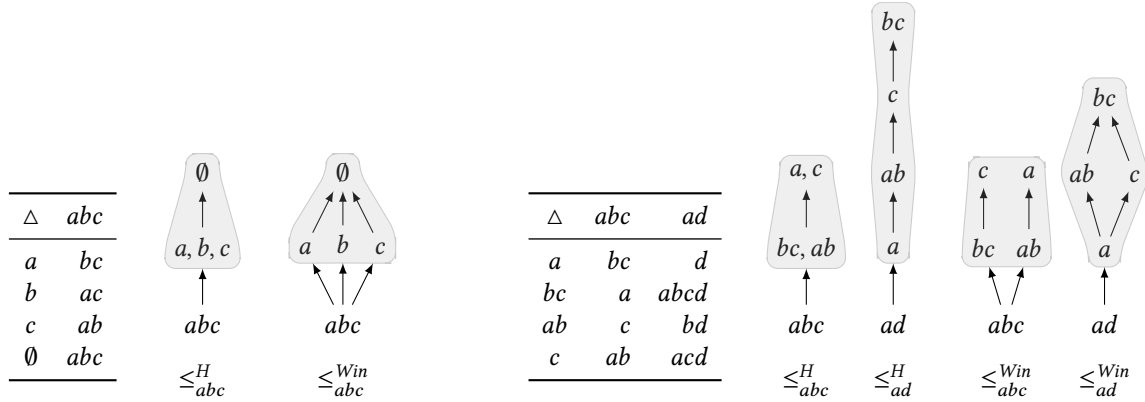
Conversely, suppose that \diamond satisfies postulates U_8 and U_4 . We show, by induction on the number of models of ψ , that $\psi \diamond \mu \equiv \bigvee_{w \in [\psi]} (\psi_w \diamond \mu)$. If ψ has a single model w , then $\psi \equiv \psi_w$ and, by postulate U_4 , it holds that $\psi \diamond \mu \equiv \psi_w \diamond \mu$. Next, suppose the claim holds if for formulas with n models, and take $[\psi] = \{w_1, \dots, w_{n+1}\}$. If ψ' is a formula such that $[\psi'] = \{w_1, \dots, w_n\}$, then $\psi \equiv \psi' \vee \psi_{w_{n+1}}$, and the following equivalences hold:

$$\begin{aligned} \psi \diamond \mu &\equiv (\psi' \vee \psi_{w_{n+1}}) \diamond \mu && // \text{ by } U_4 \\ &\equiv (\psi' \diamond \mu) \vee (\psi_{w_{n+1}} \diamond \mu) && // \text{ by } U_8 \\ &\equiv \bigvee_{w \in [\psi']} (\psi_w \diamond \mu) \vee (\psi_{w_{n+1}} \diamond \mu) && // \text{ by induction hypothesis} \\ &\equiv \bigvee_{w \in [\psi]} (\psi_w \diamond \mu) \end{aligned}$$

This shows that postulate U_8^{id} is satisfied. \square

\mathcal{L}_ -compliance.* The second challenge arises when thinking of update as the result of a semantic, preorder-based procedure. That is, we wish to use the usual semantic procedure and define the update of \mathcal{L}_* -formula ψ by \mathcal{L}_* -formula μ as an \mathcal{L}_* -formula $\psi \diamond \mu$ such that $[\psi \diamond \mu] = \bigcup_{w \in [\psi]} \min_{\leq_w} [\mu]$, where \leq_w is generated by some pointwise faithful assignment. The problem, of course, is that the set of interpretations $\bigcup_{w \in [\psi]} \min_{\leq_w} [\mu]$ might not be representable by an \mathcal{L}_* -formula. This is not a mere theoretical possibility: indeed, it afflicts all the popular update operators.

Example 6. Consider, first, the formula $\psi_1 = a \wedge b \wedge c$, to be updated by $\mu_1 = (\neg a \vee \neg b) \wedge (\neg b \vee \neg c) \wedge (\neg a \vee \neg c)$, with $[\mu_1] = \{a, b, c, \emptyset\}$. Note that ψ_1 and μ_1 are in both the Horn and Krom fragment, and are thus valid inputs to either a Horn- or Krom-update operator. As illustrated in Figure 2a, the standard operators \diamond^H and \diamond^{Win} yield $[\psi_1 \diamond^H \mu_1] = [\psi_1 \diamond^{Win} \mu_1] = \{a, b, c\}$. Since $cl_{Horn}(\{a, b, c\}) = cl_{Krom}(\{a, b, c\}) = \{\emptyset, a, b, c\}$, the output of \diamond^H and \diamond^{Win} in this case does not correspond to any Horn, or Krom, formula.



(a) Preorders for Horn/Krom formulas $[\psi_1] = \{abc\}$ and $[\mu_1] = \{a, b, c, \emptyset\}$. (b) Preorders for affine formulas $[\psi_2] = \{abc, ad\}$ and $[\mu_2] = \{a, bc, ab, c\}$.

Fig. 2. Standard operators \diamond^H and \diamond^{Win} do not stay within the Horn, Krom and affine fragments.

Consider, next, affine formulas ψ_2 and μ_2 such that $[\psi_2] = \{abc, ad\}$ and $[\mu_2] = \{a, bc, ab, c\}$. Such formulas exist, since their sets of models are closed under \oplus_3 . As illustrated in Figure 2b, the operators give $[\psi_2 \diamond^H \mu_2] = [\psi_2 \diamond^{Win} \mu_2] = \{a, ab, bc\}$. Since $cl_{affine}(\{a, ab, bc\}) = \{a, ab, bc, c\}$, there is no affine formula whose models are $\{a, ab, bc\}$.

In response, we focus our attention to fragment-friendly assignments. We say that an assignment is \mathcal{L}_* -compliant if for any \mathcal{L}_* -formulas ψ and μ , the set of interpretations $\bigcup_{w \in [\psi]} \min_{\leq_w} [\mu]$ is representable by an \mathcal{L}_* -formula. Note that \mathcal{L}_* -compliance in this context is stronger than the similar notion used in fragment-based revision (Delgrande and Peppas 2015; Delgrande, Peppas, and Woltran 2018), which requires that the minimal models of μ are \mathcal{L}_* -representable only across individual preorders. This fixes the expressibility problem by postulating it away, but shifts the burden towards ensuring that such assignments can actually be found. As Example 6 shows, the first casualties of focusing exclusively on \mathcal{L}_* -compliant assignments are the existing update operators \diamond^H and \diamond^{Win} , which are not guaranteed to output a formula in most fragments of interest to us.

With the preliminary issues of how to write the postulates and what kind of assignments to focus on out of the way, we can get on to the main concern: representation results.

3.1 Total Preorders

The main issue raised by working in fragments of propositional logic is that they are typically less expressive than full propositional logic (e.g., the Horn fragment is not closed under disjunction). As a consequence, in such fragments the standard postulates turn out to characterize relations on interpretations one would like to avoid, i.e., due to the limited expressiveness of the fragment, it suddenly becomes possible to generate reasonably looking \mathcal{L}_* -update operators from undesirable assignments, which would not be acceptable in the full propositional setting. One such type of structure admits non-transitive cycles between interpretations, which violates the intuitive idea that the relations on interpretations should correspond to some kind of distance notion. The following example illustrates the problem.

Example 7. Consider a Horn formula ψ such that $[\psi] = \{\emptyset\}$ and the binary relation \leq_{\emptyset} depicted in Figure 3. We assume that \leq_{\emptyset} behaves like a total preorder everywhere, except with respect to ab, ac and bc , which it ranks in




Fig. 3. A binary relation \leq_0 that contains a non-transitive cycle between interpretations ab , ac and bc . The cycle goes undetected by the postulates, since the set of interpretations $\{ab, ac, bc\}$ is not representable by a Horn formula.

the non-transitive cycle $ab <_0 ac <_0 bc <_0 ab$. Thus, \leq_0 is not a preorder we want to accept in an assignment representing a Horn update operator. Surprisingly, however, \leq_0 can actually be used in the context of such an assignment. To see this, note that for any Horn formula μ , $\min_{\leq_0}[\mu]$ is representable by a Horn formula:

- If $\emptyset \in [\mu]$, then $\min_{\leq_0}[\mu] = \{\emptyset\}$.
- If $\emptyset \notin [\mu]$ but $ab \in [\mu]$ and $ac, bc \notin [\mu]$, then $\min_{\leq_0}[\mu] = \{ab\}$. Similarly for ac and bc .
- If $\emptyset \notin [\mu]$ but $\{ab, ac\} \subseteq [\mu]$ and $bc \notin [\mu]$, then $\min_{\leq_0}[\mu] = \{ab\}$. Similarly for the other two pairs in $\{ab, bc, ac\}$.
- If $\{ab, ac, bc\} \subseteq [\mu]$, then $\min_{\leq_0}[\mu] = \{\emptyset\}$. This is because μ must be closed under intersection and thus contains \emptyset .
- If none of ab , bc and ac are in $[\mu]$, then $\min_{\leq_0}[\mu]$ is either $\{\emptyset\}$, if any two of $\{a, b, c\}$ or \emptyset are in μ , or $[\mu]$ otherwise.

In all of these cases, $\min_{\leq_0}[\mu]$ is representable by a Horn formula, so \leq_0 can be plugged into a Horn-compliant assignment, and can thus be used to define a Horn update operator \diamond . What is more, it can be checked that \diamond , defined on top of \leq_0 , satisfies all postulates U_{1-9}^{Horn} .

Making matters worse, there is no alternative (legitimate) preorder \leq_0^* that represents the same operator \diamond . The reason is that \leq_0^* would have to agree with \leq_0 on the result of updating ψ with any formula, and in particular with formulas $[\mu_0] = \{ab, ac, a\}$, $[\mu_1] = \{ac, bc, c\}$ and $[\mu_2] = \{ab, bc, b\}$. The preorder \leq_0^* would thus need to guarantee that $[\psi \diamond \mu_0] = \{ab\}$, $[\psi \diamond \mu_1] = \{ac\}$ and $[\psi \diamond \mu_2] = \{bc\}$, which can only happen if the non-transitive cycle is also present in \leq_0^* .

What goes wrong in Example 7, and what does this have to do with the Horn fragment? Note that in propositional logic the undesirable non-transitive cycle from Figure 3 cannot show up: take $[\mu] = \{ab, ac, bc\}$; then, by postulate U_3 , $\psi \diamond \mu$ must have a non-empty set of models, which immediately contradicts \leq_0 , since $\min_{\leq_0}[\mu] = \emptyset$. This simple argument is not available in the Horn fragment, where there is no formula corresponding to $\{ab, ac, bc\}$ and, as shown in Example 7, the cycle goes undetected by the Horn-specific postulates.

To address this problem, we first single out the semantic property involved in generating such a counter-intuitive counterexample. A binary (not necessarily transitive, or total) relation \leq on \mathcal{U} is *Suzumura consistent* (Suzumura 1976) if it satisfies the following condition:

- (SC) If $w_1 \leq \dots \leq w_n$, then $w_n \not\leq w_1$.

Intuitively, Suzumura consistency says that it is not possible to form a chain of comparisons that starts with w_1 , ends with w_n , such that every interpretation along the way is at least as good as the next one, but the last one, w_n , ends up being strictly better than the first one, w_1 .




Fig. 4. Three preorders: \leq_0 is not Suzumura consistent, while \leq'_0 and \leq''_0 (both incomplete) are.

Example 8. Take the binary relation \leq_0 from Example 7, whose restriction to $\{ab, ac, bc\}$ is depicted on the left in Figure 4. Since $ab <_0 ac$, it follows that $ab \leq_0 ac$. Similarly, $ac \leq_0 bc$. But, since $bc <_0 ab$, the relation \leq_0 is not Suzumura consistent.

The other two preorders in Figure 4 are Suzumura consistent: \leq'_0 flattens the cycle by making all three interpretations equally plausible, while \leq''_0 abandons transitivity by leaving ab and bc incomparable. Note that \leq'_0 and \leq''_0 do not represent the same update operator as \leq_0 .

Formally, Suzumura consistency lies in between transitivity and acyclicity, in the sense that transitivity implies Suzumura consistency, which in turn implies acyclicity. It is of crucial importance to us because it turns out to be both a sufficient and necessary condition for the possibility of extending a binary relation to a total preorder. To make this precise, we first introduce a bit of extra terminology. If \leq and \leq' are binary (again, not necessarily transitive, or total) relations on \mathcal{U} , then \leq' *extends* \leq if \leq' contains all the comparisons while preserving all the strict comparisons of \leq . Furthermore, \leq' is an *ordering extension* of \leq if \leq' extends \leq and is a total preorder. The following result establishes the connection between Suzumura consistency and ordering extensions.

Theorem 1 ((Suzumura 1976)). *if \leq is a binary relation on a (potentially infinite) set X , then there exists an ordering extension \leq' of \leq on X if and only if \leq is Suzumura consistent.*

Theorem 1 is the reason Suzumura consistency has been singled out in the rational choice literature as a safe fall-back option when transitivity, for whatever reason, cannot be guaranteed (Bossert and Suzumura 2010).

Returning to belief update, we can rephrase our problem as a fragment-specific failure of postulates U_{1-9}^* to enforce Suzumura consistency of the preorders used to represent the operators: thus, what goes wrong in Example 7 is that \leq_0 is not Suzumura consistent (while, at the same time, representing a postulate-abiding Horn update operator).

To fix this we simply add a new postulate, tailor-made for closed fragments, to formalize the idea that the plausibility rankings on interpretations should be Suzumura consistent. We follow previous usage in formulating the postulate (more precisely, postulate *schema*) in terms of complete formulas (Haret, Rümmele, et al. 2017). Before introducing it recall that, given interpretations w_i and w_j , we denote by ε_i an \mathcal{L}_* -formula such that $[\varepsilon_i] = \{w_i\}$ and $\varepsilon_{i,j}$ an \mathcal{L}_* -formula such that $[\varepsilon_{i,j}] = cl_*(w_i, w_j)$.

(U_{SC}^*) If ψ is a complete \mathcal{L}_* -formula, then, for any $n \geq 1$ and complete \mathcal{L}_* -formulas $\varepsilon_1, \dots, \varepsilon_n$ such that, for all $i \in \{1, \dots, n-1\}$, it holds that $(\psi \diamond \varepsilon_{i,i+1}) \wedge \varepsilon_i$ is consistent, then it does not hold both that $(\psi \diamond \varepsilon_{n,1}) \wedge \varepsilon_n$ is consistent and $(\psi \diamond \varepsilon_{n,1}) \wedge \varepsilon_1$ is inconsistent.

Returning to our problematic example, note that postulate U_{SC}^* detects that there is something suspicious about the non-transitive cycle in Figure 3.

Example 9. Consider the binary relation \leq_0 in Example 7 and complete Horn formulas $[\psi] = \{abc\}$, $[\varepsilon_1] = \{ab\}$, $[\varepsilon_2] = \{ac\}$ and $[\varepsilon_3] = \{bc\}$. We then have that $[\varepsilon_{1,2}] = \{a, ab, ac\}$, $[\psi \diamond \varepsilon_{1,2}] = \min_{\leq_0} [\varepsilon_{1,2}] = \{ab\}$ and thus $(\psi \diamond \varepsilon_{1,2}) \wedge \varepsilon_1$ is consistent. Similarly, we get that $(\psi \diamond \varepsilon_{2,3}) \wedge \varepsilon_2$ and $(\psi \diamond \varepsilon_{3,1}) \wedge \varepsilon_3$ are consistent, but $(\psi \diamond \varepsilon_{3,1}) \wedge \varepsilon_1$ is inconsistent, which means that U_{SC}^{Horn} is not satisfied by this example.

It seems, thus, that we are on the right track. But before moving any further, we pause to note that in the baseline case of propositional logic postulate U_{SC} does not contain any information not already provided by the standard postulates.

Proposition 2. In propositional logic, identified as the closed fragment $\mathcal{L}_{id} = \mathcal{L}$, postulate U_{SC}^{id} follows from postulates U_{1-9}^{id} .

PROOF. We start by using induction to prove a related statement: for any $k \geq 1$, if $w_i \in [\psi \diamond \varepsilon_{i,i+1}]$, for all $i \in \{1, \dots, k\}$, then postulates U_{1-9}^{id} imply that $w_1 \in [\psi \diamond \varepsilon_{1,k+1}]$. Note that the claim is trivially satisfied for $k = 1$. Assume, next, that the claim is true for all $i \leq k$. We show that it also holds for $k + 1$. Towards this end, consider the propositional formula $\varepsilon_{1,k,k+1}$, with $[\varepsilon_{1,k,k+1}] = \{w_1, w_k, w_{k+1}\}$.¹ By postulates U_1^{id} and U_3^{id} , we know that $[\psi \diamond \varepsilon_{1,k,k+1}]$ is a non-empty subset of $\{w_1, w_k, w_{k+1}\}$. Suppose, towards a contradiction, that $w_1 \notin [\psi \diamond \varepsilon_{1,k,k+1}]$, in which case either w_k or w_{k+1} has to be in $[\psi \diamond \varepsilon_{1,k,k+1}]$. This means that $(\psi \diamond \varepsilon_{1,k,k+1}) \wedge \varepsilon_{k,k+1}$ is consistent, and therefore:

$$\begin{aligned} (\psi \diamond \varepsilon_{1,k,k+1}) \wedge \varepsilon_{k,k+1} &\equiv \psi \diamond (\varepsilon_{1,k,k+1} \wedge \varepsilon_{k,k+1}) && // \text{ by } U_5^{id} \text{ and } U_9^{id} \\ &\equiv \psi \diamond \varepsilon_{k,k+1}. && // \text{ by } U_4^{id} \end{aligned}$$

Since, by assumption, $w_k \in [\psi \diamond \varepsilon_{k,k+1}]$, this implies that $w_k \in [\psi \diamond \varepsilon_{1,k,k+1}]$, which means that $(\psi \diamond \varepsilon_{1,k,k+1}) \wedge \varepsilon_{1,k}$ is consistent, and thus:

$$\begin{aligned} (\psi \diamond \varepsilon_{1,k,k+1}) \wedge \varepsilon_{1,k} &\equiv \psi \diamond (\varepsilon_{1,k,k+1} \wedge \varepsilon_{1,k}) && // \text{ by } U_5^{id} \text{ and } U_9^{id} \\ &\equiv \psi \diamond \varepsilon_{1,k}. && // \text{ by } U_4^{id} \end{aligned}$$

Since, by the inductive hypothesis, $w_1 \in [\psi \diamond \varepsilon_{1,k}]$, we infer that $w_1 \in [\psi \diamond \varepsilon_{1,k,k+1}]$. We have obtained the contradiction we were looking for, which allows us to conclude that, indeed, $w_1 \in [\psi \diamond \varepsilon_{1,k,k+1}]$. We now have that:

$$\begin{aligned} (\psi \diamond \varepsilon_{1,k,k+1}) \wedge \varepsilon_{1,k+1} &\models \psi \diamond (\varepsilon_{1,k,k+1} \wedge \varepsilon_{1,k+1}) && // \text{ by } U_5^{id} \\ &\equiv \psi \diamond \varepsilon_{1,k+1}, && // \text{ by } U_4^{id} \end{aligned}$$

which allows us to conclude that $w_1 \in [\psi \diamond \varepsilon_{1,k+1}]$, our target statement.

Turning now to postulate U_{SC}^{id} , take, as per its precondition, complete propositional formulas $\psi, \varepsilon_1, \dots, \varepsilon_n$, with $[\varepsilon_i] = \{w_i\}$, such that $(\psi \diamond \varepsilon_{1,2}) \wedge \varepsilon_1, \dots, (\psi \diamond \varepsilon_{n-1,n}) \wedge \varepsilon_{n-1}$ are all consistent. Note that this implies that $w_i \in [\psi \diamond \varepsilon_{i,i+1}]$, for $i \in \{1, \dots, n-1\}$. If, in addition, we assume that $w_n \in [\psi \diamond \varepsilon_{n,1}]$, we can use the statement just proven to conclude that $w_1 \in [\psi \diamond \varepsilon_{n,1}]$. Thus, it cannot be the case that $(\psi \diamond \varepsilon_{n,1}) \wedge \varepsilon_1$ is inconsistent, which shows that U_{SC}^{id} is satisfied. \square

To sum up where we are: Example 7 showed us that postulates U_{1-9}^* are not strong enough to rule out non-transitive cycles in fragments weaker than propositional logic. Adding postulate U_{SC}^* to the mix seems to prevent such cycles, which is promising. We now show postulates U_{1-9}^* , together with U_{SC}^* , allow us to actually characterize any \mathcal{L}_* -update operator in terms of total, faithful pointwise preorders.

Theorem 2. If \mathcal{L}_* is a closed fragment of propositional logic, an \mathcal{L}_* -update operator \diamond satisfies postulates U_{1-9}^* and U_{SC}^* if and only if there exists an \mathcal{L}_* -compliant, pointwise faithful assignment that maps every interpretation w to a total pre-order \leq_w and represents the operator \diamond .

¹Recall that in propositional logic it holds that $cl_{id}(M) = M$, for any set of interpretations M , and hence $[\varepsilon_M] = M$, for any set of interpretations M .

PROOF. We start with the simpler direction, which involves taking an \mathcal{L}_* -compliant, pointwise faithful assignment with total preorders, and showing that the update operator \diamond represented by it satisfies postulates U_{1-9}^* and U_{SC}^* . This holds for the simple reason that an assignment with the stated properties is just a special case of a general pointwise faithful assignment: thus, the \mathcal{L}_* -update operator \diamond represented by it satisfies the *propositional* postulates U_{1-9}^{id} , i.e., the postulates for full propositional logic. Since most of the fragment-based update postulates are restrictions of the propositional ones to the fragment \mathcal{L}_* , it follows that \diamond also satisfies these postulates.

Satisfaction of postulates U_7^* and U_8^* is straightforward. For postulate U_{SC}^* , take an \mathcal{L}_* -formula ψ such that $[\psi] = \{w\}$ and suppose, towards a contradiction, that $(\psi \diamond \varepsilon_{1,2}) \wedge \varepsilon_1, \dots, (\psi \diamond \varepsilon_{n-1,n}) \wedge \varepsilon_{n-1}$ are all consistent, but $(\psi \diamond \varepsilon_{n,1}) \wedge \varepsilon_n$ is consistent and $(\psi \diamond \varepsilon_{n,1}) \wedge \varepsilon_1$ is inconsistent, where $\varepsilon_{i,j}$ is the \mathcal{L}_* -envelope of the pair of interpretations $\{w_i, w_j\}$. It follows from this that $w_1 \leq_w w_2 \leq_w \dots \leq_w w_n <_w w_1$, which contradicts the transitivity of \leq_w .

For the other direction, we start with an \mathcal{L}_* -update operator \diamond that satisfies all the stated postulates and construct a suitable assignment that represents it. This construction proceeds in a number of steps and uses the \mathcal{L}_* -envelope $\varepsilon_{1,2}$ of interpretations w_1 and w_2 .

Step 1. Start by defining the *revealed relation* \trianglelefteq_w , for any interpretations w_1 and w_2 , as follows:

$$w_1 \trianglelefteq_w w_2 \text{ if } w_1 \in [\psi_w \diamond \varepsilon_{1,2}].$$

We write $w_1 \triangleleft_w w_2$ for the strict part of \trianglelefteq_w . By postulates U_1^* and U_3^* , it holds that $[\psi_w \diamond \varepsilon_1] = \{w_1\}$, which implies that \trianglelefteq_w is reflexive. The revealed relation \trianglelefteq_w is not guaranteed to be either total or transitive. However, postulate U_{SC}^* straightforwardly implies that \trianglelefteq_w is Suzumura consistent, which, by Theorem 1, implies that \trianglelefteq_w can be extended to a total preorder. We do this in the following step.

Step 2. Take any ordering extension \leq_w of \trianglelefteq_w , i.e., \leq_w is a total preorder that contains the transitive closure of \trianglelefteq_w and preserves the strict part of \trianglelefteq_w . From Theorem 1, we know that such an extension is guaranteed to exist. We show here that \leq_w already represents the update operator, that is, $[\psi \diamond \mu] = \bigcup_{w \in [\psi]} \min_{\leq_w} [\mu]$. We do this by first showing that $[\psi_w \diamond \mu] = \min_{\leq_w} [\mu]$, for some interpretation w and \mathcal{L}_* -formula $[\psi_w] = \{w\}$.

For the inclusion $[\psi_w \diamond \mu] \subseteq \min_{\leq_w} [\mu]$ suppose, towards a contradiction, that there exists $w_1 \in [\psi_w \diamond \mu]$ such that $w_1 \notin \min_{\leq_w} [\mu]$. This implies that there is some interpretation $w_2 \in [\mu]$ such that $w_2 <_w w_1$. At the same time, we have:

$$\begin{aligned} w_1 &\in [(\psi_w \diamond \mu) \wedge \varepsilon_{1,2}] \\ &\subseteq [\psi \diamond (\mu \wedge \varepsilon_{1,2})] && // \text{ by } U_5^* \\ &= [\psi \diamond \varepsilon_{1,2}]. && // \text{ by } U_4^* \end{aligned}$$

Hence, it also holds that $w_1 \trianglelefteq_w w_2$, and thus that $w_1 \leq_w w_2$, which provides a contradiction.

For the inclusion $\min_{\leq_w} [\mu] \subseteq [\psi_w \diamond \mu]$ suppose, towards a contradiction, that there exists an interpretation $w_1 \in \min_{\leq_w} [\mu]$ such that $w_1 \notin [\psi_w \diamond \mu]$. By postulate U_3^* we have that $[\psi_w \diamond \mu] \neq \emptyset$, so there must exist an interpretation $w_2 \in [\psi_w \diamond \mu]$. This means that $(\psi_w \diamond \mu) \wedge \varepsilon_{1,2}$ is consistent, and thus:

$$\begin{aligned} \psi_w \diamond \varepsilon_{1,2} &\equiv \psi_w \diamond (\mu \wedge \varepsilon_{1,2}) && // \text{ by } U_4^* \\ &\models (\psi_w \diamond \mu) \wedge \varepsilon_{1,2}, && // \text{ by } U_9^* \end{aligned}$$

which allows us to infer that $w_1 \notin [\psi_w \diamond \varepsilon_{1,2}]$. At the same time, postulates U_1^* and U_3^* guarantee that $[\psi_w \diamond \varepsilon_{1,2}]$ contains some interpretations found in $[\varepsilon_{1,2}]$. Assume $w_3 \in [\psi_w \diamond \varepsilon_{1,2}]$, with the understanding that w_3 can be identical to w_2 , but it can also be a different interpretation, added in the process of computing the \mathcal{L}_* -closure of $\{w_1, w_2\}$. We now switch our attention to the \mathcal{L}_* -envelope of w_1 and w_3 , the \mathcal{L}_* -formula $\varepsilon_{1,3}$, and apply the same sequence of arguments from above, but to w_3 . Note that $\varepsilon_{1,3} \models \varepsilon_{1,2}$, so we can use postulate U_5^* to infer that




Fig. 5. The revealed plausibility relation \preceq_w for a Horn update operator \diamond : a dashed arrow from w_1 to w_2 indicates that $w_1 \in [\psi_w \diamond \varepsilon_{1,2}]$ and $w_2 \notin [\psi_w \diamond \varepsilon_{1,2}]$. This relation can be extended in multiple ways to a total preorder, e.g., by making c equivalent either to ab , or to a . As long as the strict comparisons in \preceq_\emptyset are preserved, the resulting total preorder still represents \diamond .

$w_3 \in [\psi \diamond \varepsilon_{1,3}]$, and postulate U_9^* again to get that $\psi_w \diamond (\varepsilon_{1,2} \wedge \varepsilon_{1,3}) \models (\psi_w \diamond \varepsilon_{1,2}) \wedge \varepsilon_{1,3}$. Since $\varepsilon_{1,3} \models \varepsilon_{1,2}$, we infer that $\psi_w \diamond (\varepsilon_{1,2} \wedge \varepsilon_{1,3}) \equiv \psi_w \diamond \varepsilon_{1,3}$, and then that $w_1 \notin [\psi_w \diamond \varepsilon_{1,3}]$. This, now, implies that $w_3 \triangleleft_w w_1$, shorthand for the fact that $w_3 \preceq_w w_1$ but not $w_1 \preceq_w w_3$. We can immediately infer from this that $w_3 <_w w_1$ since \leq_w is an ordering extension of \preceq_w . This now contradicts the fact that $w_1 \in \min_{\leq_w} [\mu]$. Thus, our original assumption that $w_1 \notin [\psi_w \diamond \mu]$ must be wrong.

Putting these two inclusions together we obtain, as announced, that $[\psi_w \diamond \mu] = \min_{\leq_w} [\mu]$. Using postulate U_8^* we get, in addition, that $[\psi \diamond \mu] = \bigcup_{w \in [\psi]} \min_{\leq_w} [\mu]$, for any \mathcal{L}_* -formula ψ . In other words, the partial assignment defined so far already represents the \mathcal{L}_* -update operator \diamond .

Finally, we show that, for any interpretation w , \leq_w is faithful. Consider an interpretation w' distinct from w and note that $\psi_w \models \varepsilon_{w,w'}$, hence postulate U_2^* gives us that $[\psi_w \diamond \varepsilon_{w,w'}] = \{w\}$. This, in turns, implies that $w \triangleleft_w w'$, which implies that $w <_w w'$. This concludes the proof, by showing that \leq_w , as defined through Steps 1 and 2, is the faithful, total preorder we are looking for. \square

Theorem 2 shows that, for the sensitive direction that requires finding a suitable assignment to represent the \mathcal{L}_* -update operator, the preorder that does the trick is one (of possibly many) ordering extension of the revealed relation, i.e., the relation \preceq_w such that $w_1 \preceq_w w_2$ if $w_1 \in [\psi_w \diamond \varepsilon_{1,2}]$, for any interpretations w_1 and w_2 and complete formula ψ_w . Note that the final step of this proof works for *any* ordering extension, and in this respect our result ends up saying something slightly stronger than similar results on fragment-based belief change (Delgrande and Peppas 2015; Delgrande, Peppas, and Woltran 2018), since these rely on a *particular* ordering extension. An example makes this clearer.

Example 10. Consider a $\mathcal{L}^{\text{Horn}}$ -update operator that induces the revealed relation \preceq_w depicted in Figure 5, i.e., for which $[\psi_w \diamond \varepsilon_{a,\emptyset}] = \{\emptyset\}$, $[\psi_w \diamond \varepsilon_{a,ab}] = \{a\}$, and so on. Note that $[\psi_w \diamond \varepsilon_{a,c}] = \{\emptyset\}$, which implies that $\emptyset \triangleleft_\emptyset a$ and $\emptyset \triangleleft_\emptyset c$, but a and c are incomparable. As a side-effect of working in the Horn fragment, the operator \diamond cannot compare a and c directly, because the set of interpretations $\{a, c\}$ is not representable by a Horn formula. For similar reasons, ab and c end up being incomparable. The upshot is that the revealed relation \preceq_w is partial, with incomparabilities between those pairs of interpretations that cannot be represented in the Horn fragment. However, \preceq_w still manages to represent the operator \diamond : precisely because a and c are never compared directly, the exact relationship between them does not matter and can be completed in any way that leaves the revealed part untouched. In other words, the fact that \preceq_w is Suzumura consistent ensures that it can be extended to a total preorder that still represents \diamond , and Step 3 of the proof of Theorem 2 shows that any ordering extension is




Fig. 6. An update operator cannot distinguish between these two partial preorders.

suitable. For instance, we can complete \leq_w either with $ab \approx_w c$, or with $a \approx_w c$. By contrast, the construction in (Delgrande and Peppas 2015) applied to this case would make c equivalent to ab .

We conclude this section by noting that the original problem, highlighted in Example 7, occurs in an entirely similar way as for Horn revision and Horn merging, where it is also handled by employing additional, special postulates (Delgrande and Peppas 2015; Haret, Rümmele, et al. 2017). These postulates are similar to U_{SC}^* and are employed to the same effect. The added value of the present contribution lies in highlighting how the common link underlying all the constructions is Suzumura consistency.

3.2 Partial Preorders

Recall that in the case of full propositional logic postulates U_{1-8} are weaker than the full set U_{1-9} and characterize update operators representable with *partial*, rather than total, preorders (Katsuno and Mendelzon 1991a). Does a similar result hold in fragments? The answer turns out to be yes, but only under a set of particular restrictions. The first restriction is on the underlying fragment.

As already apparent from Section 3.1, working in a less expressive formalism, i.e., an \mathcal{L}_* -fragment, means less grip on the underlying ranking. Things get even worse when further weakening the postulates for the broader class of partial preorders.

Example 11. Consider partial preorders \leq_w and \leq'_w that rank interpretations a, b and $\emptyset = a \cap b$ according to Figure 6. Take, now, the Horn update operators \diamond and \diamond' that represent these rankings. These operators are perfectly equivalent with respect to all Horn formulas relevant to the interpretations: $[\psi_w \diamond \varepsilon_{a,b}] = [\psi_w \diamond' \varepsilon_{a,b}] = \{\emptyset, a\}$, $[\psi_w \diamond \varepsilon_{\emptyset,a}] = [\psi_w \diamond' \varepsilon_{\emptyset,a}] = \{\emptyset, a\}$, and $[\psi_w \diamond \varepsilon_{\emptyset,b}] = [\psi_w \diamond' \varepsilon_{\emptyset,b}] = \{\emptyset\}$. Since the pair $\{a, b\}$ is not representable in the Horn fragment (the closest we can get is through the \mathcal{L}_{Horn} -closure $\{\emptyset, a, b\}$), the direct ranking between a and b is invisible to the represented operator.

Example 11 presents two different preorders (one that considers a better than b , one leaves them incomparable) that are equivalent modulo the operator representing them. Put differently, a Horn update operator cannot tell whether a and b in the example are incomparable or not. And, lest it be thought that this problem only affects the Horn fragment, note that it applies in any fragment when a new interpretation w_0 is added when taking the closure of $\{w_1, w_2\}$. Since this type of situation creates problems for a representation result, we restrict our attention to the ranking among the two that carries more information, singled out by the following property.

We say that a partial preorder \leq is *maximally consistent with respect to \mathcal{L}_** if, for any two interpretations w_1 and w_2 , it holds that:

$$\text{if } w_1 \in \min_{\leq} cl_*(\{w_1, w_2\}) \text{ and } w_2 \notin \min_{\leq} cl_*(\{w_1, w_2\}), \text{ then } w_1 < w_2.$$

Intuitively, maximal consistency states that if w_1 gets chosen from a pool of interpretations that also contains w_2 while w_2 gets rejected, then a choice needs to be made on whether w_1 is considered more plausible than




Fig. 7. Partial preorder \leq_w , together with its maximally consistent extension \leq'_w , with \leq'_w containing a non-transitive cycle. Both \leq_w and \leq'_w represent the same \mathcal{L}_{Horn} -update operator; keeping maximal consistency will require ruling this operator out with an additional postulate.

w_2 . In such a case we assume it is sensible to stipulate, in a manner maximally consistent with an \mathcal{L}_* -operator representing $<$, that w_1 is more plausible than w_2 .

Example 12. For preorders \leq_w and \leq'_w from Example 11 and the Horn fragment, we have that $cl_{Horn}(\{a, b\}) = \{a, b, \emptyset\}$. Note that $a \in \min_{\leq_w} cl_{*}(\{a, b\})$ and $b \notin \min_{\leq_w} cl_{*}(\{a, b\})$ and the same holds if we replace \leq_w with \leq'_w . In other words, both \leq_w and \leq'_w model an operator that chooses a from $\{a, b, \emptyset\}$ and rejects b . The reason b is rejected, of course, is because \emptyset is chosen before it. Thus, if tasked with reconstructing the plausibility relation given only the choices of the operator, \emptyset is always considered more plausible than b . However, for the comparison between b and a a Horn operator does not give us enough information, and we are stuck having to choose between \leq_w and \leq'_w . Our proposal here is to choose \leq'_w , i.e., assume that a is also more plausible than b . In other words, \leq_w is not maximally consistent, whereas \leq'_w is.

Indeed, in the following we will work with maximally consistent preorders.

Our problems do not end here: isolated comparisons between pairs of interpretations live inside the broader ecosystem of the overall ranking on all interpretations, and maximally consistent preorders are still vulnerable to non-transitive cycles.

Example 13. Consider the partial preorder \leq_w depicted in Figure 7, and note that \leq_w is not maximally consistent: focusing on ab and ac , we have that the Horn closure of $\{ab, ac\}$ is $cl_{Horn}(\{ab, ac\}) = \{ab, ac, a\}$, with $\min_{\leq_w}(\{ab, ac, a\}) = \{a, ab\}$. Thus, while maximal consistency requires ab to be more plausible than ac , the two interpretations are incomparable in \leq_w . A similar situation occurs with respect to the pairs $\{ac, bc\}$ and $\{bc, ab\}$, since $\min_{\leq_w}(\{ac, bc, c\}) = \{ac, c\}$ and $\min_{\leq_w}(\{ab, bc, c\}) = \{bc, b\}$. A routine check shows that the Horn update operator represented by \leq_w satisfies all the core postulates U_{1-8}^{Horn} . At the same time, the maximally consistent extension of \leq_w , depicted in Figure 7 as \leq'_w , contains a cycle between ab , ac and bc .

The issue in Example 13 is similar to the problem in Section 3.1, where a non-transitive cycle goes undetected by an operator satisfying all the core postulates. In this case, a maximally consistent preorder ends up representing an update operator that satisfies postulates U_{1-8}^* . Commitment to maximal consistency requires, as in Section 3.1, an additional postulate.

(U_{wSC}^*) If ψ is a complete \mathcal{L}_* -formula, then, for any $n \geq 1$ and complete \mathcal{L}_* -formulas $\varepsilon_1, \dots, \varepsilon_n$ such that, for all $i \in \{1, \dots, n-1\}$, it holds that $(\psi \diamond \varepsilon_{i,i+1}) \wedge \varepsilon_i$ is consistent and $(\psi \diamond \varepsilon_{i,i+1}) \wedge \varepsilon_{i+1}$ is inconsistent, then $(\psi \diamond \varepsilon_{1,n}) \wedge \varepsilon_n$ is inconsistent.

Postulate U_{wSC}^* , short for *weak Suzumura consistency*, encodes the same intuition as postulate U_{SC}^* from Section 3.1, adapted here to work with the weaker set of postulates U_{1-8}^* , which encode update on partial preorders. The intuition, briefly, is that if interpretation w_1 gets chosen and w_2 explicitly rejected, \dots , w_{n-1} gets chosen and w_n rejected, then w_n does not get chosen from a menu that contains both w_1 and w_n . In other words, U_{wSC}^* encodes Suzumura consistency of a strict, partial relation on interpretations. This requirement turns out to prevent situations such as the one in Example 13.

Example 14. For the maximally consistent ranking \leq'_w in Figure 7 and the Horn update operator \diamond representing it, we obtain first that $[(\psi \diamond \varepsilon_{ab,ac}) \wedge \varepsilon_{ab}] = \{ab\}$ and $[(\psi \diamond \varepsilon_{ab,ac}) \wedge \varepsilon_{ac}] = \emptyset$, and then that $[(\psi \diamond \varepsilon_{ac,bc}) \wedge \varepsilon_{ac}] = \{ac\}$ and $[(\psi \diamond \varepsilon_{ac,bc}) \wedge \varepsilon_{bc}] = \emptyset$. Thus, the preconditions of $U_{\text{wSC}}^{\text{Horn}}$ are satisfied for interpretations $w_1 = ab$, $w_2 = ac$ and $w_3 = bc$. At the same time, it holds that $[(\psi \diamond \varepsilon_{ab,bc}) \wedge \varepsilon_{bc}] = \{bc\}$, contradicting the conclusion of $U_{\text{wSC}}^{\text{Horn}}$.

As expected, postulate U_{wSC}^* is redundant in propositional logic.

Proposition 3. In propositional logic, identified as the closed fragment $\mathcal{L}_{id} = \mathcal{L}$, postulate U_{wSC}^{id} follows from postulates U_{1-8}^{id} .

PROOF. The proof is similar to the proof of Proposition 2, the main difference being that in this case we cannot use postulate U_9^{id} , and must make do with the weaker postulate U_6^{id} .

We start, as before, by proving a related statement: for any $k \geq 1$, if $[\psi \diamond \varepsilon_{i,i+1}] = \{w_i\}$, for all $i \in \{1, \dots, k\}$, then postulates U_{1-8}^{id} imply that $[\psi \diamond \varepsilon_{1,k+1}] = \{w_1\}$. Note, first, that the claim is trivially satisfied for $k = 1$. Assume, next, that the claim is true for all k . We show that it also holds for $k + 1$. Consider, again, the propositional formula $\varepsilon_{1,k,k+1}$, with $[\varepsilon_{1,k,k+1}] = \{w_1, w_k, w_{k+1}\}$. Postulates U_1^{id} and U_3^{id} imply that $[\psi \diamond \varepsilon_{1,k,k+1}]$ is a non-empty subset of $\{w_1, w_k, w_{k+1}\}$. With two quick-fire applications of U_5^{id} , we can settle immediately what $\psi \diamond \varepsilon_{1,k,k+1}$ has to be. We first have:

$$\begin{aligned} [(\psi \diamond \varepsilon_{1,k,k+1}) \wedge \varepsilon_{k,k+1}] &\subseteq [\psi \diamond (\varepsilon_{1,k,k+1} \wedge \varepsilon_{k,k+1})] && // \text{ by } U_5^{id} \\ &= [\psi \diamond \varepsilon_{k,k+1}] && // \text{ by } U_4^{id} \\ &= \{w_k\}, && // \text{ by assumption} \end{aligned}$$

which implies that $w_{k+1} \notin [(\psi \diamond \varepsilon_{1,k,k+1})]$. We then have that:

$$\begin{aligned} [(\psi \diamond \varepsilon_{1,k,k+1}) \wedge \varepsilon_{1,k}] &\subseteq [\psi \diamond (\varepsilon_{1,k,k+1} \wedge \varepsilon_{1,k})] && // \text{ by } U_5^{id} \\ &= [\psi \diamond \varepsilon_{1,k}] && // \text{ by } U_4^{id} \\ &= \{w_1\}, && // \text{ by inductive hypothesis} \end{aligned}$$

which implies that $w_k \notin [(\psi \diamond \varepsilon_{1,k,k+1})]$, so the only remaining option is $[(\psi \diamond \varepsilon_{1,k,k+1})] = \{w_1\}$. This means that $\psi \diamond \varepsilon_{1,k,k+1} \models \varepsilon_{1,k+1}$, which, when coupled with the fact that $\psi \diamond \varepsilon_{1,k+1} \models \varepsilon_{1,k,k+1}$ and postulate U_6^{id} , gives us that $[\psi \diamond \varepsilon_{1,k+1}] = [\psi \diamond \varepsilon_{1,k,k+1}] = \{w_1\}$, what we set out to prove.

We can now turn to postulate U_{wSC}^{id} . Take, as per its precondition, complete propositional formulas $\psi, \varepsilon_1, \dots, \varepsilon_n$, with $[\varepsilon_i] = \{w_i\}$, such that $(\psi \diamond \varepsilon_{i,i+1}) \wedge \varepsilon_i$ is consistent and $(\psi \diamond \varepsilon_{i,i+1}) \wedge \varepsilon_{i+1}$ is inconsistent, for all $i \in \{1, \dots, n-1\}$. In propositional logic this implies that $[\psi \diamond \varepsilon_{i,i+1}] = \{w_i\}$, for all $i \in \{1, \dots, n-1\}$. Thus, the statement proven above implies that $[\psi \diamond \varepsilon_{1,n}] = \{w_1\}$, which shows that $(\psi \diamond \varepsilon_{1,n}) \wedge \varepsilon_n$ is inconsistent, as desired, and U_{wSC}^{id} is satisfied. \square

Finally, we must add a restriction on the fragment itself. We say a closed fragment \mathcal{L}_* is *tight* if:

$$|cl_*(\{w_1, w_2\})| \leq 3, \text{ for any interpretations } w_1, w_2.$$

In other words, a fragment is tight if for any two interpretations w_1 and w_2 , the \mathcal{L}_* -closure of $\{w_1, w_2\}$ adds at most one other interpretation distinct from w_1 and w_2 . This restriction sounds more demanding than it is, and almost all the fragments of interest to us here (i.e., the \mathcal{L}_{Horn} , \mathcal{L}_{Krom} and \mathcal{L}_{affine} fragments) are all tight. Indeed, in the Krom and affine fragments it holds that $cl_*(\{w_1, w_2\}) = \{w_1, w_2\}$, for any pair of interpretations w_1, w_2 , while in the Horn fragment we have that $cl_*(w_1, w_2) = \{w_1, w_2, w_1 \cap w_2\}$. The only exception is the 1CNF fragment.

Example 15. For the pair of interpretations \emptyset and abc , the \mathcal{L}_{1CNF} -closure is $cl_{1CNF}(\emptyset, abc) = \{\emptyset, a, b, c, ab, ac, bc, abc\}$, and hence \mathcal{L}_{1CNF} is not tight.

We will have more to say on the 1CNF fragment shortly, but the main representation result we present now assumes tight fragments.

Theorem 3. If \mathcal{L}_* is a tight, closed fragment of propositional logic, an \mathcal{L}_* -update operator \diamond satisfies postulates U_{1-8}^* and U_{wSC}^* if and only if there exists an \mathcal{L}_* -compliant, pointwise faithful assignment that maps every interpretation w to a maximally consistent partial pre-order \leq_w and represents the operator \diamond .

PROOF. For one direction, start with an \mathcal{L}_* -compliant assignment satisfying all mentioned conditions and note, first, that the condition of \mathcal{L}_* -compliance ensures that the \mathcal{L}_* -update operator represented by it is well defined. The argument that \diamond satisfies postulates U_{1-8}^* is familiar from Theorem 2: \diamond satisfies the propositional postulates U_{1-8}^{id} , and hence it also satisfies the restriction of these postulates to the fragment \mathcal{L}_* .

It remains to be shown that \diamond satisfies postulate U_{wSC}^* . To that end, take a complete \mathcal{L}_* -formula ψ_w and assume that there exist complete \mathcal{L}_* -formulas $\varepsilon_1, \dots, \varepsilon_n$ such that $(\psi_w \diamond \varepsilon_{i,i+1}) \wedge \varepsilon_i$ is consistent and $(\psi_w \diamond \varepsilon_{i,i+1}) \wedge \varepsilon_{i+1}$ is inconsistent, for all $i \in \{1, \dots, n-1\}$. It thus holds that $[\psi_w \diamond \varepsilon_{i,i+1}]$, which is equal to $min_{\leq_w} cl_*(\{w_i, w_{i+1}\})$, contains w_i but not w_{i+1} . By assumption \leq_w is maximally consistent, which allows us to infer that that $w_i <_w w_{i+1}$. Putting all these facts together, we get that $w_1 <_w \dots <_w w_n$ and, by the transitivity of \leq_w , that $w_1 <_w w_n$. It thus follows that $w_n \notin min_{\leq_w}[\varepsilon_{1,n}]$, and hence $(\psi_w \diamond \varepsilon_{1,n}) \wedge \varepsilon_n$ is inconsistent.

For the other direction, take an \mathcal{L}_* -update operator \diamond that satisfies postulates U_{1-8}^* and U_{wSC}^* . We proceed to construct the desired assignment in steps. The overall strategy is similar to Theorem 2, but key details differ due to the fact that we use different postulates and work with different types of preorders (partial, not total).

Step 1. For any interpretation w and complete formula $[\psi_w] = \{w\}$, the *revealed plausibility relation* \preceq_w is defined as follows:

$$w_1 \preceq_w w_2 \text{ if } w_1 = w_2, \text{ or } w_1 \in [\psi_w \diamond \varepsilon_{1,2}] \text{ and } w_2 \notin [\psi_w \diamond \varepsilon_{1,2}],$$

for any two interpretations w_1 and w_2 .

The definition of \preceq_w is slightly different from the one in Section 3.1 in that it does not allow for ties between distinct interpretations. The relation \preceq_w is reflexive by definition, but the definition implies that \preceq_w is asymmetric for distinct w_1 and w_2 . That is, if both w_1 and w_2 are in $[\psi_w \diamond \varepsilon_{1,2}]$, then w_1 and w_2 are constructed as being incomparable, and not equivalent. In this respect, \preceq_w behaves like a strict relation for distinct interpretations.

Additionally, postulate U_{wSC}^* forbids chains of comparisons $w_1 \preceq_w \dots \preceq_w w_n \preceq_w w_1$, for distinct interpretations w_1, \dots, w_n . To see why, note that in the presence of such a chain the first $n-1$ comparisons imply, by postulate U_{wSC}^* , that $(\psi_w \diamond \varepsilon_{1,n}) \wedge \varepsilon_n$ is inconsistent, whereas $w_n \preceq_w w_1$ implies, by definition of \preceq_w , that $w_n \in [\psi_w \diamond \varepsilon_{1,n}]$ —a contradiction. This feature allows us to meaningfully extend \preceq_w to a transitive relation, as follows.

Step 2. Define the partial preorder \leq_w as the transitive closure of \preceq_w . The relation \leq_w , we now show, delivers the desired assignment.

We start by noting a fact that will prove useful throughout the rest of the proof: \leq_w extends \preceq_w while preserving its strict part, i.e., if $w_1 \preceq_w w_2$ but not $w_2 \preceq_w w_1$, then $w_1 <_w w_2$. The argument is simple: if it were the case that $w_2 \leq_w w_1$, then there would need to be a chain $w_2 \preceq_w \dots \preceq_w w_1$, which, together with the assumption that $w_1 \preceq_w w_2$, contradicts the basic consequence of applying U_{wSC}^* described above.

Moving on, we claim that the assignment given by the partial preorders \leq_w represents the operator \diamond , i.e., $[\psi \diamond \mu] = \bigcup_{w \in [\psi]} \min_{\leq_w} [\mu]$, for any \mathcal{L}_* -formula ψ . We do this by first showing that $[\psi_w \diamond \mu] = \min_{\leq_w} [\mu]$, for some interpretation w and complete \mathcal{L}_* -formula $[\psi_w] = \{w\}$.

For one direction of the double inclusion, take $w_1 \in [\psi_w \diamond \mu]$ and suppose that $w_1 \notin \min_{\leq_w} [\mu]$. There exists, then, an interpretation $w_2 \in [\mu]$ such that $w_2 \neq w_1$ and $w_2 <_w w_1$, which implies that there exists a chain of $w_2 \triangleleft_w \dots \triangleleft_w w_1$ of comparisons, of length at least two, which, by postulate U_{wSC}^* , implies that $w_1 \notin [\psi_w \diamond \varepsilon_{1,2}]$. However, since $w_1 \in [(\psi_w \diamond \mu) \wedge \varepsilon_{1,2}]$, postulate U_5^* implies that $w_1 \in [\psi_w \diamond \varepsilon_{1,2}]$, which is a contradiction.

For the other direction, take $w_1 \in \min_{\leq_w} [\mu]$ and an arbitrary interpretation $w_2 \in [\mu]$. The key, here, is to show that $w_1 \in [\psi_w \diamond \varepsilon_{1,2}]$, but we will have to work a bit for this conclusion. Assume, towards a contradiction, that $w_1 \notin [\psi_w \diamond \varepsilon_{1,2}]$. By our assumption that \mathcal{L}_* is tight (this is the precise point where the assumption is needed), we can reason as follows: since $[\varepsilon_{1,2}]$ contains at most three elements, we have either that $[\varepsilon_{1,2}] = \{w_1, w_2\}$, or $[\varepsilon_{1,2}] = \{w_1, w_2, w_3\}$, with w_3 added in the process of taking the \mathcal{L}_* -closure of $\{w_1, w_2\}$. Note that, by postulates U_1^* and U_3^* , $[\psi_w \diamond \varepsilon_{1,2}]$ is a non-empty subset of $[\varepsilon_{1,2}]$. Recalling the assumption that $w_1 \notin [\psi_w \diamond \varepsilon_{1,2}]$, we are faced with three possibilities: $[\psi_w \diamond \varepsilon_{1,2}]$ is either $\{w_2\}$, $\{w_3\}$, or $\{w_2, w_3\}$. Let us analyze the possibilities. If $[\psi_w \diamond \varepsilon_{1,2}] = \{w_2\}$, then $w_2 \triangleleft_w w_1$ but not $w_1 \triangleleft_w w_2$. Thus, from the fact that the strict part of \triangleleft_w carries over to \leq_w , established above, we infer that $w_2 <_w w_1$. This, however, contradicts the \leq_w -minimality of w_1 in $[\mu]$. If $[\psi_w \diamond \varepsilon_{1,2}] = \{w_2, w_3\}$ we conclude the same thing, in the same way. If $[\psi_w \diamond \varepsilon_{1,2}] = \{w_3\}$, then note that $\psi_w \diamond \varepsilon_{1,2} \models \varepsilon_{1,3}$. By the monotonicity feature of the closure operator of \mathcal{L}_* we have that $\varepsilon_{1,3} \models \varepsilon_{1,2}$, and, using postulate U_1^* , we can infer that $\psi_w \diamond \varepsilon_{1,3} \models \varepsilon_{1,3} \models \varepsilon_{1,2}$. Then, by postulate U_6^* , it follows that $[\psi_w \diamond \varepsilon_{1,3}] = [\psi_w \diamond \varepsilon_{1,2}] = \{w_3\}$. From this we conclude that $w_3 \triangleleft_w w_1$ but not the other way around, and hence that $w_3 <_w w_1$: the final contradiction.

All these contradictions stem from the assumption that $w_1 \notin [\psi_w \diamond \varepsilon_{1,2}]$, so we can safely conclude that $w_1 \in [\psi_w \diamond \varepsilon_{1,2}]$. Recall, now, that w_2 was chosen as an arbitrary model of μ . Using postulate U_7^* over all models of μ , and invoking Lemma 1 to guarantee that $[\mu] = \bigcup_{w_i \in [\mu]} [\varepsilon_{1,i}]$, we infer:

$$\bigwedge_{w_i \in [\mu]} (\psi_w \diamond \varepsilon_{1,i}) \models \psi_w \diamond \mu,$$

which implies that $w_1 \in [\psi_w \diamond \mu]$, the fact we originally set out to prove.

Thus, we have shown that $[\psi_w \diamond \mu] = \min_{\leq_w} [\mu]$, for a complete \mathcal{L}_* -formula $[\psi_w] = \{w\}$. The last step in this part of the proof involves applying postulate U_8^* to conclude that, if ψ is an arbitrary \mathcal{L}_* -formula, then $[\psi \diamond \mu] = \bigcup_{w \in [\psi]} \min_{\leq_w} [\mu]$, as desired.

Third, we show that \leq_w is maximally consistent. For this, we need to assume that $w_1 \in \min_{\leq_w} cl_*(\{w_1, w_2\})$ and $w_2 \notin \min_{\leq_w} cl_*(\{w_1, w_2\})$, for some interpretations w_1 and w_2 . Note that $cl_*(\{w_1, w_2\}) = [\varepsilon_{1,2}]$ and, as a direct consequence of the claim shown above, $\min_{\leq_w} [\varepsilon_{1,2}] = [\psi_w \diamond \varepsilon_{1,2}]$. We thus have that $w_1 \in [\psi_w \diamond \varepsilon_{1,2}]$ and $w_2 \notin [\psi_w \diamond \varepsilon_{1,2}]$, which implies that $w_1 \triangleleft_w w_2$, though, of course, not the other way around. This, then, implies that $w_1 <_w w_2$.

Fourth, the assignment defined in this way is pointwise faithful. To see this, take an interpretation w , a complete \mathcal{L}_* -formula $[\psi_w] = \{w\}$, and an interpretation $w' \neq w$. Since $\psi_w \models \varepsilon_{w,w'}$, postulate U_2^* implies that $[\psi_w \diamond \varepsilon_{w,w'}] = [\psi_w] = w$. Hence, $w \triangleleft_w w'$ but not the other way around and, again using the fact that \leq_w extends the strict part of \triangleleft , we conclude that $w <_w w'$.

We conclude the proof by noting that the assignment thus defined is out-of-the-box \mathcal{L}_* -compliant. \square

Theorem 3 applies to tight fragments, which covers the Horn, Krom and affine fragment, but not the 1CNF fragment. This turns out not be a problem, since, as we show in the next section, update in the 1CNF fragment is less problematic: the drastic- and Hamming-distance operators work just fine (see Propositions 4 and 8 below).

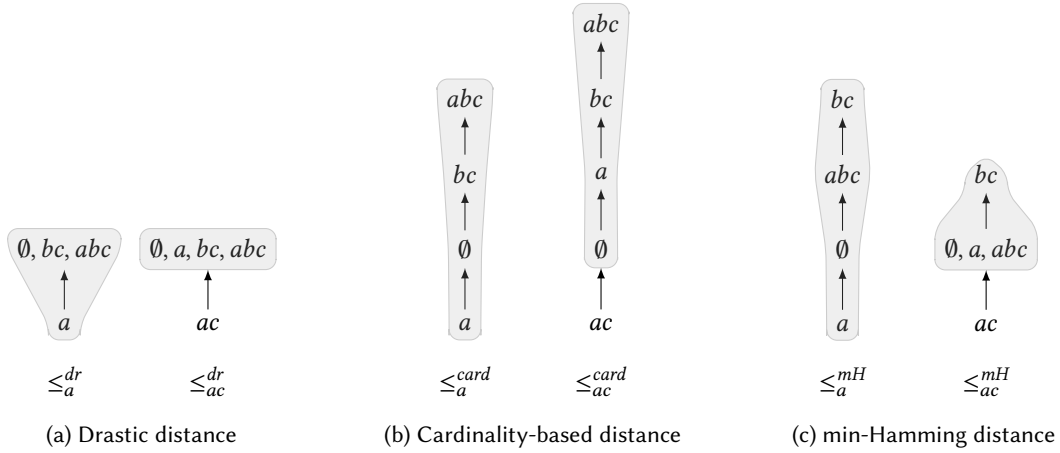


Fig. 8. Fragment-based update operators

4 Concrete Fragment-Based Operators

With the representation results in place, there is still a question of whether there do, actually, exist \mathcal{L}_* -operators that satisfy the desired properties: as Section 3 highlighted, \mathcal{L}_* -compliance rules out most of the existing operators. In this section we present a series of update operators to show that the representation results are not vacuous. Our strategy is to define operators through assignments, i.e., by finding methods for ranking interpretations, typically through some kind of distance function. The representation results in Section 3 prove their worth by showing us the type of preorders to look for.

We start with what is perhaps the simplest type of operator, defined by the drastic distance d_{dr} introduced in Section 2.

Example 16. We return to the running example from Example 1, with formulas $\psi = a \wedge \neg b$ and $\mu = b \leftrightarrow c$. Note that ψ and μ are also expressible in the Horn fragment. As already highlighted in Example 3, the drastic operator \diamond^{dr} gives us $\psi \diamond^{dr} \mu \equiv \mu$. The corresponding preorders \leq_w^{dr} are depicted in Figure 8a.

The assignments generated with the drastic distance turn out to fit all the requirements of a suitable \mathcal{L}_* -update operator.

Proposition 4. If \mathcal{L}_* is a closed fragment, the drastic assignment is \mathcal{L}_* -compliant, pointwise faithful, and \leq_w^{dr} is a total preorder for any interpretation w .

PROOF. The proof relies on the simple observation about how plausibility rankings \leq_w^{dr} look like. Note that in a ranking \leq_w^{dr} the interpretation w is at the bottom, at a distance of 0 (to itself), while every other interpretation w' sits at a distance of 1 (to w). Thus, for any \mathcal{L}_* -formulas ψ and μ , we obtain that:

$$\bigcup_{w \in [\psi]} \min_{\leq_w^{dr}} [\mu] = \begin{cases} [\psi], & \text{if } \psi \models \mu, \\ [\mu], & \text{otherwise.} \end{cases}$$

In other words, the interpretations offered up by the drastic assignment always correspond to an \mathcal{L}_* -formula, hence the assignment defined by the drastic distance is \mathcal{L}_* -compliant. It is, then, straightforward to see that the assignment is also pointwise faithful, with \leq_w^{dr} being a total preorder. \square

It follows straightforwardly from Proposition 4 and Theorem 2 that the drastic operator works in all closed fragments, and satisfies the strongest set of postulates.

Corollary 1. For any closed fragment \mathcal{L}_* , the \mathcal{L}_* -drastic operator \diamond^{dr} satisfies postulates U_{1-9}^* and U_{SC}^* . Furthermore, it is equivalent to:

$$\psi \diamond^{dr} \mu \equiv \begin{cases} \psi, & \text{if } \psi \models \mu, \\ \mu, & \text{otherwise,} \end{cases}$$

This shows, in one swoop, that the class of update operators satisfying the set of postulates we are interested in is non-empty, for all fragments. The obvious caveat is that the drastic operator is relatively unexciting: it would be disappointing if there were no other operators that fit the bill.

In the following we provide some more operators for the Horn fragment. The first is built upon the *cardinality-based distance* d_{card} defined, relative to an interpretation w , as:

$$d_{card}(w, w') = \begin{cases} 0, & \text{if } w = w', \\ |w'| + 1, & \text{otherwise,} \end{cases}$$

for any interpretation w' . Intuitively, d_{card} prioritizes interpretations of smaller cardinality, while also making sure that \leq_w^{card} is pointwise faithful, i.e., by arranging that $w <_w^{card} w'$ for any $w' \neq w$.

Example 17. Consider, again, formulas $\psi = a \wedge \neg b$ and $\mu = b \leftrightarrow c$ from Example 1. Note that $\psi \not\models \mu$, but $\psi \wedge \mu$ is consistent, with $[\psi \wedge \mu] = \{a\}$. The minimal model of μ is \emptyset , and the cardinality-based operator correspondingly gives us $[\psi \diamond^{card} \mu] = \{a\} \cup \{\emptyset\}$, with the preorders \leq_w^{card} depicted in Figure 8b.

Proposition 5. The cardinality-based assignment is \mathcal{L}_{Horn} -compliant, pointwise faithful, and \leq_w^{card} is a total preorder, for any interpretation w .

PROOF. Consider a Horn formula μ . Since μ is closed under intersection, the interpretation $w_{min} = \bigcap_{w_i \in [\mu]} w_i$ is also a model of $[\mu]$, and it has the property that $w_{min} \subseteq w_i$, for any $w_i \in [\mu]$. Since w_{min} is also the interpretation with the smallest cardinality out of all the models of μ , w_{min} will be the most plausible model of μ in any cardinality-based ranking \leq_w (other than w itself), i.e., $w_{min} \leq_w^{card} w_i$, for any $w_i \in [\mu]$ with $w_i \neq w$. We leverage this fact to conclude that the cardinality-based assignment produces the following results, for any Horn formulas ψ and μ :

$$\bigcup_{w \in [\psi]} \min_{\leq_w^{card}} [\mu] = \begin{cases} [\psi], & \text{if } [\psi] \subseteq [\mu], \\ ([\psi] \cap [\mu]) \cup \{w_{min}\}, & \text{if } [\psi] \not\subseteq [\mu] \text{ but } [\psi] \cap [\mu] \neq \emptyset, \\ \{w_{min}\}, & \text{otherwise.} \end{cases}$$

Clearly, $[\psi]$ is representable by a Horn formula. Then, note that $[\psi] \cap [\mu]$ is closed under intersection. Furthermore, it is a subset of $[\mu]$, and thus $w_{min} \subseteq w$, for any $w \in [\psi] \cap [\mu]$. Hence, adding w_{min} to it creates a set of interpretations that is also closed under intersection, and hence representable by a Horn formula. Lastly, the singleton $\{w_{min}\}$ is also representable by a Horn formula, as any singleton set is. Thus, the cardinality based assignment is Horn compliant. It is also straightforward to see that the assignment is pointwise faithful and total. \square

The immediate upshot, as per Theorem 2, is that the cardinality-based operator is suitable as a Horn operator.

Corollary 2. The operator \diamond^{card} satisfies postulates U_{1-9}^{Horn} and U_{SC}^{Horn} . Furthermore, it is equivalent to:

$$\psi \diamond^{card} \mu \equiv \begin{cases} \psi, & \text{if } \psi \models \mu, \\ (\psi \wedge \mu) \vee \varepsilon_{w_{min}}, & \text{if } \psi \not\models \mu \text{ but } \psi \wedge \mu \text{ is consistent,} \\ \varepsilon_{w_{min}}, & \text{otherwise, where } w_{min} = \bigcap_{w_i \in [\mu]} w_i. \end{cases}$$

The main drawback of the cardinality-based distance is that it disregards most information contained in w when it calculates \leq_w . At the same time, we have seen in Example 6 that Hamming distance does not deliver Horn-compliant outcomes. As a middle-ground, we propose the *min-Hamming distance* d_{mH} defined, for any interpretations w and w_0 , as follows:

$$d_{mH}(w, w_0) = \min\left(\{d_H(w, w_0)\} \cup \{d_H(w, w_1) \mid w_1 \notin \{w, w_0\} \text{ and } w_1 \cap w_2 = w_0, \text{ for some } w_2 \notin \{w, w_0\}\}\right).$$

The min-Hamming distance between a reference interpretation w and an interpretation w_0 looks at whether there exist interpretations w_1 and w_2 , both distinct from w and w_0 , such that $w_1 \cap w_2 = w_0$. If so, $d_{mH}(w, w_0)$ is set to the minimal Hamming distance to w over all such pairs of interpretations. Intuitively, the idea is to make sure that w_0 does not end up higher in the ranking than any distinct interpretations w_1 and w_2 such that $w_1 \cap w_2 = w_0$, since these are the kinds of situations that lead to failures of Horn compliance. If w_0 cannot be obtained as the intersection of any two other interpretations, then the distance defaults to its Hamming distance to w .

Example 18. Consider, as before, Horn formulas $\psi = a \wedge \neg b$ and $\mu = b \leftrightarrow c$ from the motivating Example 1. The preorders \leq_w^{mH} generated with the d_{mH} distance are depicted in Figure 8c. Note, for instance, that the default Hamming distance between ac and bc is $d_H(ac, bc) = 2$. There are, then, no other distinct interpretations w_1 and w_2 such that $w_1 \cap w_2 = bc$, so bc 's distance to a does not need to be corrected. On the other hand, $d_H(ac, \emptyset) = 2$, but $\emptyset = a \cap b$ and a and b are closer to ac , at a distance of 1: we thus adjust \emptyset 's distance to 1. In the end, we obtain that $[\psi \diamond^{mH} \mu] = \{\emptyset, a, abc\} = [\mu]$.

Proposition 6. The min-Hamming-based assignment is \mathcal{L}_{Horn} -compliant, pointwise faithful, and \leq_w^{mH} is a total preorder, for any interpretation w .

PROOF. It is straightforward to see that \leq_w^{mH} is pointwise faithful and total, for any interpretation w . The crucial detail is that the assignment is also Horn compliant. To see why this holds, note that d_{mH} ensures that, for any interpretations w_1 and w_2 such that $w_1 \not\subseteq w_2$ and $w_2 \not\subseteq w_1$ we have that $w_0 \leq_w^{mH} w_1$ and $w_0 \leq_w^{mH} w_2$, where $w_0 = w_1 \cap w_2$. Thus, it is not possible to have that $w_1 \approx_w^{mH} w_2 <_w^{mH} w_0$, and thus $\min_{\leq_w^{mH}}[\mu]$ is representable in the Horn fragment, for any Horn formula μ . To show that the assignment as a whole is Horn compliant, assume that $\bigcup_{w \in [\mu]} \min_{\leq_w^{mH}}[\mu]$ is not Horn-closed, which means it contains some interpretations w_1 and w_2 , but not $w_0 = w_1 \cap w_2$. Since the outcome is the union of minimal models of μ across several preorders, this means that w_0 is not among the minimal models of μ in any of these preorders, whereas w_1 and w_2 are, at least once. This means that there must exist some preorder \leq_w^{mH} , for a model $w \in [\mu]$, in which $w_1 <_w^{mH} w_0$, or $w_2 <_w^{mH} w_0$, which contradicts the conclusion derived earlier. \square

With this, it follows immediately that the \diamond^{mH} operator satisfies the strongest set of postulates.

Corollary 3. The operator \diamond^{mH} satisfies postulates U_{1-9}^{Horn} and U_{SC}^{Horn} .

We also introduce a Horn operator based on partial preorders: the *basic subset-based relation* \leq_w^{b-sub} , for an interpretation w , is defined as by taking $w_1 \leq_w^{b-sub} w_2$ if $w_1 = w$, or, in case $w_1 \neq w$, if $w_1 \subseteq w_2$. Note that \leq_w^{b-sub} is a partial preorder: it is clearly transitive, reflexive and maximally consistent, but considers interpretations incomparable if one is not a subset of the other. The resulting ranking is also pointwise faithful, maximally

consistent, and the assignment is Horn compliant, so Theorem 3 implies that the represented update operator $\diamond^{b\text{-sub}}$ satisfies postulates U_{1-8}^* and U_{wSC}^* . In fact, something even stronger holds: $\diamond^{b\text{-sub}}$ also favors interpretations with lower cardinality and turns out to be equivalent to the cardinality-based operator \diamond^{card} .

Proposition 7. For any Horn formulas ψ and μ , it holds that $\psi \diamond^{b\text{-sub}} \mu \equiv \psi \diamond^{card} \mu$.

PROOF. Note that, for any Horn formula μ , its model $w_{min} = \bigcap_{w_i \in [\mu]} w_i$ is the only $\leq_w^{b\text{-sub}}$ -minimal element in $[\mu]$ (if $w_{min} \neq w$, of course), it is also the only cardinality-minimum model of μ . Therefore, $\min_{\leq_w^{b\text{-sub}}} [\mu] = \min_{\leq_w^{card}} [\mu]$, for any interpretation w . The conclusion follows immediately. \square

Unfortunately, these operators do not work for either the Krom or affine fragments. We leave the search for interesting operators for these fragments to future work.

Finally, we point out that the standard Hamming operator, which does not work for most of the other fragments, works without any problems in the 1CNF fragment where it is equivalent to a simple syntactic updating procedure: updating ψ by μ amounts to accepting the literals of μ , even if they conflict with the literals of ψ , together with the literals of ψ untouched by μ . More precisely, recall that a 1CNF formula is a conjunction of literals that can be positive (atoms) or negative (negated atoms). In this respect, we can think of each 1CNF formula as a set of literals, and we write $\ell \in \psi$ to denote the fact that ℓ is a literal occurring in ψ , and $\ell \in (\psi \setminus \mu)$ to denote the fact that ℓ occurs in ψ but not in μ .

Proposition 8. The assignment generated by the Hamming distance is \mathcal{L}_{1CNF} -compliant. Furthermore, the \diamond^H operator restricted to the 1CNF fragment is equivalent to:

$$\psi \diamond^H \mu \equiv \bigwedge_{\ell \in \psi \setminus \mu} \ell \wedge \bigwedge_{\ell \in \mu} \ell.$$

PROOF. We show that $\psi \diamond^H \mu$ is equivalent to the literals of μ together with the literals of ψ untouched by μ , which also implies that $\psi \diamond^H \mu$ is a valid 1CNF formula.

Note that we can ignore the literals outside ψ and μ , so the task is to show that $[\psi \diamond^H \mu] = \{w^*\}$, where $w^* = \{\ell \mid \ell \in \psi \setminus \mu\} \cup \{\ell \mid \ell \in \mu\}$, i.e., w^* is the interpretation that contains all the literals of μ along with the literals of ψ untouched by μ . This involves showing that w^* is the most plausible model of μ in every preorder \leq_w^H , where $w \in [\psi]$.

To this end note, first, that w^* is clearly a model of μ . Take, now, a model w of ψ . If $w \in [\psi \wedge \mu]$, then w satisfies all the literals in both μ and ψ , i.e., it is identical to w^* and thus $d_H(w, w^*) = 0$. This means that w^* is the most plausible interpretation in \leq_w . If $w \notin [\psi \wedge \mu]$, then this means that w satisfies all the literals in ψ and flips some of the literals of μ . Let $k \geq 1$ be the exact number of literals of μ that are flipped in w , and observe that $d_H(w, w^*) = k$: this is because w and w^* agree on all the literals untouched by μ (because they end up being literals in ψ), and disagree only on the literals of μ that get flipped. Consider, now, another model w' of μ , with $w' \neq w^*$. Note that w' contains the same literals of μ as w^* (i.e., all of them), which means that w' and w disagree on at least the k literals of μ that w flips, i.e., $d_H(w, w') \geq k$. Furthermore, since $w' \neq w^*$, it must be the case that w' flips at least one literal in ψ . This implies that $d_H(w, w') \geq k + 1 > k$, and thus $w^* <_w w'$. \square

Thus, \diamond^H can be used as a 1CNF operator with the knowledge that it stays within the fragment and satisfies postulates U_{1-9}^{1CNF} and U_{SC}^{1CNF} . Note that the \diamond^H operator is not the only suitable 1CNF operator satisfying the postulates: as mentioned earlier, the drastic operator \diamond^{dr} also fits the framework.

Example 19. Consider 1CNF formulas $\psi = a \wedge \neg b$ and $\mu = b \wedge c$. We have that $\psi \diamond^{dr} \mu \equiv \mu$ and $\psi \diamond^H \mu \equiv a \wedge b \wedge c$. That is, the drastic operator \diamond^{dr} copies the information in μ , while \diamond^H copies the information of μ while also preserving the literals of ψ that do not conflict with μ .

5 Revision in Closed Fragments

Revision and update are closely connected operators (more on this connection below), with much the same problems occurring for fragment-based revision operators as occur for update operators. We thus introduce, without further ado, the formal machinery needed to model revision operators under the blanket statement that the underlying motivation for this machinery is the same as for update. An \mathcal{L}_* -revision operator is a function $\circ: \mathcal{L}_* \times \mathcal{L}_* \rightarrow \mathcal{L}_*$. The postulates we present below consist of the core Katsuno-Mendelzon set (Katsuno and Mendelzon 1991a,b), together with two additional postulates encoding the varieties of Suzumura consistency needed, and hold for any \mathcal{L}_* -formulas:

- (R₁^{*}) $\psi \circ \mu \models \mu$.
- (R₂^{*}) If $\psi \wedge \mu$ is consistent, then $\psi \circ \mu \equiv \psi \wedge \mu$.
- (R₃^{*}) If μ is consistent, then $\psi \circ \mu$ is consistent.
- (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^{*}) $(\psi \circ \mu_1) \wedge \mu_2 \models \psi \circ (\mu_1 \wedge \mu_2)$.
- (R₆^{*}) If $\psi \circ \mu_1 \models \mu_2$ and $\psi \circ \mu_2 \models \mu_1$, then $\psi \circ \mu_1 \equiv \psi \circ \mu_2$.
- (R₇^{*}) If μ_1, μ_2 are such that $\mu \equiv \mu_1 \vee \mu_2$, then $(\psi \circ \mu_1) \wedge (\psi \circ \mu_2) \models \psi \circ \mu$.
- (R₈^{*}) If $(\psi \circ \mu_1) \wedge \mu_2$ is satisfiable, then $\psi \circ (\mu_1 \wedge \mu_2) \models (\psi \circ \mu_1) \wedge \mu_2$.
- (R_{SC}^{*}) For any $n \geq 1$, if $\varepsilon_1, \dots, \varepsilon_n$ are complete formulas such that for all $i \in \{1, \dots, n-1\}$, $(\psi \circ \varepsilon_{i,i+1}) \wedge \varepsilon_i$ is consistent, then it does not hold both that $(\psi \circ \varepsilon_{n,1}) \wedge \varepsilon_n$ is consistent and that $(\psi \circ \varepsilon_{n,1}) \wedge \varepsilon_1$ is inconsistent.
- (R_{wSC}^{*}) For any $n \geq 1$, if $\varepsilon_1, \dots, \varepsilon_n$ are complete formulas such that for all $i \in \{1, \dots, n-1\}$ it holds that $(\psi \circ \varepsilon_{i,i+1}) \wedge \varepsilon_i$ is consistent and $(\psi \circ \varepsilon_{i,i+1}) \wedge \varepsilon_{i+1}$ is inconsistent, then it does not hold that $(\psi \circ \varepsilon_{n,1}) \wedge \varepsilon_n$ is consistent.

As with update, postulates R₁₋₅^{*} and R₈^{*} form the strongest set of postulates available in propositional logic, i.e., R₁₋₅^{id} and R₈^{id} imply R₆₋₇^{*}, as well as R_{SC}^{id} and R_{wSC}^{id} in propositional logic. We omit the proofs here, as they are entirely similar to the ones from Section 3. The similarity between the revision and update postulates goes even further: postulate R₂^{*} implies U₂^{*}, and is equivalent to it if ψ is complete; postulates R₁^{*}, R₃₋₇^{*}, R_{SC}^{*} and R_{wSC}^{*} are almost identical to postulates U₁^{*}, U₃₋₇^{*}, U_{SC}^{*} and U_{wSC}^{*}, respectively; and, finally, postulate R₈^{*} corresponds to U₉^{*}. The only difference between these sets is the omission of completeness of ψ as a pre-condition for the revision postulates. There is no revision equivalent to postulate U₈^{*}, which explains this omission: the prior belief ψ needs not be complete in any of the revision postulates, hence there is no need for instructions on how to combine results from such scenarios, which is the role of postulate U₈^{*} in update.

The connection between the two types of operators becomes clearer when looking at them through the angle of choice functions on rankings of interpretations, though we need a bit more notation to make this formal. A *faithful assignment* maps any formula ψ to a preorder, possibly partial, \leq_ψ over interpretations such that: (i) if $w_1 \in [\psi]$ and $w_2 \notin [\psi]$, then $w_1 <_\psi w_2$, (ii) if $w_1, w_2 \in [\psi]$, then it does not hold that $w_1 <_\psi w_2$, and (iii) if $\psi_1 \equiv \psi_2$, then $\leq_{\psi_1} = \leq_{\psi_2}$. Faithful assignments make sure models of ψ are the most plausible elements, and rank everything else as less plausible, and they are a generalization of the pointwise faithful assignments used in Section 3. The key to understanding the workings of the revision postulates lies in how a revision operator \circ is *represented* by an assignment (alternatively, we say the assignment *represents* \circ): this happens if $[\psi \circ \mu] = \min_{\leq_\psi} [\mu]$, for any formulas ψ and μ . Naturally enough, an assignment is \mathcal{L}_* -compliant if $\min_{\leq_\psi} [\mu]$ is representable by an \mathcal{L}_* -formula, for any \mathcal{L}_* -formula μ .

Note, now, that revision occurs by taking the minimal models of the new information μ in the preorder \leq_ψ assigned to the prior belief ψ : the revision postulates ensure that this choice is done in orderly fashion. An update operator does the same thing, but across the preorders \leq_w assigned to the models w of μ , with all the postulates except U₈^{*} responsible for making this choice coherent. In the case of update, postulate U₈^{*} then specifies how to combine the results of all the mini-revision events into one outcome. Revision dispenses with this last step, and does not need an equivalent to postulate U₈^{*}.

In summary, revision operators choose from one order (\leq_ψ), while update operators compile the outcome by gathering choices over many orders (the orders \leq_w , for $w \in [\psi]$). The standard representation result for revision says that a propositional revision operator \circ satisfies postulates R_{1-8}^{id} if and only if there exists a faithful assignment mapping each formula ψ to a total preorder \leq_ψ that represents \circ (Katsuno and Mendelzon 1991b). Postulates R_{1-7}^* are similarly represented using partial preorders.

In fragments we obtain a similar set of representation results, but only with additional postulates and restrictions on the assignments.

Theorem 4. If \mathcal{L}_* is a closed fragment, an \mathcal{L}_* -revision operator \circ satisfies postulates R_{1-8}^* and R_{SC}^* iff there exists an \mathcal{L}_* -compliant, faithful assignment mapping each formula ψ to a total preorder \leq_ψ that represents \circ .

This result can be found in previous work (Delgrande and Peppas 2015; Delgrande, Peppas, and Woltran 2018), though using a different version of the R_{SC}^* postulate. We include it here to highlight the usefulness of thinking in terms of Suzumura consistency.

For partial preorders, we bring back maximally consistent assignments and tight fragments from Section 3.2.

Theorem 5. If \mathcal{L}_* is a closed, tight fragment, an \mathcal{L}_* -revision operator \circ satisfies postulates R_{1-7}^* and R_{wSC}^* iff there exists an \mathcal{L}_* -compliant faithful assignment mapping each formula ψ to a maximally consistent partial preorder \leq_ψ that represents \circ .

The proofs are entirely similar to the proofs of Theorems 2 and 3, so we do not restate them here.

Connections between revision and update that allow transfer of results between the two operators have been explored in greater detail in previous work (Aravanis 2025; Peppas et al. 1996). It would be useful to explore their relationship in fragments, in light of the representation results obtained here.

6 Conclusion

We have identified a series of constraints, both on the postulates and the preorders, that allow us to represent fragment-based belief update operators using total preorders (Theorem 2) and partial preorders (Theorem 3). For partial preorders we also require the fragment to be *tight*, but, as shown in Section 4, the non-tight fragment we care about (1CNF) has such a particular structure as to make update in it non-problematic. A pleasant consequence of setting up the results for update is that their lessons could then be adapted with minimal effort to revision in closed fragments (Section 5). Fragment-based revision with total preorders is well understood (Delgrande and Peppas 2015; Delgrande, Peppas, and Woltran 2018), but similar results for belief update, and revision with partial preorders, have been lacking. We have also presented a series of concrete update operators (Section 4), though more work is needed to find interesting operators for the Krom and affine fragments.

As apparent from Section 3, representation results in fragments rely on finding a way to balance the intention of the postulates with the expressive capabilities of the fragments, and are generally non-trivial. A significant contribution of our work lies in highlighting the importance of Suzumura consistency for the representation results: though known in the social choice literature, its importance in belief change scenarios has so far been mostly implicit. As our results suggest, the manner in which Suzumura consistency is enforced will vary depending on the expressive capabilities of the fragment we are working with, though, as the minimal property needed to ensure that a relation can be extended to a total preorder (Suzumura 1976), it will need to show up in one way or another in any axiomatization. Thus, an interesting avenue for future work is the relationship between a fragment's expressive capabilities and the types of representation results available: in Section 3.2 we appealed to a tightness condition on the fragments, but it would be useful to understand the minimal conditions on the closure operator that still yield nice representation results.

Acknowledgments

The authors would like to thank the anonymous reviewers for their helpful comments and suggestions. This research was supported by the Austrian Science Fund (FWF) grant 10.557766/COE12 and by a Leverhulme Trust International Professorship Grant (LIP-2022-001).

References

- M. Aiguier, J. Atif, I. Bloch, and C. Hudelot. 2018. “Belief revision, minimal change and relaxation: a general framework based on satisfaction systems, and applications to description logics.” *Artificial Intelligence*, 256, 160–180.
- T. I. Aravanis. 2025. “On the consistency between belief revision and belief update.” *Journal of Artificial Intelligence Research*, 82, 1743–1771.
- S. Benferhat, S. Lagrue, and O. Papini. 2005. “Revision of partially ordered information: axiomatization, semantics and iteration.” In: *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI 2005)*, 376–381.
- G. Bonanno. 2009. “Rational choice and AGM belief revision.” *Artificial Intelligence*, 173, 12–13, 1194–1203.
- R. Booth, T. Meyer, I. Varzinczak, and R. Wassermann. 2011. “On the link between partial meet, kernel, and infra contraction and its application to Horn logic.” *Journal of Artificial Intelligence Research*, 42, 31–53.
- R. Booth, T. Meyer, and I. Varzinczak. 2009. “Next steps in propositional Horn contraction.” In: *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, 702–707.
- D. Borchmann, T. Hanika, and S. Obiedkov. 2020. “Probably approximately correct learning of Horn envelopes from queries.” *Discrete Applied Mathematics*, 273, 30–42.
- W. Bossert and K. Suzumura. 2010. *Consistency, Choice, and Rationality*. Harvard University Press.
- N. Creignou, A. Haret, O. Papini, and S. Woltran. 2018. “Belief update in the Horn fragment.” In: *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI 2018)*, 1781–1787.
- N. Creignou, R. Ktari, and O. Papini. 2022. “Belief contraction and erasure in fragments of propositional logic.” *Journal of Logic and Computation*, 32, 7, 1436–1468.
- N. Creignou, R. Ktari, and O. Papini. 2018. “Belief update within propositional fragments.” *Journal of Artificial Intelligence Research*, 61, 807–834.
- N. Creignou, O. Papini, R. Pichler, and S. Woltran. 2014. “Belief revision within fragments of propositional logic.” *Journal of Computer and System Sciences*, 80, 2, 427–449.
- N. Creignou, O. Papini, S. Rümmele, and S. Woltran. 2016. “Belief merging within fragments of propositional logic.” *ACM Transactions on Computational Logic*, 17, 3, 1–28.
- R. Dechter and J. Pearl. 1992. “Structure identification in relational data.” *Artificial Intelligence*, 58, 1, 237–270.
- J. Delgrande. 2008. “Horn clause belief change: contraction functions.” In: *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR 2008)*, 156–165.
- J. Delgrande and P. Peppas. 2015. “Belief revision in Horn theories.” *Artificial Intelligence*, 218, 1–22.
- J. Delgrande, P. Peppas, and S. Woltran. 2018. “General belief revision.” *Journal of the ACM*, 65, 5, 1–34.
- J. Delgrande, T. Schaub, H. Tompits, and S. Woltran. 2013. “A model-theoretic approach to belief change in answer set programming.” *ACM Transactions on Computational Logic*, 14, 2, 14:1–14:46.
- J. Delgrande and R. Wassermann. 2013. “Horn clause contraction functions.” *Journal of Artificial Intelligence Research*, 48, 457–511.
- J. Delgrande and R. Wassermann. 2010. “Horn clause contraction functions: belief set and belief base approaches.” In: *Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning (KR 2010)*.
- J. Delobelle, A. Haret, S. Konieczny, J.-G. Mailly, J. Rossit, and S. Woltran. 2016. “Merging of abstract argumentation frameworks.” In: *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning (KR 2016)*, 33–42.
- M. Diller, A. Haret, T. Linsbichler, S. Rümmele, and S. Woltran. 2018. “An extension-based approach to belief revision in abstract argumentation.” *International Journal of Approximate Reasoning*, 93, 395–423.
- R. Fagin, J. Ullman, and M. Vardi. 1983. “On the semantics of updates in databases.” In: *Proceedings of the 2nd ACM SIGACT-SIGMOD Symposium on Principles of Database Systems (PODS 1983)*, 352–365.
- F. Falakh, S. Rudolph, and K. Sauerwald. 2025. “AGM belief revision, semantically.” *ACM Transactions on Computational Logic*, 26, 4, 1–54.
- K. Forbus. 1989. “Introducing actions into qualitative simulation.” In: *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI 1989)*, 1273–1278.
- A. Haret. 2020. “Choosing what to believe: belief change through the lens of rational choice.” Ph.D. Dissertation. Technische Universität Wien.
- A. Haret, S. Rümmele, and S. Woltran. 2017. “Merging in the Horn fragment.” *ACM Transactions on Computational Logic*, 18, 1, 6:1–6:32.
- A. Haret and J. Wallner. 2022. “An axiomatic approach to revising preferences.” *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI 2022)*, 36, 5, 5676–5683.

- A. Haret, J. Wallner, and S. Woltran. 2018. “Two sides of the same coin: Belief revision and enforcing arguments.” In: *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI 2018)*, 1854–1860.
- A. Herzig and O. Rifi. 1999. “Propositional belief base update and minimal change.” *Artificial Intelligence*, 115, 1, 107–138.
- H. Katsuno and A. Mendelzon. 1991a. “On the difference between updating a knowledge base and revising it.” In: *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR 1991)*, 230–237.
- H. Katsuno and A. Mendelzon. 1991b. “Propositional knowledge base revision and minimal change.” *Artificial intelligence*, 52, 3, 263–294.
- H. Kautz, M. Kearns, and B. Selman. 1995. “Horn approximations of empirical data.” *Artificial Intelligence*, 74, 1, 129–145.
- E. Kharlamov, D. Zheleznyakov, and D. Calvanese. 2013. “Capturing model-based ontology evolution at the instance level: The case of DL-Lite.” *Journal of Computer and System Sciences*, 79, 6, 835–872.
- R. Ktari. 2016. “Changement de croyances dans des fragments de la logique propositionnelle.” Ph.D. Dissertation. Aix-Marseille Université.
- J. Ma, S. Benferhat, and W. Liu. 2012. “Revising partial pre-orders with partial pre-orders: a unit-based revision framework.” In: *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR 2012)*, 633–637.
- P. Peppas, A. Nayak, M. Pagnucco, N. Foo, R. Kwok, and M. Prokopenko. 1996. “Revision vs. update: taking a closer look.” In: *Proceedings of the 12th European Conference on Artificial Intelligence (ECAI 1996)*, 95–99.
- G. Qi and J. Du. 2009. “Model-based revision operators for terminologies in description logics.” In: *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI 2009)*, 891–897.
- G. Qi, W. Liu, and D. Bell. 2006. “Knowledge base revision in description logics.” In: *Proceedings of the 10th European Conference on Logics in Artificial Intelligence (JELIA 2006)*. Vol. 4160, 386–398.
- H. Rott. 2001. *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford University Press.
- T. Schaefer. 1978. “The complexity of satisfiability problems.” In: *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing (STOC 1978)*, 216–226.
- A. Sen. 1971. “Choice functions and revealed preference.” *The Review of Economic Studies*, 38, 3, 307.
- M. Sérayet, P. Drap, and O. Papini. 2009. “Encoding the revision of partially preordered information in answer set programming.” In: *Proceedings of the 10th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2009)*, 421–433.
- K. Suzumura. 1976. “Remarks on the theory of collective choice.” *Economica*, 43, 172, 381–390.
- M. Winslett. 1990. *Updating Logical Databases*. Cambridge University Press.
- Z. Zhuang and M. Pagnucco. 2012. “Model based Horn contraction.” In: *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR 2012)*.
- Z. Zhuang, J. Delgrande, A. Nayak, and A. Sattar. 2016. “Reconsidering AGM-style belief revision in the context of logic programs.” In: *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI 2016)*, 671–679.
- Z. Zhuang and M. Pagnucco. 2014. “Entrenchment-based Horn contraction.” *Journal of Artificial Intelligence Research*, 51, 227–254.
- Z. Zhuang, M. Pagnucco, and Y. Zhang. 2013. “Definability of Horn revision from Horn contraction.” In: *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*, 1205–1212.
- Z. Zhuang, M. Pagnucco, and Y. Zhang. 2017. “Inter-definability of Horn contraction and Horn revision.” *Journal of Philosophical Logic*, 46, 3, 299–332.
- Z. Zhuang, Z. Wang, K. Wang, and J. Delgrande. 2019. “A generalisation of AGM contraction and revision to fragments of first-order logic.” *Journal of Artificial Intelligence Research*, 64, 147–179.
- Z. Zhuang, Z. Wang, K. Wang, and J. Delgrande. 2015. “Extending AGM contraction to arbitrary logics.” In: *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, 3299–3307.
- Z. Zhuang, Z. Wang, K. Wang, and G. Qi. 2016. “DL-Lite contraction and revision.” *Journal of Artificial Intelligence Research*, 56, 329–378.

Received 05 August 2025; accepted 31 March 2026