dbai

TECHNICAL REPORT

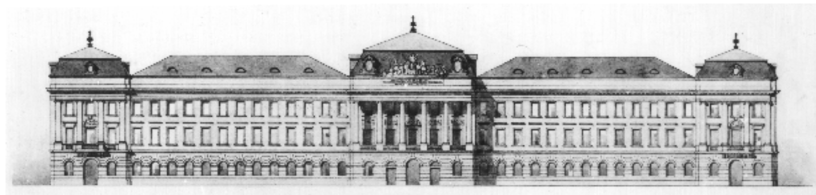

INSTITUT FÜR INFORMATIONSSYSTEME

ABTEILUNG DATENBANKEN UND ARTIFICIAL INTELLIGENCE

Merging in the Horn Fragment

DBAI-TR-2015-91

Adrian Haret, Stefan Rümmele, Stefan Woltran

DBAI TECHNICAL REPORT
2015

Institut für Informationssysteme
Abteilung Datenbanken und Artificial Intelligence
Technische Universität Wien
Favoritenstr. 9
A-1040 Vienna, Austria
Tel: +43-1-58801-18403
Fax: +43-1-58801-18493
sekret@dbai.tuwien.ac.at
www.dbai.tuwien.ac.at

TECHNISCHE UNIVERSITÄT WIEN
Vienna University of Technology

DBAI Technical Report
DBAI Technical Report DBAI-TR-2015-91, 2015

Technical Report

Adrian Haret [1] **Stefan Rümmele** [1] **Stefan Woltran** [1]

Abstract. Belief merging is a central operation within the field of belief change and addresses the problem of combining multiple, possibly mutually inconsistent knowledge bases into a single, consistent one. A current research trend in belief change is concerned with tailored representation theorems for fragments of logic, in particular Horn logic. Hereby, the goal is to guarantee that the result of the change operations stays within the fragment under consideration. While several such results have been obtained for Horn revision and Horn contraction, merging of Horn theories has been neglected so far. In this paper, we provide a novel representation theorem for Horn merging by strengthening the standard merging postulates. Moreover, we present a concrete Horn merging operator satisfying all postulates.

[1]TU Wien. E-mail: {haret, ruemmele, woltran}@dbai.tuwien.ac.at

Acknowledgements: This work has been supported by the Austrian Science Fund (FWF) under grants P25518 and P25521.

Contents

1 Introduction

Belief merging uses a logical approach to study how information coming from multiple, possibly mutually inconsistent knowledge bases should be combined to form a single, consistent knowledge base. Merging shares a common methodology with other belief change operators such as revision [1, 7], contraction [1] and update [6]. This methodology consists, first of all, in formulating logical postulates which any rational operator should satisfy. For merging, the IC-merging postulates [8, 9] are commonly used. In a further step, a representation result is usually derived: this shows that all (merging) operators satisfying the postulates can be characterized using rankings on the possible worlds described by the underlying language, which is typically taken to be full propositional logic.

Recently, the restriction of belief change formalisms to different fragments of propositional logic has become a vivid research branch. There are pragmatic reasons for focusing on fragments, especially Horn logic. Firstly, Horn clauses are a natural way of formulating basic facts and rules, and thus are useful to encode expert knowledge. Secondly, Horn logic affords very efficient algorithms. Thus the computational cost of reasoning in this fragment is comparatively low.

While revision [4, 11, 13] and contraction [2, 5, 12] have received a lot of attention in this direction, belief merging has yet remained unexplored, with the notable exception of [3]. We aim to fill this gap and investigate the problem of merging in the Horn fragment of propositional logic. We find that restricting the underlying language poses a series of non-trivial challenges, as representation results which work for full propositional logic break down in the Horn case.

Firstly, we find that we cannot rely on the same types of rankings as the ones used for merging in the case of full propositional logic. The reason is that such rankings lead to outputs that cannot be expressed as Horn formulas. We fix this problem by adding the restriction of *Horn compliance*: this narrows down the notion of ranking in a way that is coherent with the semantics of Horn formulas. Since standard merging operators are found not to be Horn compliant (hence useless for our needs), we also give a concrete operator that exhibits this property. This is remarkable, as previous research [3] resulted only in Horn merging operators that do not satisfy all postulates. Secondly, Horn merging operators that satisfy the standard postulates turn out to represent more rankings than was expected, some of which are undesirable. We interpret this as an inadequacy of the standard postulates to capture the intended intuitive behaviour of a merging operator. Hence, we propose an alternative formulation of some key postulates, which allows us to derive an appealing representation result for the case of Horn merging. Our approach here is inspired by existing work on Horn revision [4], though we go significantly beyond it to tackle the problems posed by merging.

The rest of the paper is organized as follows. In Section 2 we introduce the background to merging. In Section 3 we argue that standard model-based merging operators are inappropriate for Horn merging and introduce the property of Horn compliance. In Section 4 we argue that a subset of the IC-merging postulates should be replaced by a strengthened version, and introduce a representation result using the strengthened postulates. Finally, in Section 5 we describe a concrete Horn merging operator satisfying all postulates.

2 Preliminaries

Propositional logic. We work with the language $\mathcal{L}$ of propositional logic over a fixed *alphabet* $\mathcal{P} = \{p_1, \ldots, p_n\}$ of propositional atoms. We use standard connectives $\vee$, $\wedge$, $\neg$ and the logical constants $\top$ and $\bot$. A literal is an atom or a negated atom. A clause is a disjunction of literals. A clause is called *Horn* if at most one of its literals is positive. The Horn fragment $\mathcal{L}_H \subset \mathcal{L}$ is the set of all formulas in $\mathcal{L}$ that are conjunctions of Horn clauses. An interpretation is a set $w \subseteq \mathcal{P}$ of atoms. The set of all interpretations is denoted by $\mathcal{W}$. We will typically represent an interpretation by its corresponding bit-vector of length $|\mathcal{P}|$. As an example, if $|\mathcal{P}| = 3$, then 101 is the interpretation $\{p_1, p_3\}$. A pre-order $\leq$ on $\mathcal{W}$ is a reflexive, transitive binary relation on $\mathcal{W}$. If $w_1, w_2 \in \mathcal{W}$, then $w_1 < w_2$ denotes the strict part of $\leq$, i.e., $w_1 \leq w_2$ but $w_2 \nleq w_1$. We write $w_1 \approx w_2$ to abbreviate $w_1 \leq w_2$ and $w_2 \leq w_1$. If $\mathcal{M}$ is a set of interpretations, then the set of *minimal elements of $\mathcal{M}$ with respect to* $\leq$ is defined as $min_{\leq}\mathcal{M} = \{w_1 \in \mathcal{M} \mid \nexists w_2 \in \mathcal{M} \text{ s.t. } w_2 < w_1\}$. If interpretation w satisfies formula φ, we call *w a model of φ*. We denote the set of models of φ by $[\varphi]$. Given a set $\mathcal{M}$ of interpretations, we define $\mathrm{Cl}_\cap(\mathcal{M})$, the closure of $\mathcal{M}$ under intersection, as the smallest superset of $\mathcal{M}$ that is closed under $\cap$. By closure under $\cap$ we mean that if $w_1, w_2 \in \mathrm{Cl}_\cap(\mathcal{M})$ then also $w_1 \cap w_2 \in \mathrm{Cl}_\cap(\mathcal{M})$. We recall here a classic result concerning Horn formulas and their models (see, for example, [10]).

Proposition 1. *A set of interpretations $\mathcal{M}$ is the set of models of a Horn formula φ if and only if $\mathcal{M} = Cl_\cap(\mathcal{M})$.*

Proposition 1 highlights the expressive limitations of the Horn fragment. It shows, for instance, that the set of interpretations $\{01, 10\}$, which is easily representable in full propositional logic by the formula $(p_1 \wedge \neg p_2) \vee (\neg p_1 \wedge p_2)$, cannot be represented by a Horn formula. However, we can represent the closure $Cl_\cap(\{01, 10\}) = \{00, 01, 10\}$ by the Horn formula $\neg p_1 \vee \neg p_2$.

This expressive limitations of the Horn fragment motivate the following notions. A formula is called *complete* if it has exactly one model. If w_i is an interpretation, we write σ_{w_i} or σ_i to denote a complete formula that has w_i as a model. In other words, σ_i is a propositional formula such that $[\sigma_i] = \{w_i\}$. As an observation, it is easy to see that for any singleton $\{w_i\}$ we can find a complete propositional Horn formula σ_i such that $[\sigma_i] = \{w_i\}$. For instance, the singleton $\{01\}$ is representable by the Horn formula $\neg p_1 \wedge p_2$.

If σ_i and σ_j are complete formulas, then $\sigma_{i,j}$ is a formula such that $[\sigma_{i,j}] = \{w_i, w_j\}$. If we are working in the Horn fragment, we take $\sigma_{i,j}$ to be a Horn formula such that $[\sigma_{i,j}] = Cl_\cap(\{w_i, w_j\})$.

Belief Merging. A *knowledge base* is a finite set of propositional formulas. A *profile* is a non-empty finite multi-set $E = \{K_1, \ldots, K_n\}$ of consistent knowledge bases. *Horn knowledge bases* and *Horn profiles* contain only Horn formulas and Horn knowledge bases, respectively. The sets of all knowledge bases, Horn knowledge bases, profiles and Horn profiles are denoted by $\mathcal{K}$, $\mathcal{K}_H$, $\mathcal{E}$ and $\mathcal{E}_H$, respectively. If E_1 and E_2 are profiles, then $E_1 \sqcup E_2$ is the multi-set union of E_1 and E_2. Interpretation w is a model of knowledge base K if it is a model of every formula in K. Interpretation w is a model of profile E if it is a model of every $K \in E$. We denote by $[K]$ and $[E]$ the set of models of K and E, respectively. We write $\bigwedge E$ for $\bigwedge_{K \in E} \bigwedge_{\varphi \in K} \varphi$. This reduces a profile to a single propositional formula. Clearly, $[\bigwedge E] = [E]$.

Profiles E_1 and E_2 are *equivalent*, written $E_1 \equiv E_2$, if there exists a bijection $f \colon E_1 \to E_2$ such that for any $K \in E_1$ we have $[K] = [f(K)]$.

A merging operator is a function $\Delta \colon \mathcal{E} \times \mathcal{L} \to \mathcal{K}$. It maps a profile E and a formula μ, typically referred to as the constraint, onto a knowledge base. We write $\Delta_\mu(E)$ instead of $\Delta(E, \mu)$. As is common in the belief change literature, logical postulates are employed to set out properties which any merging operator Δ should satisfy. An operator satisfying the following postulates is called IC-merging operator [8, 9]:

(IC$_0$) $\Delta_\mu(E) \models \mu$.

(IC$_1$) If μ is consistent, then $\Delta_\mu(E)$ is consistent.

(IC$_2$) If $\bigwedge E$ is consistent with μ, then $\Delta_\mu(E) \equiv \bigwedge E \wedge \mu$.

(IC$_3$) If $E_1 \equiv E_2$ and $\mu_1 \equiv \mu_2$, then $\Delta_{\mu_1}(E_1) \equiv \Delta_{\mu_2}(E_2)$.

(IC$_4$) If $K_1 \models \mu$ and $K_2 \models \mu$, then $\Delta_\mu(\{K_1, K_2\}) \wedge K_1$ is consistent iff $\Delta_\mu(\{K_1, K_2\}) \wedge K_2$ is consistent.

(IC$_5$) $\Delta_\mu(E_1) \wedge \Delta_\mu(E_2) \models \Delta_\mu(E_1 \sqcup E_2)$.

(IC$_6$) If $\Delta_\mu(E_1) \wedge \Delta_\mu(E_2)$ is consistent, then $\Delta_\mu(E_1 \sqcup E_2) \models \Delta_\mu(E_1) \wedge \Delta_\mu(E_2)$.

(IC$_7$) $\Delta_{\mu_1}(E) \wedge \mu_2 \models \Delta_{\mu_1 \wedge \mu_2}(E)$.

(IC$_8$) If $\Delta_{\mu_1}(E) \wedge \mu_2$ is consistent, then $\Delta_{\mu_1 \wedge \mu_2}(E) \models \Delta_{\mu_1}(E) \wedge \mu_2$.

Though these postulates tell what properties $\Delta_\mu(E)$ should have, they do not spell out how to actually construct $\Delta_\mu(E)$, given E and μ. Towards this latter task, it is useful to consider what happens at the interpretation level when merging occurs. The notion of a syncretic assignment proves crucial here.

Definition 1. A *syncretic assignment* is a function mapping each $E \in \mathcal{E}$ to a pre-order $\leq_E$ on $\mathcal{W}$ such that, for any $E, E_1, E_2 \in \mathcal{E}$, $K_1, K_2 \in \mathcal{K}$ and $w_1, w_2 \in \mathcal{W}$ the following conditions hold:

(s_1) If $w_1 \in [E]$ and $w_2 \in [E]$, then $w_1 \approx_E w_2$.

(s_2) If $w_1 \in [E]$ and $w_2 \notin [E]$, then $w_1 <_E w_2$.

(s_3) If $E_1 \equiv E_2$, then $\leq_{E_1} = \leq_{E_2}$.

(s_4) If $w_1 \in [K_1]$, then there is $w_2 \in [K_2]$ such that $w_2 \leq_{\{K_1,K_2\}} w_1$.

(s_5) If $w_1 \leq_{E_1} w_2$ and $w_1 \leq_{E_2} w_2$, then $w_1 \leq_{E_1 \sqcup E_2} w_2$.

(s_6) If $w_1 \leq_{E_1} w_2$ and $w_1 <_{E_2} w_2$, then $w_1 <_{E_1 \sqcup E_2} w_2$.

We define syncretic assignments in a way that allows the pre-orders $\leq_E$ to be partial, as we will make use of partial pre-orders in our own results on Horn operators (Theorems 3 and 4). In the context of full propositional logic, however, the classical result below characterizes all IC-merging operators in terms of syncretic assignments with total pre-orders.

Theorem 1 ([8, 9])**.** *A merging operator Δ is an IC-merging operator if and only if there exists a syncretic assignment mapping each $E \in \mathcal{E}$ to a total pre-order $\leq_E$ such that $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any $\mu \in \mathcal{L}$.*

When this equation holds we will say that the assignment represents *the operator.*

It is useful to think of a profile $E = \{K_1, \ldots, K_n\}$ as a multi-set of agents, represented by their sets of beliefs K_i. Each agent is equipped with a pre-order $\leq_{K_i}$ on $\mathcal{W}$ which can be thought of as the way in which the agent ranks possible worlds in terms of their plausibility. Merging is then the task of finding a common ranking that approximates, as best as possible, the individual rankings. Proposition 1 tells us that if this process is done using syncretic assignments, we are in agreement with postulates IC_0 – IC_8.

Two parts need to be filled out to get a concrete merging operator: how to compute the individual rankings and how to aggregate them. For computing the individual rankings, the common approach in the literature is to use some notion of distance between interpretations, such as Hamming distance d_H or the drastic distance d_D. The minimal distance between interpretations and models of K_i is used to construct $\leq_{K_i}$. For aggregating the rankings, common functions used are the sum Σ or $GMAX$, giving

w	$d(w, K_1)$	$d(w, K_2)$	$d(w, K_3)$	Σ	$GMAX$
000	3	1	1	5	(3,1,1)
001	2	0	0	2	(2,0,0)
010	2	0	1	3	(2,0,1)
011	1	1	0	2	(1,1,0)
100	2	2	1	5	(2,2,1)
101	1	1	0	2	(1,1,0)
110	1	1	2	4	(2,1,1)
111	0	2	1	3	(2,1,0)

Table 1: Example of how $\Delta_{\mu}^{d_H,\Sigma}(E_1)$ is computed.

us operators $\Delta^{d_H,\Sigma}$, $\Delta^{d_H,GMAX}$, $\Delta^{d_D,\Sigma}$, $\Delta^{d_H,GMAX}$. See Example 1 to see how $\Delta^{d_H,\Sigma}$ and $\Delta^{d_H,GMAX}$ are computed in a concrete case, and [8, 9] for more details.

Example 1. Take $K_1 = \{p_1, p_2, p_3\}$, $K_2 = \{\neg p_1,\ p_2 \leftrightarrow \neg p_3\}$, $K_3 = \{p_3,\ p_1 \to \neg p_2\}$, $E_1 = \{K_1, K_2, K_3\}$ and $\mu = p_3 \wedge (p_1 \to p_2)$. To compute $\Delta_{\mu}^{d_H,\Sigma}(E)$ and $\Delta_{\mu}^{d_H,GMAX}(E)$ we first use (in this case) Hamming distance d_H to derive a ranking for each of the knowledge bases K_i. The rankings are based on the distance between interpretations w and the knowledge bases K_i, denoted by $d(w, K_i)$ and shown in Table 1. The distance $d(w, K_i)$ between w and K_i is computed by taking the minimal distance (in this case, Hamming distance d_H) between w and all $w' \in [K_i]$. For instance:

$$\begin{aligned} d(010, K_3) &= \min\{d_H(010, w') \mid w' \in [K_3]\} \\ &= \min\{d_H(010, 001), d_H(010, 011), d_H(010, 101)\} \\ &= \min\{2, 1, 3\} \\ &= 1. \end{aligned}$$

Thus 1 is the level of 010 in the ranking for K_3, and is the entry in Table 1 for the $d(010, K_3)$ field. Once all the numbers are in, we use an aggregation function (in this case Σ and $GMAX$) to obtain the final ranking $\leq_E$. The aggregation function Σ simply adds the numbers interpretation-wise, and the final ranking $\leq_E^{\Sigma}$ is determined by the order of the final levels for each interpretation. The aggregation function $GMAX$ takes the ordered vector of levels for each interpretation, and then determines $\leq_E^{GMAX}$ by ordering the vectors lexicographically. Thus, $GMAX$ gives us that $001 \leq_E^{GMAX} 111$, since $(2, 0, 0)$ is lexicographically smaller than $(2, 1, 0)$.

Finally, we look at the models of the constraint μ (highlighted in grey in Table 1) and pick out the ones that have the minimal levels in the final ranking. In our case, we get:

$$[\Delta_{\mu}^{d_H,\Sigma}(E)] = min_{\leq_E^{\Sigma}}[\mu] = \{001, 011\},$$
$$[\Delta_{\mu}^{d_H,GMAX}(E)] = min_{\leq_E^{GMAX}}[\mu] = \{011\}.$$

w	$d(w,K)$	Σ	$GMAX$
00	2	2	(2)
01	1	1	(1)
10	1	1	(1)
11	0	0	(0)

Table 2: $\Delta_{\mu}^{d_H,\Sigma}(E_1)$ and $\Delta_{\mu}^{d_H,GMAX}(E_1)$ do not stay in the Horn fragment.

Thus, the merging operators $\Delta^{d_H,\Sigma}$ and $\Delta^{d_H,GMAX}$ produce a set of models: the result, we see, is different depending on the aggregation function used. We can then take that set of models and represent it by a propositional knowledge base: $\{\neg p_1, p_3\}$ in the case of $\Delta_{\mu}^{d_H,\Sigma}(E)$, and $\{\neg p_1, p_2, p_3\}$ in the case of $\Delta_{\mu}^{d_H,GMAX}(E)$.

3 Restricting assignments: Horn compliance

A Horn merging operator is a function $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$. Our aim is to characterize the class of such operators in the manner of Proposition 1 and to exhibit a concrete operator. However, we encounter a first problem when we try to apply standard merging operators to Horn profiles and formulas: it turns out the output of standard merging operators cannot always be represented by a Horn formula.

Examples 2 and 3 show how standard merging operators fail in the Horn fragment. In both cases we choose Horn profiles over the 2-letter alphabet and construct rankings using the Hamming distance and the drastic distance. We then aggregate the rankings with Σ and $GMAX$. The rankings and the result of their aggregation are shown in Tables 2 and 3. Each row displays one possible interpretation over $\{p_1, p_2\}$ (denoted in the first column). The second column displays the minimal distance of the interpretation to any model of K. The third and fourth column show the aggregation of the distances (in our case only one distance) according to the Σ as well as the $GMAX$ function. The output of the merging operator is the set of those interpretations that are models of μ (marked grey) and have the smallest aggregated value.

Example 2. Take $K = \{p_1,\ p_2\}$, $E_1 = \{K\}$ and $\mu = \neg p_1 \vee \neg p_2$ (all of them Horn). We compute $\Delta_{\mu}^{d_H,\Sigma}(E_1)$ and $\Delta_{\mu}^{d_H,GMAX}(E_1)$, keeping in mind that $[K] = \{11\}$ and $[\mu] = \{00, 10, 01\}$. Table 2 displays how the output of these two mergings is computed.

We get $[\Delta_{\mu}^{d_H,\Sigma}(E_1)] = min[\mu] = \{10, 01\}$, and the same result is obtained for $\Delta_{\mu}^{d_H,GMAX}(E_1)$. It holds, then, that $\Delta_{\mu}^{d_H,\Sigma}(E_1)$ and $\Delta_{\mu}^{d_H,GMAX}(E_1)$ cannot be expressed as Horn formulas.

Example 3. Take $K_1 = \{p_1,\ \neg p_2\}$, $K_2 = \{\neg p_1,\ p_2\}$, $E_2 = \{K_1,\ K_2\}$ and $\mu = \neg p_1 \vee \neg p_2$: We get (see Table 3) $[\Delta_{\mu}^{d_D,\Sigma}(E_2)] = min[\mu] = \{10, 01\}$, and the same result is obtained

w	$d(w,K_1)$	$d(w,K_2)$	Σ	$GMAX$
00	1	1	2	(1,1)
01	1	0	1	(1,0)
10	0	1	1	(1,0)
11	1	1	2	(1,1)

Table 3: $\Delta_\mu^{d_D,\Sigma}(E_2)$ and $\Delta_\mu^{d_D,GMAX}(E_2)$ do not stay in the Horn fragment.

for $\Delta_\mu^{d_D,GMAX}(E_2)$. Again, $\Delta_\mu^{d_D,\Sigma}(E_2)$ and $\Delta_\mu^{d_D,GMAX}(E_2)$ cannot be represented by Horn formulas.

As we see in Examples 2 and 3, standard distance and aggregation functions are not adequate for the Horn fragment. To circumvent this problem we adopt a solution used for Horn revision [4], which is to impose an extra condition on the pre-orders.

Definition 2. A pre-order $\leq$ is *Horn compliant* if for any $\mu \in \mathcal{L}_H$, $min_{\leq}[\mu]$ can be represented by a Horn formula.

Example 4. The computed pre-orders for E_1 and E_2 in Examples 2 and 3 are not Horn compliant, as when $\mu = \neg p_1 \vee \neg p_2$ we get that $min[\mu] = \{01, 10\}$ in both cases, and $\{01, 10\}$ is not closed under intersection.

Adding Horn compliance makes it possible to define a merging operator for the Horn fragment, and this gives us one direction of a representation theorem.

Theorem 2. *If there exists a syncretic assignment mapping every $E \in \mathcal{E}_H$ to a Horn compliant total pre-order $\leq_E$, then we can define an operator $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$ by taking $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any $\mu \in \mathcal{L}_H$, and Δ satisfies postulates* $\mathsf{IC}_0 - \mathsf{IC}_8$.

Proof. If $\leq_E$ is Horn compliant then $min_{\leq_E}[\mu]$ can be represented by a Horn formula, for any Horn formula μ. Then $\Delta_\mu(E)$ is well defined as a Horn operator and we can use the proof for full propositional logic to show that Δ satisfies $\mathsf{IC}_0 - \mathsf{IC}_8$ (see [8]), as Horn compliant pre-orders are just a special type of pre-orders on $\mathcal{W}$. $\square$

4 Strengthening the postulates

Conversely, we want to show that for any Horn merging operator Δ there exists a syncretic assignment which represents it. This is true when the language is not restricted (see Proposition 1), but interesting problems arise as soon as we restrict ourselves to the Horn case. In the following, we assume we are given a Horn merging operator Δ satisfying the merging postulates.

4.1 Non-transitive cycles and Acyc

We know from existing work on Horn revision [4] that we can find non-syncretic assignments to represent a Horn operator Δ. Such assignments are non-syncretic in the sense that they contain non-transitive cycles between interpretations. Yet, we can still define an operator on top of these rankings, which satisfies postulates $\mathsf{IC}_0 - \mathsf{IC}_8$. Furthermore, it can be shown that there are no non-cyclic pre-orders that represent the same operator Δ. The solution proposed in [4] is to add an extra postulate, called Acyc, specifically to eliminate cycles:

(Acyc**)** If for every $n \geq 1$ and $i \in \{0, \ldots, n-1\}$, $\mu_i \wedge \Delta_{\mu_{i+1}}(E)$ and $\mu_n \wedge \Delta_{\mu_0}(E)$ are all consistent, then $\mu_0 \wedge \Delta_{\mu_n}(E)$ is also consistent.

Acyc provably follows from postulates $\mathsf{IC}_0 - \mathsf{IC}_8$ in full propositional logic (see [4]), so it only makes a difference when we restrict the language to the Horn fragment. Here we employ the same strategy of adding an extra postulate to deal with non-transitive cycles, but we propose a postulate formulated in terms of complete formulas:

(Acyc$'$**)** For any complete formulas $\sigma_0, \ldots, \sigma_n$, $n \geq 1$ and $i \in \{0, \ldots, n-1\}$, it holds that if $\sigma_i \wedge \Delta_{\sigma_{i,i+1}}(E)$ and $\sigma_n \wedge \Delta_{\sigma_{n,0}}(E)$ are all consistent, then $\sigma_0 \wedge \Delta_{\sigma_{0,n}}(E)$ is also consistent.

There is clearly a strong similarity between Acyc and Acyc$'$, though we prefer the latter here for its intuitive appeal. Moreover, it can be shown that Acyc and Acyc$'$ are equivalent modulo the merging postulates.

Proposition 2. *Given postulates* $\mathsf{IC}_0 - \mathsf{IC}_8$ *and full propositional logic,* Acyc *and* Acyc$'$ *are equivalent.*

Proof. We prove each direction in turn.

"$\Rightarrow$" We show first that Acyc implies Acyc$'$. Take complete formulas $\sigma_0, \ldots, \sigma_n$ such that for $i \in \{0, \ldots, n-1\}$, $\sigma_i \wedge \Delta_{\sigma_{i,i+1}}(E)$ and $\sigma_n \wedge \Delta_{\sigma_{n,0}}(E)$ are all consistent. Now take $\mu_0 = \sigma_{0,1}$, $\mu_1 = \sigma_{1,2}$, $\ldots$, $\mu_n = \sigma_{n,0}$.

If $[\sigma_i] = \{w_i\}$, then for $i \in \{1, \ldots, n-1\}$ we get that $w_{i-1} \in [\sigma_{i-2,\ i-1} \wedge \Delta_{\sigma_{i-1,\ i}}(E)]$ and $w_n \in [\sigma_{n-1,\ n} \wedge \Delta_{\sigma_{n,0}}(E)]$. This shows that for $i \in \{0, \ldots, n-1\}$, we have that all of $\mu_i \wedge \Delta_{\mu_{i+1}}(E)$ and $\mu_n \wedge \Delta_{\mu_0}(E)$ are consistent. Applying Acyc, we get that $\mu_0 \wedge \Delta_{\mu_n}(E)$ is also consistent, which is the same as saying that $\sigma_{0,1} \wedge \Delta_{\sigma_{n,0}}(E)$ is consistent. This implies that $w_0 \in [\sigma_{0,1} \wedge \Delta_{\sigma_{n,0}}(E)]$, so $\sigma_0 \wedge \Delta_{\sigma_{n,0}}(E)$ is consistent. Hence Acyc$'$ holds.

"$\Leftarrow$" Next we show that Acyc$'$ implies Acyc. Take formulas $\mu_0, \ldots, \mu_n$ such that for $i \in \{0, \ldots, n-1\}$, it holds that $\mu_i \wedge \Delta_{\mu_{i+1}}(E)$ and $\mu_n \wedge \Delta_{\mu_0}(E)$ are all consistent. Take, now, complete formulas $\sigma_0, \ldots, \sigma_n$ such that $[\sigma_i] = \{w_i\}$ for $i \in \{0, \ldots, n\}$. Since $w_0 \in \Delta_{\mu_0}(E)$, we get that $\sigma_0 \wedge \Delta_{\mu_0}(E)$ is consistent. We also have that $w_1 \in [\mu_0]$, hence $\sigma_{0,1} \models \mu_0$. It follows that $\sigma_{0,1} \wedge \mu_0 \equiv \sigma_{0,1}$, which by IC_3 implies that $\Delta_{\sigma_{0,1} \wedge \mu_0}(E) \equiv \Delta_{\sigma_{0,1}}(E)$.

Next, note that $w_0 \in [\sigma_0 \wedge \Delta_{\mu_0}(E)]$ implies that $w_0 \in [\sigma_{0,1} \wedge \Delta_{\mu_0}(E)]$, hence the latter formula is consistent. Applying IC_7, IC_8 and what we have already deduced, we get that:

$$\sigma_{0,1} \wedge \Delta_{\mu_0}(E) \equiv \Delta_{\sigma_{0,1} \wedge \mu_0}(E) \equiv \Delta_{\sigma_{0,1}}(E).$$

Since $w_0 \in [\sigma_{0,1} \wedge \Delta_{\mu_0}(E)]$, it follows that $w_0 \in [\Delta_{\sigma_{0,1}}(E)]$. This implies that $\sigma_0 \wedge \Delta_{\sigma_{0,1}}(E)$ is consistent. By the same argument, we can show that for $i \in \{0, \dots, n-1\}$, it holds that $\sigma_i \wedge \Delta_{\sigma_{i,\ i+1}}(E)$ and $\sigma_n \wedge \Delta_{\sigma_{n,0}}(E)$ are all consistent. Applying Acyc', it follows that $\sigma_0 \wedge \Delta_{\sigma_{0,n}}(E)$ is consistent, which means that $w_0 \in [\Delta_{\sigma_{0,n}}(E)]$. From this we can conclude that $\sigma_{0,n} \wedge \Delta_{\mu_n}(E)$ is consistent. Thus $\sigma_{0,n} \models \mu_n$ and (by IC_7, IC_8 and IC_3), we get that:

$$\sigma_{0,n} \wedge \Delta_{\mu_n}(E) \equiv \Delta_{\sigma_{0,n} \wedge \mu_n}(E) \equiv \Delta_{\sigma_{0,n}}(E).$$

Putting these last things together, it follows that $w_0 \in [\Delta_{\mu_n}(E)]$ and hence $\mu_0 \wedge \Delta_{\mu_n}(E)$ is consistent. □

The formal justification of Acyc' is that we need it for our representation result: without Acyc' (or an equivalent of it) we cannot ensure that every Horn merging operator is represented by a syncretic assignment. Acyc' is featured in the proof of Theorem 4.

Intuitively, Acyc' prevents non-transitive cycles between chains of interpretations of arbitrary length. Suppose $n = 2$ and the antecedent of Acyc' is true: then from the fact that $\sigma_0 \wedge \Delta_{\sigma_{0,1}}(E)$ is consistent we conclude that $w_0 \in [\Delta_{\sigma_{0,1}}(E)]$, where $[\sigma_i] = \{w_i\}$. This means that w_0 is among the models of $\sigma_{0,1}$ that are 'preferred', or considered more plausible, by Δ. Thus, in the pre-order $\leq_E$ that represents Δ, it should hold that $w_0 \leq_E w_1$. By the same token, we get that $w_1 \leq_E w_2 \leq_E w_0$ should hold. Since we want $\leq_E$ to be transitive, it should also hold that $w_0 \leq_E w_2$, or $w_0 \in [\Delta_{\sigma_{0,2}}(E)]$— and this is exactly what Acyc' requires at this point. Thus, we need the extra postulate Acyc' (or something equivalent) to ensure that the pre-orders representing a given Horn merging operator Δ preserve intuitive properties such as (in this case) transitivity.

4.2 s_4 – s_6 and IC_4 – IC_6

Introducing Acyc' is not enough, as even after we add it to the set of postulates we can still find non-syncretic assignments to represent a Horn merging operator Δ. Non-syncreticity, in this case, occurs because properties s_4 – s_6 are not enforced by IC_4 – IC_6. The following examples make this clearer.

Example 5. Consider Horn knowledge bases $[K_1] = \{01\}$, $[K_2] = \{10\}$ and an assignment that works as in Figure 1 when restricted to K_1 and K_2. Figure 1 shows the rankings associated with K_1 and K_2 and the result of merging them into the new ranking $\leq_{\{K_1,K_2\}}$.[1] Notice that s_4 is not true: s_4 requires that $01 \approx_{\{K_1,K_2\}} 10$, whereas we have $01 <_{\{K_1,K_2\}} 10$.

However, we can define a (Horn) merging operator Δ on top of this assignment in the usual way, by taking $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any Horn formula μ. The operator Δ will satisfy postulates IC_0 – IC_8 + Acyc': the argument for this is not difficult, though it is tedious to spell out in detail. It is straightforward to check, by direct inspection,

[1] It is worth noting that $\leq_{K_1}$, $\leq_{K_2}$ and $\leq_{\{K_1,K_2\}}$ are *not* generated using any familiar notion of distance—the rankings were hand-picked.

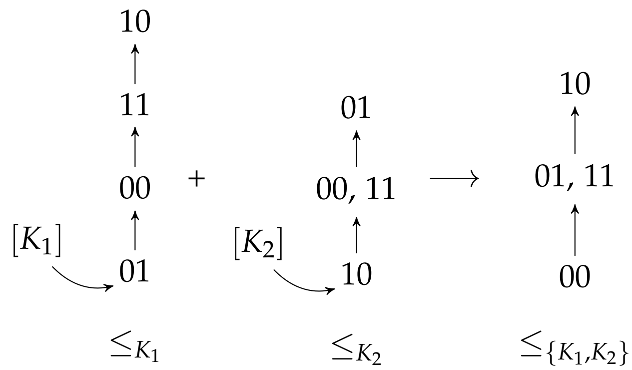

Figure 1: s_4 does not hold.

that the merging postulates hold for this example, so let us focus here only on IC_4. The problematic interpretations are the models 01 and 10 of K_1 and K_2, respectively. Notice that there is no Horn formula that represents exactly the set $\{01, 10\}$: the best we can do is take a Horn formula μ such that $[\mu] = [\sigma_{01,10}] = \{00, 01, 10\}$. Obviously, $K_1 \models \mu$ and $K_2 \models \mu$, so the antecedent of IC_4 holds. We notice that the consequent also holds: $[\Delta_\mu(\{K_1, K_2\})] = \{00\}$, and $[\Delta_\mu(\{K_1, K_2\}) \wedge K_1] = [\Delta_\mu(\{K_1, K_2\}) \wedge K_2] = \emptyset$, so IC_4 is satisfied for this particular μ.

Example 5 is significant because it shows that an assignment which does not satisfy s_4 may still represent an operator Δ obeying IC_4. And given the standard formulation of IC_4, this type of situation turns out to be unavoidable, as Proposition 3 shows.

Proposition 3. *There is no syncretic assignment representing Δ from Example 5 that assigns to $\{K_1, K_2\}$ a pre-order $\leq^\star_{\{K_1,K_2\}}$ where* $01 \approx^\star_{\{K_1,K_2\}} 10$.

Proof. Suppose $01 \approx^\star_{\{K_1,K_2\}} 10$. From Figure 1 we know that $[\Delta_{\sigma_{10,11}}(\{K_1, K_2\})] = \{11\}$, hence $min_{\leq^\star}\{10, 11\} = \{11\}$ and thus $11 <^\star_{\{K_1,K_2\}} 10$. Similarly, we obtain $01 \approx^\star_{\{K_1,K_2\}} 11$. By transitivity of $\leq^\star_{\{K_1,K_2\}}$ it follows that $01 <^\star_{\{K_1,K_2\}} 10$. This creates a contradiction. □

The following example shows how s_5 fails to be enforced by IC_5 in the case of Horn logic.

Example 6. Assume there exists an assignment which for two profiles E_1 and E_2 behaves as in Figure 2, and is otherwise Horn compliant and syncretic. Property s_5 does not hold: $010 \approx_{E_1} 100$ and $010 \approx_{E_2} 100$, but $010 <_{E_1 \sqcup E_2} 100$. However, as in Example 5, we can define a (Horn) merging operator Δ on top of this assignment and Δ will satisfy postulates IC_0 – IC_8 + Acyc'. Let us check here that IC_5 holds. The problematic interpretations are 010 and 100 (for which s_5 does not hold). In this case we have that $\Delta_{\sigma_{010,100}}(E_1) \wedge \Delta_{\sigma_{010,100}}(E_2)$ is consistent, and $[\Delta_{\sigma_{010,100}}(E_1) \wedge \Delta_{\sigma_{010,100}}(E_2)] = [\Delta_{\sigma_{010,100}}(E_1 \sqcup E_2)] = \{000\}$. This shows that for the case we are interested in, which is $\mu = \sigma_{010,100}$, IC_5 is true.

$$\begin{array}{c}110\\ \uparrow\\ 010, 111, 100\\ \uparrow\\ 000\\ \uparrow\\ 011\\ \uparrow\\ 001\\ \uparrow\\ 101\\ \\ \leq_{E_1}\end{array} \quad + \quad \begin{array}{c}010, 011, 001, 101, 100\\ \uparrow\\ 000\\ \uparrow\\ 110\\ \uparrow\\ 111\\ \\ \leq_{E_2}\end{array} \quad \longrightarrow \quad \begin{array}{c}100\\ \uparrow\\ 110\\ \uparrow\\ 010\\ \uparrow\\ 000\\ \uparrow\\ 011\\ \uparrow\\ 001\\ \uparrow\\ 101, 111\\ \\ \leq_{E_1 \sqcup E_2}\end{array}$$

Figure 2: s_5 does not hold for 010 and 100.

Similarly as for IC_4, we can show that such a counter-example to s_5 is unavoidable.

Proposition 4. *There is no syncretic assignment representing Δ from Example 6 that assigns to $E_1 \sqcup E_2$ a pre-order $\leq^\star_{E_1 \sqcup E_2}$ where $100 \approx^\star_{E_1 \sqcup E_2} 010$.*

Proof. Assume $100 \approx^\star_{E_1 \sqcup E_2} 010$. Looking at $\leq_{E_1 \sqcup E_2}$, we get that:

$$[\Delta_{\sigma_{100,110}}(E_1 \sqcup E_2)] = \{110\},$$
$$[\Delta_{\sigma_{110,010}}(E_1 \sqcup E_2)] = \{010\}.$$

This leads us to conclude that $010 <^\star_{E_1 \sqcup E_2} 110 <^\star_{E_1 \sqcup E_2} 100$, and by transitivity $010 <^\star_{E_1 \sqcup E_2} 100$, which creates a contradiction. $\square$

It is perhaps surprising to see that IC_5 can be satisfied in an assignment where s_5 does not hold, but closer thought shows this is to be expected: since in the Horn fragment we cannot represent the set $\{100, 010\}$ with a formula, it becomes harder to control the order in which 100 and 010 appear. Without any additional constraints on Δ, one cannot prevent it from varying the order of 100 and 010 in ways that directly contradict s_5. A similar counter-example can be constructed for s_6.

Examples 5 and 6 show that a syncretic assignment with total pre-orders is not the most natural way to represent a Horn merging operator. Hence, we introduce the following notion.

Definition 3. A pre-order $\leq$ on $\mathcal{W}$ is *Horn connected* if:

Figure 3: $\leq_1$ ✓

Figure 4: $\leq_2$ ✓

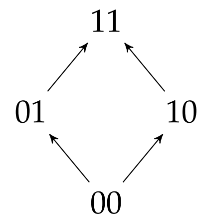

Figure 5: $\leq_3$ ✓

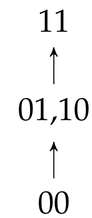

Figure 6: $\leq_4$ $\times$

(h_1) $\leq$ is Horn compliant,

(h_2) any $w_i, w_j \in \mathcal{W}$ that are in the subset relation are in $\leq$, and

(h_3) for any $w_i, w_j \in \mathcal{W}$ such that $w_i \not\subseteq w_j$ and $w_j \not\subseteq w_i$, it holds that if $w_i \leq w_j$ then:

- ($h_{3.1}$) $w_i \in min_{\leq} Cl_{\cap}(\{w_i, w_j\})$, or
- ($h_{3.2}$) for some $n > 2$, there exist interpretations $w_1, \ldots, w_n$, pair-wise distinct, such that $w_1 = w_i$, $w_n = w_j$ and $w_1 \leq \cdots \leq w_n$.

A Horn connected pre-order $\leq_E$ is not necessarily total. Example 7 illustrates this.

Example 7. Consider the following pre-orders on the 2-letter alphabet (see also Figures 3-6): (a) $11 <_1 01 <_1 10 <_1 00$, (b) $00 <_2 01 <_2 11 <_2 10$, (c) $00 <_3 01 <_3 11$, $00 <_3 10 <_3 11$, $01 \not\lesssim_3 10$, $10 \not\lesssim_3 01$, (d) $00 <_4 01 \approx_4 10 <_4 11$. It is immediately visible that all pre-orders are Horn compliant (h_1) and that they satisfy h_2. To check h_3, let us focus on interpretations 01 and 10.

In $\leq_1$ we have $01 \in min_{\leq_1} Cl_{\cap}(\{01, 10\})$. Thus $h_{3.1}$ is satisfied, and $\leq_1$ is Horn connected. In $\leq_2$ we do not have $01 \in min_{\leq_2} Cl_{\cap}(\{01, 10\})$, but there is interpretation 11 such that $01 <_2 11 <_2 10$. Thus $h_{3.2}$ is satisfied and $\leq_2$ is Horn connected. Pre-order $\leq_3$ is partial, as 01 and 10 are not in $\leq_3$, and thus h_3 is vacuously true. In $\leq_4$ we have $01 \approx_4 10$ though none of $h_{3.1}$ and $h_{3.2}$ holds, thus $\leq_4$ is not Horn connected.

Next, for every Horn operator Δ and Horn profile E, we define a (partial) pre-order on complete formulas of $\mathcal{L}_H$.[2]

Definition 4. Given a Horn operator Δ, then for any Horn profile E and complete Horn formulas σ_i, σ_j, we say that $\sigma_i \preceq_E \sigma_j$ if there exist complete Horn formulas $\sigma_1, \ldots, \sigma_n$ such that $\sigma_1 = \sigma_i$, $\sigma_n = \sigma_j$, and for $i \in \{1, \ldots, n-1\}$, $\sigma_i \wedge \Delta_{\sigma_{i,i+1}}(E)$ are all consistent.

[2]The pre-order on Horn complete formulas $\preceq_E$ is not to be confused with the pre-order $\leq_E$ on interpretations, though it is meant to mirror it.

It is straightforward to check that $\preceq_E$ is reflexive and transitive, and thus a pre-order on complete Horn formulas. We write $\prec_E$ for the strict part of $\preceq_E$. It is also worth noting that $\preceq_E$ is total when the underlying language is full propositional logic, since $[\sigma_{i,j}] = \{w_i, w_j\}$ and we can take the sequence $\sigma_1, \ldots, \sigma_n$ to be just σ_i, σ_j or σ_j, σ_i. This does not necessarily hold in the case of the Horn fragment, where $\preceq_E$ can be partial.

We now reformulate IC_4, IC_5 and IC_6 for $K_1, K_2 \in \mathcal{K}_H$, $E_1, E_2 \in \mathcal{E}_H$ and complete Horn formulas σ_i, σ_j as follows:

(IC_4') For any $\sigma_i \models K_1$, there exists $\sigma_j \models K_2$ such that $\sigma_j \preceq_{\{K_1, K_2\}} \sigma_i$.

(IC_5') If $\sigma_i \preceq_{E_1} \sigma_j$ and $\sigma_i \preceq_{E_2} \sigma_j$, then $\sigma_i \preceq_{E_1 \sqcup E_2} \sigma_j$.

(IC_6') If $\sigma_i \preceq_{E_1} \sigma_j$ and $\sigma_i \prec_{E_2} \sigma_j$, then $\sigma_i \prec_{E_1 \sqcup E_2} \sigma_j$.

These postulates make a difference only in the Horn fragment, while in full propositional logic they are redundant.

Proposition 5. *In the case of full propositional logic,* IC_4', IC_5' *and* IC_6' *follow from the standard* $\mathsf{IC}_0 - \mathsf{IC}_8$ *postulates.*

Proof. We start with IC_4' and show that it follows from $\mathsf{IC}_0 - \mathsf{IC}_8$. Take a complete formula $\sigma_i \models K_1$ and $\mu = K_1 \vee K_2$. We show first that $\Delta_\mu(\{K_1, K_2\}) \wedge K_2$ is consistent. Suppose, on the contrary, that it is inconsistent. Since $\Delta_\mu(\{K_1, K_2\})$ is consistent (by IC_1), it follows that $\Delta_\mu(\{K_1, K_2\}) \not\models K_2$. From IC_0 we have that $\Delta_\mu(\{K_1, K_2\}) \models \mu$ and this, together with the previous result implies that $\Delta_\mu(\{K_1, K_2\}) \models K_1$. From here we conclude that $\Delta_\mu(\{K_1, K_2\}) \wedge K_1$ is consistent, and (by IC_4) this implies that $\Delta_\mu(\{K_1, K_2\}) \wedge K_2$ is consistent, which contradicts our assumption.

Thus, we have that $\Delta_\mu(\{K_1, K_2\}) \wedge K_2$ is consistent. Take, then, a complete formula σ_j such that $\sigma_j \models \Delta_\mu(\{K_1, K_2\}) \wedge K_2$. It follows that $\sigma_j \models \Delta_\mu(\{K_1, K_2\})$, and as a result $\Delta_\mu(\{K_1, K_2\}) \wedge \sigma_{i,j}$ is consistent. By IC_7 and IC_8, this implies that $\Delta_\mu(\{K_1, K_2\}) \wedge \sigma_{i,j} \equiv \Delta_{\mu \wedge \sigma_{i,j}}(\{K_1, K_2\})$. Since $\sigma_i \models K_1$ and $\sigma_j \models K_2$, we get that $\sigma_{i,j} \models K_1 \vee K_2$. Thus $\mu \wedge \sigma_{i,j} \equiv \sigma_{i,j}$ and by IC_3 this implies that $\Delta_{\mu \wedge \sigma_{i,j}}(\{K_1, K_2\}) \equiv \Delta_{\sigma_{i,j}}(\{K_1, K_2\})$. Plugging this into our previous result we get that $\Delta_\mu(\{K_1, K_2\}) \wedge \sigma_{i,j} \equiv \Delta_{\sigma_{i,j}}(\{K_1, K_2\})$. As a consequence $\sigma_j \models \Delta_{\sigma_{i,j}}(\{K_1, K_2\})$, hence $\sigma_j \wedge \Delta_{\sigma_{i,j}}(\{K_1, K_2\})$ is consistent and $\sigma_j \preceq_{\{K_1, K_2\}} \sigma_i$.

The next step is to show that IC_5' follows from $\mathsf{IC}_0 - \mathsf{IC}_8$. We first prove the following two lemmas.

Lemma 1. *Let Δ be a merging operator satisfying postulates* $\mathsf{IC}_0 - \mathsf{IC}_8$ *and let E be a profile. In the case of full propositional logic, if σ_1, σ_2, σ_3 are complete formulas such that $\sigma_1 \wedge \Delta_{\sigma_{1,2}}(E)$ and $\sigma_2 \wedge \Delta_{\sigma_{2,3}}(E)$ are both consistent, then $\sigma_1 \wedge \Delta_{\sigma_{1,3}}(E)$ is also consistent.*

Proof. Suppose $[\sigma_1] = \{w_1\}$, $[\sigma_2] = \{w_2\}$ and $[\sigma_3] = \{w_3\}$. Let $\sigma_{1,2,3}$ be a formula such that $[\sigma_{1,2,3}] = \{w_1, w_2, w_3\}$. We first show that $\sigma_1 \wedge \Delta_{\sigma_{1,2,3}}(E)$ is consistent: by IC_0 and IC_1 we have that $[\Delta_{\sigma_{1,2,3}}(E)]$ is a non-empty subset of $[\sigma_{1,2,3}] = \{w_1, w_2, w_3\}$. This implies that at least one of $\sigma_{1,2} \wedge \Delta_{\sigma_{1,2,3}}(E)$ and $\sigma_{2,3} \wedge \Delta_{\sigma_{1,2,3}}(E)$ is consistent. Let us do a case analysis and show that both cases imply our intermediary conclusion.

Case 1. If $\sigma_{1,2} \wedge \Delta_{\sigma_{1,2,3}}(E)$ is consistent, then by IC_8 we get that $\Delta_{\sigma_{1,2,3} \wedge \sigma_{1,2}}(E) \models \sigma_{1,2} \wedge \Delta_{\sigma_{1,2,3}}(E)$. Since $\sigma_{1,2} \models \sigma_{1,2,3}$, it follows that $\sigma_{1,2,3} \wedge \sigma_{1,2} \equiv \sigma_{1,2}$ and, by IC_3 we get that $\Delta_{\sigma_{1,2,3} \wedge \sigma_{1,2}}(E) \equiv \Delta_{\sigma_{1,2}}(E)$. From the hypothesis we have $w_1 \in [\Delta_{\sigma_{1,2}}(E)]$, and using the previous results we get that $w_1 \in [\sigma_{1,2} \wedge \Delta_{\sigma_{1,2,3}}(E)]$ and, by consequence, $w_1 \in [\Delta_{\sigma_{1,2,3}}(E)]$. Clearly, then, $w_1 \in [\sigma_1 \wedge \Delta_{\sigma_{1,2,3}}(E)]$ so $\sigma_1 \wedge \Delta_{\sigma_{1,2,3}}(E)$ is consistent.

Case 2. If $\sigma_{2,3} \wedge \Delta_{\sigma_{1,2,3}}(E)$ is consistent, then by IC_8 we get that $\Delta_{\sigma_{1,2,3} \wedge \sigma_{2,3}}(E) \models \sigma_{2,3} \wedge \Delta_{\sigma_{1,2,3}}(E)$. Similarly to before, we get that $\sigma_{1,2,3} \wedge \sigma_{2,3} \equiv \sigma_{2,3}$ and, by IC_3 we have $\Delta_{\sigma_{1,2,3} \wedge \sigma_{2,3}}(E) \equiv \Delta_{\sigma_{2,3}}(E)$. From the hypothesis we get that $w_2 \in [\Delta_{\sigma_{2,3}}(E)]$. Our previous results imply that $w_2 \in [\sigma_{2,3} \wedge \Delta_{\sigma_{1,2,3}}(E)]$ and, by consequence, $w_2 \in [\Delta_{\sigma_{1,2,3}}(E)]$. It follows that $w_2 \in [\sigma_{1,2} \wedge \Delta_{\sigma_{1,2,3}}(E)]$ and we apply the reasoning from *Case 1* to get the conclusion that $\sigma_1 \wedge \Delta_{\sigma_{1,2,3}}(E)$ is consistent. This concludes our case analysis

We have shown, then, that $\sigma_1 \wedge \Delta_{\sigma_{1,2,3}}(E)$ is consistent. Thus $w_1 \in [\sigma_1 \wedge \Delta_{\sigma_{1,2,3}}(E)]$, which implies that $w_1 \in [\sigma_{1,3} \wedge \Delta_{\sigma_{1,2,3}}(E)]$. Now, by IC_7, we get that $\sigma_{1,3} \wedge \Delta_{\sigma_{1,2,3}}(E) \models \Delta_{\sigma_{1,3} \wedge \sigma_{1,2,3}}(E)$. Since $\sigma_{1,3} \wedge \sigma_{1,2,3} \equiv \sigma_{1,3}$ and applying IC_3, we have that $\Delta_{\sigma_{1,3} \wedge \sigma_{1,2,3}}(E) \equiv \Delta_{\sigma_{1,3}}(E)$. Using all these results we infer that $w_1 \in [\Delta_{\sigma_{1,3}}(E)]$. Hence $w_1 \in [\sigma_1 \wedge \Delta_{\sigma_{1,3}}(E)]$ and $\sigma_1 \wedge \Delta_{\sigma_{1,3}}(E)$ is consistent. $\square$

Lemma 2. *In the case of full propositional logic, if σ_i and σ_j are complete formulas such that $\sigma_i \preceq_E \sigma_j$, then $\sigma_i \wedge \Delta_{\sigma_{i,j}}(E)$ is consistent.*

Proof. By definition, there exist complete formulas $\sigma_1, \dots, \sigma_n$ such that $\sigma_i = \sigma_1$, $\sigma_j = \sigma_n$ and for $i \in \{1, \dots, n-1\}$, all of $\sigma_i \wedge \Delta_{\sigma_{i,i+1}}(E)$ are consistent. By induction on n and using Lemma 1, it follows immediately that $\sigma_1 \wedge \Delta_{\sigma_{1,n}}(E)$ is consistent. Since $\sigma_i = \sigma_1$ and $\sigma_j = \sigma_n$, we get the conclusion. $\square$

Now we can show that IC_5' follows from the standard postulates. Take complete formulas σ_i and σ_j such that $\sigma_i \preceq_{E_1} \sigma_j$ and $\sigma_i \preceq_{E_2} \sigma_j$. Suppose $[\sigma_i] = \{w_i\}$ and $[\sigma_j] = \{w_j\}$. From Lemma 2 it follows that $\sigma_i \wedge \Delta_{\sigma_{i,j}}(E_1)$ and $\sigma_i \wedge \Delta_{\sigma_{i,j}}(E_2)$ are both consistent, which implies that $w_i \in [\Delta_{\sigma_{i,j}}(E_1)]$ and $w_i \in [\Delta_{\sigma_{i,j}}(E_2)]$, so $w_i \in [\Delta_{\sigma_{i,j}}(E_1) \wedge \Delta_{\sigma_{i,j}}(E_2)]$. By IC_5, it follows that $w_i \in [\Delta_{\sigma_{i,j}}(E_1 \sqcup E_2)]$, so $\sigma_i \wedge \Delta_{\sigma_{i,j}}(E_1 \sqcup E_2)$ is consistent. This immediately implies that $\sigma_i \preceq_{E_1 \sqcup E_2} \sigma_j$, since we can just take $\sigma_1 = \sigma_i$ and $\sigma_2 = \sigma_j$.

For IC_6', take complete formulas σ_i, σ_j such that $\sigma_i \preceq_{E_1} \sigma_j$, $\sigma_i \prec_{E_2} \sigma_j$. Using the result for IC_5', it follows immediately that $\sigma_i \preceq_{E_1 \sqcup E_2} \sigma_j$. We want to show that $\sigma_i \prec_{E_1 \sqcup E_2} \sigma_j$. Suppose, on the contrary, that $\sigma_j \preceq_{E_1 \sqcup E_2} \sigma_i$. By Lemma 2, our assumptions give us that $\sigma_i \wedge \Delta_{\sigma_{i,j}}(E_1)$, $\sigma_i \wedge \Delta_{\sigma_{i,j}}(E_2)$ and $\sigma_j \wedge \Delta_{\sigma_{i,j}}(E_1 \sqcup E_2)$ are all consistent. Thus, $\Delta_{\sigma_{i,j}}(E_1) \wedge \Delta_{\sigma_{i,j}}(E_2)$ is consistent, and by IC_6 we get that $\Delta_{\sigma_{i,j}}(E_1 \sqcup E_2) \models \Delta_{\sigma_{i,j}}(E_1) \wedge \Delta_{\sigma_{i,j}}(E_2)$. This means that $\sigma_j \wedge \Delta_{\sigma_{i,j}}(E_2)$ is consistent, and thus $\sigma_j \preceq_{E_2} \sigma_i$, which creates a contradiction. $\square$

With postulates IC_4', IC_5' and IC_6' we can derive a representation result for syncretic assignments with Horn connected pre-orders. The result is split across two theorems: Theorem 3 shows that Horn connected pre-orders can be used to construct a Horn merging operator satisfying our amended set of postulates. Its converse, Theorem 4, shows that any Horn merging operator satisfying the amended postulates is represented by a syncretic assignment with Horn connected pre-orders.

Theorem 3. *If there exists a syncretic assignment mapping every $E \in \mathcal{E}_H$ to a Horn connected total pre-order $\leq_E$, then we can define an operator $\Delta \colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$ by taking $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any $\mu \in \mathcal{L}_H$, and Δ satisfies postulates $\mathsf{IC}_0 - \mathsf{IC}_3 + \mathsf{IC}_4' + \mathsf{IC}_5' + \mathsf{IC}_6' + \mathsf{IC}_7 - \mathsf{IC}_8 + \mathsf{Acyc}'$.*

Proof. We start from a syncretic assignment with Horn connected pre-orders and show that a Horn merging operator defined on top of it satisfies the postulates. We first prove the following lemma.

Lemma 3. *If $\leq$ is a Horn connected pre-order, then for any Horn formula μ and $w_i, w_j \in [\mu]$, it is the case that if $w_i \in min_{\leq}[\mu]$, then $w_i \leq w_j$.*

Proof. Take $w_i, w_j \in [\mu]$ such that $w_i \in min_{\leq}[\mu]$ and let us show, by a quick case analysis, that $w_i \leq w_j$.

Case 1. If $w_i \subseteq w_j$ or $w_j \subseteq w_i$, then because $\leq$ is Horn connected, it holds that $w_i \leq w_j$ or $w_j \leq w_i$. Since $w_i \in min_{\leq}[\mu]$ we clearly cannot have $w_j < w_i$, so $w_i \leq w_j$.

Case 2. If $w_i \nsubseteq w_j$ and $w_j \nsubseteq w_i$, take $w_{i,j}$ to be the interpretation such that $w_{i,j} = w_i \cap w_j$. Since μ is a Horn formula, we have that $w_{i,j} \in [\mu]$. Also, $w_{i,j} \subseteq w_i$ and $w_{i,j} \subseteq w_j$, so the pairs of interpretations $w_i, w_{i,j}$ and $w_j, w_{i,j}$ must be connected by $\leq$. Clearly, since $w_i \in min_{\leq}[\mu]$, it holds that $w_i \leq w_{i,j}$. This fixes the first pair. What about the second? Let us do a quick case analysis here as well.

Case 2.1. If $w_{i,j} \leq w_j$, then from this and $w_i \leq w_{i,j}$, we get (by transitivity) that $w_i \leq w_j$.

Case 2.2. If $w_j < w_{i,j}$, then it cannot be the case that $w_{i,j} \in min_{\leq}[\mu]$, so $w_i < w_{i,j}$ as well. We claim now that w_i and w_j must be in $\leq$, and the argument goes as follows. Supposing w_i and w_j were not in $\leq$, then consider a Horn formula σ_{w_i,w_j} such that $[\sigma_{w_i,w_j}] = \{w_{i,j}, w_i, w_j\}$. We get that $min_{\leq}[\sigma_{w_i,w_j}] = \{w_i, w_j\}$, which contradicts the fact that $\leq$ is Horn compliant. It follows, therefore, that w_i and w_j must be in $\leq$, and since $w_i \in min_{\leq}[\mu]$ and $w_j \in [\mu]$, it can only be the case that $w_i \leq w_j$.[3] $\square$

Now, Horn compliance guarantees that for any Horn formula μ, we have that $min_{\leq_E}[\mu]$ is a set of models representable as a Horn formula. Thus $\Delta_\mu(E)$ is well-defined. Using this and Lemma 3, it follows immediately that Δ satisfies $\mathsf{IC}_0 - \mathsf{IC}_8$: the proof is essentially the same as in the case of full propositional logic (see Theorem 2). To show that Δ satisfies Acyc, the proof in [4] works here with no modifications. Since Acyc and Acyc$'$ are equivalent modulo $\mathsf{IC}_0 - \mathsf{IC}_8$, it follows that Δ also satisfies Acyc$'$. And, since postulates IC_4, IC_5 IC_6 imply IC_4', IC_5' and IC_6' in the case of full propositional logic, Δ also satisfies IC_4', IC_5' and IC_6'. $\square$

The next step is to show that every Horn merging operator satisfying our strengthened postulates is represented by an assignment using Horn connected pre-orders.

Theorem 4. *If a Horn operator $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$ satisfies postulates $\mathsf{IC}_0 - \mathsf{IC}_3 + \mathsf{IC}_4' + \mathsf{IC}_5' + \mathsf{IC}_6' + \mathsf{IC}_7 - \mathsf{IC}_8 +$ Acyc$'$, then there exists a syncretic assignment mapping every Horn profile E to a Horn connected pre-order $\leq_E$, such that, for any Horn formula μ, it holds that $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$.*

Proof. Assume we are given a Horn merging operator Δ satisfying the postulates. For any Horn profile E, we define a Horn connected pre-order $\leq_E$ on $\mathcal{W}$ in two steps, as follows.

Step 1. First we define a relation $\leq_E'$ on $\mathcal{W}$. For any two interpretations w_i and w_j, say that:[4]

$w_i \leq_E' w_j$ iff $w_i \in [\Delta_{\sigma_{i,j}}(E)]$.

[3] In fact in this particular sub-case we can draw the even stronger conclusion that $w_i < w_j$—since if $w_i \approx w_j$ we get again that $min_{\leq}[\sigma_{w_i,w_j}] = \{w_i, w_j\}$, which contradicts the fact that $\leq$ is Horn compliant.

[4] As a reminder, if w_i and w_j are two interpretations, then in the context of the Horn fragment $\sigma_{i,j}$ is a Horn formula such that $[\sigma_{i,j}] = \{w_i, w_j, w_i \cap w_j\}$.

Relation $\leq'$ as defined above is reflexive: for any interpretation w, $[\Delta_{\sigma_w}(E)]$ is (by IC_0 and IC_1) a non-empty subset of $[\sigma_w] = \{w\}$, thus $w \leq'_E w$. Relation $\leq'_E$ is not, however, necessarily total.

Step 2. Define $\leq_E$ as the transitive closure of $\leq'_E$. In other words:

$$w_i \leq_E w_j \text{ iff } \quad \text{there exist } w_1, \ldots, w_n \text{ such that } w_i = w_1, w_j = w_n \text{ and } w_1 \leq'_E \cdots \leq'_E w_n.$$

The relation $\leq_E$ thus defined extends $\leq'_E$ and is therefore reflexive. By construction, $\leq_E$ is transitive, so it is a pre-order on $\mathcal{W}$.[5] We show now that pre-orders defined in this way give rise to an assignment that satisfies our requirements. There are several parts to this.

First, we show that $\leq_E$ is a Horn compliant pre-order such that $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$ and s_1 – s_3 hold. The following lemmas prove useful.[6]

Lemma 4. *For any Horn formula μ and interpretations w_1, w_2 such that $w_1 \in [\mu]$, $w_2 \in [\Delta_\mu(E)]$ and $w_1 \leq'_E w_2$, it holds that $w_1 \in [\Delta_\mu(E)]$.*

Proof. Using our assumption we have that $w_2 \in [\Delta_\mu(E) \wedge \sigma_{1,2}]$, so $\Delta_\mu(E) \wedge \sigma_{1,2}$ is consistent. From IC_7 – IC_8 we have $\Delta_\mu(E) \wedge \sigma_{1,2} \equiv \Delta_{\mu \wedge \sigma_{1,2}}(E)$. We also have that $\sigma_{1,2} \models \mu$, so $\mu \wedge \sigma_{1,2} \equiv \sigma_{1,2}$ and by IC_3 it holds that $\Delta_{\mu \wedge \sigma_{1,2}}(E) \equiv \Delta_{\sigma_{1,2}}(E)$. By definition of $\leq'_E$ we have that $w_1 \in [\Delta_{\sigma_{1,2}}(E)]$, which together with the previous results imply that $w_1 \in [\Delta_\mu(E)]$. □

Lemma 5. *For any Horn formula μ, $min_{\leq'_E}[\mu] = [\Delta_\mu(E)]$.*

Proof. We prove the lemma by double inclusion.

"$\subseteq$" Take $w_1 \in min_{\leq'_E}[\mu]$. This implies, by IC_0, that $w_1 \in [\mu]$ and thus $[\mu]$ is consistent. By IC_1, $[\Delta_\mu(E)]$ is non-empty. Take, then, $w_2 \in [\Delta_\mu(E)]$ and consider a Horn formula $\sigma_{1,2}$ such that $[\sigma_{1,2}] = \mathrm{Cl}_\cap(\{w_1, w_2\})$. We have that $w_2 \in [\sigma_{1,2} \wedge \Delta_\mu(E)]$ is consistent, and thus by IC_7 – IC_8 it holds that $\sigma_{1,2} \wedge \Delta_\mu(E) \equiv \Delta_{\mu \wedge \sigma_{1,2}}(E)$. Since $w_1, w_2 \in [\mu]$ and μ is a Horn formula, it follows that $\mathrm{Cl}_\cap(\{w_1, w_2\}) \subseteq [\mu]$ and thus $\sigma_{1,2} \models \mu$ and, further, $\sigma_{1,2} \wedge \mu \equiv \sigma_{1,2}$. By IC_3 this implies that $\Delta_{\mu \wedge \sigma_{1,2}}(E) \equiv \Delta_{\sigma_{1,2}}(E)$. Putting all these results together, it follows that $w_2 \in [\Delta_{\sigma_{1,2}}(E)]$ and thus by the definition of $\leq'_E$ we have $w_2 \leq'_E w_1$. Since w_1 and w_2 are both models of μ and $w_1 \in min_{\leq'_E}[\mu]$, it follows that $w_1 \leq'_E w_2$. From this and the fact that $w_2 \in [\Delta_\mu(E)]$ it follows by Lemma 4 that $w_1 \in [\Delta_\mu(E)]$.

"$\supseteq$" Take $w_1 \in [\Delta_\mu(E)]$ and an arbitrary interpretation $w_2 \in [\mu]$. Consider a Horn formula $\sigma_{1,2}$ such that $[\sigma_{1,2}] = \mathrm{Cl}_\cap(\{w_1, w_2\})$. We have that $w_1 \in [\sigma_{1,2} \wedge \Delta_\mu(E)]$ is consistent, and thus by IC_7 – IC_8 it holds that $\sigma_{1,2} \wedge \Delta_\mu(E) \equiv \Delta_{\mu \wedge \sigma_{1,2}}(E)$. Since $w_1, w_2 \in [\mu]$ and μ is a Horn formula, it follows that $\mathrm{Cl}_\cap(\{w_1, w_2\}) \subseteq [\mu]$ and thus $\sigma_{1,2} \models \mu$ and, further, $\sigma_{1,2} \wedge \mu \equiv \sigma_{1,2}$. By IC_3 this implies that $\Delta_{\mu \wedge \sigma_{1,2}}(E) \equiv \Delta_{\sigma_{1,2}}(E)$. Putting all these results together, it follows that $w_1 \in [\Delta_{\sigma_{1,2}}(E)]$ and thus by the definition of $\leq'_E$ we have $w_1 \leq'_E w_2$. Hence $w_1 \in min_{\leq'_E}[\mu]$. □

Lemma 6. *If $w_1 \leq'_E w_2 \leq'_E \cdots \leq'_E w_n \leq'_E w_1$, then $w_1 \leq'_E w_n$.*

[5] This is essentially the same construction as in [4], except that here we stop at Step 2, and do not perform the extra step of defining a total order on top of the transitive closure of $\leq'_E$.

[6] Lemmas 4 – 7 are essentially similar to Lemmas 1 – 4 in [4].

Proof. This Lemma follows as a straightforward applications of Acyc'. Take complete Horn formulas $\sigma_1, \dots, \sigma_n$ such that $[\sigma_i] = \{w_i\}$, for $i \in \{1, \dots, n\}$. As per usual, $\sigma_{i,j}$ is a Horn formula such that $[\sigma_{i,j}] = \mathrm{Cl}_\cap(\{w_i, w_j\})$. From $w_1 \leq'_E w_2$ and the definition of $\leq'_E$ we get that $w_1 \in [\Delta_{\sigma_{1,2}}(E)]$ and thus $\sigma_1 \wedge \Delta_{\sigma_{1,2}}(E)$ is consistent. By the same token, we get that for $i \in \{1, \dots, n-1\}$, $\sigma_i \wedge \Delta_{\sigma_{i,i+1}}(E)$ and $\sigma_n \wedge \Delta_{\sigma_{n,1}}(E)$ are all consistent. By Acyc' we get that $\sigma_1 \wedge \Delta_{\sigma_{n,1}}(E)$ is also consistent, from which it follows that $w_1 \in [\Delta_{\sigma_{n,1}}(E)]$ and thus $w_1 \leq'_E w_n$. □

Lemma 7. *For any two interpretations w_1, w_2, if $w_1 <'_E w_2$, then $w_1 <_E w_2$.*

Proof. Consider two interpretations w_1, w_2 such that $w_1 <'_E w_2$. From this it already follows that $w_1 \leq_E w_2$. What we still have to show is that $w_2 \not\leq_E w_1$. Suppose, on the contrary, that $w_2 \leq_E w_1$. Then there exist interpretations $u_0, \dots, u_n$ such that $w_2 = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n = w_1$ Because $w_1 \leq'_E w_2$ (since, by assumption, $w_1 <'_E w_2$), we have that $w_2 \leq'_E u_1 \leq'_E \cdots \leq'_E u_n \leq'_E w_1 \leq'_E w_2$ Using this and Lemma 6 we get that $w_2 \leq'_E w_1$, which together with the assumption that $w_1 <'_E w_2$ leads to a contradiction. □

Lemma 8. *For any Horn formula μ, $min_{\leq_E}[\mu] = min_{\leq'_E}[\mu]$.*

Proof. "$\subseteq$" Take $w_i \in min_{\leq_E}[\mu]$ and suppose there is some interpretation $w_j \in [\mu]$ such that $w_j <'_E w_i$. But then we get that $w_j <_E w_i$ by Lemma 7, which is a contradiction.

"$\supseteq$" Take $w_i \in min_{\leq'_E}[\mu]$ and an arbitrary interpretation $w_j \in [\mu]$. Also, consider a Horn formula $\sigma_{i,j}$ such that $[\sigma_{i,j}] = \mathrm{Cl}_\cap(\{w_i, w_j\})$. By IC_0 and IC_1, $[\Delta_{\sigma_{i,j}}(E)]$ is a non-empty subset of $[\sigma_{i,j}]$. We want to show that $w_i \leq_E w_j$. Let us do a case analysis.

Case 1. If $w_i \subseteq w_j$ or $w_j \subseteq w_i$, then $\mathrm{Cl}_\cap(\{w_i, w_j\}) = \{w_i, w_j\}$. It cannot be the case that $[\Delta_{\sigma_{i,j}}(E)] = \{w_j\}$ since that would imply that $w_j <'_E w_i$, which contradicts the hypothesis that $w_i \in min_{\leq'_E}[\mu]$. It follows that $w_i \in [\Delta_{\sigma_{i,j}}(E)]$ and so $w_i \leq'_E w_j$, which implies that $w_i \leq_E w_j$.

Case 2. If $w_i \nsubseteq w_j$ and $w_j \nsubseteq w_i$, take $w_{i,j} = w_i \cap w_j$. Since μ is a Horn formula, $w_{i,j} \in [\mu]$. Clearly, the sets $\{w_i, w_{i,j}\}$ and $\{w_j, w_{i,j}\}$ are representable by Horn formulas, and thus both pairs of interpretations must be in $\leq'_E$. What is more: it must be the case that $w_i \leq'_E w_{i,j}$, otherwise we would get a contradiction with $w_i \in min_{\leq'_E}[\mu]$. Let us see, next, what happens depending on how w_j and $w_{i,j}$ stand in relation to each other in $\leq'_E$.

Case 2.1. If $w_{i,j} \leq'_E w_j$ we get that $w_i \leq'_E w_{i,j} \leq'_E w_j$. When we take the transitive closure of $\leq'_E$ we get that $w_i \leq_E w_j$.

Case 2.2. If $w_j <'_E w_{i,j}$, then $w_{i,j} \notin [\Delta_{\sigma_{i,j}}(E)]$. But remember that $[\Delta_{\sigma_{i,j}}(E)]$ is a non-empty subset of $\mathrm{Cl}_\cap(\{w_i, w_j\}) = \{w_{i,j}, w_i, w_j\}$, and we cannot have that $[\Delta_{\sigma_{i,j}}(E)] = \{w_i, w_j\}$ or $[\Delta_{\sigma_{i,j}}(E)] = \{w_j\}$, because the first case this contradicts the fact that $\Delta_{\sigma_{i,j}}(E)$ is a Horn formula and the second case implies that $w_j <'_E w_i$, which contradicts the fact that $w_i \in min_{\leq'_E}[\mu]$. It only remains that $[\Delta_{\sigma_{i,j}}(E)] = \{w_i\}$, which implies that $w_i \leq'_E w_j$ and hence $w_i \leq_E w_j$. □

Next we show that $\leq_E$ is Horn connected:

Lemma 9. *The pre-order $\leq_E$ is Horn connected.*

Proof. Since, by assumption, $\Delta_\mu(E)$ is a Horn formula, $[\Delta_\mu(E)]$ is always closed under intersection. In other words, $\min_{\leq_E}[\mu]$ can always be represented by a Horn formula. This shows that $\leq_E$ is Horn compliant, thereby satisfying h_1. For h_2 and h_3, there are two cases to consider.

Case 1. If $w_i \subseteq w_j$ or $w_j \subseteq w_i$, then $\{w_i, w_j\}$ is representable by a Horn formula $\sigma_{i,j}$, which implies that $w_i \leq'_E w_j$ or $w_j \leq'_E w_i$. It follows that $w_i \leq_E w_j$ or $w_j \leq_E w_i$.

Case 2. If $w_i \not\subseteq w_j$, $w_j \not\subseteq w_i$, then take $w_{i,j} = w_i \cap w_j$ and let us assume $w_i \leq_E w_j$. By the definition of $\leq_E$, we know that there must be $w_1 \leq'_E \cdots \leq'_E w_n$ such that $w_1 = w_i$, $w_n = w_j$, for some $n \geq 2$.

Case 2.1. If $n = 2$, then $w_i \leq'_E w_j$, which means that $w_i \in [\Delta_{\sigma_{i,j}}(E)]$. By Lemmas 5 and 8 we have that $[\Delta_{\sigma_{i,j}}(E)] = min_{\leq'_E}[\sigma_{i,j}] = min_{\leq_E}[\sigma_{i,j}] = min_{\leq_E} Cl_\cap(\{w_i, w_j\})$, so property $\mathsf{h}_{3.1}$ is satisfied.

Case 2.2. If $n > 2$, we have that $w_1 \leq'_E w_2 \leq'_E \cdots \leq'_E w_n$. When we take the transitive closure of $\leq'_E$, we get that $w_1 \leq_E w_2 \leq_E \cdots \leq_E w_n$. Thus, property $\mathsf{h}_{3.2}$ is satisfied. □

Lemma 10. *If $w_1, w_2 \in [E]$, then $w_1 \approx_E w_2$.*

Proof. Consider a Horn formula $\sigma_{i,j}$ such that $[\sigma_{i,j}] = \mathrm{Cl}_\cap(\{w_i, w_j\})$. We have that $\sigma_{1,2} \wedge \bigwedge E$ is consistent and therefore by IC_2 it holds that $\Delta_{\sigma_{1,2}}(E) \equiv \bigwedge E \wedge \sigma_{1,2}$. Since $w_1, w_2 \in [\bigwedge E \wedge \sigma_{1,2}]$, it follows that $w_1, w_2 \in [\Delta_{\sigma_{1,2}}(E)]$ and therefore $w_1 \approx'_E w_2$ which implies that $w_1 \approx_E w_2$. □

Lemma 11. *If $w_1 \in [E]$ and $w_2 \notin [E]$, then $w_1 <_E w_2$.*

Proof. Consider a Horn formula $\sigma_{i,j}$ such that $[\sigma_{i,j}] = \mathrm{Cl}_\cap(\{w_i, w_j\})$. We have that $\sigma_{1,2} \wedge \bigwedge E$ is consistent and therefore by IC_2 it holds that $\Delta_{\sigma_{1,2}}(E) \equiv \bigwedge E \wedge \sigma_{1,2}$. Since $w_1 \in [\bigwedge E \wedge \sigma_{1,2}]$ and $w_2 \notin [\bigwedge E \wedge \sigma_{1,2}]$, it follows that $w_1 \in [\Delta_{\sigma_{1,2}}(E)]$ and $w_2 \notin [\Delta_{\sigma_{1,2}}(E)]$. Therefore $w_1 <'_E w_2$ and by Lemma 7 it holds that $w_1 <_E w_2$. □

In order to show that the assignment satisfies $\mathsf{s}_4 - \mathsf{s}_6$, the following lemma will prove very useful.

Lemma 12. *If w_i, w_j are two interpretations, σ_i and σ_j are two complete Horn formulas such that $[\sigma_i] = \{w_i\}$, $[\sigma_j] = \{w_j\}$ and $\leq_E$ is a Horn connected pre-order constructed as above, then $w_i \leq_E w_j$ if and only if $\sigma_i \preceq_E \sigma_j$.*

Proof. Consider a Horn formula $\sigma_{i,j}$ such that $[\sigma_{i,j}] = \mathrm{Cl}_\cap(\{w_i, w_j\})$. We will show each direction in turn.

"$\Rightarrow$" Since $w_i \leq_E w_j$ and $\leq_E$ is the transitive closure of relation $\leq'_E$ constructed using Δ, there must exist interpretations $w_1, \ldots, w_n$, for some $n \geq 2$, such that $w_i = w_1$, $w_j = w_n$ and $w_1 \leq'_E \cdots \leq'_E w_n$. Thus, by the definition of $\leq'_E$ and for $k \in \{1, \ldots, n-1\}$, it holds that $w_k \in [\Delta_{\sigma_{k,k+1}}(E)]$. This implies that $\sigma_i \preceq_E \sigma_j$.

"$\Leftarrow$" From $\sigma_i \preceq_E \sigma_j$ it follows that for $n \geq 1$ there exist $\sigma_1, \ldots, \sigma_n$ such that $\sigma_1 = \sigma_i$, $\sigma_n = \sigma_j$ and $\sigma_k \wedge \Delta_{\sigma_{k,k+1}}(E)$ is consistent, for $k \in \{1, \ldots, n-1\}$. This means that $w_k \in [\Delta_{\sigma_{k,k+1}}(E)]$, for $k \in \{1, \ldots, n-1\}$ and thus $w_1 \leq'_E \cdots \leq'_E w_n$. Since $\leq_E$ is the transitive closure of $\leq'_E$, we get that $w_i \leq_E w_j$. □

Lemma 12 establishes a neat correspondence between the partial pre-order $\preceq_E$ on complete Horn formulas and the partial pre-order $\leq_E$ on interpretations constructed according to the steps outlined above. It is easy to see, further, that the strict parts of $\preceq_E$ and $\leq_E$ coincide as well. We make this observation explicit as a corollary which we will make use of shortly.

Corollary 1. If w_i, w_j are two interpretations, σ_i and σ_j are two complete Horn formulas such that $[\sigma_i] = \{w_i\}$, $[\sigma_j] = \{w_j\}$ and $\leq_E$ is a Horn connected pre-order constructed as above, then $w_i <_E w_j$ if and only if $\sigma_i \prec_E \sigma_j$.

Proof. If $w_i <_E w_j$, then by Lemma 12 it follows that $\sigma_i \preceq_E \sigma_j$. If we also had $\sigma_j \preceq_E \sigma_i$, this would imply by Lemma 12 that $w_j \leq_E w_i$, which would be a contradiction. The converse is entirely similar. □

Given Lemma 12 and Corollary 1, $\mathsf{s}_4 - \mathsf{s}_6$ follow immediately.

Lemma 13. *The assignment constructed according to steps 1 and 2 outlined above satisfies* $\mathsf{s}_4 - \mathsf{s}_6$.

Proof. In the following we will assume that if w_i is an interpretation, then σ_i is a complete Horn formula such that $[\sigma_i] = \{w_i\}$.

For s_4, take $w_i \in [K_1]$. Then $\sigma_i \models K_1$ and by IC'_4 there exists a complete Horn formula σ_j such that $\sigma_j \models K_2$ and $\sigma_j \preceq_{\{K_1,K_2\}} \sigma_i$. By Lemma 12, this implies that $w_j \leq_{\{K_1,K_2\}} w_i$.

For s_5, take two interpretations w_i, w_j such that $w_i \leq_{E_1} w_j$ and $w_i \leq_{E_2} w_j$. By Lemma 12, we get that $\sigma_i \preceq_{E_1} \sigma_j$ and $\sigma_i \preceq_{E_2} \sigma_j$. By IC'_5 this implies that $\sigma_i \preceq_{E_1 \sqcup E_2} \sigma_j$. It follows, by Lemma 12 again, that $w_i \leq_{E_1 \sqcup E_2} w_j$.

For s_6, take two interpretations w_i, w_j such that $w_i \leq_{E_1} w_j$ and $w_i <_{E_2} w_j$. By Lemma 12 and Corollary 1, we get that $\sigma_i \preceq_{E_1} \sigma_j$ and $\sigma_i \prec_{E_2} \sigma_j$. By IC'_6 this implies that $\sigma_i \prec_{E_1 \sqcup E_2} \sigma_j$. It follows, by Corollary 1, that $w_i <_{E_1 \sqcup E_2} w_j$. □

We can now gather all the results and state our conclusion. The relation $\leq_E$ is reflexive and transitive, thus a pre-order. Lemma 9 gives us that $\leq_E$ is Horn connected. From Lemmas 5 and 8 it follows that $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$. Lemmas 10 and 11 imply that s_1 and s_2 hold. Since $\leq_E$ is constructed only with regard to the models of E, s_3 also holds. Finally, Lemma 13 gives us properties $\mathsf{s}_4 - \mathsf{s}_6$. □

In Theorem 4, the strengthened postulates IC'_4, IC'_5 and IC'_6 rule out Horn merging operators Δ represented by non-syncretic assignments such as the ones in Examples 5 and 6, and thus justify their presence. Our focus on Horn connected pre-orders, on the other hand, should not be seen as a restriction: we can translate any Horn compliant pre-order $\leq_E$ into a Horn connected one $\leq^\star_E$ such that the overall assignment (1) represents the same (Horn) merging operator and (2) remains syncretic. This can be done simply by 'uncoupling' pairs w_i and w_j which are not in the subset relation and do not satisfy either of the properties $\mathsf{h}_{3.1}$ and $\mathsf{h}_{3.2}$. Since w_i and w_j do not appear together in any set of the type $min_{\leq_E}[\mu]$, for $\mu \in \mathcal{L}_H$, the Horn merging operators represented by $\leq_E$ and $\leq^\star_E$ are the same. As an example consider Figure 7, where s_5 does not hold because of interpretations 100 and 001. Notice that in this case IC_5 is satisfied—not only that, but IC'_5 is also satisfied, since 100 and 001 are 'alone' on the same level in $\leq_{E_1}$ and $\leq_{E_2}$, and thus σ_{001} and σ_{100} are in neither of $\preceq_{E_1}$ and $\preceq_{E_2}$. This case would be particularly problematic if we insisted that $\leq_{E_1}$ and $\leq_{E_2}$ be total, since there seems to be no way to express that 100 and 001 are equally preferred in $\leq_{E_1}$ and $\leq_{E_2}$ using just the means of Horn logic

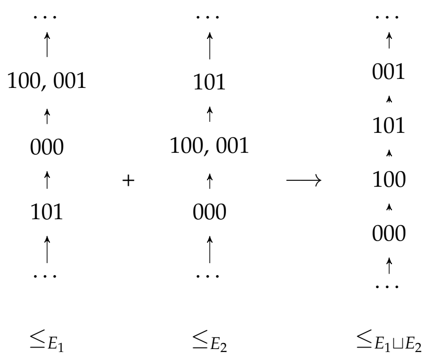

Figure 7: s_5 does not hold

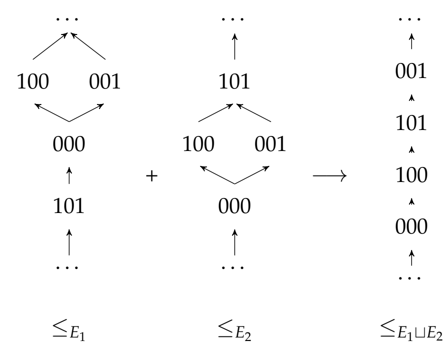

Figure 8: s_5 holds

and Δ. But our problem disappears when we uncouple 100 and 001 in $\le_{E_1}$ and $\le_{E_2}$ (see Figure 8), as now the antecedent of s_5 is not triggered any more by 100 and 001, and thus s_5 is trivially satisfied. This solution is also *natural*, in the sense that allowing 100 and 001 to be incomparable reflects the fact that the Horn operator Δ does not have any opinion on their relative ranking.

Conversely, going from Horn connected pre-orders to total Horn compliant pre-orders is not as straightforward: for any Horn connected pre-order there exist more than one Horn compliant pre-orders representing the same Horn merging operator: any interpretations w_i and w_j that are not in $\le_E^\star$ can be related in several ways if we care about making $\le_E^\star$ total (we could have $w_i <_E^\star w_j$, or $w_j <_E^\star w_i$, etc.), and some of the configurations give rise to non-syncretic assignments. Our point is that w_i and w_j do not need to be related as long as the represented merging operator Δ stays the same. Indeed, the main motivation for formulating the representation result with partial pre-orders is that if w_i and w_j do not satisfy $\mathsf{h}_{3.1}$ and $\mathsf{h}_{3.2}$ then a Horn merging operator Δ does not give us any information on what the order between them should be. It makes sense, in this case, to not include w_i and w_j in the pre-order representing Δ.

5 A concrete Horn merging operator

By Theorem 2, we can find a Horn merging operator simply by exhibiting a Horn compliant, syncretic assignment. As in Examples 2 and 3, we can think of a pre-order $\le_K$ as being generated by the distances $d(w, K)$, for any interpretation w: these distances are just positive integers, and their order determines the ranking of interpretations in $\le_K$. In the rest of this section pre-orders will always be specified in terms of integers. We write $l_K(w)$ to denote the number assigned to w with respect to the knowledge base K. If K has exactly one model w', we simply write $l_{w'}(w)$. This is a slight abuse of notation, but it makes sense, as the pre-order $\le_K$ depends only on the models of K and not its syntactic structure.

One difficulty in finding a Horn merging operator this way is that there is no obvious candidate for an off-the-shelf assignment which satisfies all the required properties: the requirement of Horn compliance rules out standard approaches that are built with familiar distances between interpretations (e.g. Hamming or drastic distance). Therefore, we start by describing some general conditions sufficient to guarantee that the resulting assignment satisfies $\mathsf{s}_1 - \mathsf{s}_6$ and is Horn compliant.

5.1 General conditions

We take $l_K(w) \geq 0$, for any knowledge base K and any $w \in \mathcal{W}$, with $l_K(w) = 0$ if and only if $w \in [K]$. This guarantees the assignment satisfies $\mathsf{s}_1 - \mathsf{s}_3$. We use the sum Σ to aggregate individual pre-orders, and this guarantees $\mathsf{s}_5 - \mathsf{s}_6$. The next conditions spell out what is needed for an assignment to satisfy s_4.

Definition 5. The *distance between knowledge bases K_1 and K_2* is defined as $d(K_1, K_2) = min\{l_{K_1}(w) \mid w \in [K_2]\}$.

We are interested in knowledge bases that satisfy the following property.

Definition 6. Knowledge bases K_1 and K_2 are *symmetric* if $d(K_1, K_2) = d(K_2, K_1)$.

Symmetry is important because it guarantees s_4.

Proposition 6 ([8])**.** *If an assignment satisfies $\mathsf{s}_1 - \mathsf{s}_3$, then it satisfies s_4 iff any two knowledge bases are symmetric.*

We can further simplify the symmetry condition by noticing that we can focus on a particular subset of pre-orders $\leq_K$.

Definition 7. An *initial assignment* is an assignment for knowledge bases that have exactly one model.

In an n letter alphabet, the initial assignment can be visualized as a $2^n \times 2^n$ matrix, with the entries representing the numbers assigned to each interpretation in a pre-order. We shall call this *the initial matrix*. Properties $\mathsf{s}_1 - \mathsf{s}_3$ mean that the matrix has positive entries and 0 on the main diagonal. The symmetry condition means that the matrix must be symmetric (see Example 8).

Definition 8. A *basic assignment* is an assignment that satisfies $\mathsf{s}_1 - \mathsf{s}_3$ for consistent knowledge bases.

Interestingly, as Lemma 14 shows, it turns out that if we fix the pre-orders for every knowledge base K with exactly one model, then pre-orders for knowledge bases K with more than one model are completely determined by this initial assignment (see Example 8).

	00	01	10	11	{10, 11}
00	0	1	2	3	2
01	1	0	3	5	3
10	2	3	0	8	0
11	3	5	8	0	0

Table 4: An initial assignment determines the remaining rankings by symmetry.

Lemma 14. *In a symmetric assignment, the basic assignment is completely determined by the initial assignment.*

Proof. If $E = \{K\}$, we may identify $\leq_E$ with $\leq_K$. Now, if K is a knowledge base having exactly one model, the pre-order $\leq_K$ is assumed to be given. Let us suppose, now, that $[K] = \{w_1, \ldots, w_n\}$, for $n > 1$. Take an interpretation $w_i \in \mathcal{W}$. We denote by K_i a knowledge base such that $[K_i] = \{w_i\}$. By symmetry, we have that $l_K(K_i) = l_{K_i}(K)$. Unpacking this, we get:

$$min\{l_K(w) \mid w \in K_i\} = min\{l_{K_i}(w) \mid w \in [K]\}.$$

Since $[K_i] = \{w_i\}$, we get that the left-hand term is equal to $min\{l_K(w_i)\}$, which is just $l_K(w_i)$. Our problem boils down to showing that this number is determined by the assignment for knowledge bases having exactly one model. This is immediate when we look at the right-hand term: remember that K_i has exactly one model and therefore $l_{K_i}(w)$ is assumed to be given for any interpretation w in the initial assignment. Therefore $min\{l_{K_i}(w) \mid w \in [K]\}$ is completely determined by the initial assignment. □

Example 8. Table 4 shows the initial matrix for the 2 letter alphabet, plus an additional ranking obtained through symmetry. Each column represents a ranking: for instance the first column represents the ranking for a knowledge base that has 00 as its sole model. The number assigned to 00 in this ranking is 0, the number assigned to 01 is 1, etc. The ranking for a knowledge base K that has $\{10, 11\}$ as its set of models is computed from the initial assignment matrix with symmetry. For example, consider interpretation 00. By symmetry, we have that $l_K(00) = l_{00}(K)$. Thus, we obtain $l_{00}(K) = min\{l_{00}(10), l_{00}(11)\} = min\{2, 3\} = 2$.

The nice thing is that if a basic assignment is defined from a symmetric initial assignment by symmetry, then the entire basic assignment is symmetric.

Proposition 7. *If a basic assignment is defined from a symmetric initial assignment by symmetry, then the basic assignment is symmetric.*

Proof. We have to show that for any consistent knowledge bases K_1, K_2, it is the case that $l_{K_1}(K_2) = l_{K_2}(K_1)$. Let us do a case analysis and see that in a couple of cases the conclusion falls out easily.

Case 1. If either of K_1 or K_2 has exactly one model, then $l_{K_1}(K_2) = l_{K_2}(K_1)$ by definition.

	w_1	$\dots$	w_m	w'_1	$\dots$	w'_n
w_1	0			$\bullet$	$\dots$	$\bullet$
$\dots$				$\bullet$	$\dots$	$\bullet$
w_m			0	$\bullet$	$\dots$	$\bullet$
w'_1	$\bullet$	$\dots$	$\bullet$	0		
$\dots$	$\bullet$	$\dots$	$\bullet$			
w'_n	$\bullet$	$\dots$	$\bullet$			0

Table 5: Symmetry

Case 2. If $[K_1] \cap [K_2] \neq \emptyset$, then $l_{K_1}(K_2) = l_{K_2}(K_1) = 0$.

Case 3. The only case left to analyse is when K_1 and K_2 are consistent knowledge bases that each have more than one model and they share no models between them. Suppose, then, that $[K_1] = \{w_1, \dots, w_m\}$ and $[K_2] = \{w'_1, \dots, w'_n\}$, with $m, n > 1$ and $[K_1] \cap [K_2] = \emptyset$. Then:

$$\begin{aligned} l_{K_1}(K_2) &= min\{l_{K_1}(w'_1), \dots, l_{K_1}(w'_n)\} \\ &= min\{min\{l_{w'_1}(w_1), \dots, l_{w'_1}(w_m)\}, \dots, min\{l_{w'_n}(w_1), \dots, l_{w'_n}(w_m)\}\}. \end{aligned}$$

The last step there was taken by applying symmetry. To visualize what this statement says, consult Table 5 and focus on the square of dots in the upper right corner: $min\{l_{w'_1}(w_1), \dots, l_{w'_1}(w_m)\}$ takes the minimum of the dotted elements in the w'_1-column, while $min\{l_{w'_n}(w_1), \dots, l_{w'_n}(w_m)\}$ takes the minimum of the dotted elements in the w'_n-column. We then have to take the minimum of all these minima, which essentially means that $l_{K_1}(K_2)$ takes the minimum element from the upper right dotted square.

Completely analogously, we get that $l_{K_2}(K_1)$ takes the minimum element in the lower left dotted square of Table 5. Crucially, remember that the initial assignment matrix is symmetric—hence, the sub-matrix we have selected in Table 5 is also symmetric. It follows that the dotted squares contain the same elements, and so they have the same minima. This proves that $l_{K_2}(K_1) = l_{K_1}(K_2)$. □

All that is left is Horn compliance. The first thing we show is that Horn compliance in a pre-order reduces to the fact that any triple $\{w_0, w_1, w_2\}$ where $w_1 \not\subseteq w_2$, $w_2 \not\subseteq w_1$ and $w_0 = w_1 \cap w_2$ has to be Horn compliant.

Lemma 15. *A total pre-order $\leq$ on interpretations is Horn compliant if and only if for every triple of interpretations $\{w_0, w_1, w_2\}$, where $w_1 \not\subseteq w_2$, $w_2 \not\subseteq w_1$ and $w_0 = w_1 \cap w_2$, it is the case that $min_{\leq}\{w_0, w_1, w_2\}$ is representable by a Horn formula.*

Proof. As a clarification, for a triple $\{w_0, w_1, w_2\}$ that satisfies the conditions above, $min_{\leq}\{w_0, w_1, w_2\}$ is not representable by a Horn formula exactly when $w_1 \approx w_2 < w_0$. So what we need to show is that Horn compliance means that this arrangement never occurs in $\leq$.

"$\subseteq$" If $\leq$ is Horn compliant then (by definition) the minimal elements of any subset of interpretations closed under intersection are representable by a Horn formula, and this includes triples $\{w_0, w_1, w_2\}$.

"$\supseteq$" Assume that any triple $\{w_0, w_1, w_2\}$ is Horn compliant, but that $\leq$ is not. This means that there exists a Horn formula μ such that $min_{\leq}[\mu]$ is not representable by a Horn formula. Clearly, $min_{\leq}[\mu]$ must contain at least 2 elements (otherwise it would be representable in the Horn fragment). Since $min_{\leq}[\mu]$ is not representable by a Horn formula, this means that $min_{\leq}[\mu]$ is not closed under intersection. In other words, there are two interpretations $w_1, w_2 \in min_{\leq}[\mu]$ such that $w_1 \cap w_2 = w_0 \notin min_{\leq}[\mu]$. Because μ is a Horn formula and $w_1, w_2 \in [\mu]$, then $w_0 \in [\mu]$. It follows, therefore, that $w_1 \approx w_2 < w_0$. But this implies that $\{w_0, w_1, w_2\}$ is not Horn compliant, which contradicts our starting assumption. $\square$

Next we propose the following notion.

Definition 9. A pre-order $\leq$ is *well-behaved* if and only if for any interpretations w_0, w_1, w_2 such that $w_1 \nsubseteq w_2$, $w_2 \nsubseteq w_1$ and $w_0 = w_1 \cap w_2$, both of the following properties hold:

(wb_1) $w_0 \leq w_1$ or $w_0 \leq w_2$ and

(wb_2) $|min\{l(w_1), l(w_2)\} - l(w_0)| \leq |max\{l(w_1), l(w_2)\} - l(w_0)|$.

Notice that a well-behaved pre-order $\leq$ is also Horn compliant, as all the problematic triples of interpretations are Horn compliant and this implies, by Lemma 15, that $\leq$ is also Horn compliant. What makes well-behavedness suitable for our needs, however, is that it is transmitted through Σ-aggregation. We prove this in Proposition 8.

Proposition 8. *If $\leq_1$ and $\leq_2$ are well-behaved, then the pre-order obtained by Σ-aggregating $\leq_1$ and $\leq_2$ is well-behaved.*

Proof. We denote by $\leq_{1+2}$ the pre-order obtained by Σ-aggregating pre-orders $\leq_1$ and $\leq_2$ and by $l_i(w)$ the level of w in pre-order $\leq_i$. Our goal is to show that $\leq_{1+2}$ is well-behaved. Take, then, a triple of interpretations $\{w_0, w_1, w_2\}$ such that $w_1 \nsubseteq w_2$, $w_2 \nsubseteq w_1$ and $w_0 = w_1 \cap w_2$.

First of all notice that, due to property wb_1, the possible ways in which w_0, w_1 and w_2 can be arranged in $\leq_1$ are $w_0 \leq_1 w_1 \leq_1 w_2$, or $w_0 \leq_1 w_2 \leq_1 w_1$, or $w_1 \leq_1 w_0 \leq_1 w_2$, or $w_2 \leq_1 w_0 \leq_1 w_1$. Using property wb_2, it is easy to see that in all these cases it holds that $2l_1(w_0) \leq l_1(w_1) + l_1(w_2)$. In the cases $w_0 \leq_1 w_1 \leq_1 w_2$ and $w_0 \leq_1 w_2 \leq_1 w_1$ we have that $l_1(w_0) \leq l_1(w_1)$ and $l_1(w_0) \leq l_1(w_2)$, thus the conclusion is immediate. In the case $w_1 \leq_1 w_0 \leq_1 w_2$, property wb_2 gives us that $l_1(w_0) - l_1(w_1) \leq l_2(w_2) - l_2(w_0)$. Adding $l_1(w_1)$ and $l_1(w_0)$ to both sides gives us the conclusion. The case $w_2 \leq_1 w_0 \leq_1 w_1$ is analogous.

Similarly, we get that in $\leq_2$ it holds that $2l_2(w_0) \leq l_2(w_1) + l_2(w_2)$. Thus, when we add up the two equalities, it follows that $2(l_1(w_0) + l_2(w_0)) \leq l_1(w_1) + l_1(w_2) + l_2(w_1) + l_2(w_2)$.

Suppose, now, that in $\leq_{1+2}$, property wb_1 does not hold. This means that we have $w_1 <_{1+2} w_0$ and $w_2 <_{1+2} w_0$. Adding up these two inequalities, it follows that $l_{1+2}(w_1) + l_{1+2}(w_2) < 2l_{1+2}(w_0)$. But $\leq_{1+2}$ is obtained by Σ-aggregating $\leq_1$ and $\leq_2$, thus $l_{1+2}(w_i) = l_1(w_i) + l_2(w_i)$ for $i \in \{0, 1, 2\}$. Plugging this into the inequality just derived leads to a contradiction.

Similarly, suppose that $\leq_{1+2}$ does not satisfy property wb_2. Since we have just shown that $\leq_{1+2}$ satisfies property wb_1, the possible arrangement of w_0, w_1, w_2 in $\leq_{1+2}$ is one of $w_0 \leq_{1+2} w_1 \leq_{1+2}$

w_2, or $w_0 \leq_{1+2} w_2 \leq_{1+2} w_1$, or $w_1 \leq_{1+2} w_0 \leq_{1+2} w_2$, or $w_2 \leq_{1+2} w_0 \leq_{1+2} w_1$. Since the levels are always positive numbers, property wb_2 holds in the first two cases. Thus the only remaining possibilities are the last two: let us assume, without loss of generality, that we are in the third case (the remaining case is completely analogous). The fact that property wb_2 does not hold means that $l_{1+2}(w_0) - l_{1+2}(w_1) > l_{1+2}(w_2) - l_{1+2}(w_0)$. This implies that $2l_{1+2}(w_0) > l_{1+2}(w_1) + l_{1+2}(w_2)$. Using again the fact that $l_{1+2}(w_i) = l_1(w_i) + l_2(w_i)$ for $i \in \{0, 1, 2\}$, we get a contradiction. □

5.2 The summation assignment

Using all this knowledge, we can now define a specific Horn compliant syncretic assignment, which we will call *the summation assignment*. We define this assignment for the general case of an alphabet of size n. As suggested by the previous discussion, we give the initial matrix and use symmetry to determine pre-orders $\leq_K$, when $|[K]| > 1$. Since the matrix for the initial assignment has to itself be symmetric and have 0 on the main diagonal, we will only define the entries in the matrix below the main diagonal, with the understanding that the entries above the main diagonal are fixed by symmetry. Also, the order in which interpretations appear in the rows and columns is fixed by the number of 1's in the corresponding bit-vector. For instance, the matrix for the 3-letter alphabet has its rows and columns ordered as follows: 000, 001, 010, 100, 011, 101, 110, 111. We refer to these interpretations as $w_0, w_1, \ldots, w_7$, respectively.

The definition of the bottom half of the initial assignment matrix is recursive. First, put:

$$l_{w_0}(w_i) = i, \text{ for } i \in \{0, \ldots, 2^n - 1\}.$$

Hence, the levels on the first column are $0, 1, 2, \ldots, 2^n - 1$ (see Table 6).

Second, for $i \in \{1, \ldots, \leq 2^n - 1\}$, put:

$$l_{w_i}(w_{i+1}) = l_{w_{i-1}}(w_i) + l_{w_{i-1}}(w_{i+1}).$$

Roughly, this means that the number in a particular cell under the main diagonal is the sum of its two neighbours to the left. In Table 6: if $l_{w_{i-1}}(w_i) = a$, $l_{w_{i-1}}(w_{i+1}) = b$, then $l_{w_i}(w_{i+1}) = a + b$. This is simpler than it sounds, and Table 4 shows the matrix that we get for the 2-letter alphabet.

To be in accordance with the existing literature on merging, we can see the summation assignment as generated by a custom-defined distance function, call it d_S. Thus, for any interpretation w and Horn knowledge base K, $d_S(w, K)$ is just $l_K(w)$, the level of w with respect to K, as defined above.

As motivation for why the levels are assigned the way they are, consider the example of the 2-letter alphabet and what we get when we want to Σ-aggregate the rankings corresponding to models 01 and 10 (see Table 7). We assume $l_{00}(01) = a$ and $l_{00}(10) = b$. By symmetry, $l_{01}(10) = l_{10}(01) = c$. When we compute the aggregated pre-order $\leq_{01+10}$

	w_0	$\dots$	w_{i-1}	w_i	w_{i+1}	$\dots$
w_0	0	$\dots$	$i-1$	i	$i+1$	$\dots$
$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$
w_{i-1}	$i-1$	$\dots$	0			$\dots$
w_i	i	$\dots$	a	0		$\dots$
w_{i+1}	$i+1$	$\dots$	b	$a+b$	0	$\dots$
$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$

Table 6: The recursive relation for levels.

	01	10	$\leq_{01+10}$
00	a	b	$a+b$
01	0	c	c
10	c	0	c
11	$\dots$	$\dots$	$\dots$

Table 7: Horn compliance forces us to assign increasing levels.

we get that $l_{01+10}(00) = a+b$ and $l_{01+10}(01) = l_{01+10}(10) = c$. Since 01 and 10 are on the same level in $\leq_{01+10}$, Horn compliance requires that 00 is assigned a lower numbers, or in other words $c \geq a+b$. Roughly, the idea behind the summation assignment is simply to take, in this situation, $c = a+b$.

We now show that the summation assignment is Horn compliant and stays Horn compliant through repeated Σ-aggregations.

Proposition 9. *The summation assignment is well-behaved.*

Proof. We begin by noting a couple of key aspects of our assignment.

Observation 1. Notice that, because of the way we order the vector of interpretations that make up the column and row heads of the matrix, an interpretation $w_1 \cap w_2$ always comes 'before' w_1 and w_2: that is to say, the column for $w_1 \cap w_2$ is always to the left of both w_1 and w_2, and the row for it is always above the row for w_1 and w_2.

Observation 2. Notice that as we traverse the initial assignment matrix from left to right and from top to bottom, the levels keep increasing. More precisely, say we select two interpretations w_x and w_y which appear in this order in the matrix, and we extract their columns from the initial assignment matrix (see Figure 8). We get a sub-matrix of the original one as in Figure 8. Suppose, now, that we select some subset of interpretations $\{w_{i_m}, \dots, w_{i_n}\}$, which also appear in this order. Take:

$$min\{l_{w_a}(w_i) \mid w_i \in \{w_{i_m}, \dots, w_{i_n}\} \text{ and } l_{w_x}(w_i) \neq 0\} = a,$$

that is to say: a is the smallest level on the column for w_x, except 0. We represent this by writing $a+$ in the places for all levels except 0, to show that they are all at least as great as a. Similarly, let us say that the smallest element on the column for w_y, is b. We represent this by writing $b+$.

	w_x	w_y
$\dots$	$\dots$	$\dots$
w_{i_m}	$a+$	$b+$
$\dots$	$a+$	$b+$
w_{i_n}	$a+$	$b+$
$\dots$	$\dots$	$\dots$

Table 8: Extracting a sub-matrix from the initial assignment matrix

	w_0	w_1	w_2	$w_{i_1}, \dots, w_{i_k}$
w_0	0			a
w_1		0		b
w_2			0	c
w_{i_1}	$a+$	$b+$	$c+$	0
$\dots$				0
w_{i_k}	$a+$	$b+$	$c+$	0

Table 9: The level of w_0 in $\leq_K$ has to be smaller than the levels of w_1 and w_2.

The crucial thing to see here is that $a \leq b$. This follows from the way we defined the initial assignment, and the fact that levels keep increasing as we go from left to write (or top to bottom).

Using these observations, we show now that the summation assignment is well-behaved. Take a triple of interpretations $\{w_0, w_1, w_2\}$ such that $w_1 \nsubseteq w_2$, $w_2 \nsubseteq w_1$ and $w_0 = w_1 \cap w_2$. We want to show that properties wb_1 and wb_2 hold in every pre-order $\leq_K$ of the basic assignment. First, notice that the assignment for knowledge bases with exactly one model is well-behaved. This is because, from the way the assignment is defined, w_0 always has a lower level than both w_1 and w_2. Next let us look at knowledge bases K that have more than one model. We will do a case distinction.

Case 1. First, suppose neither of w_0, w_1, w_2 is in $[K]$. We claim that $w_0 \leq_K w_1$ and $w_0 \leq_K w_2$. Suppose $[K] = \{w_{i_1}, \dots, w_{i_k}\}$. By symmetry, we have that $l_K(w_0) = l_{w_0}(K)$, which means that $l_K(w_0) = min\{l_{w_0}(w_{i_1}), \dots, l_{w_0}(w_{i_k})\}$. Similarly, we get that $l_K(w_1) = min\{l_{w_1}(w_{i_1}), \dots, l_{w_1}(w_{i_k})\}$. To see that $l_K(w_0) \leq l_K(w_1)$, let us extract the sub-matrix with pre-orders for w_0, w_1 and $[K] = \{w_{i_1}, \dots, w_{i_k}\}$ (see Table 9, where we have also included w_2). Suppose $l_K(w_0) = a$ and $l_K(w_1) = b$. Then, by symmetry, we must have that $min\{l_{w_0}(w_{i_1}), \dots, l_{w_0}(w_{i_k})\} = a$ and $min\{l_{w_1}(w_{i_1}), \dots, l_{w_1}(w_{i_k})\} = b$. Using Observations 1 and 2, it follows that $a \leq b$. Similarly, it follows that $a \leq c$, which implies the conclusion.

Case 2. Second, suppose $w_1 \in [K]$, and that $[K] = \{w_1, w_{i_1}, \dots, w_{i_k}\}$. Since $l_K(w_1) = 0$, the condition for well-behavedness amounts to showing that $l_K(w_0) \leq l_K(w_2)$. An argument similar to the one before shows why this holds. By consulting Table 10, which was also completed using symmetry, and by Observations 1 and 2, we conclude that $a \leq b$ and hence that the level of w_0 in $\leq_K$ is smaller than the level of w_2. The case when $w_2 \in [K]$ is completely analogous.

Together, these considerations show that the pre-order is well-behaved, for any knowledge base that has more than one model. □

	w_0	w_1	w_2	$w_1, w_{i_1}, \dots, w_{i_k}$
w_0	0			a
w_1	$a+$	0	$b+$	0
w_2			0	b
w_{i_1}	$a+$		$b+$	0
$\dots$				0
w_{i_k}	$a+$		$b+$	0

Table 10: The level of w_0 in $\leq_K$ has to be smaller than the level of w_2.

	000	001	010	100	011	101	110	111	$d(w, K_1)$	$d(w, K_2)$	Σ
000	0	1	2	3	4	5	6	7	3	1	4
001	1	0	3	5	7	9	11	13	5	0	5
010	2	3	0	8	12	16	20	24	8	3	11
100	3	5	8	0	20	28	36	44	0	5	5
011	4	7	12	20	0	48	64	80	20	7	27
101	5	9	16	28	48	0	112	144	0	0	0
110	6	11	20	36	64	112	0	256	36	11	47
111	7	13	24	44	80	144	256	0	0	0	0

Table 11: Summation assignment example

This is the last piece of information needed. We can now assert the following theorem.

Theorem 5. *The summation assignment represents a Horn merging operator. We call this operator* $\Delta^{d_S, \Sigma}$.

Proof. The initial matrix has positive entries and 0 on the main diagonal, thus s_1 – s_3 are satisfied. It is also symmetric, which guarantees s_4. Proposition 9 guarantees that the pre-orders in the basic assignment are well-behaved and that they stay well-behaved through Σ-aggregation, which is sufficient for Horn compliance. Thus the summation assignment is Horn compliant and syncretic, and by Theorem 2 it represents a Horn merging operator. □

We illustrate the summation assignment on an example.

Example 9. Take $K_1 = \{p_1, \neg p_2 \vee p_3\}$, $K_2 = \{p_3, \neg p_2 \vee p_1\}$, $E = \{K_1, K_2\}$ and $\mu = \neg p_1 \vee \neg p_3$. Notice that K_1 and K_2 are Horn theories and μ is a Horn formula. Merging E under constraint μ with $\Delta_\mu^{d_H, \Sigma}$ or $\Delta_\mu^{d_D, \Sigma}$ would give us $[\Delta_\mu^{d_H, \Sigma}(E)] = [\Delta_\mu^{d_D, \Sigma}(E)] = \{001, 100\}$, thus not representable by a Horn formula. Clearly, in this case, the problem is that the pre-order $\leq_E$ is not Horn compliant. However, with the custom distance function encoded with our summation assignment, we generate Horn compliant pre-orders for K_1 and K_2 that, when aggregated with Σ, produce a Horn compliant pre-order $\leq_E$. Table 11 shows the initial matrix of the summation assignment for the 3 letter alphabet, together with the pre-orders for K_1 and K_2 and the Σ-aggregated pre-order $\leq_E$.The pre-orders for K_1 and K_2 are computed using the symmetry property and data from the initial matrix. For instance:

$$\begin{aligned} l_{K_1}(110) &= \min\{l_{110}(w) \mid w \in [K_1]\} \\ &= \min\{l_{110}(100),\ l_{110}(101),\ l_{110}(111)\} \\ &= \min\{36, 112, 256\} \\ &= 36. \end{aligned}$$

Thus, the entry in Table 11 for $d(110, K_1)$ is 36. In the end, we obtain $[\Delta_{\mu}^{d_S,\Sigma}(E) = \{000\}]$. Thus $\Delta_{\mu}^{d_S,\Sigma}(E)$ can be represented by a Horn formula, in this case $\neg p_1 \wedge \neg p_2 \wedge \neg p_3$.

6 Conclusion and future work

In this paper, we provided a novel representation theorem for Horn merging by strengthening the standard merging postulates. Belief change operators for the Horn fragment have attracted increasing attention over the last years, in particular revision and contraction, while merging in the Horn fragment remained rather unexplored so far. An exception is [3], where the authors propose to *adapt* known merging operators by means of a certain post-processing and study the limits of this approach in terms of satisfaction of the merging postulates. One of the main results of that paper is that in their framework it is not possible to keep all postulates satisfied. In our work, we have presented a novel concrete Horn merging operator satisfying *all* postulates.

The moral of the present work is that, while going from syncretic assignments to Horn merging operators is relatively easy (Horn compliance is sufficient, by Theorem 2), going from Horn merging operators to syncretic assignments requires considerably more machinery (in particular, stronger postulates). Thus, all the work in Section 4 is needed to obtain a full representation result. Even so, Section 5 highlights that the easiness of the first direction is only relative, as finding concrete syncretic assignments that are also Horn compliant requires some conceptual work, and there is no obvious trivial operator that does the job. The main difficulty here lies in making sure that if two pre-orders $\leq_1$ and $\leq_2$ are Horn compliant, then the pre-order resulted from Σ-aggregating them is also Horn compliant. Our well-behavedness property guarantees this.

Future work on merging in the Horn fragment would have to consider extending the family of Horn merging operators. This requires seeing how Horn compliance interacts, on the model side, with other aggregation functions (such as *GMAX*) and exploring the range of conditions guaranteeing Horn compliance of an assignment. We may add to this the study of other merging postulates (e.g., majority and arbitration), considered in the merging literature [8, 9] but not touched upon here. Finally, we would like to extend our approach to other fragments of propositional logic (e.g., Krom or dual Horn), where similar problems arise and for which tailored notions of compliance and strengthened postulates are likely needed.

References

[1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.*, 50(2):510–530, 1985.

[2] Richard Booth, Thomas A. Meyer, Ivan J. Varzinczak, and Renata Wassermann. On the link between partial meet, kernel, and infra contraction and its application to Horn logic. *J. Artif. Intell. Res. (JAIR)*, 42:31–53, 2011.

[3] Nadia Creignou, Odile Papini, Reinhard Pichler, and Stefan Woltran. Belief revision within fragments of propositional logic. *J. Comput. Syst. Sci.*, 80(2):427–449, 2014.

[4] James P. Delgrande and Pavlos Peppas. Belief revision in Horn theories. *Artif. Intell.*, 218:1–22, 2015.

[5] James P. Delgrande and Renata Wassermann. Horn clause contraction functions. *J. Artif. Intell. Res. (JAIR)*, 48:475–511, 2013.

[6] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proc. KR 1991*, pages 387–394. Morgan Kaufmann, 1991.

[7] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artif. Intell.*, 52(3):263–294, 1992.

[8] Sébastien Konieczny and Ramón Pino Pérez. Merging information under constraints: A logical framework. *J. Log. Comput.*, 12(5):773–808, 2002.

[9] Sébastien Konieczny and Ramón Pino Pérez. Logic Based Merging. *J. Philosophical Logic*, 40(2):239–270, 2011.

[10] Thomas J. Schaefer. The Complexity of Satisfiability Problems. In *Proc. STOC 1978*, pages 216–226. ACM, 1978.

[11] Frederik Van De Putte. Prime implicates and relevant belief revision. *J. Log. Comput.*, 23(1):109–119, 2013.

[12] Zhi Qiang Zhuang and Maurice Pagnucco. Model Based Horn Contraction. In *Proc. KR 2012*. AAAI Press, 2012.

[13] Zhi Qiang Zhuang, Maurice Pagnucco, and Yan Zhang. Definability of Horn revision from Horn contraction. In *Proc. IJCAI 2013*, 2013.